PUBLIC PERCEPTIONS OF GLOBAL WARMING

RICHARD A. BERK and DANIEL SCHULMAN
Center for the Study of the Environment and Society Department of Sociology and Interdivisional Program in Statistics, UCLA, U.S.A.

Abstract. In this paper, we examine the way public opinion responds to the prospect of global warming. In particular, we focus on the public's "willingness to pay" in order to prevent various hypothetical climate scenarios from transpiring. To this end, fractional factorial survey methods are employed with a sample of over 600 residents of Southern California. By and large, the public is able to understand and evaluate rather complicated hypothetical climate scenarios, but the public appreciates some features of climate far better than others. In this context, the contingent valuation estimates we provide, while promising, are clearly not ready of consideration by policy makers.

1. Introduction

The massive outpouring of greenhouse gases over the last 150 years has launched an anthropogenic experiment of unprecedented proportions. While the precise nature of the experiment is unclear and its exact implications largely unknown (MacCracken *et al.*, 1990), serious proposals are being entertained to alter human activity around the world so that the release of greenhouse gases is curtailed.

It is not surprising that proposals to reduce greenhouse gas emissions are controversial. The economic and social stakes are enormous. But to date, most of the controversy has been generated by scientists, environmental groups, industry representatives and government officials. Global warming has not produced the same level of citizen involvement as other environmental problems such as air quality, nuclear waste disposal, and toxics in the water supply. Part of the explanation is that global warming is not yet evident in people's daily lives. Rather, global warming is a process that some time in the future may have rather specific, but as yet unknown, effects. Another explanation is that the science on global warming is almost certainly confusing to many and therefore, not easily transformed into a clear social problem (Schnaiberg, 1980). Yet, if the weight of scientific opinion is to be believed, important changes in how fossil fuels are used must soon begin, well before the possible impacts of global warming are observed. That is, important changes must be initiated when public opinion is still unformed.

In this paper, we report on research undertaken to measure factors that may affect what form public opinion will take. Factorial survey methods (Rossi and Nock, 1982) are used to provide survey respondents with a variety of future global warming scenarios. The key response is willingness to pay (WTP): how much respondents are willing to pay to keep a given global warming scenario from happening. Other responses are also examined. Insofar as these response variables

Climatic Change **29:** 1–33, 1995.

vary with specific features of global warming scenarios, it is possible to infer something about the factors that may ultimately shape public opinion.

Theoretical links are made to the literature on the value of environmental goods. It is well known that as a particular kind of public good, environmental goods do not have clear property rights (Baumol and Oates, 1988). As a result, there is no market and no observable price: it is very difficult to place an explicit value on environmental goods (Sargoff, 1988). Climate, as an environmental good, shares these problems, but at a more fundamental level, it is not clear which aspects of climate people value. Thus, conceptual issues raised in the paper are related to the long-standing interest among social scientists on how value is determined. We also consider theoretical and technical issues of interest to survey researchers who worry about response errors. Willingness to pay, and the policy implications of numbers elicited, raise well-known problems about which we will have some findings and constructive suggestions.

2. Theoretical Perspectives on Valuing Environmental Goods

In market-oriented societies, the value of conventional economic goods is normally ascertained by observing the behavior of buyers and sellers. Under ideal market circumstances, a good's market price can be taken as a measure of value at the margin. As public goods, environmental goods are not traded on a market, and obtaining measures of value is very difficult. Thus, there is no market for clean air or litter-free parks.

Yet, it is clear that environmental goods have value and that decisions have to be made taking their value into account (Pearce, 1993). Billions of dollars are spent each year in the United States alone to improve air quality and to maintain public parks. But, whether that spending is too much or too little clearly depends on the value of the environmental goods in question.

2.1. A THEORY OF VALUE FOR ENVIRONMENTAL GOODS

How then, does one think about environmental goods? To begin, the value of environmental goods needs to be placed in a larger structure of economic relations. Drawing heavily on an excellent review paper by Cropper and Oates (1992), it is common in environmental economics to approach the issue of value through equations that capture: (1) the utility of the representative consumer; (2) a production function for the level of goods consumed; and (3) the level of environmental quality. In abbreviated form:

$$U = \mathrm{U}(\mathbf{X},\ Q), \tag{1}$$

$$\mathbf{X} = \mathrm{X}(\mathbf{L},\ E,\ Q), \tag{2}$$

$$Q = \mathrm{Q}(E). \tag{3}$$

U denotes the utility of the representative consumer, $\mathbf{X}$ is a vector of consumer goods, $\mathbf{L}$ is a vector of conventional production inputs, Q is the level of environmental quality, and E is the level of wate emissions from some set of industrial or agricultural processes.

From Equation 1, one learns that the well-being of consumers depends on both the goods consumed and environmental quality. From Equation 2, it is apparent that waste emissions are just another factor of production. From Equation 3, one can see that environmental quality, Q, is a function of emissions (e.g., hydrocarbons transformed into smog) and is likewise, an argument in Equation 2, the production function.

Reducing the level of emissions is assumed to reduce output (at least in the short term). For example, reductions in the emissions of nitrogen compounds into the air, other things equal, may mean less oil will be burned, and less electricity produced. Reductions in environmental quality are assumed to decrease output. For example, minerals from water used in irrigation may slowly build up in the soil and reduce agricultural production.

Maximizing Equation 1, subject to Equations 2 and 3, and resource constraints, yields several interesting results.* For example, ideally each polluting agent should pursue pollution control measures until the marginal benefits from reduced pollution equal the marginal abatement costs. But, to address the question of value, several more steps are required.

In particular, a damage function is defined to replace Q:

$$S = \mathrm{S}(Q). \tag{4}$$

S replaces Q because it is, after all, the consequences of environmental quality that ultimately matter, not environmental quality itself. It follows that there are two ways to think about value. If the S (i.e., damage) is beyond an actor's control, the value of a change in S may be directly measured, from which it is possible to work backwards to determine the value of Q. For example, insofar as the removal of lead paint from the market by federal law reduced the number of new lead poisoning cases, one might take the medical costs saved as a measure of the incremental value of lead-free paint compared to paint with lead.**

Sometimes, however, it may be possible to mitigate the effects of Q by introducing other inputs in the damage function. For example, consumers might install water filters on their home faucets to remove certain contaminants. Letting Z represent such inputs, one now writes

$$S = \mathrm{S}(Q, Z). \tag{5}$$

Including Z in the damage function complicates efforts to empirically estimate costs because while Z can affect S, S can now also affect Z. Therefore, direct observation

* The maximization will in practice depend on the functional forms assumed. And whether the assumed forms are the correct forms can be a troubling question.

** In practice, one would want to also consider additional cost saving such as lifetime earnings.

of S cannot be used to infer the value of Q. As an alternative, consider the situation where the consumer is being asked to pay for an improvement in environmental quality. It follows that the most a consumer should be willing to pay is equal to the reduction in expenditures necessary to achieve the original utility. Suppose a consumer pays a $ 1000 for 100 utils at some existing level of environmental quality and prices. Now suppose that consumer could get those 100 utils for $ 800 because there is higher environmental quality. The consumer's willingness to pay (WTP) should be $ 200, and the value of the environmental improvement is $ 200.

Now, consider a situation where a consumer is faced with a reduction in environmental quality. Then, the smallest amount of money a consumer should be willing to accept as compensation is the additional amount needed to achieve the original level of utility. Suppose that a consumer's utility would be reduced from 1000 to 900 utils because of a decrease in environmental quality. The consumer should be willing to accept (WTA) at a minimum the amount of money necessary to achieve 1000 utils again, which, in turn, equals the cost of the decline in environmental quality.

Thus, willingness to pay and willingness to accept provide two ways to value environmental goods. In general, a consumer's willingness to pay will depend on income and the price of available substitutes (Hoehn and Loomis, 1993), which implies that a consumer's willingness to pay will not necessarily be the same as a consumer's willingness to accept. For willingness to accept, the substitutes for the environmental goods foregone may be more costly. However, since value is tied to willingness to pay or accept and since, at least in principle, willingness to pay or accept can be measured, it is possible to obtain measures of value.

2.2. APPLICATIONS TO USE AND NONUSE VALUE

Much of the conventional theory and research on WTP and WTA has focused on environmental goods that people experience directly and, therefore, use. Such environmental goods have use value, much like conventional market goods. There are other environmental goods that have value to people although these goods are not directly used in any sense. A Midwestern farmer may value the preservation of whales, although that farmer may never see a live whale. An urban resident from California may value the restoration of prairie grassland, although the closest that urban resident may get to a Midwestern prairie is 30,000 feet, while flying coast-to-coast. Large set-asides for wilderness areas are supported by millions of Americans even when those lands will be effectively closed to them. In other words, some environmental goods have nonuse value (Krutilla, 1967). Mitchell and Carson (1989: 59–67) favor instead the term existence value and then divide existence value into vicarious consumption (i.e., consumption by others) and stewardship (i.e., preservation). Other divisions are possible (Stern *et al.*, 1993; Diamond *et al.*, 1993; Milgrom, 1993).

One must be very clear on what is contained within existence value. Our discussion of value began with a conventional set of economic relationships and the goal of allocating material societal resources (capital and labor) in an economically efficient manner. How much should be paid for litter-free parks or clean air? But, insofar as existence value is meant to inform *pareto optimal* societal resources allocations through cost-benefit analysis, existence value cannot include altruism.* That is, among the reasons for valuing some environmental good cannot be moral satisfaction from stewardship or from other peoples' enjoyment of it. Milgrom (1993) shows that when altruism is included as a part of the value of some environmental good, there is double (or more) counting of the social benefits derived from that environmental good. The consequence is that society will invest too little in the environmental goods compared to a pareto optimal outcome.

Milgrom also points out that similar reasoning about existence value excludes the value of how benefits are distributed (e.g., supporting Native American business enterprises), the value of the process by which an environmental good is improved or damaged, (e.g., a forest fire caused by lightning versus a campfire), and the value of whether some moral obligation is being upheld or challenged (e.g., the moral obligation of humans to all living things). While no one would deny the importance of such concerns, there is no provision for them within a pareto optimal framework. If these kinds of nonuse values are included, economic theory will not provide proper guidance in how resources should be allocated. In short, if one wishes to use conventional economic theory to help society allocate scarce resources between various public goods, the content of the existence value must be very narrowly circumscribed.

However, even circumscribed definitions of existence value are controversial (Rosenthal and Nelson, 1992; Kopp, 1992). For an environmental good to have existence value, the individuals doing the valuing must know that it exists. Consequently, if there are animal or plant species that no one has yet identified, for example, they cannot have existence value. In a similar fashion, if only a small number of people know about the existence of some pristine wilderness, the maximum existence value that wilderness can have is the sum of existence values of the people who know about it. And it is entirely possible for an environmental good to have more existence value after it is damaged insofar as publicity about the damage has advertised its existence (Daum, 1993). In short, since existence depends on the information that consumers have, as the information changes, so will existence value. There is no requirement that the information be complete or even accurate. For example, consumers probably value the existence of seals far more than bees, although bees are clearly much more important for human food production.

* In the hands of some, a very ambitious agenda lies behind the goal of pareto optimal societal allocations (Mishan, 1976: pp. 382–402). Many decisions that are usually defined as political are redefined as technical. Also, whether pareto optimality is the most desirable social objective certainly can be argued (Blaug, 1980: pp. 140–152) along whether the underlying behavioral and institutional assumptions are realistic (Russell and Wilkinson, 1979).

If existence value were the concern solely of social psychologists, such complications would be of interest in their own right, and the stimulus for research on how people think about the value of environmental goods. If in the minds of consumers, certain environmental goods have more vlaue after they are damaged then before they are damaged, for example, so be it. But if the purpose of existence value is to help in the allocation of scarce resources, the difficulties we have summarized are very unsettling.

2.3. APPLICATIONS TO GLOBAL WARMING

The conceptual difficulties with the existence (nonuse) value of environmental goods must be kept in mind especially when new and difficult environmental concerns, such as global warming, are identified by science. The prospect of global warming and its possible implications for all life on the planet, implies the need to value our current climate against what a future climate might be like. With that done, choices in principle can be made between a variety of possible responses. However, there is clearly no market for climate change and just as clearly no market price. While people, loosely speaking, use the climate of today, they have not used the possible climates of the future. Indeed, a large fraction of the people living today will be dead before a significantly changed climate could materialize. So, one cannot estimate the value people place on different climates by observing the choices they make between their use of today's climate compared to their use of the climate of the future. By default, therefore, we are back into the willingness to pay and the willingness to accept framework.*

There are perhaps two ways to think about valuing global warming within this framework. If climate has use value, a comparison between the climate of today and an uncertain prospect about a future climate can, in principle, generate a willingness to pay or willingness to accept. In effect, the consumer has a choice between the status quo and placing a bet on one or more climate horses. However, matters are complicated by the fact that change is effectively irreversible and more information about the costs and benefits will continue to be available over time. These conditions make current decisions more costly (Arrow and Fisher, 1974). In effect, there is a penalty for not having the information that will be available in the future. Nevertheless, consumers will, in principle, incorporate these costs in the willingness to pay or accept, and both concepts would seem to be viable when the use value of climate is the issue. If climate also has existence value, however, all of the earlier difficulties resurface. They are also magnified by the incomplete and rapidly changing information base on which existence value of a particular climate is determined. To summarize, climate probably has use value and existence value. But the conceptual problems with existence value would seem

* As one reviewer suggested, however, it might be possible to observe behavior in response to existing approximations of climate change. For example, if a family moved from Seattle to Phoenix, the price they paid to purchase air conditioners might provide a useful proxy for willingness to accept an average temperature increase of, say, 15 degrees.

to argue for focusing on the use value only. Then, willingness to pay may prove to be an instructive concept.

3. Measuring the Value of Environmental Goods

Over the past several decades, researchers have tried to estimate the use value of environmental goods using three 'indirect' strategies (Cropper and Oates, 1992).* Studies of 'averting behavior' measure the costs of techniques adopted to avoid the damage in question. For example, the increased cost of using oxygenated gasoline in winter months can be used to provide estimates of the value of particular carbon monoxide reductions. Studies using 'weak complimentarity' estimate costs by observing the costs consumers are prepared to incur to use a better rather than worse environmental good. For example, the costs of the additional trips to a lake that has had its water quality improved, provide a measure of the value of those improvements. Finally, 'hedonic market' studies compare the price of market goods, such as homes, that are comparable except for the quality of the environmental good in question. For example, if two neighborhoods have similar amenities but differ in their air quality, the average difference in the market value of homes in those two neighborhoods provides a measure of the value of the difference in air quality.

Studies of existence value are unable to capitalize on indirect methods of measuring value; there is no behavior to observe. This has led to widespread use of the method of 'contingent valuation' (CV) in which people are asked in surveys to report on their hypothetical willingness to pay or their hypothetical willingness to accept. Implicit is the assumption that consumers have a coherent set of preferences for both market and non-market goods, that these preferences could in principle be revealed by market behavior, and that these preferences can be elicited in a contingent valuation survey (Kahneman, 1986). Beginning in the early 1960's economists took to the field with such surveys (Davis, 1963), and they are now routine (Mitchell and Carson, 1988; Cropper and Oates, 1992).

In the typical CV survey, respondents are presented with a detailed description of the environmental goods in question. These descriptions can sometimes cover several pages and sometimes include maps and figures of various kinds. Following each description is one or more questions about willingness to pay or willingness to accept. Questions asking for a dollar amount were popular at first, but now there seems to be a general consensus that 'referendum-style' questions are superior (Cameron, 1991). Respondents are presented with a dollar figure which they can accept or reject, as if they were voting on a referendum. Sometimes there is a follow-up question using a higher or lower figure.

While it seems to have come as some surprise to the economists who conducted the early surveys, CV items appear to have all the usual sort of response errors with

* The three strategies are 'indirect' because the prices of environmental goods are not observed. Prices are inferred from other information.

which survey researchers are familiar (Groves, 1989). Also present are the variety of cognitive anomalies demonstrated by experimental psychologists (Fischhoff, 1991; Loomis *et al.*, 1993). Finally, respondents seem to provide WTP or WTA figures that include a great variety of values beyond existence value (Kahneman and Knetsch, 1992; Schkade and Payne, 1993, 1994; Stern *et al.*, 1993).

There is little doubt that WTP and WTA figures elicited from surveys currently provide such crude measures of value that the point estimates are of little help in setting up public policy (Arrow *et al.*, 1993). That is, the precise dollar figures cannot be taken seriously. However, the dollar values may provide a ratio scale quantitatively ranking environmental goods. At a minimum, a qualitative ranking should be possible.

Perhaps, therefore, CV techniques could be usefully applied to our problem of valuing climate. All that we need is a response metric with which to learn about the relative ranking of various climate features. Dollars are a good metric because they are units with which people have lots of experience. But we stress that the point estimates should not be taken literally as the amount of money respondents are willing to pay or accept.

4. Research Design

The research design consists of three parts, the construction of the instrument, the selection of the sample, and the interviewing procedures. We will address these in order, but since the instrument is by far the most complex issue, we will focus primarily on it.

4.1. INSTRUMENT DESIGN

In order to consider the use value of climate, our research design addresses climate as people *experience* it. In people's daily lives, climate is potentially composed of a number of features: temperature, precipitation, wind, cloud cover, humidity and others. For purposes of this study of global warming, we will focus on temperature and precipitation. Both have been a central focus of scientific research on the enhanced greenhouse effect, and both clearly figure in what people experience day-to-day.* Climate as experienced is in some sense an aggregation of daily experience. The need to represent many day-to-day experiences over time implies that temperature and precipitation are best conceptualized as *distributions*. We will explore in this paper the role of four 'parameters' of these distributions:

1. mean;
2. spread;
3. extreme values; and
4. clustering of extreme values.

* We are ducking a discussion of the scientific definition of climate. There is some disagreement over whether climate is nothing more than the temporal average of weather or if climate is a distinct phenomenon.

For example, a future climate could be warmer on the average, have more variability in temperature, could have more very hot days, and these hot days could come in sets. Likewise, a future climate could be drier on the average, with greater variation between wet periods and dry periods, and with more severe storms that tend to come in waves. And, of course, various combinations of these characteristics are possible in principle. Our research goal is to determine the impact of temperature and precipitation on willingness to pay, with temperature and precipitation characterized by variation in the four distributional parameters.

To this end, factorial survey methods were used (Rossi and Nock, 1982). The strategy was to give respondents a brief summary of their climate and pose to them a number of hypothetical climates of the future, eliciting their responses to each. For temperature and precipitation, each of the four parameters became a dimension. Each dimension was represented by a number of levels to be sampled at random as the hypothetical climate scenarios were assembled. In effect, each respondent received *a random sample* of all possible climate scenarios: a *fractional* factorial design was employed.

Since the scenarios were to be linked to daily experience, the scenarios had to be situated in a particular geographical area. Scenarios for Seattle would be very different from scenarios for Phoenix. The Los Angeles area was chosen for our study for reasons that will be briefly discussed below.

Our instrument was refined through extensive pilot testing in the Los Angeles area. We first conducted 40 telephone interviews with lengthy debriefing of respondents. Substantial revisions of the instrument followed. We then conducted 150 computer aided telephone interviews (CATI) as a 'dress rehearsal' for the study. The 150 interviews were statistically analyzed, which lead to a number of more modest revisions. We eventually settled on an instrument with four dimensions designed as follows.

I. Temperature
1. Values for the Mean (in degrees) – 60, 70, 80, 90, 100, 110, 120;
2. Values for the Spread (in degrees) – plus or minus 5, 10, 15, 20;
3. Values for the Extreme (in degrees) – a few days when the temperature was 40, 60, 80, 100, 110, 120, 130, 140;
4. Values for the Clustering of Extremes (in degrees) – no consecutive days over 100, 2 consecutive days over 100, an entire week over 100, 2 straight weeks over a 100, 3 straight weeks over 100, an entire month over 100;

II. Precipitation
1. Values for the Mean (in inches of rain per year) – 5, 10, 15, 20, 25, 30, 35;
2. Values for the Spread (when most of the rain fell) – within a single month, over two months, over three months, over six months, throughout the year;
3. Values for the Extremes (intensity of rainfall) – very light rain, light rain, steady rain, heavy rain, very heavy rain;

4. Values for Clustering of Extremes – no consecutive weeks of rain, 2 consecutive weeks of rain, 4 consecutive weeks of rain, 6 consecutive weeks of rain, 8 consecutive weeks of rain.

While there is some arbitrariness in the particular values and language assigned to each of the levels within dimensions, we had rather specific causal effects in mind. For the dimensions of average, spread, and extremes, we anticipated at least a monotonic increase in undesirability with greater departures from the norm, perhaps at an increasing rate. That is, one might find a V-shaped or U-shaped functional relationship. For the dimensions of clustering of extremes movement toward more clustering of extremes would produce at least a monotonic increase in undesirability and perhaps an increase in the rate of increase as well. That is, one might find approximately a linear function or a power function.

Through the pilot testing, it became apparent that some respondents had great difficulty with climate scenarios in which all eight dimensions were included. There was too much information to retain and process over the phone. So, we reduced the information by: (a) allowing only mean temperature and mean rainfall to be included in each scenario; and then, (b) selecting at random for inclusion (with equal probability) one of the remaining three temperature dimensions and one of the remaining three rainfall dimensions. It was clear that at least for context, the means for temperature and rainfall had to be in each scenario. Sampling from the other dimensions meant that three-way interaction effects could not be studied and that there would be some loss of statistical power for estimates of the relevant main effects and two-way interactions.

The need to elicit responses based on personal experiences meant that each climate scenario had to be compared to each respondent's current micro-climate. The Los Angeles region provides an interesting site because there are a number of different micro-climates. In particular, the communities along the coast have more moderate temperatures and a greater number of overcast days then communities on the floor of the San Fernando Valley. However, there is more rainfall on the average in the Valley because of the west-facing coastal range. Consequently, there is the possibility of testing our assumption that baseline micro-climates affect how respondents react to future climate scenarios. Further justifications for using the Los Angeles area follow shortly.

Before the climate scenarios were presented to each respondent, the baseline micro-climate was described. While all residents knew in broad terms what their local climate was like, we wanted to provide a correct and standardized understanding of key details. For residents living on the floor of the San Fernando Valley, the climate over the past 10 years was described as follows.

If you took the highest temperature on each day during the summer months and averaged them, the number would be 90 degrees. We call this number the summer average high temperature. While the average has been 90 degrees, we often see daily highs range from 85 to 95 degrees, with some extreme days over 100. These extreme days tend to occur back-to-back with as many as five in a row.

While there have been differences from year to year, this area has averaged 15 inches of rain per year over the last 10 years. Typically, the rainy season has been from November to February with up to three weeks in a row where there has been at least some rain.

For residents of the coastal communities, the climate over the past 10 years was described as follows.

If you took the highest temperature on each day during the summer months and averaged them, the number would be 75 degrees. We call this number the summer average high temperature. While the average has been 75 degrees, we often see daily highs range from 70 to 80 degrees, with some extreme days over 90 degrees. These extreme days tend to occur back-to-back with as many as five in a row.

While there have been differences from year to year, this area has averaged 13 inches of rain per year over the last 10 years. Typically, the rainy season has been from November to February with up to three weeks in a row where there has been at least some rain.

Given the number of dimension, the number of levels and how the sampling of each was done, there were a total of 13,965 scenarios. The following is one example of the 13,965, which in principle would have followed one of the baseline scenarios above.

Now, suppose over the next ten years the average high temperature during the summer months was 100 degrees, and the daily high temperature generally ranged between 80 to 120 degrees. Also suppose that the Los Angeles area received on the average 20 inches of rain per year and that most of the rain fell within a two month period during the winter months.

Note that in addition to the mean temperature, a measure of the spread in temperature was randomly selected. And in addition to the mean rainfall, a measure of spread in rainfall was randomly selected. By the luck of the draw, other dimensions could have appeared, as well as one randomly selected level within each.

After each climate scenario, there were a number of questions.

1. How uncomfortable or comfortable would all this make you feel, say, in and around your home? (four possible responses from very comfortable to very uncomfortable);
2. How concerned would this make you feel about the greenhouse effect? (four possible responses from very concerned to very unconcerned);
3. If these conditions really happened, how likely would it be that you might move out of the Los Angeles area? (Contingent Behavior – four possible responses from very likely to very unlikely);
4. Would you be willing to pay X dollars per year for things you normally buy in order to prevent the situation just described? (Contingent Valuation – with X a random value between 25 and 500 in $25 increments, and a yes or no response);
5. Two possible regulations, sometimes drawn from the regulations promulgated or proposed by the South Coast Air Quality Management District, chosen at

random from a set of 37, which the respondent could support or oppose. For example, after a scenario a respondent might be asked whether he/she would support or oppose banning all oil based paints or placing a city tax on gasoline. These will not be discussed in this paper.*

Twelve scenarios were given to each respondent with the baseline information read before the first scenario. However, the baseline information could be read again if requested by the respondent. The instrument also included the usual sorts of biographical questions and several more conventional attitude questions. We will make limited use of those items here.

4.2. SAMPLING AND INTERVIEWS

The population was defined as all adults living in either designated coastal communities or designated communities on the floor (as opposed to the foothills) of the San Fernando Valley.** The study may be usefully conceptualized, therefore, as two randomized experiments; the two populations are of little inherent interest, and statistical generalization, should that be desirable, is best accomplished through replications in other sites and with other kinds of respondents (Ehrenberg and Bound, 1993).

Respondents were selected by random digit dialing and by selecting a random respondent within households. The goal was to obtain approximately 600 respondents split evenly between the two regions. In fact, the effective sample includes 308 residents from the Valley floor and 312 respondents from the coastal strip. Since each rated 12 scenarios, the number of scenarios to be analyzed is 7,440. The response rate was approximately 60%.***

The interviewing was done on a CATI system using the program CASES. Scenarios were constructed on the fly as needed. Overall, the interviews seemed to go quite well. The interviewers reported that 95% of the respondents appeared to understand the questions and that in only 9% of the interviews did they (interviewers) have any difficulty probing for answers. Thirty seven percent of the respondents asked questions during the course of the interview, but nearly 90% of the time, interviewers felt they had the information required to answer. Twenty percent of the respondents complained about at least one question; the most common complaint was that there were too many climate scenarios. We will pursue their concerns in the data analysis.

* The use of these regulations in the study was another reason for choosing the Los Angeles area. The regulations are among the most demanding in the country and a source of public controversy. If one wants to study public opinion and environmental regulations, Los Angeles is an ideal place.

** The coastal strip was defined by zip codes adjacent to the ocean from Santa Monica to Redondo Beach. The Valley floor was defined by zip codes that excluded communities in the foothills. Included communities were parts of Van Nuys, North Hollywood, Northridge, Conoga Park and others.

*** We say 'approximately' because the response rate depends on what gets counted in the numerator and denominator.

4.3. DATA ANALYSIS

Given the randomized design, the impact of the scenario dimensions is relatively easy to analyze. Unbiased estimates of main effects may be obtained by considering each dimension in isolation. However, at least a few of the respondent biographical characteristics are of interest along with some interaction effects between scenario dimensions and respondent characteristics. There is also the complication introduced by the fact that four of the eight dimensions are 'missing' from each scenario. Finally, since scenarios are nested with respondents, one risks correlated disturbances within respondents.

All of these difficulties can be addressed, at least in principle, with a properly formulated hierarchal model (Wong and Mason, 1985; Hoox *et al.*, 1991); scenarios are the micro level and respondents are the macro level. Given the binary outcome, hierarchal logistic regression is a natural choice. The main technical advantage of the hierarchal model is proper estimates of the standard errors. For our analysis, correlated disturbance within respondents (across the 12 scenarios per respondent) implies that there is effectively a smaller sample than the number of scenarios. If these within respondent correlations are not addressed, standard errors will be too optimistic.

However, the model's estimated structural parameters (typically regression coefficients) will be at least consistent and in some cases, unbiased as well. Moreover, the improved estimates of the standard errors only follow if the model specified contains the appropriate functional forms and appropriate explanatory variables. Finally as a computational matter, the response variable of the main concern is binary, and the 620 macro units (i.e., respondents) and the 12 within-person micro units outstrip any publicly available software of which we are aware for hierarchal logistic regression.*

We report below, therefore, results produced by a conventional generalized linear model, but with a heavy reliance on graphics and smoothers to summarize key relationships. The problem of falsely optimistic standard errors was considered by first computing the average within-person correlation. That correlation was approximately 0.30 with some small variation depending on the model estimated. This is a modest figure suggesting that the usual standard errors were not grossly misleading. We then tried replicating the full sample analysis (i.e., $n = 7440$) with random subsets of the data using samples of different sizes. We found that the overall story was essentially unchanged until more than about a third of the data were discarded randomly (i.e., $n = 5000$). Some of the larger effects persisted with far smaller samples.** We eventually concluded that a sufficient margin of error

* We consulted quite extensively with our statistical colleagues at UCLA, several of whom have published in the hierarchical modeling literature (e.g., William Mason, Jan De Leeuw).

** We sampled by selecting less than the 12 scenarios for each individual: 11, 10, 9 and so on. Since the problem was the within-person correlations, each sample retained the full compliment of 620 respondents and randomly discarded scenarios.

existed for the reported standard errors; the correlation of 0.30 does not imply that our effective sample is less than two-thirds of the actual sample.

The response variable of initial concern was whether or not the respondent was "willing to pay *X* amount of dollars for things you normally buy in order to prevent the situation just described". Recall that *X* is a random number drawn between the values of 25 and 500 in increments of 25, which means that *X* must be included as an explanatory variable and, in effect, becomes a part of each climate scenario. We treated the dollar variable as a set of binary variables, since we had no theoretical rationale for specifying any particular functional form beyond a monotonically decreasing one.

Explanatory variables from the climate scenarios were simply representations of the design matrix with presence of any given level (within a given dimension) coded as '1' and absence coded as '0'. The design matrix included all of the levels within all 8 dimensions. Recall that since only four dimensions were used in any given scenario, half the dimensions were 'missing' (completely at random) from any given scenario. 'Missing' was simply included as one of the levels in the design matrix. When these were added to the other variables in the design matrix, a total of 63 scenario explanatory were initially defined. Later, we briefly consider smoother versions of these variables.

We were concerned that wherever in the sequence of twelve a given scenario was evaluated it could affect WTP. The scenarios were demanding, and it is unlikely that any of the respondents previously had thought much about global warming in the manner in which it was presented. Thus, there was a reason to expect learning to occur in the first couple of scenarios and perhaps some fatigue and/or frustration to surface near the end. To capture such patterns, we included eleven binary variables for order: 2nd, 3rd, ..., 12th.* The use of the binary variables means that no functional form is being assumed.

We initially included as macro explanatory variables biographical indicators which might be related to WTP (Stern *et al.*, 1993): income, education, gender, age, ethnicity, in which of the two micro-climates the respondent lived and whether the respondent was employed. We did not include any attitudinal measures because we could not with confidence specify the causal direction.

5. Findings

5.1. DOLLAR FIGURE OFFERED

Perhaps the most important issue before the details of the results are considered is what evidence there is that respondents took the scenario exercise seriously and then were able to make sense of the task. An indicator is how they responded to the various dollar amounts in their WTP assessments. One would have to be very skeptical of any substantive conclusions from the analysis if the probability of

* We deleted the first one to prevent linear dependence.

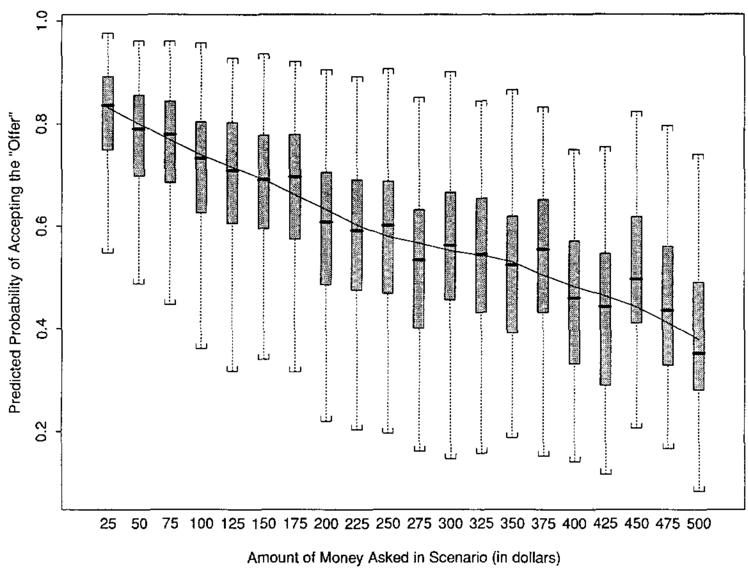

Fig. 1. Predicted probabilities by category of amount of money asked.

agreeing to pay was unrelated to the dollar amount offered. Rather, one should find a monotonic decrease in WTP as the dollar amount increased.

Figure 1 shows side-by-side boxplots of the predicted probability of WTP for different dollar amounts. For example, the first boxplot shows the distribution of predicted values for all scenarios in which WTP was assessed at $25. The dispersion reflects the impact of *other variables in the model*; the dispersion is *not* a representation of the residuals.

The goal of Figure 1 is to display the impact of the set of categorical variables for the dollar amount offered to respondents. A loess smooth through the medians of each distribution is included as a visual aside (Hastie and Tibshirani, 1990: 29–31). Clearly, respondents reacted as anticipated despite the novel nature of the task and the amount of information they were being asked to evaluate. A likelihood ratio test for the set of binary variables easily rejects the null hypothesis of no impact, and a linear fit treating the dollar amount as continuous has a statistically significant negative slope. While Figure 1 is hardly proof that all is well, the findings are encouraging.*

Figure 1 also provides a general sense of the importance of the dollar amount offered. The dollar amount offered can on the average shift the probability that a respondent will accept the dollar amount from approximately 0.45 to approximately 0.85. That is, much of the middle segment of the logistic curve is reproduced. However, one can also see from Figure 1 that while the dollar amount offered is important in WTP, it is only a small part of the overall story. The dispersions

* The conclusions are the same when one-third of the data are randomly dropped as described above, and this will be true for all the statistically significant effects we report.

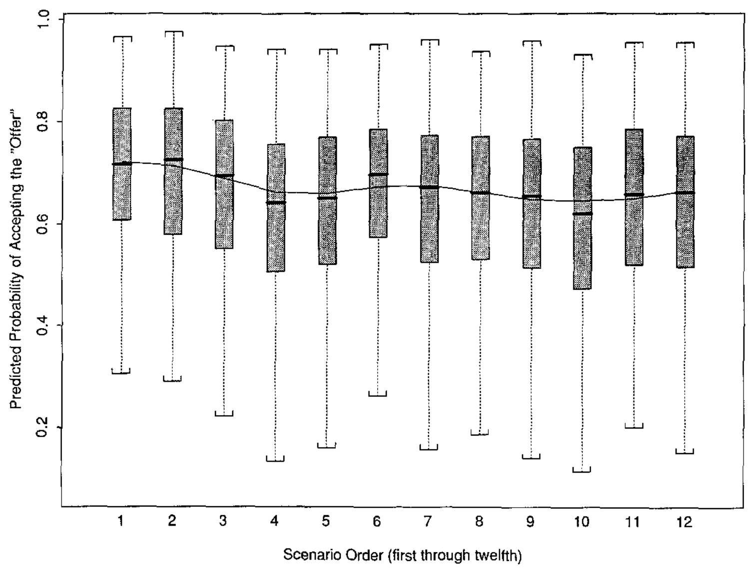

Fig. 2. Predicted probabilities by category of scenario order.

represented are the effects of the *other variables in the model.* For example, the offering of $300 has a median probability of acceptance of about 0.60. But other variables can move that probability from about 0.20 to about 0.90.

More generally, the wide dispersions apparent in Figure 1 are another indication that the scenarios may have worked well. Figure 1 shows systematic variation produced by the scenarios and respondent biography; various combinations of scenario characteristics were able to produce systematic variation over much the range between 0.0 and 0.1 for the willingness to pay.

5.2. SCENARIO ORDER EFFECTS

Figure 2 addresses possible effects from the random ordering of the scenarios. The vertical axis is again the predicted probability and the horizontal axis is random placement (1st, 2nd, ..., 12th) of the scenario. Recall that we were interested in the possibility of the respondent learning and/or respondent fatigue. There is virtually no evidence that the probability of accepting the offered dollar amount varies systematically with order.* This is good news and a bit of further support that on the average the instrument worked roughly as planned. More evidence will be considered as we proceed.

* Because of the randomization, order is on the average unrelated to the other properties of the scenarios and respondents. As a result, ignoring additive order effects will not bias the results. And in any case, we include order as a set of explanatory variables in the multivariate analysis.

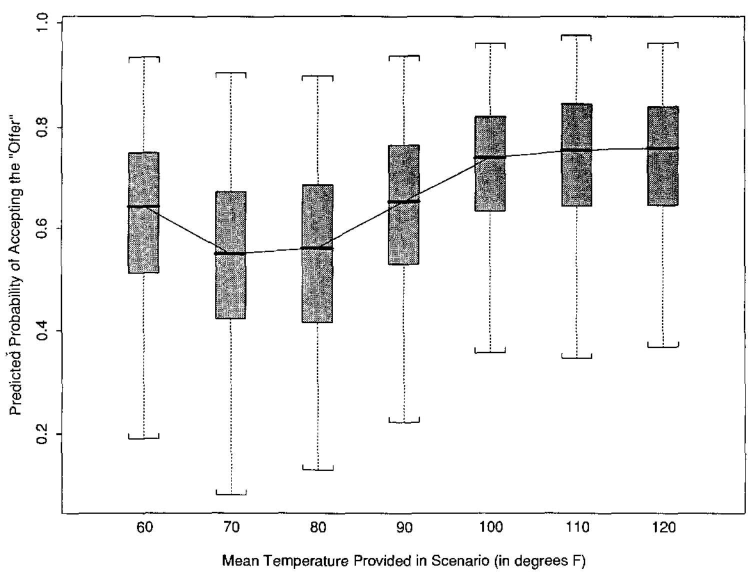

Fig. 3. Predicted probabilities by category of mean temperature.

5.3. SCENARIO CLIMATE EFFECTS

A useful way to consider initially the role of the different scenario dimensions is through likelihood ratio tests of the various clusters of binary variables. For these, it is readily apparent that for rainfall, only the means affect WTP. Only for the mean can one reject the null hypothesis of no contribution to the deviance. For temperature, the mean, extreme values, and the clustering of extreme values all are apparently related to WTP. For all three, the null hypothesis of no contribution to the deviance is easily rejected. Perhaps the major conclusion is that the spread for temperature and rainfall by themselves (and within the ranges we examined) do not seem to affect WTP. *Variability in the climate at the local level does not seem to concern our respondents.* As a scientific matter, of course, climate variability by itself can have enormous implications for agriculture and wildlife.

Figure 3 shows side-by-side boxplots of the predicted probability of WTP for different mean high temperatures. (Recall, these are mean daily highs, not the mean overall). A loess smooth of the medians is undertaken to summarize the pattern. As expected, WTP increases with distance from a 'comfort zone' of between 70 and 80 degrees. Any cooler or any hotter, and WTP increases. We wondered whether WTP might increase at an increasing rate, but if anything, the rate of increase slows down where there are enough data points draw to any conclusion. It may be that when mean high daily temperatures exceed 100 degrees, some respondents may see the situation as catastrophic. Willingness to pay then becomes irrelevant. Another possibility is that at these very high mean temperatures, some respondents found the exercise artificial and stopped making considered judgements.

While from Figure 3 respondents may have behaved roughly as anticipated on the average, it is difficult to know from Figure 3 alone if the impact of the mean high daily temperature has a large substantive impact on WTP. The median predicted probabilities range from a low of about 0.50 to a high of about 0.70. On their face, these figures seem to imply a non-trivial impact. In addition, if despite our earlier warnings the dollar values are taken literally, the differences in dollars are considerable. At a mean high temperature of 80 degrees, the respondents are willing to pay on the average about $100. At a mean high temperature of 100 degrees, respondents are willing to pay on the average about $140.* The increment represents a 40% gain in WTP as a function of a 20 degree increase in mean high temperature, averaging over other scenario characteristics.

However, a 20 degree (Fahrenheit) increase in mean temperature is well outside the range of current scientific estimates of greenhouse warming. It follows that more realistic mean temperature increases would produce much smaller increments in WTP. Moreover, our scenarios asked about mean *high* temperatures, and we know of no scientific estimates for that climate parameter. Depending on the future shape of the temperature distribution, a 20 degree increase in mean high temperature may or may not be credible. In future work, therefore, it would be useful to undertake separate analyses for the particular scenarios that are more realistic.**

In Figure 2, we considered the possible additive effects of order. No order effects of any importance was found. In Figures 4 and 5 we explore the possible interaction effects between order and mean high temperature. For example, if respondents were becoming tired or frustrated as they worked through their twelve scenarios, the impact of mean temperature on WTP might be apparent for the early scenarios and not for the later scenarios. Alternatively, respondents may come to believe after the first few scenarios that there are combinations of climate conditions that can be anticipated. Consequently, later scenarios would be evaluated partly as a function of earlier scenarios. In fact, the response patterns for the first four scenarios (Figure 4) and the last four scenarios (Figure 5) are about the same. Respondents seem to be reacting to variation in climate in a similar fashion whether the scenario comes early or late. Figures 4 and 5 are still another bit of evidence that the scenarios on the average worked approximately as advertised.

Figure 6 shows the impact of extremes in temperature. The effect of the category for omitting the dimension from the scenario altogether is included at the far left. With a median probability of WTP of about 0.65, WTP for 'omitted' falls roughly in the middle of the other distributions shown. This suggests that respondents were implicitly (and almost certainly unconsciously) doing some averaging over the

* These figures were computed by multiplying the mean of the dollar amount offered by the probability of accepting the offer.

** We have recently begun such work collaboratively with atmospheric scientists. A critical initial step is to learn what scenarios are in fact more realistic. For example, our climate dimensions were orthogonal by design. In the real world, our climate dimensions would not be. But what are plausible correlations between the climate dimensions? For instance, is mean temperature positively or negatively correlated with mean precipitation?

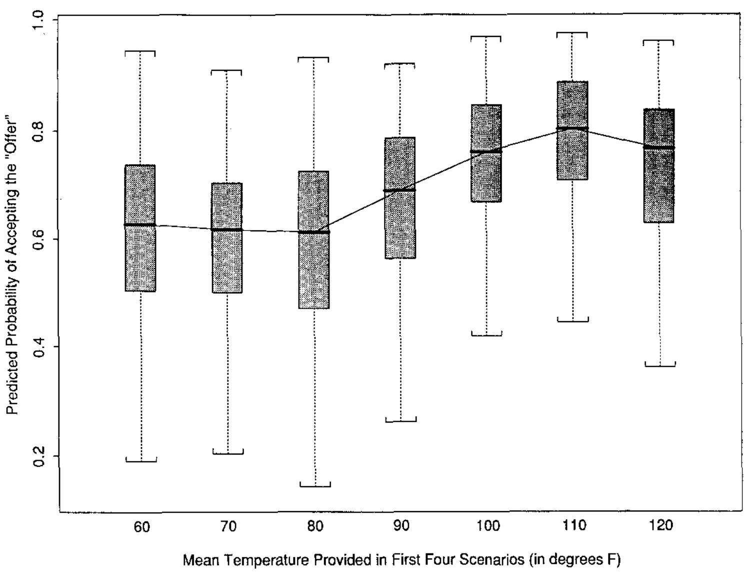

Fig. 4. Predicted probabilities by category of mean temperature, first four scenarios only.

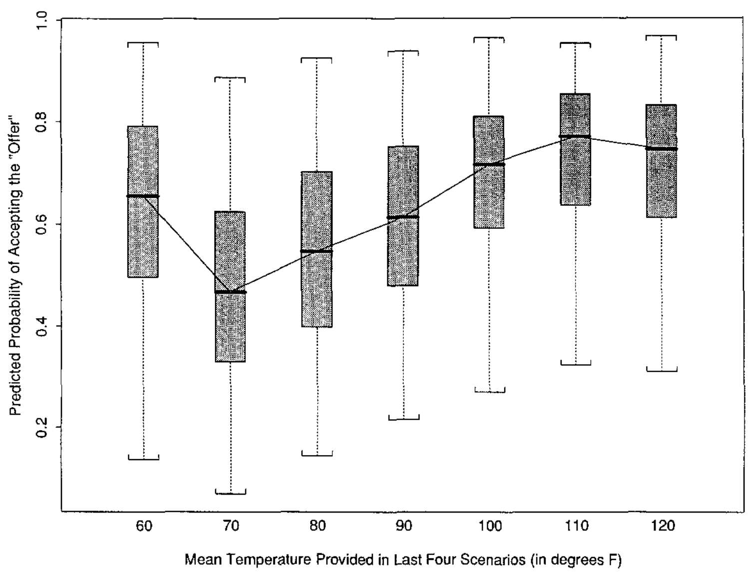

Fig. 5. Predicted probabilities by category of mean temperature, last four scenarios only.

scenarios they evaluated and then inserting that average when extreme values were not mentioned.

Scenario information on extreme temperature values produce roughly a V-shaped response curve. Extreme values of 90 degrees seem to produce the least WTP. Higher or lower extremes produce a greater WTP. Still, there is no evidence of acceleration in WTP at the tails. The increase in WTP is relatively constant.

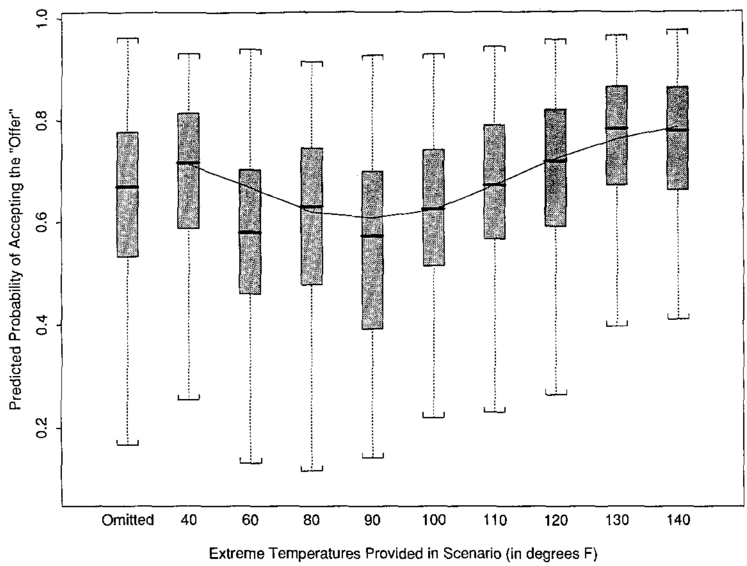

Fig. 6. Predicted probabilities by category of extreme temperature.

Finally, the size of the effects produced by the extreme temperatures is about the same size as the effects of the mean of high temperatures. The extreme temperatures can shift the median probability of accepting the dollar offer from about 0.55 to 0.75.

Figure 7 shows in the same way the impact of the clustering of extreme temperatures. Again, the impact of the variable for omitting the dimensions is roughly the average of the other effects. And again, an inference is that respondents insert some typical value when a particular dimension is not in the presented scenario.

The impact of clustering is much smaller than the impact of either mean high temperature or extreme temperature. The difference between the average of the highest and the average of the lowest probability is about 0.10. Perhaps more important, it is difficult to see any consistent association between the degree of clustering and WTP. When there is no clustering, WTP is the lowest. But the two highest probabilities are for 2 days and for a month of extreme temperature. Values in between have lower probabilities.

Two interpretations may be constructed. First, any clustering whatsoever increases willingness to pay, compared to no clustering at all. So, the clustering of extremes seems to matter, but more crudely than one might expect. Second, the average predicted probability for 2 days and 7 days of extreme temperatures are the major exception to a monotonic increase associated with the number of days of clustering. Sampling error could easily account for the anomaly.

Turning to rainfall, recall that only mean rainfall affected WTP. The spread, extremes and clustering of extremes had no discernible impact on WTP. Figure 8 shows the impact of mean rainfall.

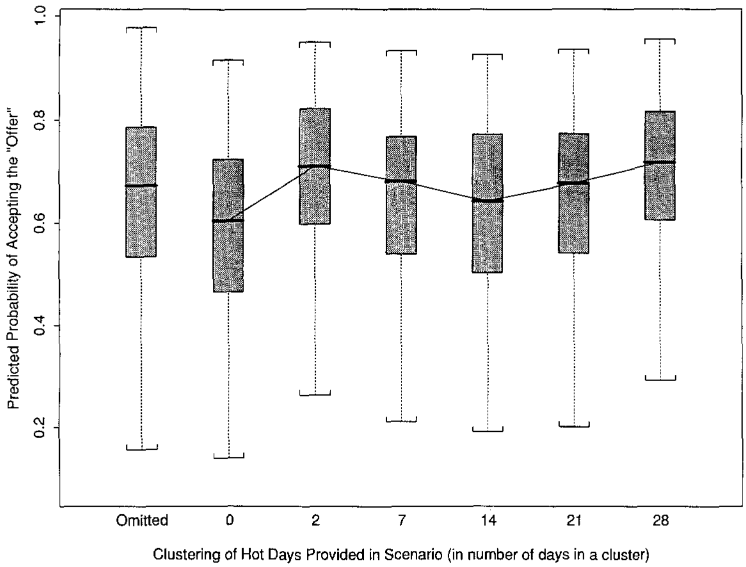

Fig. 7. Predicted probabilities by category of temperature clustering.

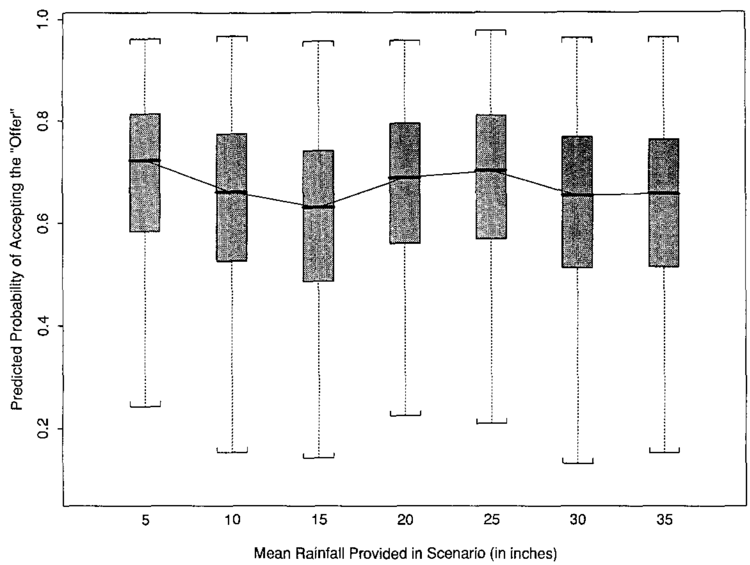

Fig. 8. Predicted probabilities by category of mean rainfall.

Overall, the effect of mean rainfall is modest. The difference in the probability of accepting the dollar amount offered between the most desirable and least desirable average amount of rain is only about 0.10. However, there is again a suggestion of the V-shaped curve. Mean rainfall of about 15 inches has the lowest predicted probability. The probabilities increase in both directions until the mean of 25 inches is reached. The probabilities then decline. However, all adjacent differences

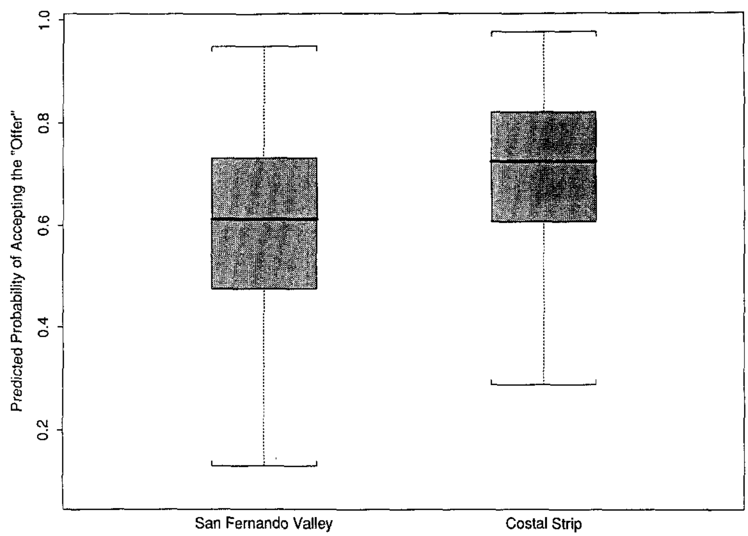

Fig. 9. Predicted probabilities by micro-climate.

between the average predicted probabilities are really quite small and within the conventional boundaries of sampling error. We are inclined, therefore, not to take the drop too seriously.

5.4. EFFECTS OF RESPONDENTS' BACKGROUNDS

The estimated impact of micro-climate is relatively modest. As Figure 9 shows, respondents living along the coastal strip had a median probability of about 0.70 of accepting the dollar amount offered, while respondents living on the floor of the San Fernando Valley had a probability of about 0.60. Note that in Figure 9 statistical adjustments have been made for all other biographical variables in the model, such as education and income; these are net effects for micro-climate.

We have no compelling explanation for the micro-climate impact. One possibility is that coastal residents are simply living up to their green reputation. Coast communities such as Santa Monica and Malibu have long been associated with such local environmental groups as Heal the Bay and American Oceans' Campaign. Another possibility is that Valley residents have more experience with hot summers and may feel less threatened by the prospects of global warming. In other words, how people evaluate the prospect of global warming may depend on the current climate that they experience day-to-day.

Figures 10–15 show similarly modest effects, respectively, for ethnicity, income, employment, gender, age, and education. WTP is higher for Whites (compared to all minorities), males, respondents who are employed, respondents who earn more than $60,000 a year, younger respondents, and more educated respondents. The impacts of employment and income on WTP are predicted by economic theory and common sense: with greater financial resources come a greater ability to pay.

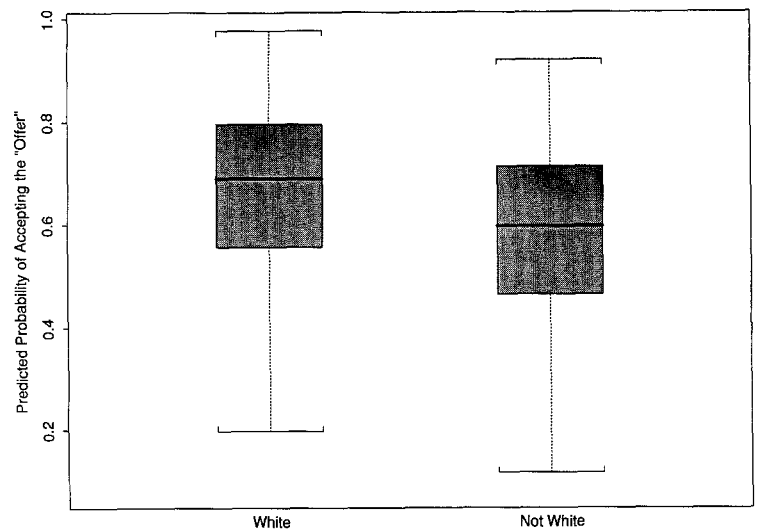

Fig. 10. Predicted probabilities by ethnicity.

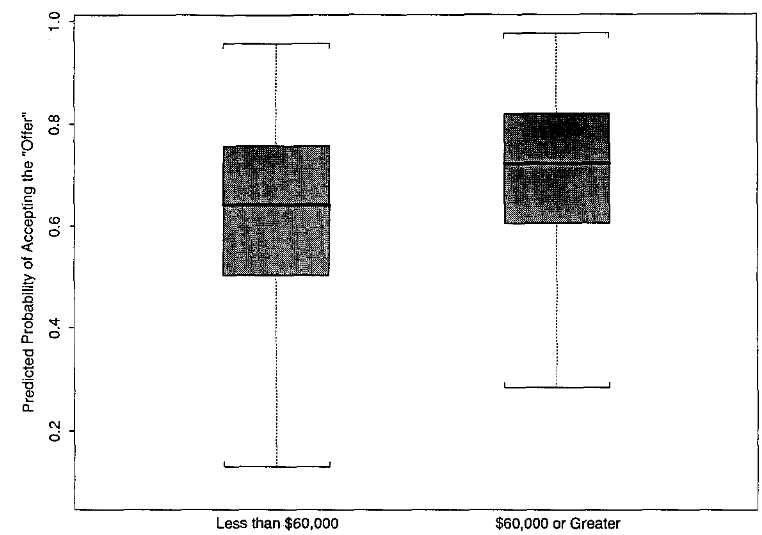

Fig. 11. Predicted probabilities by total household income.

Although the literature on environmental attitudes offers many possibilities (Stern *et al.*, 1993), the impacts of ethnicity, gender, age, and education are far more difficult to explain. Part of the problem is that all of these variables represent proxies for life experiences and circumstance that need to be measured directly. For example, what is it about the experiences and circumstances of Whites, *net all the other biographical variables in the model*, that make them more willing to pay?

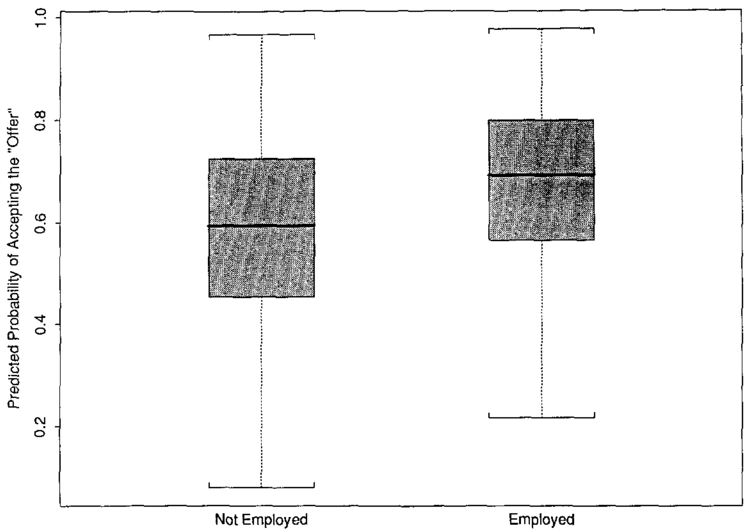

Fig. 12. Predicted probabilities by employment.

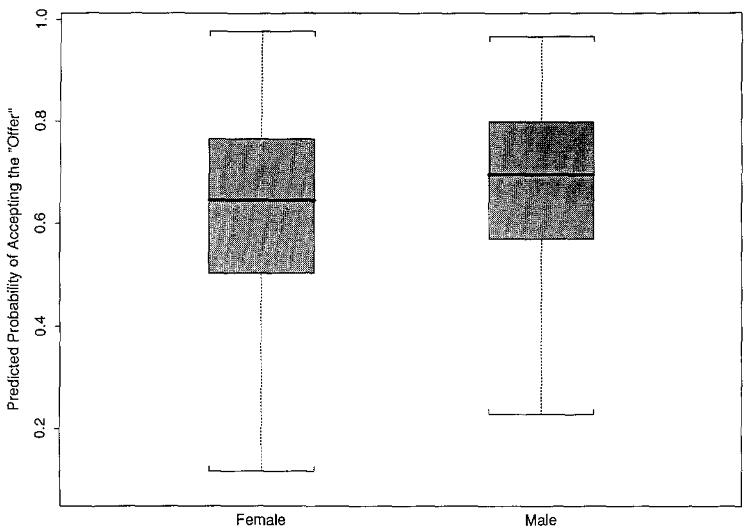

Fig. 13. Predicted probabilities by gender.

In short, respondents brought to the scenarios some predispositions, but the sources of these predispositions typically remain obscure. Moreover, WTP was affected greatly by scenario content so that the predispositions that people brought to the task were hardly definitive.

5.5. SMOOTHING RESULTS

All of the scenario results reported in the first eight figures are non-parametric in the sense that no functional form is being assumed between the given scenario

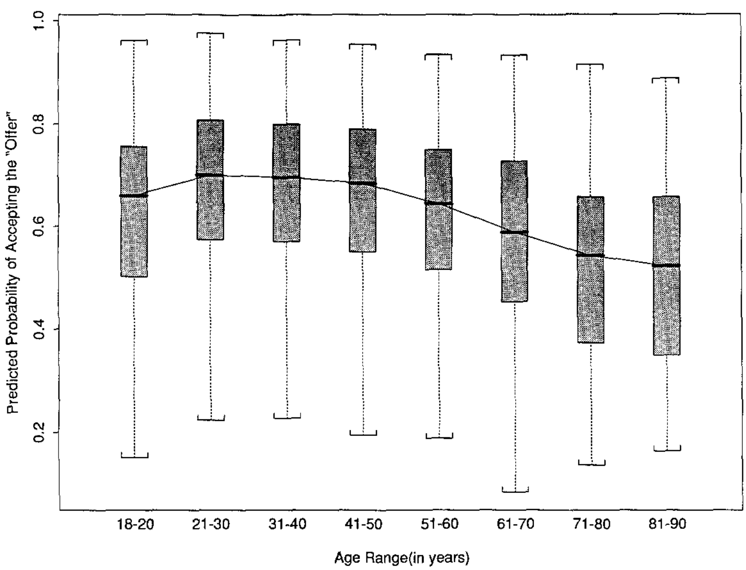

Fig. 14. Predicted probabilities by age.

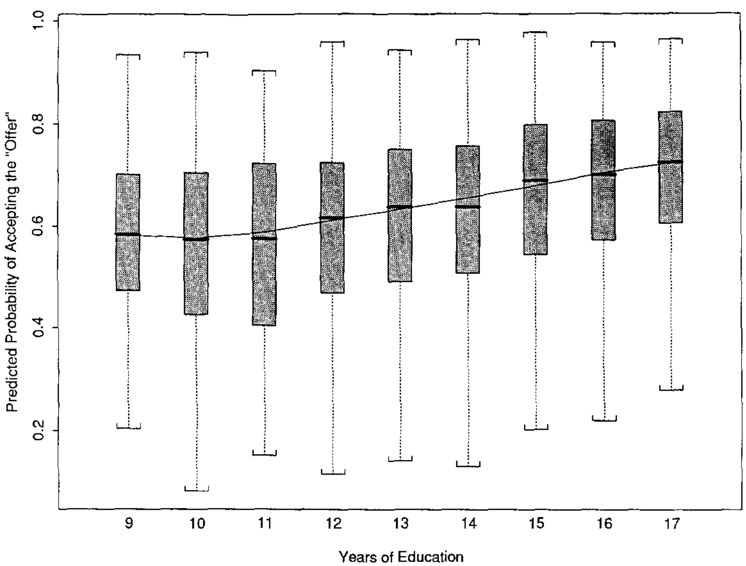

Fig. 15. Predicted probabilities by education.

dimension and the response. It is also possible to move toward a functional form inductively having examined Figures 1–8.

Consider as example, Figure 16, where a plot of a smooth version of the impact of the mean temperature and its 95% confidence interval are shown. The smoother is loess, estimated as a part of a larger generalized additive model (Hastie and Tibshirani, 1990). That is, the same multivariate model as reported above is once

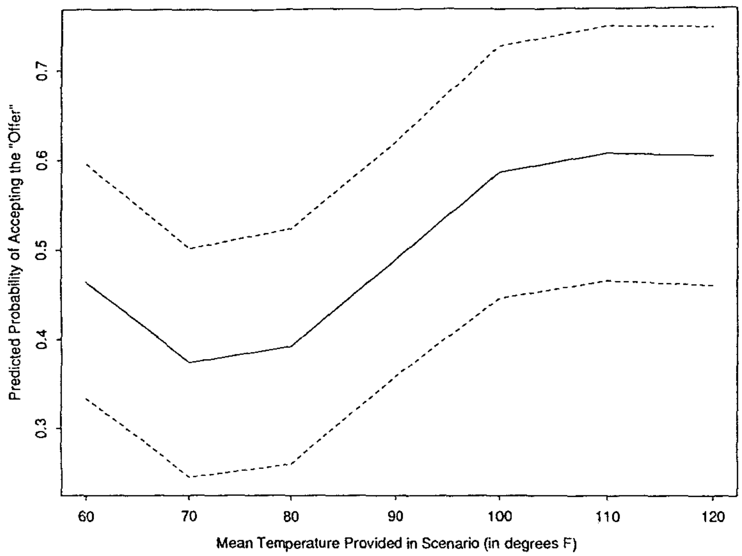

Fig. 16. Lowess smooth of mean temperature using a generalized additive model with a 95% confidence interval.

again used, but with mean temperature treated as an equal interval variable and with a loess smooth used only for that variable.* The story is much the same; the now-familiar (drooping) V-shaped pattern reappears.

From Figures 3 and 15, it would seem that a third degree polynomial would fit well the data for mean high temperature. In fact, when the loess smooth was replaced by a third degree polynomial, the fit was excellent. A likelihood ratio test failed to reject the null hypothesis that the polynomial fit the data as well as the loess smooth. This confirms that the rate of increase does indeed seem to taper off at high mean temperature values.

We had the same experience with each of the figures; in each case, a linear, quadratic or cubic polynomial fit the data as well as a loess smooth. And in each case, the substantive story from the original figures did not change. In short, it is possible to simplify the parameterization of the model with no real loss of information. The simplification was not important for our substantive story, in the smooths may be hints of functional forms useful for future theoretical work.

5.6. INTERACTION EFFECTS

We experimented with several interaction effects. Perhaps the most potentially interesting was the interaction between micro-climate and temperature. Figure 17 shows the impact of mean high daily temperature for the two micro climates. Consistent with Figure 9, the difference in the probability of accepting the dollar

* Dimensions not included in a given scenario were dropped from the analysis. Treating only one dimension at a time as an equal interval minimized the amount of data that were deleted. Mean temperature and mean rainfall were included in all scenarios and therefore, never deleted.

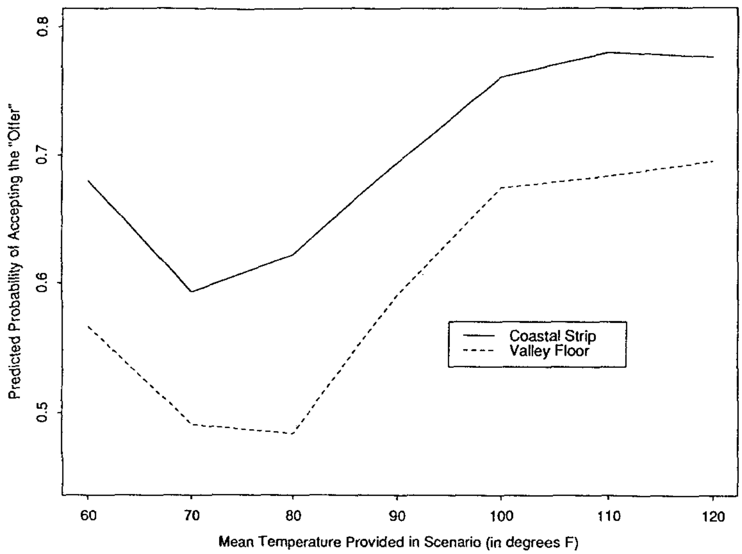

Fig. 17. Predicted probabilities by mean temperature for each micro-climate.

amount averages about 0.10. It is clear that the residents of the coastal communities are generally more willing to pay overall. However, the gap between the impacts of the two micro-climates is rather constant over the full range of mean high temperature. There is no evidence of an interaction effect. Similar null findings surfaced for the few other interaction effects we explored. We are skeptical that large two-way interaction effects exist in the data, but we leave a more complete exploration to future work.

5.7. USING OTHER RESPONSES AS EXPLANATORY VARIABLES

As discussed at length earlier, a key issue in the contingent valuation literature is what survey questions on willingness to pay or willingness to accept really measure. We can provide some insight by examining the impact of our other responses measures used as *explanatory variables*. We begin with the model that has provided most of the findings so far (basically, the design matrix plus biographical variables) and then add: (1) degree of comfort; (2) degree of concern about global warming; and (3) contingent behavior (moving from Los Angeles). If the contingent valuation response to the original set of explanatory variables is little more than a restatement of these three items, the impact of the scenario content and respondent background on WTP should disappear when the three variables are added to the model.

When all three alternative responses are included in the model, there is *no statistically significant impact of scenario content on willingness to pay*. Whatever content WTP has beyond the content of the three responses, it is not related to characteristics of the scenario. That is, the scenarios affect WTP and a linear combination of the three other responses in a very similar fashion. However, the predisposing effects from respondent background remained. That is, respondent biography has somewhat different patterns of association with the contingent valuation responses; for the impact of respondent background, WTP is *not* redundant with a linear combination of the other three responses variables. In short, respondent predispositions distinguished between the different response variables much more than did scenario content. For example, employment has no impact on degree of comfort or concern, but a large impact on willingness to move from Los Angeles; not surprisingly, unemployed people were more inclined to move. In general, respondent background had about the same impact on degree of comfort and degree of concern, but a rather different impact on willingness to move and still different impact on willingness to pay. With 20–20 hindsight, this makes sense. However, a more thorough exploration of impact of respondent background on the different response variables is underway and will be the subject of another paper.

More important for this paper, when each of the alternative response variables was added by itself to the full model, small but statistically significant effects remain for the impact of climate scenarios. That is, no single alternative response variable fully eliminates the impact of the scenarios. However, degrees of comfort or discomfort is by far the most effective of the three in reducing the role of scenario content. One implication is that of the three, scenario impact on WTP and on comfort are the most alike. This is encouraging because comfort is clearly one important aspect of use value.

Degree of concern and contingent behavior (i.e. moving) were rather less effective in eliminating scenario effects. It is encouraging that there is some overlap between the impact of the scenarios on WTP and on contingent behavior. Moving from Los Angeles entails costs that properly should be captured by WTP. However, the overlap in scenario impact on WTP and on concern about global warming may be troubling for contingent valuation advocates. Concern about global warming would seem to capture not just the immediate implications of climate change for personal well-being, but commitments to stewardship of the planet and altruistic motives about the well-being of others. We earlier reviewed the difficulties with these sources of value.

In summary, our findings suggest that WTP captures some genuine aspects of use value. At the same time, our findings suggest that some aspects of stewardship and altruistic value are being captured as well. In our data at least, a combination of the two seem to account fully for the impact of the climate scenarios on WTP.

6. Summary and Conclusions

6.1. INSTRUMENT ARTIFACTS

Whatever our findings, everything depends on how respondents undertook the task set before them. On the one hand, there is ample reason to be concerned that respondents either could not or would not properly evaluate the information provided. References scattered throughout the paper are sufficient grounds of skepticism. On the other hand, all of the evidence we have is consistent with an alternative view that respondents evaluated the scenarios in a serious and sensible manner. The substantive findings seem plausible and our efforts to document response errors of various sorts failed.

Only if we had data on how each respondent evaluated the information provided could we definitively decide between these conflicting views. Unfortunately, the only data available comes from early pre-testing when respondents were debriefed. From that information, we suspect that respondents proceeded by first scanning a given scenario for its most salient features. Extreme values on one or more climate dimensions were perhaps especially important. Then, the reported WTP was based solely on those salient features. The other features were ignored. One implication is that respondents may not have used simultaneously all of the climate information provided. However, since by design each climate dimension evidenced extreme values at one time or another, our aggregate analyses would reveal the impact of each climate dimension. If these speculations are correct, there is perhaps no conflict between those who believe that climate scenarios such as ours are too demanding and those who accept our aggregate effects. In short, our findings rest on averages across respondents and scenarios. How those averages are related to the evaluations of individual respondents cannot be effectively determined from our data.*

What, then, are our major conclusions about possible response errors? To begin, the respondents appear to have been up to the task put before them. They reacted to the dollar figures offered in sensible ways and largely as anticipated to the climate scenarios. In addition, respondents appeared to handle well the omitted scenario dimensions and were seemingly unaffected by the order in which the scenarios were presented. With respect to the omitted scenario dimensions, recall that respondents seemed to implicitly impute a rough average for each missing scenario dimension. This is certainly one sensible heuristic in the face of missing information. With respect to order effects, recall that when order was included as an explanatory variable, no additive effects or interaction effects were found. One inference is that respondents may have been able to get up to speed quickly and not appreciably tire.

An important implication is that studies of attitudes about environmental issues need not be restricted to simple one-liners couched in the social rhetoric of the day (unless the goal is to study that rhetoric). The public seems capable of considering environmental issues in some of the same complexity as informed policy makers. This is good news for advocates of contingent valuation.

But, there is bad news as well. Willingness to pay, as elicited by at least our questionnaire, seems to confound several sources of value. Stewardship and altruism are two important possibilities. Consequently, the dollar figures respondents generate cannot be used in conventional cost-benefit analyses. However, we see

* This is a vexing problem in economics more generally (Stoker, 1993).

some promise in the design of questionnaires with a set of response variables that would far more directly tap different kinds of value. It might then be statistically possible to isolate use value from other kinds of value with multivariate techniques much like we employed.

6.2. MAJOR FINDINGS

Turning to the impact of scenario dimensions, our respondents were more likely to be moved by changes in temperature than by changes in precipitation. This may result in part from the use of the term global warming in our introductory material and the more general publicity that global warming has received. Or, Southern California residents may simply be more attuned to variation in temperature than variation in rainfall. In fact, most of the water used in Southern California is imported from hundreds of miles away. However, any such conclusions are no doubt influenced by the range of values built into our scenarios and by the baseline micro-climates. It is difficult to know what would have happened had we allowed for a wider range of precipitation futures or had our respondents lived in substantively wetter or drier climates.

It is also interesting that spread had no impact on WTP for either temperature or precipitation: variation in climate by itself did not seem to concern our respondents. We find this surprising because as a factual matter, variation in climate can have dramatic impacts on wildlife and agriculture. It may be, however, that such concerns are not sufficiently salient and that whatever the potential impact on personal comfort, it is easily muted. Spread refers to climate variation around some central tendency, and one can simply wear a coat or take it off. One can simply carry an umbrella or not. It is easy for people to respond to variation per se.

The impact of micro-climate underscores that how people respond to the hypothetical events depends upon their current circumstances. If the issue is change, a critical aspect is change from what; the baseline is vital. Future studies of environmental studies might benefit from a substantive investment in measures of the environmental status quo.

Consistent with most studies using factorial survey methods, elicited responses seem to vary much more with characteristics of the scenarios than with particular biographical characteristics. Yet, it is difficult to draw any general substantive conclusions because variation in the scenarios is designed; the variation in our scenarios may or may not represent naturally occurring variation. Our climate scenarios were designed with plausible futures in mind (insofar as that is known), but with more emphasis on extreme values for all of the dimensions than is likely to materialize. We were following the common practice of anchoring stimuli at the tails. Were the climate scenarios anchored at more likely values, the impact of the scenarios might have been more comparable to the impact of respondent biography. On the other hand, climate change may be a true leveler in which concerns about large scale perturbations in temperature and precipitation are shared by people from very different backgrounds. That is, there may be substantial similarities among people in the structure of how they value climate (Rossi and Berk, 1985). It is also possible, however, that with increasing exposure to the prospect of global warming, public opinion will become more differentiated.

6.3. IMPLICATIONS FOR FUTURE WORK

Looking ahead to the future work, we have three general suggestions. First, despite some genuine concerns about what respondents are doing with the scenarios, we continue to favor factorial survey methods over conventional single dimension questions. Perhaps most important, the scenarios provide information that would otherwise be invisibly introduced by respondents. A sensible answer to a conventional single dimension question may be 'it depends'. If, nevertheless, an answer is coerced, respondents will likely construct the needed information themselves. While scenarios certainly do not preclude such complications, we suspect that when faced with relatively rich scenarios, respondents will be less tempted to fill in the blanks.

Second, it is possible to do a much better job isolating the kinds of value WTP is really capturing. We reiterate our suggestion above that a wider range of response variables be included in future surveys. We also think there is promise in designing scenarios so that different sources of value could be explicitly manipulated. For example, a scenario might include various adaptations to global warming that minimized any impact on human well-being, but which did not protect certain plant and animal species.

Third, we strongly favor research that would better reveal how respondents employ the information provided in the scenarios. Verbal protocol studies of the sort undertaken by Schkade and Payne (1992, 1993) have considerable promise. In addition, the scenarios themselves can be designed to address important cognitive questions. For example, one can randomly drop entire scenario dimensions to see how the reduction in information affects the response. But, until there is far better understanding about how respondents arrive at their contingent valuations, WTP estimates will not rest on a sound foundation.

6.4. IMPLICATIONS FOR SOCIAL POLICY

For reasons that should now be apparent, we do not believe that our WTP estimates can be properly used to place values on climate. Nevertheless, there are in the results some general implications for social policy. To begin, policy makers should not underestimate the public's capacity to understand the key scientific issues. Efforts to 'educate' the public through sound bites and other baby talk may do more harm than good. Our results suggest that given the proper setting, people with little scientific training can digest and evaluate rather complex factual material.

At the same time, however, responses to the climate scenarios highlight potential gaps in the public's ability to understand climate change. It appears that very large changes in climate are required before WTP is significantly affected. We suspect that the public does not fully appreciate the consequences of seemingly small climate changes. In addition, the null findings for climate variability suggest that the effects of variation in climate also is not fully appreciated.

The impact of micro-climate on WTP underscores that insofar as global climate change occurs, it will be experienced and evaluated locally. Climate modelers appreciate that global averages convey little about regional effects. Policy makers must appreciate that public opinion about climate change will vary substantially by locale.

Finally, our skepticism about the current state of contingent valuation should not be misconstrued as denying the need to better value climate and other environmental goods. Nor are we saying that contingent valuation is a scientific dead end. Given the central role that tradeoffs play in environmental policymaking, it is essential that we learn to more accurately place a value on the environment. Contingent valuation may one day help in that process.

Acknowledgements

Funding for this research was provided by the Program on Global Change within the National Science Foundation. Thanks go to Eric Hanley for his assistance in the pre-test phase and to Trudy Cameron and Hilary Sigman for suggestions on our use of the environmental economics.

Finally, many of the ideas found in this paper were first discussed at a workshop under the NATO Special Programme on the Science of Global Environmental Change. We are deeply indebted to the international group of scientists who participated.

References

Arrow, K. J. and Fisher, A. C.: 1974, 'Environmental Preservation, Uncertainty, and Irreversibility', *Quart. J. Econom.* **87** (May), 312–319.

Arrow, K. J., Solow, R., Leamer, E., Portney, P., Radner, R., and Schuman, H.: 1993, *Report on the NOAA Panel on Contingent Valuation*, January, 12.

Baumol, W. J. and Oates, W. E.: 1988, *The Theory of Environmental Policy* (second edition), Cambridge University Press, Cambridge.

Cameron, T. A.: 1991, 'A New Paradigm for Valuing Non-Market Good Using Referendum Data: Maximum Likelihood Estimation by Censored Logistic Regression', *J. Environm. Econom. Managem.* **76**, 413–421.

Cropper, M. L. and Oates, W. E.: 1992, 'Environmental Economics: A Survey', *J. Econom. Literat.* **30**, 675–740.

Davis, R. K.: 1963, 'Recreational Planning as an Economic Problem', *Natural Resourc. J.* **3**, 239–249.

Daum, J. F.: 1993, 'Some Legal and Regulatory Aspects of Contingent Valuation', in Hausman, J. A. (ed.), *Contingent Valuation: A Critical Assessment*, North Holland, New York.

Diamond, P. A., Hausman, J. A., Leonard, G. K., and Denning, M. A.: 1993, 'Does Contingent Valuation Measure Preferences? Experimental Evidence', in Hausman, J. A. (ed.), *Contingent Valuation: A Critical Assessment*, North Holland, New York, pp. 41–85.

Ehrenberg, A. S. C. and Bound, J. A.: 1993, 'Predictability and Prediction (with discussion)', *Statist. Soc.* – Series A 156 (part 2), 167–206.

Fischhoff, B.: 1991, 'Value Elicitation: Is There Anything There?', *Amer. Psychol.* **46**, 835–847.

Groves, R. M.: 1989, *Survey Errors and Survey Costs*, John Wiley, New York.

Hastie, T. J. and Tibshirani, R. J.: 1990, *Generalized Additive Models*, John Wiley, New York.
Hausman, J. A.: 1993, *Contingent Valuation: A Critical Assessment*, North Holland, New York.
Hoehn, J. P. and Loomis, J. B.: 1993, 'Substitution Effects in the Valuation of Multiple Environmental Programs', *J. Environm. Econom. Managem.* **25**, 56–75.
Howe, C. W.: 1979, *Natural Resource Economics*, John Wiley, New York.
Hoox, J. J., Kreft, G. G., and Hermkens, P. L. J.: 1991, 'The Analysis of Factorial Surveys', *Sociolog. Methods Res.* **4(19)**: 493–510.
Kahneman, D.: 1986, 'Comments', in Cummings, R. D., Brookshire, S., and Schultz, W. D. (eds.), *Valuing Environmental Goods: An Assessment of the Contingent Valuation Method*, Totowa, N. J.: Rowman and Allen, 1986.
Kahneman, D. and Knetsch, J. L.: 1992, 'Valuing Public Goods: The Purchase of Moral Satisfaction', *J. Environm. Econom. Managem.* **22**, 57–70.
Kopp, R. J.: 1992, 'Why Existence Value Should Be Used in Cost-Benefit Analysis', *J. Policy Analys. Managem.* **11(1)**, 123–130.
Krutilla, J. V.: 1967, 'Conservation Reconsidered', *Amer. Econom. Rev.* **57**, 787–796.
Loomis, J., Lockwood, M., and Delacy, T.: 'Some Empirical Evidence on Embedding Effects in Contingent Valuation of Forest Protection', *J. Environm. Econom. Managem.* **24**, 45–55.
MacCracken, M. C., Budyko, M. I., Hecht, A. D., and Israel, Y. A.: 1990, *Prospects for Future Climate*, Lewis Publisher, Chelsea, Michigan.
Milgrom, P.: 1993, 'Is Sympathy an Economic Value? Philosophy, Economics, and the Contingent Valuation Method', in Hausman, J. A. (ed.), *Contingent Valuation: A Critical Assessment*, North Holland, New York, pp. 415–435.
Michan, E. J.: 1976, *Cost-Benefit Analysis* (second edition), Praeger, New York.
Mitchell, R. C. and Carson, R. T.: 1989, *Using Surveys to Value Public Goods: The Contingent Valuation Method*, Resources for the Future, Washington, D.C.
Pearce, D. W.: 1993, *Economic Values and the Natural World*, MIT Press, Cambridge, MA.
Rosenthal, D. H. and Nelson, R. H.: 1992, 'Why Existence Value Should Not Be Used in Cost-Benefit Analysis', *J. Policy analys. Managem.* **11(1)**, 116–122.
Rossi, P. H. and Berk, R. A.: 1985, 'Varieties of Normative Consensus', *Amer. Sociolog. Rev.* **50(3)**, 333–346.
Rossi, P. H. and Nock, S. L.: 1982, *Measuring Social Judgements*, Sage Publications, Beverly Hills.
Russell, R. R. and Wilkinson, M.: 1979, *Microeconomics: A Synthesis of Modern and Neoclassical Theory*, John Wiley, New York.
Sargoff, R. M.: 1988, *The Economy of Earth*, Cambridge University Press, Cambridge.
Schkade, D. A. and Payne, J. W.: 1993, 'Where Do the Numbers Come from? How People Respond to Contingent Valuation Questions', in Hausman, J. A. (ed.), *Contingent Valuation: A Critical Assessment*, North Holland, New York, pp. 271–293.
Schkade, D. A. and Payne, J. W.: 1994, 'How People Respond to Contingent Valuation Questions: A Verbal Protocol Analysis of Willingness to Pay for an Environmental Regulation', *J. Environm. Econom. Managem.* **26**, 88–109.
Stern, P. C., Dietz, T., and Kalof, L.: 1993, 'Value Orientation, Gender, and Environmental Concern', *Environm. Behav.* **25(3)**, 322–348.
Stoker, R. M.: 1993, 'Empirical Approaches to the Problem of Aggregation over Individuals', *J. Econom. Literat.* **31** (December), 1827–1874.
Wong, G. Y. and Mason, W. M.: 1985, 'The Hierarchical Logistic Regression Model for Multilevel Data', *J. Amer. Statist. Assoc.* **80(391)**, 512–524.

(Received 22 November, 1993; in revised form 13 June, 1994)