PREFERENCE LEARNING AND CONTINGENT VALUATION METHODS

Thomas D. Crocker Jason F. Shogren

1991
Book Chapter WWRC-91-19

Chapter 5

in

Environmental Policy and the Economy

F. Dietz, F. van der Ploeg and
J. van der Straaten (Editors)
Elsevier Science Publishers B.V. 1991

Environmental Policy and the Economy
F. Dietz, F. van der Ploeg and J. van der Straaten (Editors)

5. PREFERENCE LEARNING AND CONTINGENT VALUATION METHODS*

Thomas D. Crocker and Jason F. Shogren

1 Introduction

Standard rhetoric in the guidelines for implementing the contingent valuation method (CVM) insists that respondents be "familiar" with the decision problem that the investigator poses.[1] Familiarity does not breed contempt; it improves the reliability of responses. Reliability is properly said to be enhanced by investigator clarification of what is to be valued (the contingent commodity) and the means by which it is to be valued (the exchange medium). In spite of the intuitive appeal of clarification, only Hoehn and Randall (1987) have investigated its analytical implications for *CVM* survey designs. They refer to it as the value formulation problem.

Hoehn and Randall (1987) suggest that the value formulation problem originates in "the time and resource constraints of the *CVM* context..." (p. 229). Repetition and review constraints inhibit the investigator's abilities to communicate complex commodity attributes, thus making what is being valued appear ambiguous to respondents. Moreover, even in the absence of ambiguities about what is being valued, the respondent's time and decision resources inhibit his value formulation process. Hoehn and Randall (1987) model value formulation by comparing the results obtained with a standard consumer optimisation problem to those produced when elements of the posited constraint system are ambiguous or when the respondent must engage in a "time-constrained search and decision process" (p. 231) that results in

*Comments received from an anonymous referee and in presentations at the 1990 meetings of the Southern Regional Science Association, Weber State University, and Tilburg University have been helpful.

[1]See Cummings et al. (1986) and Mitchell and Carson (1989), for example. "Realism" and "scenario plausibility" are sometimes substituted for "familiarity." The semantics of these guidelines are confusing. Realism is a necessary but not a sufficient condition for a *CVM* respondent to be familiar (experienced) with a trade-off.

incomplete optimisation. [2] They demonstrate that both imperfect communication and incomplete optimisation cause the Hicksian compensating measure of value to be weakly less than the same measure formulated under ideal situations. They also show that this compensating measure is non-decreasing with time or with other resources devoted to respondent decision-making.

Throughout their demonstration, Hoehn and Randall (1987, p. 230) presume "...that the respondent ... knows his/her initial level of well being." In short, they assume that the respondent's utility function is static and invariant and that he suffers no doubts about the utility associated with any particular lottery or outcome. They thus follow Stigler and Becker (1977) in proposing that it is more useful and valid to treat individual preferences as constant and to seek a constraint system explanation for any observed changes in commodity demands. Such changes can arise only from shifts in the shadow prices of the individual's resources or changes in household technologies. They cannot arise from the individual's lack of knowledge about how a particular commodity enters his utility function or from changes in his preference ordering (March, 1978; Cohen and Axelrod, 1984).

In spite of the *CVM* guidelines on familiarity, numerous applications involve nonmarketed environmental commodities with which respondents are unfamiliar. Frequently, respondents plausibly have had little or no actual experience with the commodity, e.g. atmospheric visibility in the Grand Canyon or acid mine drainage in Colorado mountain streams. In other cases, respondents may daily experience the commodity, e.g. trace metal exposures. However, they may view efforts to influence these experiences as futile and have therefore devoted little effort to understanding how the commodity affects their well-being. In order to forge values for such commodities, they must first explore their preference orderings.

In the next section, we construct a model in which an individual who is unconstrained by time and decision resources must form conjectures and accumulate experience about the effect that a well-specified and clearly communicated commodity has on his well-being. We presume that the individual does not know everything about his preferences, not that he has changing

[2]Imperfect communication implies misperception of his opportunity set such that for given time and decision resources, the respondent may select a point within rather than on the boundary of the set. Limited time and decision resources or bounded rationality resulting from the inability to store, retrieve, and process information imply that more such resources would cause some part or all of the opportunity set boundary to shift outward. See Colantoni et al. (1976) for further discussion of the welfare implications of the distinction. The impacts of imperfect communication and bounded rationality would differ with the questionnaire structure, the survey method (e.g., face-to-face, telephone, or mail), and other factors.

tastes. We develop two propositions: (1) if one does not know his preference for a commodity, his willingness-to-pay for a given quantity will be weakly greater than when he has better preference knowledge; and (2) an initial lack of knowledge about one's preferences causes willingness-to-pay to be non-increasing with time or with the application of other decision-making resources. We also argue that although the sequence of beliefs about own-preferences converges to a limit belief, incomplete preference learning may nevertheless prove optimal. A third section reports empirical tests of our two basic propositions. The results from a series of experiments involving a pair of hypothetical and nonhypothetical markets do not refute either proposition. A concluding section summarises our results, and discusses their implications for the reliability of *CVM*.

2 Preference Learning

Intuitively, the frequency with which one hears the phrase "I don't know whether I'll like this or not" makes evident the reality of the preference learning phenomenon. We follow a Bayesian approach to describe the process of learning one's utility function in a multiperiod world. Initially, the individual has some expectation of the utility he will derive from a particular commodity but he does not know the exact utility that he will receive, not because of any properties of the commodity but because he has not yet established the intensity of his preference for it. Each round of trading constitutes a test of his conjectures about his preference for the commodity. Below, we demonstrate that this individual need not initially choose the commodity amount that will maximize his first-period utility. Instead, he may choose an amount that initially has smaller expected marginal utility from consumption but which yields utility information that will enable him to attain a greater utility level in subsequent periods. As he invests more money, time, or effort and his consumption horizon unfolds, his uncertainty about his utility function declines. Moreover, the greater this initial investment, the more quickly can he determine this utility function because he can have greater confidence that the utility of his sample represents true utility. In short, he has both a consumption and an investment problem. He deliberately invests in conjectured suboptimal commodity bundles in order to ascertain whether his currently conjectured preferences are true or in error. He is therefore engaging in a form of self-protection.[3]

[3]The trade-off between current information and present consumption has also been examined in the context of sequential experimentation in statistics. See Berry and Fristedt (1985), and Easley and Keifer (1989) on the bandit problem, for example.

Consider an individual who is concerned only with a single period and who does not know his utility function, U. He must allocate his total wealth, M, among n nonstorable commodities, $x_1, ..., x_n$, subject to the wealth constraint $\sum_{i=1}^{n}, p_i x_i = M$, where price, p_i, and M are predetermined. Thus his instantaneous expected utility can be written as $V(x_1, ..., x_n) = E[U(x_1, ..., x_n)]$, where E is the expectations operator calculated with respect to the individual's prior distribution for U. Though the demand functions associated with V may embody option values, consumption is the singular focus of this problem.

Following Cyert and De Groot (1975), now presume that this individual who does not know his utility function must allocate his wealth each period among the nonstorable $x_1, ..., x_n$ over a number of periods. He must exhaust his wealth in each period. The values of $x_1, ..., x_n$ that solve his instantaneous allocation problem could result in any one of a variety of realised utilities in subsequent periods. He learns something about his likes and dislikes from these realisations and is able to control what he learns by his choices of the x. Each $x_1, ..., x_n$ causes formation of a posterior distribution of utilities. The individual calculates the discounted expected utility, W^*, associated with this posterior. Alternative values of $x_1, ..., x_n$ lead to different posteriors and different expected utility outcomes. Thus the individual's prior distribution for his utility function, along with the different values of $x_1, ..., x_n$ that he could choose, induce a set of posteriors, each member of which has an associated expected utility, W^*. These expected utilities are themselves distributed according to the probability function, Ω. By behaving as a Bayesian statistician, the individual can calculate the discounted expected utility function $W(x_1, ..., x_n) = E[W|\Omega(x_1, ..., x_n)]$. He uses his experiences to update his priors and thus chooses an allocation $x_1, ..., x_n$ for which $V(\cdot) + W(\cdot)$ is maximized subject to the wealth constraint.

We develop the key implications of this structure by using simplifying devices found in the dependent learning versions of the quasi-option value literature (Miller and Lad, 1984; Freeman, 1984). Assume that the individual can choose between two goods, x_1 and x_2. He knows the utility he will derive from x_1 but does not know how, if at all, x_2 will offer utility. For example, as in Rowe et al. (1980), x_1 might be electricity and x_2 might be rarely experienced, visibility-impairing levels of air pollution. For any particular value of x_1, the information the individual amasses about his utility function is a nondecreasing concave function of the amount of x_2 that he consumes in the first period. This information allows him to increase his expected utility, $W(\cdot)$, in future periods. It is crucial to assume that $W(\cdot)$ is nondecreasing in

x_2 since it implies that more valuable information is generated by greater consumption of x_2. In the sense of Blackwell (1951), one can reach a higher level of expected utility if beliefs are formed from a more informative experiment. See theorems 1 and 3, lemmas 1 and 2 in Grossman et al. (1977) for proof that $W(\cdot)$ is nondecreasing in x_2.

Given the two goods, choose the appropriate pecuniary measurement units so that $M = p_1 = p_2 = 1$. Then $d_1 = x_1$ and $d_2 = 1 - x_1$ can represent the proportions of M spent on each good. In a single period setting, the individual will choose x_1 and x_2 such that $\partial V/\partial x_1 = \partial V/\partial x_2$. Let $(\hat{d}_1, \hat{d}_2)$ be this optimal solution to the one period problem. It then follows from the concavity of the utility function that

$$\left.\begin{array}{l} d_1 < \hat{d}_1 \\ d_2 > \hat{d}_2 \end{array}\right\} \quad \text{if } \partial V/\partial x_1 > \partial V/\partial x_2. \tag{1}$$

and

$$\left.\begin{array}{l} d_1 > \hat{d}_1 \\ d_2 < \hat{d}_2 \end{array}\right\} \quad \text{if } \partial V/\partial x_1 < \partial V/\partial x_2. \tag{2}$$

In a multiperiod problem, the individual will choose x_1 and x_2 in the first period such that $\partial V/\partial x_1 + \partial W/\partial x_1 = \partial V/\partial x_2 + \partial W/\partial x_2$. Let $(\overline{d}_1, \overline{d}_2)$ be this optimal solution to the multiperiod problem. By previous assumption, $\partial W/\partial x_1 = 0$, and $\partial W/\partial x_2 > 0$, which implies that $\partial V/\partial x_1 > \partial V/\partial x_2$ if $(\overline{d}_1, \overline{d}_2)$ is to be attained. It then follows that $\overline{d}_1 < \hat{d}_1$ and $\overline{d}_2 > \hat{d}_2$. This result can be summarised in the following proposition.

Proposition 1: *In the first period of a multiperiod problem with dependent learning of his utility function, the individual will be willing to pay relatively less for the commodity with the unambiguous utility impact and relatively more for the commodity with the ambiguous utility impact than he would in a single period problem or in a problem where he knows his utility function.*

Because $\overline{d}_1 < \hat{d}_1$ and $\overline{d}_2 > \hat{d}_2$ in the first period of the multiperiod problem, it follows that $V(\hat{d}_1, \hat{d}_2) - V(\overline{d}_1, \overline{d}_2) > 0$; that is, the individual sacrifices consumption utility in the first period in order to acquire information about the utility that the ambiguous good might offer in future periods. The first period utility sacrifice measures the value of this information. Given his limited income, the more the individual spends on the ambiguous good in the first period, the greater the chance that he will suffer a shortfall in that period in his realised utility level for the bundle of goods. However, he must trade off this potential loss against the opportunity to learn more about the future potential of the ambiguous good.[4]

In a multiperiod setting, the greater the discrepancy between the expected and the realised utility of a good in any period, the greater the value of any information garnered by consuming the good in that period. If the individual's accumulated information about his likes and dislikes is nondecaying and if its marginal value is decreasing, then $V(\hat{d}_1, \hat{d}_2) - V(\overline{d}_1, \overline{d}_2)$ must approach zero from above as time passes. The individual's expected and realised utilities will then coincide and he will no longer have an incentive to consume in order to acquire utility information.

However, and more formally, define learning as a convergence of the sequence of beliefs $\{\Omega_t\}, t = 1, ..., m$, to a distribution which is point mass at true values (Easley and Keifer, 1988). When limit beliefs are point mass, then consumption of the ambiguous commodity generates no new information about its ability to generate utility. To see how beliefs converge, note that Bayes rule implies that a stochastic process of beliefs is a Martingale

$$E[\Omega_{t+1}|\Omega_t] = \Omega_t. \tag{3}$$

A Martingale implies that beliefs are not expected to evolve in any predictable way given new information. By the Martingale convergence theorem, beliefs converge to a limit belief Ω_∞ such that

$$Pr\{\lim_{t\to\infty} \Omega_t - \Omega_\infty\} = 1 \tag{4}$$

and

$$\lim_{t\to\infty} E[\Omega_t - \Omega_\infty] = 0. \tag{5}$$

[4] An individual who is practicing self-control in the sense of Thaler and Shefrin (1981) would reverse these results. Self-control theory allows the individual to have two sets of conflicting preferences simultaneously, those of a myopic doer and of a long-term planner. In our notation, the doer's utility function is $V(\cdot)$ and the planner's is $W(\cdot)$. If the utility effects of *not* having the commodity are ambiguous, increased consumption will provide a less informative experiment. The following corollary can therefore be stated. **Corollary 1:** Given a multiperiod problem, dependent learning, and a self-imposed restriction on its consumption, the individual will be willing to pay less for a good with an ambiguous utility impact than he would in a single period problem or where he knows his utility function.

A parallel exists between this corollary and results in the quasi-option value literature on development versus preservation. As Fisher and Haneman (1987) show, if learning is dependent on development, then more development is warranted than in the absence of dependent learning. However, if the relevant information is the value of preservation (i.e. a self-imposed restriction on consumption), then efficiency requires less development than otherwise.

Easley and Keifer (1988, 1989) demonstrate that although beliefs almost surely converge to a limit distribution, complete learning may or may not be optimal. If the sequence of consumption decisions converges rapidly, insufficient information may be generated to identify the uncertain parameter. With respect to value formation, this implies

$$\lim_{t \to \infty} \frac{\partial W(\cdot|\Omega_t)}{\partial x} = K \geq 0, \tag{6}$$

where $K > 0$ implies incomplete preference learning and $K = 0$ represents complete learning. Therefore the consumption of a commodity may be insufficient to approach the level, $\hat{x}$, where its potential utility contributions are fully known, that is,

$$\lim_{t \to \infty} x(\Omega_t) - \hat{x} = \epsilon \geq 0. \tag{7}$$

These results can be summarised in a second proposition.

Proposition 2: *Given a Bayesian updating process, then the individual's consumption of a commodity with an ambiguous utility contribution* will *converge from above toward the level of consumption that would occur in a single period problem or in a problem where he knows his utility function. Note, however, that optimal preference learning may be incomplete.*

Keifer (1989) simulates a case in which beliefs converge away from the true state and remain stable at incomplete learning. Clearly, if optimal learning about one's utility function is incomplete, then value formation will never be complete.

In succeeding sections, we report experiments designed to test the above two propositions. The next section describes the experimental designs and the fourth section reports results.

3 Experimental Design and Procedures

Experimental design

Our experiments focus on markets for reductions in risk. Risk combines probability and severity (Ehrlich and Becker, 1972). Markets for risk reduction can thus involve ex ante opportunities to reduce the likelihood of an undesirable event being realised (protection) or ex ante opportunities to reduce the severity of a realised undesirable event (insurance). In order to reduce the chance that our empirical results are unique to a single commodity, we apply our experimental design to a self-protection asset and a self-insurance asset and to four levels of risk. The design of the markets for these two assets and four risk levels is described below. The assets were rivalrous.

a) Every market had 30 participants who were assigned in groups of six to five distinct replications of the market. Each participant was identifiable only by a randomly assigned number.[5] Prior to the opening of a market, a perfectly fungible endowment, M, was granted each participant. $M = \$10$ was identical among participants and between markets and their replications.

b) Each market had four binary lottery levels $(\pi, \$L; 1-\pi, \$G)$, where $\pi(0 \leq \pi \leq 1)$ is the probability of a monetary loss, $\$ L$, and $1-\pi$ is the probability of a monetary gain, $\$ G$. The lottery levels were identical between asset markets, as was $L = -\$4$ and $G = +\$1$.

c) Participants confronted 40 per cent, 20 per cent, 10 per cent, and 1 per cent probabilities of suffering wealth losses in the order 20 per cent, 10 per cent, 1 per cent and 40 per cent.

d) In every replication of each market and for every loss probability, each individual knew that he was to make 12 bids to reduce risk to zero. His first and last bids were hypothetical (inexperienced and experienced). All other bids were nonhypothetical. Each participant therefore made eight hypothetical bids and 40 nonhypothetical bids for each risk reduction asset.

e) In the self-protection market, participants bid upon an asset unit that would reduce the probability, π, of a loss to zero. In the self-insurance market, they bid upon an asset unit that would reduce the severity, $\$L$, of a loss to zero.

f) Self-protection and self-insurance nonhypothetical bids were made in Vickrey (1961) sealed bid, second price auctions.[6] Thus each subject competed in each round against the other five participants for the purchase of a protection or insurance asset that reduced risk to zero. The winner in each round was the subject with the highest bid. This winner had to pay the second highest bid. Both the randomly assigned number of the winner and the amount of the second highest bid were posted as public information at the end of each bidding round. The sealed bids in each round were submitted at

[5]The participants were undergraduate students at Appalachian State University, Boone, NC. Bennett (1987) found undergraduate student responses in experimental settings to be statistically no different from those of the general population.

[6]The Vickrey (1961) auction has well-known demand-revealing properties. The participant's dominant strategy is to reveal his full willingness-to-pay since he does not have to pay what he bid.

the same time.

g) At the end of a round, a single draw was made from an opaque urn containing 100π red chips and $(1-100\pi)$ white chips. A red chip represented a \$4 loss for all participants except for the bidder who won the round, and a white chip provided a \$1 gain to all participants, including the winning bidder. The drawn chip was replaced before the next round. Participants pocketed their gains and losses at the end of each round.

h) In each succeeding nonhypothetical round, the \$10 endowment of all participants was restored, thereby making capital gains and losses independent across the nonhypothetical rounds (see McKee, 1989). Protection and insurance purchases did not carry over from round to round.

Procedures

No participants had ever played a role in an experimental risk reduction market. No side payments were allowed among participants. Except for the posted second highest bid and the number of the winning bidder, communication among participants was forbidden. Altruistic motives were thus suppressed. Given these restrictions and demonstrations by Owens (1982) and others that collusive activities without sanctions cannot be long sustained for groups with more than four members, we believe that participants' bids within each round are independent.

Prior to the first bidding round, participants read at least once the instructions for their roles in the experiments. A monitor then read those same instructions to them once. Participants were then allowed to ask any questions they wished about their roles. No time constraints were imposed. When all questions were satisfied, the experiment commenced.

Initially, the monitor asked each participant to write down, seal, and submit hypothetical bids to reduce to zero each loss probability. With these initial hypothetical bids, participant statements of self-protection or self-insurance purchase intentions were not binding upon him, real money did not change hands, and the risk reduction did not occur. Because participants were under no time pressure to submit these initial hypothetical bids, the binding time and resource decision constraints that concern Hoehn and Randall (1987) were plausibly weak. However, participants may try harder to learn their preferences when payment depends on the unbiasedness and the precision of the estimates on which they base their decision (Thaler, 1987). The interpretation that we subsequently offer of the implications for *CVM* of our experimental results requires that Thaler's (1987) point be set aside. In particular, in the absence of dissimilar preference learning opportunities and time and decision resource constraints, we maintain the hypothesis of the identity of the hypothetical bids of *CVM* and actual market bids.

Next, the monitor stated that a 20 per cent chance of a red chip being drawn existed and that participants would bid in each of ten rounds to reduce this risk to zero. Participants were again told that their bids would be binding and would thus influence their take-home pay. They were also reminded again that gains, losses, and purchases of protection or insurance were not transferable across rounds. In each round, after participants' bids were collected, the lottery was immediately resolved. Participants averaged 30 seconds to a minute to make their bids in each round.

Third, after the ten rounds of nonhypothetical bids were completed, a final experienced hypothetical bid was solicited. Participants had unlimited time to submit this bid and they were informed that this final bid was not binding. They took about one minute to submit this bid.

Next, the monitor informed participants of the 10 per cent chance of a red chip being drawn. He then elicited in the same fashion as with the 20 per cent chance the initial hypothetical bid, the ten nonhypothetical bids, and the final hypothetical bid. This process was then repeated for the 1 per cent chance of a red chip and finally, for the 40 per cent chance.

4 Experimental Results

Sixty individuals participated in the experiments. Five experimental sessions with six subjects each were run for each of the two asset markets. Table 1 summarises the results for all experiments. The first two columns describe the two risk reduction asset markets and the four probability levels of having a red chip drawn.[7] Reported in the table are the mean, its variance, and the median for the initial hypothetical bid, the nonhypothetical bids, and the final hypothetical bid. The hypothetical bid entries are calculated with respect to the one-time bids of 30 participants; the nonhypothetical entries represent the arithmetic mean bids over 10 bidding rounds of the same 30 participants.

When a one-tailed Wilcoxon matched sample sign test is applied at the 99 per cent confidence level, the hypothesis is rejected for each asset market and for all loss probabilities that the initial hypothetical bid (round 1) and the initial nonhypothetical bid (round 2) came from the same parental distribution. Figure 1, which charts the mean bids by bidding rounds dramatically

[7] Shogren (1990) uses similar data to conclude that the bids for self-protection were significantly greater statistically than were the bids for self-insurance.

Table 1: Summary statistics of bids in $.

Asset Market[a]	Loss Probability	Mean			Variance			Median		
		A	B	C	A	B	C	A	B	C
Self-protection	1%	2.73	0.78	0.81	11.72	0.94	2.08	1.50	0.38	0.35
	10%	2.87	1.09	1.13	5.42	1.42	4.19	3.00	2.98	1.38
	20%	3.35	2.93	3.45	5.40	1.49	3.80	3.08	3.36	3.50
	40%	4.62	3.93	4.37	7.45	2.66	4.57	4.00	3.70	4.00
Self-insurance	1%	1.85	0.09	0.07	10.35	0.02	2.01	0.50	0.07	0.03
	10%	2.93	1.09	1.13	6.23	1.10	0.92	2.28	0.79	0.86
	20%	3.93	2.56	2.44	5.26	2.69	2.59	4.00	2.16	2.25
	40%	4.91	3.35	3.33	5.32	1.81	1.87	5.00	3.31	3.58

A = Initial hypothetical bid
B = Arithmetic mean nonhypothetical bid, 10 bidding rounds
C = Final hypothetical bid

[a] n = 30 for each asset market: 5 experiments with 6 participants each.

NOTE: For all entries, we use a one-tailed test to reject at the .01 level the null hypothesis that the population mean is zero.

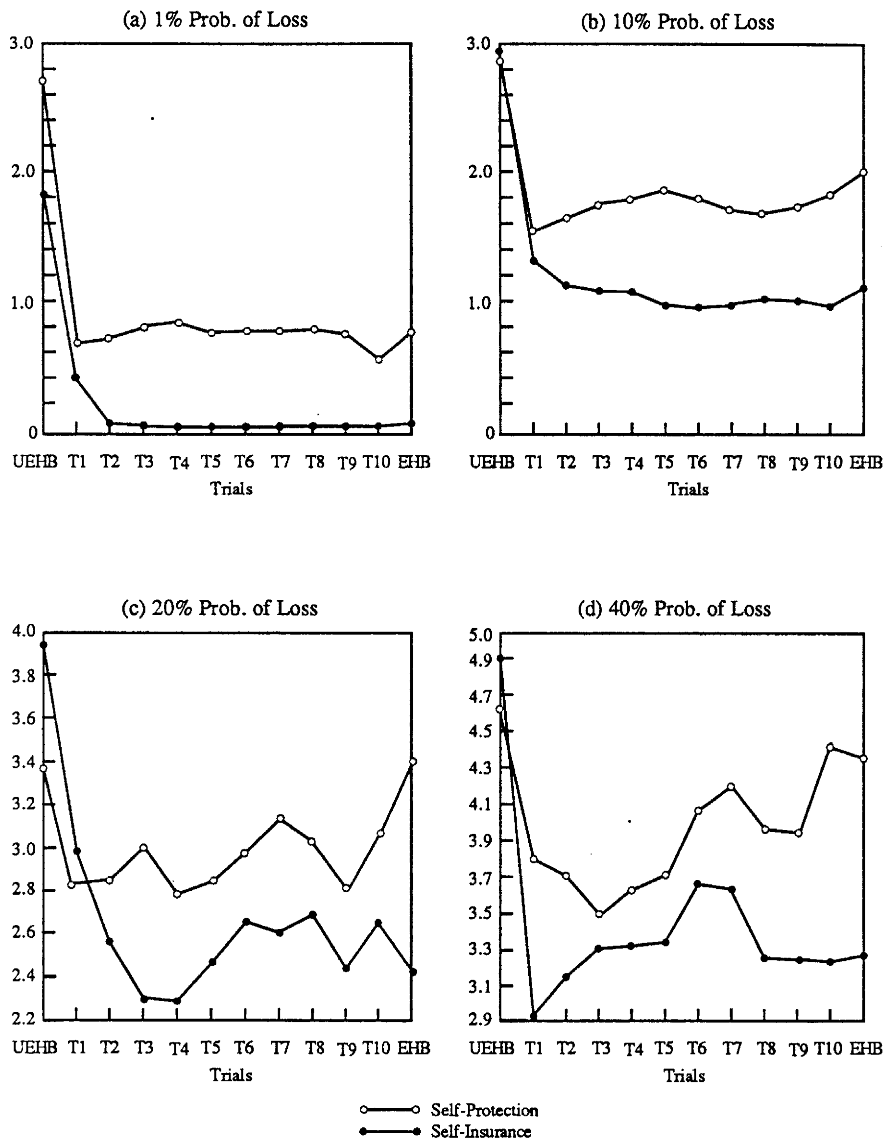

Figure 1: Self-protection and self-insurance bids (mean bids by rounds).

confirms this statistical result. The figure also shows that the round 2 mean bid is always less than the round 1 mean bid. Together, figure 1 and table 1 support our two theoretical propositions about the structure that preference learning will impose upon a multiperiod exchange activity. For each asset market and for all loss probabilities, initial bids are higher than and approach immediately subsequent bids from above.

Note, however, that table 1 and figure 1 display at least three other prominent patterns. First, with the sole exception of the 40 per cent loss probability in the self-insurance market, the relative discrepancies between the round 1 and the round 2 bids are consistently greater with the one and the 10 per cent loss probabilities than with the 20 and the 40 percent loss probabilities. Within the context of our model, this implies that participants knew relatively less about how small risks than how large risks affected their well-being. The value of consuming to learn about one's preferences was greater for the low loss probabilities than for the high loss probabilities.

Second, only the final hypothetical bid in the 20 per cent loss probability self-protection market was higher than the initial hypothetical bid in the same market. This implies that the value of preference learning declines with experience in consumption. Finally, after the decline in bid magnitudes from round 1 to rounds 2 or 3, bids usually increased or stayed constant over the next ten bidding rounds. Our model provides no insight into this third phenomenon. Nevertheless, Hoehn and Randall's (1987) approach to value formulation does offer an explanation. Suppose that the minute or less that participants used to make a nonhypothetical bid in each round somehow imposed a binding time or decision resource constraint upon them. For example, once a single participant had bid, a desire not to be viewed as dull and a laggard relative to one's peers may have operated among those who had not yet bid. If so, the bidding sequences observed in rounds 3 through 12 are broadly consistent with the propositions of Hoehn and Randall (1987).

5 Summary and Conclusions

What do value estimates established with the *CVM* mean? The model of Hoehn and Randall (1987) implies that they are weakly less than the true measure; our model implies that they are weakly greater than this measure. The empirical results reported above support both models, though the support appears at different points in the bidding sequences.

In the field settings where *CVM* surveys are performed, both the familiarity that respondents have with the good being valued and the severity of the time and decision resource constraints under which they have to act are difficult to specify. If the preference learning phenomenon dominates these settings, the amounts that individuals are willing to pay to acquire knowledge about the potential utility value of a commodity should not be mistaken for measures of the consumption value of that commodity. In our self-protection and self-insurance experiments, these knowledge values were about equal to the initial, nonhypothetical willingnesses-to-pay. Knowledge values that are this large raise serious questions about the exaggerations introduced into consumption value estimates generated by *CVM* surveys where opportunities for learning one's preferences are absent.

If the time and decision resource constraint phenomenon dominates, the amounts that individuals are willing to pay under incomplete optimisation should not be mistaken for measures of the consumption value of the commodity. There is no reason whatsoever to think that the downward biases induced by the time and decision resource constraints of the standard *CVM* format bear any resemblance to the same constraints that a respondent might confront in a nonhypothetical setting.

Most importantly for the intellectual validity and policy reliability of *CVM*, it is not at all obvious how one would determine whether and when the preference learning phenomenon or the incomplete optimisation phenomenon dominate for any particular respondent.[8] In the absence of this knowledge, one cannot know whether the individual respondent's estimates involve an upward or a downward bias. The choice appears arbitrary. If so, one cannot know what values of nonmarketed commodities are illuminated by *CVM* surveys that do not provide plentiful opportunities for preference learning and the relaxation of time and decision resource constraints.

The singular focus upon incomplete optimisation in the *CVM* value formulation problem is either a maintained hypothesis for a research agenda or a canon of interpretation. If the research agenda requires that *CVM* analysis presume that respondents know their preferences, then it is untenable in the light of our experimental results. It is also untenable in the light of introspective observations about one's own knowledge of the consumption value that he might obtain from nonmarketed goods which he has never or only rarely experienced. It would then be more plausible to treat the focus as a rule for interpreting *CVM* estimates, directing policymakers to regard differences in these estimates as reflecting differences in the completeness of

[8]Recall that we took the identity of hypothetical and actual bids as a maintained hypothesis. If the identity does not hold, our preference learning concerns and the time and decision resource constraint concerns of Hoehn and Randall (1987) are immaterial. Given that market bids are the baseline measure of value, a failure of our maintained hypothesis would invalidate *CVM*.

respondents' optimisations rather than as differences in own-preference knowledge. Especially since the directions of bias from incomplete optimisation and from preference learning are quite different, we prefer a neutral rule of interpretation which says that our experimental results and, by extension, *CVM* survey estimates are consistent with both. Unfortunately, this interpretation implies that *CVM* estimates are themselves at this time uninterpretable. Although beyond the scope of this paper, this suggests the need for criteria that define the appropriate combination of incomplete optimisation and preference knowledge for the valuation of a specific commodity. If such criteria can be defined, we will then know how to compare an individual's valuations when his situation, time and decision resource constraints, and his preference knowledge change simultaneously. Otherwise, *CVM* value estimates are inherently nonrebuttable.[9]

However, the above remarks may be too hard on *CVM*. Actual market behaviour also plausibly reflects preference learning and time and decision resource constraints. The observed market values that serve as benchmarks for evaluations of *CVM* reliability can therefore shift. Perhaps the notion of economic value itself needs refinement such that the disequilibrium process of value formulation is allowed to have a coherent role.

References

Bennett, J., (1987), Strategic Behavior, *Journal of Public Economics,* **32**, 355-368.

Berry, D. and B. Fristedt, (1985), *Bandit Problems: Sequential Allocation of Experiments*, London: Chapman and Hall.

Blackwell, D., (1951), The Comparison of Experiments, *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley: University of California Press, 93-102.

Cohen, M.D., and R. Axelrod, (1984), Coping with Complexity, *The American Economic Review,* **74**, 30-42.

Colantoni, C.S., O.A. Davis and M. Swaminuthan, (1976), Imperfect Consumers and Welfare Comparisons of Policies Concerning Information and Regulation, *Bell Journal of Economics and Management Science,* **7**, 602-615.

Cummings, R.G., D.S. Brookshire and W.D. Schulze, (1986), *Valuing Environmental Goods: An Assessment of the Contingent Valuation Method*, Totawa, New Jersey: Bowman and Allanheld.

Cyert, R.M. and M.H. de Groot, (1975), Adaptive Utility, R.H. Day and T. Groves (eds.), *Adaptive Economic Models*, New York: Academic Press, 223-246.

[9]The reliability of *CVM* is taken as a rebuttable presumption in U.S. environmental litigation brought under the Comprehensive Environmental Response, Compensation and Liability Act of 1980 (CERCLA). See Phillips and Zeckhauser (1989) for another expression of strong doubts about the reliability of *CVM*, especially when applied to compensation claims in CERCLA-type litigation.

Easley, D. and N. Keifer, (1988), Controlling a Stochastic Process with Unknown Parameters, *Econometrica,* **56**, 1045-1064.

Easley, D. and N. Keifer, (1989), Optimal Learning with Endogenous Data, *International Economic Review,* **30**, 963-978.

Ehrlich, I. and G.S. Becker, (1972), Market Insurance, Self-Insurance, and Self-Protection, *Journal of Political Economy,* **80**, 623-648.

Fisher, A. and M. Haneman, (1987), Quasi-Option Value: Same Misconceptions Dispelled, *Journal of Environmental Economics and Management,* **14**, 183-190.

Freeman, A.M. III, (1984), The Quasi-Option Value of Irreversible Development, *Journal of Environmental Economics and Management,* **11**, 292-295.

Grossman, S., R. Kihlstrom and L. Mirman, (1977), A Bayesian Approach to the Production of Information and Learning by Doing, *Review of Economic Studies,* **46**, 533-547.

Hoehn, J.P. and A. Randall, (1987), A Satisfactory Benefit Cost Indicator for Contingent Valuation, *Journal of Environmental Economics and Management,* **14**, 226-247.

Keifer, N., (1989), A Value Function Arising in the Economics of Information, *Journal of Economic Dynamics and Control,* **13**, 201-223.

March, J.G., (1978), Bounded Rationality, Ambiguity, and the Engineering of Choice, *Bell Journal of Economics and Management Science,* **9**, 587-608.

McKee, M., (1989), Intra-Experimental Income Effects and Risk Aversion, *Economic Letters,* **30**, 109-115.

Miller, J.R. and F. Lad, (1984), Flexibility, Learning, and Irreversibility in Environmental Decisions: A Bayesian Approach, *Journal of Environmental Economics and Management,* **11**, 161-172.

Mitchell, R.C. and R.T. Carson, (1989), *Using Surveys to Value Public Goods: The Contingent Valuation Method*, Washington DC: Resources for the Future, Inc..

Owens, G., (1982), *Game Theory,* New York, NY: Academic Press.

Phillips, C.V., and R.J. Zeckhauser, (1989), Contingent Valuation of Damage to Natural Resources: How Accurate? How Appropriate?, *Toxics Law Reporter*, 520-529.

Rowe, R.D., R.C. d'Arge, and D.S. Brookshire, (1980), An Experiment on the Economic Value of Visibility, *Journal of Environmental Economics and Management,* **7**, 1-19.

Shogren, J.F., (1990), The Impact of Self-Protection and Self-Insurance on Individual Response to Risk, *Journal of Risk and Uncertainty,* **3**, 191-204.

Stigler, G.J. and G.S. Becker, (1977), De Gustibus Non Est Disputandum, *The American Economic Review,* **67**, 76-90.

Thaler, R., (1987), The Psychology of Choice and the Assumption of Economics, A.E. Roth (ed.), *Laboratory Experimentation in Economics*, New York, NY: Cambridge University Press, 99-130.

Thaler, R. and H. Shefrin, (1981), An Economic Theory of Self-Control, *Journal of Political Economy,* **89**, 389-406.

Vickrey, W., (1961), Counterspeculation, Auctions and Competitive Sealed Tenders, *Journal of Finance,* **16**, 8-37.