Comparing Predictions and Outcomes: Theory and Application to Income Changes [2]

Marcel Das*, Jeff Dominitz**, and Arthur van Soest*

* Tilburg University, Department of Econometrics and CentER
** Deloitte & Touche, Dispute Consulting Services

June 1998

Abstract

Household surveys often elicit respondents' intentions or predictions of future outcomes. The survey questions may ask respondents to choose among a selection of (ordered) response categories. If panel data or repeated cross-sections are available, predictions may be compared with realized outcomes. The categorical nature of the predictions data, however, complicates this comparison. Generalizing previous findings on binary intentions data, we derive bounds on features of the empirical distribution of realized outcomes under the "best-case" hypothesis that respondents form rational expectations and that reported expectations are best predictions of future outcomes. These bounds are shown to depend on the assumed model of how respondents form their "best prediction" when forced to choose among (ordered) categories. An application to data on income change expectations and realizations illustrates how alternative response models may be used to test the best-case hypothesis.

---

1 Introduction

Subjective data on respondents' intentions or predictions are commonly used for many purposes. They are routinely used in psychology, sociology, and political science. Economists, however, tend to be skeptical of subjective data, in general, and expectations data, in particular. This skepticism may be traced to a history of negative findings on the predictive value of the data (see, for example, Tobin, 1959). More recently, it has been claimed that expectations data need not match up to future outcomes because respondents have no incentive to report expectations accurately (see, for example, Keane and Runkle, 1990).

Some examples in the recent literature suggest that this attitude is changing. Dominitz and Manski (1996, 1997a) analyze long-term income expectations of students and near-term income expectations of U.S. households. Das and Van Soest (1996, 1997) analyze income change expectations of Dutch households. Guiso et al. (1992, 1996) use expectations data to construct a measure of subjective income uncertainty which is included in models of saving and portfolio choice. In the literature on labor supply, data on desired hours of work have been used to disentangle preferences and hours restrictions (Ilmakunnas and Pudney, 1990).

If panel data or repeated cross-sections are available, data on expectations of prospective outcomes may be compared with data on realized outcomes. When qualitative rather than quantitative expectations data are to be analyzed, these comparisons may not be straightforward. Manski (1990) studied this problem for the case of a binary outcome. Under the "best-case" hypothesis that respondents have rational expectations and report best predictions of future outcomes, he showed that these expectations data bound but do not identify the probability of each possible outcome.

Say, for example, that households are asked whether or not they intend to buy a new car in the next twelve months. Given their information set, and their (subjective) distribution of relevant future variables, they will have some (subjective) probability of buying a car. A possible model for the answer to the intention question is: "yes", if this probability exceeds 0.5, and "no" otherwise. If, for some group of households, the subjective probability is 0.4, they will all answer "no". On the other hand, if the subjective distributions of the future variables are correct, and if the realizations of the future variables are independent, 40% will actually buy a car. The response rule that maps the probability onto a "yes" or "no" answer causes the discrepancy between the proportions of "yes" intentions and realizations. Although the subjective and actual distributions coincide, the intentions and realizations variables are not directly comparable.

"Yes/no" expectations about binary outcomes may be thought of as a special case of ordered-category expectations. In particular, they are 2-ordered-category expectations of a variable that takes on just two values (e.g., 0 and 1). We extend Manski's analysis to the general case of multiple-ordered-category expectations of a variable that takes on more than two values. Our empirical analysis focuses on expectations of a change in household income, which respondents report by choosing among five ordered categories.

Adopting Manski's "best-case" framework, we consider three models generating best predictions of the prospective outcome given the respondent's subjective probability distribution. Each model may be thought of as based on minimizing an expected loss function.

These behavioral models yield responses of (1) the modal category, (2) the category containing the median of the subjective distribution, or (3) the category containing the mean of the subjective distribution. For each case, we derive bounds on features of the distribution of realizations under the best-case hypothesis. In contrast to the case of yes/no expectations, different symmetric loss functions may yield different ordered-category survey responses and therefore imply different best-case bounds.

Our application focuses on income change expectations and outcomes reported in the 1984 – 1989 waves of the Dutch Socio-Economic Panel (SEP). Heads of household are asked whether they expect their income to decrease strongly, decrease, remain the same, increase, or increase strongly in the next twelve months. A similar categorical question is asked about the change in income over the past twelve months. In addition, we use a quantitative measure of income constructed from detailed data on income components of all household members.

In the majority of empirical life cycle models of consumption and savings, rational expectations of prospective income is taken for granted (see, for example, the survey of Browning and Lusardi, 1996). Our results suggest that in at least four out of the five years considered, the best-case scenario does not hold, and that, on average, people tend to underestimate future income. This finding suggests that either household expectations are not rational or macroeconomic shocks take place in a number of consecutive years or both. Persistent underestimation of household income leads to excess savings in a life-cycle model.

The outline of the paper is as follows. Section 2 discusses previous analyses of responses to ordered-category expectations questions. We present various models of survey respondent behavior and discuss the implications for previous findings in empirical research on expectations data. Section 3 derives best-case bounds on conditional probabilities of outcomes given reported expectations and the hypothesized model of survey response. In Section 4, the data of the empirical application are discussed. Expectations of income changes are compared with categorical and quantitative realizations in Sections 5 and 6, respectively. Section 7 concludes.

2   Background

Surveys regularly elicit respondent expectations of prospective economic outcomes. Such data may be collected in order to forecast realizations, to monitor public sentiment, or to increase our understanding of behaviors thought to be related to these expectations. We consider responses to qualitative expectations questions asking respondents to choose among ordered-response categories. While the number of responses may vary, this type of question is quite common. For example, the University of Michigan's Survey of Consumers asks a series of such questions, responses to which are used to generate the Index of Consumer Expectations (ICE) and the Index of Consumer Sentiment (ICS). These data are regularly analyzed to forecast realizations, monitor public sentiment, and estimate behavioral models. This section begins by discussing standard analyses of these data. Then, we present more rigorous approaches that explicitly consider how individuals may respond to such questions.

2.1 Standard approaches

A prototypical economic expectations question takes the following form:

> "Now looking ahead – do you think that a year from now you (and your family living there) will be better off financially, or worse off, or just about the same as now?"

Respondents to the Surveys of Consumers and the National Election Studies, both conducted by the University of Michigan, are asked this ordered-category expectations question. The responses are often aggregated in some way.

The ICE, for example, is based on responses to this question, in addition to two ordered-category expectations questions concerning national economic conditions. To construct the index, a "relative score" is calculated for each question based on the difference between the percentage of positive responses (e.g., "better") and the percentage of negative responses (e.g., "worse"). Each question receives equal weight in the index. Relative scores for two ordered-category retrospective assessment questions (e.g, "Are you better off or worse off than you were a year ago?") are averaged in with these values to construct the ICS. The Conference Board's Consumer Confidence Index is based on responses to a similar series of questions aggregated in a similar manner (Linden, 1982).

Despite the range of potential applications of economic expectations data, economists are generally skeptical of their use. Dominitz and Manski (1997a) trace this skepticism back to a scientific controversy that began in the 1940s and persisted until the 1960s. During that time period, when only a short time-series of observations were available, a number of studies found data of the type collected in the Michigan surveys to be of little or no predictive value in microeconomic analyses. A leading example of the firmly negative conclusions may be found in Tobin's (1959) study of the re-interview portion of the 1952-53 Survey of Consumer Finances. He estimated best linear predictors of household durable goods expenditures and of household savings given observed household attributes and a selection of variables derived from responses to ordered-category expectations questions – for example, an individual-level ICS, separate components of the ICS, and the following income change question:

> "How about a year from now – do you think you people will be making more money or less money than you are now, or what do you expect?"

Despite the negative conclusions, such questions continue to be included in University of Michigan Surveys, among others. Macroeconomists, especially those employed in the Federal Reserve System, continue to study the time-series relationship between these indexes and aggregate economic outcomes. See, for example, Carroll et al. (1994), Fuhrer (1988), Garner (1991), Otoo (1997) and Throop (1992). In addition, the ICE is a component of the Index of Leading Economic Indicators compiled by the U.S. Department of Commerce and reported in its monthly *Survey of Current Business.*

It is not just macroeconomists who use these data. A considerable amount of empirical political science research focuses on the relationship between responses to Michigan's ordered-category expectations questions and presidential voting behavior (or presidential

approval ratings). See, for example, the influential work of Kiewiet (1983) and a recent summary in Norpoth (1996). Typically, average responses are calculated and included in equations predicting presidential vote share or average approval rating. Two striking conclusions arise from this branch of research. First, the electorate is found to engage in retrospective rather than prospective voting; that is, responses to retrospective assessment questions are found to be of predictive value whereas reported expectations are not. For a notable challenge to this conclusion, see MacKuen et al. (1992). Second, the electorate is sensitive to national economic conditions rather than personal economic conditions; that is, responses to the (retrospective) national business conditions questions are of predictive value whereas responses to family finances questions are not.

2.2 Modelling ordered-category expectations

Dominitz and Manski (1997a, 1997b) discuss weaknesses of qualitative, ordered-category expectations questions but do not systematically assess the information content of such data. The purpose of their studies is to assess the merits of an alternative approach – elicitation of quantitative expectations in the form of subjective probabilities. Given the generally positive findings on probability elicitation reported there and elsewhere (e.g, Guiso et al., 1992, 1996; Hurd and McGarry, 1995), the general skepticism (among economists) of expectations data may be alleviated. Researchers may then be interested in utilizing other, more readily accessible, forms of expectations data, such as responses to ordered-category questions. It therefore seems worthwhile to present a rigorous approach to analysis of such data, making explicit the types of restrictive assumptions that must be made to appropriately interpret results. Recognition of these restrictions, as well as the inherently limited information content of responses, is crucial to understanding why previous findings on the predictive value of expectations data and related behavioral conclusions deserve to be revisited. In addition, this approach may be extended to apply to responses to any ordered-category survey question in which respondent uncertainty exists (e.g., retrospective reports and hypothetical choices).

Following on Manski's (1990) analysis of yes/no expectations, we consider respondents who attempt to report best predictions of future outcomes. In particular, we propose a model in which this best prediction is found by minimizing an expected loss function. This interpretation of ordered-category responses follows directly from Manski, but the framework is implicit in the work of Tobin (1959) and Juster (1966).

Influenced by, perhaps, the phrasing of the question, the respondent may adopt any of a variety of loss functions. If the respondent interprets the ordered-category question as one eliciting the most likely outcome, then we may assume he or she will report the category that contains the most subjective probability mass. This response rule arises when the loss function is an indicator function taking on the value 0 when the realization is in the predicted category and taking on the value of 1 otherwise.

The modal category response model appears sensible, but it is not often adopted in analyses of ordered-category expectations data. Instead, researchers typically assert that the respondent forms some point expectation and then chooses the category that contains this point expectation. In particular, respondents are often thought to interpret questions of what they "think" or "expect" to happen as questions eliciting the subjective mean of

the variable of interest. According to this model, the respondent behaves as if he or she is minimizing squared forecast errors. If, instead, the respondent were thought to report the category containing the subjective median, this model would correspond to minimizing absolute forecast errors.

We consider these three symmetric loss functions in the remainder of this paper. The survey response is invariant to the choice among symmetric loss functions in a special case of ordered-category expectations – yes/no expectations of binary outcomes (Manski, 1990). The three symmetric models, however, may yield varying responses when outcomes are not binary. Asymmetric loss functions cannot be ruled out. We therefore also consider the α-quantile assumption, corresponding to minimizing the asymmetrically weighted sum of absolute deviations.

2.3   Examples in expectations research

In an analysis of ordered-category purchase intentions data, Juster (1966) hypothesizes that respondents report best predictions of prospective outcomes. He then argues that subjective probability elicitation should yield more efficient predictors of subsequent realizations and presents empirical evidence to support this argument. Manski (1990) formalizes the argument to derive an upper bound on the information content of yes/no expectations data under the best-case hypothesis that respondents form rational expectations and report best predictions of future outcomes. If respondents behave as if they minimize a symmetric loss function, then "yes" is reported if the subjective probability of the event occurring is at least 0.5 and "no" is reported otherwise. In this case, the best case hypothesis predicts that at least half of all "yes" respondents will subsequently report that the event did occur whereas at most half of all "no" respondents will do so. Manski applies these bounds to study schooling-work expectations and realizations reported by respondents to the National Longitudinal Survey of the High School Class of 1972.

Carlson and Parkin (1975) study 3-ordered-category inflation expectations data. They adopt a model in which the respondent chooses one of the three ordered categories if that category contains at least 0.5 probability mass. Otherwise, *don't know* is reported. This study represents a rare instance in which *don't know* responses are modeled. It can be seen as a modification of both the modal and median response models. That is, in any ordered category case, if one category contains at least 0.5 probability mass, then it is both the modal category and the category that contains the subjective median. If no category satisfies this restriction, then some other response rule must be followed, such as report *don't know*.

Expectations data have often been used to test predictions of rational expectations models. For surveys of this literature, see Lovell (1986) and Maddala (1994). When ordered-category expectations data are studied, the researcher typically acts as if each respondent chooses the category that contains the subjective mean, and then the researcher attempts to quantify these qualitative responses (Maddala, 1994). Tests of unbiasedness are conducted by confronting the expectations data with subsequent realizations. For example, Nerlove (1983) confronts 3-ordered-category expectations with 3-ordered-category realizations reported by French and German firms. He chooses to "regard expectations and plans as single-valued but to recognize that the economic agent knows that they may

turn out to be wrong" (p. 1252). This example shows that the framework for this type of analysis is not always clearly specified in terms of stating (1) the feature of the subjective probability distribution that respondents are assumed to report and (2) the rational expectations implications of the assumed response model.

3   Outcome probabilities conditional on predictions

This section generalizes the framework in Manski (1990) and derives restrictions on the distribution of actual outcomes for given values of the subjective predictions in the best-case scenario. As the starting point for the analysis of responses to ordered-category expectations questions, consider a respondent who has a subjective probability density $f(y|s)$ over the support of prospective realizations of y given his or her current information captured in variables s. The expectations question asks the respondent to choose one category from K categories $C_1, ..., C_K$, which typically will be of the form $C_k = (m_{k-1}, m_k]$, with $-\infty = m_0 < m_1 < ... < m_{K-1} < m_K = \infty$. The threshold values m_k are typically not defined by the survey question; instead, they are subjectively determined (but not reported) by the respondent. The response to this question is denoted by p, where p is a best prediction in some well-defined sense.

We concentrate on the three different assumptions about the respondents' strategy for answering the subjective questions discussed in Section 2.2. The three assumptions refer to which feature of the subjective distribution is reflected by p_i, the prediction of respondent i ($p_i \in \{1, ..., K\}$). The difference between the three assumptions is illustrated in Figure 1. Section 3.1 presents the *modal category* assumption (in Figure 1, this leads to $p_i = 4$. Section 3.2 discusses the *α-quantile* assumption, which reduces to the *median category* assumption when $\alpha = 0.5$ (in Figure 1, $p_i = 3$). Section 3.3 presents the *mean* assumption ($p_i = 4$ in Figure 1).

Figure 1 : Graphical illustration of the three assumptions that refer to which feature of the subjective distribution is reflected by the prediction.

INSERT FIGURE 1 HERE

The observed prediction p_i is always a categorical variable. We distinguish, however, two cases for the realization. We either observe the exact realization (quantitative) y_i, or the (qualitative) category $c_i(\in \{1, ..., K\})$ in which y_i is contained: $c_i = k$ iff $y_i \in C_{k,i}$. If the threshold values are known, observing y_i clearly implies that c_i is also known. In the other case, to use the c_i, we will assume that they refer to the same categories as the p_i. Under this assumption, they may be more useful than the y_i.

Rational expectations means that the respondent's subjective distribution is correct, in the sense that the realization y_i is drawn from the same distribution on which the expectation p_i is based. To test the predictions of rational expectations models, we compare reported predictions with the distribution of realizations across the sample of respondents. This does not exclude common shocks, which would lead to correlation between the y_i for different respondents i. For our rational expectations tests, we need realizations to be independent across respondents. We therefore do not allow for common shocks. Thus, when we say we test the best-case scenario, we actually test the joint null hypothesis of (1) rational expectations, (2) best predictions under assumed loss function, and (3) independence of realizations (y_i or c_i) across respondents.

3.1 Modal category assumption

The rational expectations implications of the modal category assumption can be formalized as

$$\mathrm{P}\{c_i = k|s_i, p_i = k\} \geq \mathrm{P}\{c_i = j|s_i, p_i = k\}, \quad j = 1, \dots, K. \tag{1}$$

The probabilities here are computed according to the subjective distribution of respondent i, given the information s_i. As in Manski (1990), let x_i denote some component of s_i that is observed by the econometrician. Using that x_i is contained in s_i, we have

$$\mathrm{P}\{c_i = k|x_i, p_i = k\} \geq \mathrm{P}\{c_i = j|x_i, p_i = k\}, \quad j = 1, \dots, K. \tag{2}$$

Under this model, the best-case scenario implies that, for any group of respondents who report $p_i = k$, a plurality of realizations will fall in category k. For the density in Figure 1, this would imply $p_i = 2$: C_2 has the largest probability. Realizations are based upon drawings from the same distribution leading to the probabilities in (1) and (2). We can then use observations of c_i to check whether (2) holds. Consider the case that x_i is discrete. For notational convenience, assume that x_i is fixed, and define $P_{j|k} \equiv \mathrm{P}\{c_i = j|x_i, p_i = k\}$. Let $\hat{P}_{j|k}$ be the sample equivalent of $P_{j|k}$, i.e. the number of observations with $c_i = j$ and $p_i = k$ and the given value of x_i, divided by n_k, the number of observations with $p_i = k$ and the given value of x_i. Finally, define

$$P_k \equiv (P_{1|k}, ..., P_{K|k})', \quad \hat{P}_k \equiv (\hat{P}_{1|k}, ..., \hat{P}_{K|k})'.$$

If there are no macro-economic shocks, the c_i are independent (conditional on x_i and p_i) and the limiting distribution of $\sqrt{n_k}(\hat{P}_k - P_k)$ is $\mathrm{N}(0, \Sigma)$, with the i-th diagonal element of Σ given by $P_{i|k}(1-P_{i|k})$ and the (i, j)-th off-diagonal element given by $-P_{i|k}P_{j|k}$ ($i \neq j$). For each $j \neq k$ separately, we can now test the inequality in (2), i.e. the null $H_0 : P_{k|k} \geq P_{j|k}$ versus the one sided alternative $H_1 : P_{k|k} < P_{j|k}$. Under the null, we have that

$$\sqrt{n_k}(2\hat{P}_{k|k})^{-1/2}(\hat{P}_{k|k} - \hat{P}_{j|k}) \xrightarrow{\mathcal{L}} \mathrm{N}(0, 1),$$

on which we base our test.

To test the inequality (2), we need the categorical information on c_i and not the exact realizations y_i. If we observe only y_i but the threshold values are unknown, the test cannot be performed. The test does not use the ordered nature of the categories; the same procedure can be used for unordered outcomes. Note also that the categories cannot be combined (*ex post*), since this can change the modal category.

3.2 α-Quantile category assumption

One natural interpretation of p_i is that p_i is the category that contains the α-quantile of the respondent's subjective distribution of y_i. The most obvious choice is $\alpha = 0.5$, in which case p_i is the category containing the median of y_i. In Figure 1, this would lead to $p_i = 3$. Since the categories are ordered, this means that p_i is the median category. Other values of α can be relevant if respondents use their response to make actual decisions. See, for example, Leonard (1982). We see little justification for asymmetric loss functions in the case we study.

Assume, for convenience, that the subjective distribution of y_i is such that the α-quantile is uniquely defined and corresponds exactly to cumulative probability α. Let p_i^* denote this α-quantile. In the best-case scenario, the actual outcome y_i is drawn from this same subjective distribution, and thus we have

$$P\{y_i - p_i^* < 0|s_i\} = \alpha. \tag{3}$$

If the observed predicted category p_i is equal to k then $p_i^* \in C_{k,i} = (m_{k-1,i}, m_{k,i}]$, so

$$m_{k-1,i} < p_i^* \leq m_{k,i}. \tag{4}$$

This implies

$$y_i - m_{k,i} \leq y_i - p_i^* < y_i - m_{k-1,i}.$$

With (3), it follows directly that

$$P\{y_i - m_{k-1,i} < 0|s_i, p_i = k\} \leq \alpha \leq P\{y_i - m_{k,i} < 0|s_i, p_i = k\}. \tag{5}$$

If y_i itself is observed but the $m_{k,i}$ are unknown, this is of little value without further assumptions on the $m_{k,i}$. We will come back to this in Section 6. Here, we focus on the case that we observe the category c_i, with $c_i = k$ iff $y_i \in C_{k,i}$. This imposes no restrictions on the $m_{k,i}$ across individuals; all we need is that the outcome variable c_i is based on the same categories as the prediction p_i. Equation (5) can be written as

$$P\{c_i \leq k - 1|s_i, p_i = k\} \leq \alpha \leq P\{c_i \leq k|s_i, p_i = k\}.$$

This implies the following inequalities for the α-quantile category assumption:

$$P\{c_i > k|x_i, p_i = k\} \leq 1 - \alpha \tag{6}$$

$$P\{c_i < k|x_i, p_i = k\} \leq \alpha. \tag{7}$$

The best-case scenario now implies that, for any group of respondents who report $p_i = k$, the α-quantile of the distribution of realizations falls in category k. Therefore, no more than $100\alpha\%$ of realized values are in lower categories and no more than $100(1 - \alpha)\%$ are in higher categories.

Whether (6) and (7) are satisfied for given k and α can be tested straightforwardly. For example, with $P_{j|k}$ and $\hat{P}_{j|k}$ defined as in Section 3.1, a test of (6) can be based upon

$$\sqrt{n_k}(\sum_{j=k+1}^{K} \hat{P}_{j|k} - \sum_{j=k+1}^{K} P_{j|k}) \xrightarrow{\mathcal{L}} \mathrm{N}(0, (1 - \sum_{j=k+1}^{K} P_{j|k}) \sum_{j=k+1}^{K} P_{j|k}). \qquad (8)$$

Unlike the test in Section 3.1, this test uses the ordering of the categories. This suggests that the required assumptions are stronger than those used for the modal category assumption. But for the case that $\alpha = 0.5$ (median category assumption) we see that (6) and (7) for all k do not imply that (2) holds for all k and j, and vice versa. It is true, however, that for $k = 1$ (i.e., the lowest category) (6) implies (2) and for $k = K$ (i.e., the highest category) (7) implies (2). Thus the median category assumption is stronger than the modal category assumption in the sense that it imposes sharper lower bounds on the probabilities that the extreme predictions (i.e., k equals either 1 or K) are realized. The modal category assumption always requires a plurality of probability mass in the predicted category, whereas the median category requires a majority, when either the lowest or highest category is predicted.

3.3 Mean assumption

The third interpretation of what respondents may have in mind when they provide their subjective prediction is that p_i is the category that contains $\mathrm{E}\{y_i|s_i\}$, the subjective mean of y_i (in Figure 1, this leads to $p_i = 4$). As in the previous subsection, $p_i = k$ implies (4). Thus

$$\mathrm{E}\{y_i|s_i, p_i = k\} \in (m_{k-1,i}, m_{k,i}],$$

and also

$$\mathrm{E}\{y_i|x_i, p_i = k\} \in (m_{k-1,i}, m_{k,i}]. \qquad (9)$$

The best-case scenario here implies that, for any group of respondents who report $p_i = k$, the mean of the distribution of realizations falls in category k.

Under the mean assumption, categorical information on y_i cannot be used to test the best-case scenario. Actual values of y_i and information on the threshold values $m_{k,i}$ are required. If the $m_{k,i}$ are known and if independent observations y_i are available, a test of (9) can be based upon the standard asymptotic behavior of a sample mean (conditional upon x_i). If the $m_{k,i}$ are unknown but some prior information on them is available, we may still be able to carry out a test based upon a sample mean of the y_i. We come back to this in the empirical application in Section 6.1.

4   Application to income change predictions

We apply the tests developed in the previous section to data on income change predictions and realizations. The data are taken from the Dutch Socio-Economic Panel (SEP), ad-

ministered by Statistics Netherlands. The SEP is based upon a two-level sample design: about 200 local authorities are drawn in the first stage and households are drawn randomly per local authority in the second stage. The fraction of households drawn per local authority is so small that clustering problems can be ignored. The sample is designed to be representative of all Dutch households excluding those living in institutions like nursing homes etc. Due to nonresponse, some groups, such as single and elderly individuals, are underrepresented. We have no reason to believe that this type of nonresponse leads to mistaken inferences on population behavior.

Households were interviewed in October 1984, then twice a year (April and October) until 1989, and once a year since 1990. Personal interviews are held with all adult household members. In 1990, the questions related to (actual) income have changed substantially. We therefore focus on 1984 till 1989. We use the October waves which contain the information on income. The attrition rate in the panel is about 25 percent on average, and decreases over time. New households have entered the panel each year. After eliminating observations with item nonresponse, we retained a sample of 6845 households. Only 10.5% of them are in all six waves. For 14% of all households the required information is available in five waves, for 18% in four, for 16.8% in three, and for 16.4% in two waves. The remaining households (24.3%) provided information for only one wave. Most of those who are in more than one wave participate in consecutive waves. The numbers of observations per wave and in consecutive waves are presented in Table 1.

Table 1 : Numbers of observations per
wave and in consecutive waves

wave	# obs.	cons. waves	# obs.
1984	2683		
1985	2787	1984 − 1985	1647
1986	3850	1985 − 1986	2087
1987	3899	1986 − 1987	2835
1988	4059	1987 − 1988	3027
1989	4133	1988 − 1989	3219

Heads of household are asked to answer similar questions on realized income changes and future income changes. The question on the future is given by

> What will happen to your household's income in the next twelve months? Possible answers are: strong decrease (1); decrease (2); no change (3); increase (4); strong increase (5).

The answer to this question of head of household i in the sample is denoted by p_i. In each wave, heads of households are also asked what *happened* to their household income in the last twelve months. This question is formulated in the same way as the one on future income, with the same categories as possible answers. The answer is denoted by c_i. Since the questions are similar, and the question on p_i immediately follows the question on c_i, it seems reasonable to assume that the respondents use the same income concept for both answers. We compare p_i in wave t with c_i in wave $t + 1$ ($t =' 84, ...,' 88$). In the next section, we discuss the tests using the qualitative data.

We also have a quantitative measure of household income, based upon survey questions on many income components of all household members. This is used to construct a continuous measure of realized income change, which will be used in Section 6.

We lose approximately 25% of the households in the data set because of at least one missing income component for some household member. Given the households with complete records on all income components, the non-response rate of the subjective answers to the income prediction and realization question is quite low (about 1% in each wave).

5   Qualitative data on realized income

Under the additional assumption that p_i and c_i are based upon the same income concept and the same category bounds, we can test the best-case hypothesis that respondents form rational expectations and report best predictions, using the three models in Section 3 to define the best prediction. As noted previously, the tests are actually joint tests of the best-case scenario and the statistical independence of realizations.

5.1   Modal category assumption

Table 2 displays the frequencies of realized change categories c_i given predicted change categories p_i for five combinations of adjacent years. We present frequencies that do

not condition on other covariates, so x_i is just the "year of observation." Since the SEP is unbalanced, the numbers of observations varies across waves (see the final column of Table 2).

Table 2 shows that, for $k = 1$ (strong decrease predicted), the inequality (2) is not satisfied in three years: in '86-'87 the frequencies for $c = 2$ and $c = 3$ exceed the frequency for $c = 1$, in '84-'85 and '87-'88, this holds for the frequency for $c = 3$ only. None of these results, however, are statistically significant at conventional levels (nor are they when the data are pooled across years).

For $k = 2$, however, the numbers of observations are larger, and the findings are stronger. The inequalities are violated for each year: of those who predict a moderate income fall, the number of households who actually experience no change is larger than the number whose income moderately falls. This is statistically significant in four of the five years (and also if the data are pooled). The systematic violation of inequality (2) suggests that either the modal category assumption is inappropriate or the best-case scenario is not realistic. For $k = 3$, $k = 4$, and $k = 5$, (2) is never violated.

Table 2 : Estimates of $P\{c_i = c | p_i = k\}$ (in percentages), where k is *predicted* **category and c is** *realized* **category of income change**

		$c = 1$	$c = 2$	$c = 3$	$c = 4$	$c = 5$	$n_k^{*)}$
$k = 1$:	'84 - '85	29.7	26.7	31.7	10.9	1.0	101
strong decrease	'85 - '86	42.1	15.8	28.9	13.2	0.0	38
	'86 - '87	24.5	28.6	32.7	8.2	6.1	49
	'87 - '88	32.4	19.1	41.2	2.9	4.4	68
	'88 - '89	41.5	9.8	29.3	17.1	2.4	41
	pooled	32.7	21.5	33.3	9.8	2.7	297
$k = 2$:	'84 - '85	10.6	24.6	53.2	10.0	1.6	549
decrease	'85 - '86	10.6	24.7	51.6	10.6	2.4	376
	'86 - '87	12.2	35.7	42.7	7.8	1.7	361
	'87 - '88	7.5	20.3	61.4	8.7	2.0	492
	'88 - '89	9.4	21.6	53.5	13.6	1.9	361
	pooled	10.0	25.0	53.1	10.1	1.9	2139
$k = 3$:	'84 - '85	3.0	10.4	68.8	15.0	2.8	808
no change	'85 - '86	2.4	8.7	66.0	20.1	2.8	1313
	'86 - '87	3.5	13.7	64.1	16.4	2.3	1919
	'87 - '88	2.2	7.1	70.2	16.8	3.8	1944
	'88 - '89	1.7	5.5	67.9	21.0	3.9	2232
	pooled	2.5	8.8	67.3	18.2	3.2	8216
$k = 4$:	'84 - '85	3.9	7.7	28.7	48.1	11.6	181
increase	'85 - '86	0.9	3.2	34.8	50.0	11.1	342
	'86 - '87	1.8	5.7	37.8	43.9	10.8	492
	'87 - '88	1.8	4.1	37.0	44.3	12.8	508
	'88 - '89	2.1	3.6	26.0	52.8	15.5	561
	pooled	1.9	4.5	33.2	47.7	12.7	2084
$k = 5$:	'84 - '85	0.0	0.0	25.0	12.5	62.5	8
strong increase	'85 - '86	0.0	0.0	33.3	16.7	50.0	18
	'86 - '87	0.0	7.1	28.6	21.4	42.9	14
	'87 - '88	6.7	0.0	13.3	26.7	53.3	15
	'88 - '89	0.0	4.2	25.0	25.0	45.8	24
	pooled	1.3	2.5	25.3	21.5	49.4	79

*) $n_k = \#\{i : p_i = k\}$

It should also be noted here that the outcome data might contain mistakes due to limits on respondents' memory. The outcome questions require respondents to recall their income from the past year, their income from one year ago, and to compare the two. It seems reasonable to suppose that some respondents are not able to recall this information accurately and that others never take the trouble to try. These respondents may give "no change" as a kind of "can't remember" response, inflating this category relative to other response options. This might explain why so many people who expected a decrease report that their income did not change. But it would lead to the same phenomenon for people who expect an increase, and this is not confirmed by the data. Moreover, using the quantitative income data, we calculated a 90% confidence interval for the median real income change for those who reported a "no change" in last year's income. The result is

the interval (0.82%; 1.17%). Thus, those who reported "no change" in fact experienced a small increase, on average. This can not explain the underestimation of future income found in Table 2: since the realized income change is even larger than reported, the underestimation becomes even stronger.

We also calculated the estimates in Table 2 conditional on several covariates x_i, such as the level of net household income, dummies for actual income changes in the past year (lags of c_i), sex of the head of household, and dummies for the labor market state of head and spouse. For a continuous x_i it is possible to split the sample into groups (such as low and high incomes), or to use nonparametric (kernel) estimates.

The overall conclusion of the conditional analysis is that the pattern in Table 2 basically remains the same if we condition on given values of x_i. (Exact results are available from the authors upon request.) For almost all x_i and combinations of adjacent years, the estimate of $P\{c_i = 3|x_i, p_i = 2\}$ exceeds that of $P\{c_i = 2|x_i, p_i = 2\}$. Thus the violation of (2) cannot be ascribed to one specific income category, to households with a specific composition or labor market state, or to households whose income fell in the past.

5.2 Median and other quantile category assumptions

In this subsection, we first test the best-case implications [inequalities (6) and (7)] of the median response model (i.e., $\alpha = 0.5$). For the case x_i includes "year of observation" only, the tests for the best-case scenario under the median category assumption can be derived from the data in Table 2. By adding up the relevant probabilities and replacing the unknown variance in (8) with a consistent estimate, we can construct confidence intervals for the probabilities in (6) and (7). Table 3 displays (two-sided) 90% confidence intervals. We perform one-sided tests, with significance level 5%.

For $k = 1$ the hypothesis $P\{c_i > k|p_i = k\} \leq 0.5$ is rejected in three years: three confidence intervals do not contain the value 0.5, and violation of (6) is statistically significant. This also holds for the data pooled across years. For $k = 2$, four of the five probabilities are significantly larger than 0.5. For $k = 5$, inequality (7) is not violated significantly (Table 3), although for two means, the point estimates $P\{c_i < 5|p_i = 5\}$ are larger than 0.5 (Table 2). The conclusion is therefore the same as in the previous subsection: those who expect a moderate decrease appear to be too pessimistic, on average.

If we repeat the calculations conditional on certain values of covariates, the results are somewhat clearer than for the modal category assumption. Partitioning according to income level, we find that, for those who predict their income to fall, (7) is often violated significantly for the lower and intermediate income quartiles, but less so for the highest income quartile. For the lowest income quartile, we also find for two years significant violations of (6) for those who predict a moderate income rise. This group in particular seems to expect a (positive or negative) income change too often. A similar conclusion can be drawn for those who did not experience an income change in the previous year. For $k = 3$, the data respect both inequalities, indicating that the best-case hypothesis cannot be rejected for the groups who predict their income to be stable.

Table 3 also allows to test under α-quantile category assumptions for other values of α. For each separate row in the table, the confidence intervals together with (6) and (7) allow us to determine ranges of α for which the best-case scenario is not rejected. For

example, for $k = 1$, '86-'87, the best-case scenario is not rejected for $\alpha \leq 0.346$. For $k = 4$, '86-'87, however, we do not reject if $\alpha \geq 0.416$. Thus our data do not support the best-case hypothesis combined with a uniform asymmetric loss function based upon a single value of α per year. A similar result is found for '87-'88 and '88-'89.

Table 3 : 90% confidence intervals for the (cumulative) probabilities (in percentages)

		$P\{c_i < k\|p_i = k\}$		$P\{c_i > k\|p_i = k\}$		$n_k^{*)}$
		lower	upper	lower	upper	
$k = 1$:	'84 - '85	–	–	62.8	77.8	101
strong decrease	'85 - '86	–	–	44.7	71.1	38
	'86 - '87	–	–	65.4	85.6	49
	'87 - '88	–	–	58.3	77.0	68
	'88 - '89	–	–	45.9	71.2	41
	pooled	–	–	62.9	71.8	297
$k = 2$:	'84 - '85	8.4	12.7	61.5	68.2	549
decrease	'85 - '86	8.0	13.3	60.6	68.7	376
	'86 - '87	9.4	15.0	47.8	56.4	361
	'87 - '88	5.6	9.5	68.8	75.5	492
	'88 - '89	6.9	11.9	65.0	73.0	361
	pooled	8.9	11.0	63.3	66.7	2139
$k = 3$:	'84 - '85	11.4	15.3	15.6	20.0	808
no change	'85 - '86	9.6	12.5	21.0	24.8	1313
	'86 - '87	15.8	18.6	17.2	20.2	1919
	'87 - '88	8.2	10.3	19.1	22.1	1944
	'88 - '89	6.3	8.1	23.4	26.4	2232
	pooled	10.7	11.8	20.7	22.2	8216
$k = 4$:	'84 - '85	34.3	46.3	7.7	15.5	181
increase	'85 - '86	34.6	43.2	8.3	13.9	342
	'86 - '87	41.6	49.0	8.5	13.1	492
	'87 - '88	39.3	46.5	10.4	15.2	508
	'88 - '89	28.5	35.0	13.0	18.0	561
	pooled	37.8	41.3	11.5	13.9	2084
$k = 5$:	'84 - '85	9.3	65.7	–	–	8
strong increase	'85 - '86	30.6	69.4	–	–	18
	'86 - '87	35.4	78.9	–	–	14
	'87 - '88	25.5	67.9	–	–	15
	'88 - '89	37.4	70.9	–	–	24
	pooled	41.4	59.9	–	–	79

$^{*)}$ $n_k = \#\{i : p_i = k\}$

6 Quantitative data on realized income

6.1 Mean assumption

The categorical information on realized income is not enough to test the best-case hypothesis under the assumption that p_i reflects the category containing the mean. Instead, we need quantitative information on realized income. The SEP contains detailed information on income from about twenty potential sources for each household member. After tax household income is constructed by adding up all income components of all family members. The change in household income is then obtained from two consecutive waves.

The subjective questions on past and future income changes are not precise. It is not clear whether households should consider real or nominal income, absolute or percentage changes, or which threshold values $m_{k,i}$ they should use to distinguish between a strong change, a moderate change, and no change. Thus, additional assumptions on respondent behavior are now required. On the other hand, we no longer need to assume that respondents use the same concept or category bounds for predicted and actual income changes.

It appears that, whichever concept of income change is used, the income change variable suffers from enormous outliers. This has strong effects on the means for the subsamples with a given income change prediction. They are estimated inprecisely, and the tests based upon (9) do not seem meaningful.

A practical solution to this is to remove the observations in the upper and lower tails of the distribution of the income change variables. In Table 4, the 5% lowest and 5% highest observations are deleted. This is done for each income change variable and each year separately, without partitioning according to p_i.

In Table 4, we assume that households consider percentage income changes, either in nominal or in real (or inflation adjusted) terms. The table presents estimates of the mean and their standard errors for all values of p_i and all years. (Standard errors are not corrected for the trimming procedure.) As in Tables 2 and 3, the only covariate we condition on is the year of observation.

Table 4 : 5%-Trimmed sample means of the (actual) income change (in %) per prediction category k (standard errors of sample means in parentheses)

CHANGE IN TERMS OF PERCENTAGES					
		nominal		real		$\#\{i : p_i = k\}^*$
'84 - '85	$k = 1$	0.7	(2.1)	-1.8	(2.0)	90
	2	4.5	(0.7)	1.9	(0.7)	499
	3	5.5	(0.6)	2.9	(0.6)	725
	4	9.0	(1.4)	6.3	(1.3)	161
	5	0.8	(5.8)	-1.6	(5.6)	6
'85 - '86	$k = 1$	-7.1	(3.1)	-7.3	(3.1)	33
	2	-0.8	(0.9)	-1.0	(0.9)	332
	3	2.2	(0.5)	2.0	(0.5)	1190
	4	8.3	(1.0)	8.1	(1.0)	310
	5	13.2	(5.2)	13.0	(5.1)	12
'86 - '87	$k = 1$	0.7	(4.1)	0.9	(4.1)	45
	2	3.7	(1.2)	3.9	(1.2)	320
	3	4.9	(0.5)	5.1	(0.5)	1730
	4	7.2	(1.1)	7.4	(1.1)	444
	5	28.7	(4.8)	28.9	(4.8)	12
'87 - '88	$k = 1$	-5.0	(2.5)	-5.9	(2.5)	63
	2	1.8	(0.8)	0.9	(0.8)	435
	3	2.7	(0.5)	1.8	(0.5)	1748
	4	4.9	(0.8)	4.0	(0.8)	465
	5	7.2	(5.6)	6.2	(5.6)	12
'88 - '89	$k = 1$	3.7	(4.9)	2.1	(4.8)	31
	2	1.5	(1.1)	-0.1	(1.1)	312
	3	6.0	(0.4)	4.3	(0.4)	2017
	4	8.7	(0.9)	7.0	(0.9)	516
	5	20.0	(6.3)	18.1	(6.2)	21

[*] Outliers are determined for the nominal and inflation adjusted change separately. The nominal change is calculated as $\frac{1}{n} \sum_{i=1}^{n} \frac{y_{i,t+1} - y_{i,t}}{y_{i,t}}$ where $y_{i,t}$ is income of family i and n is the number of respondents in the considered category. In the formula for the real change, $y_{i,t+1}$ is replaced by $y_{i,t+1}/I_{t+1,t}$ with $I_{t+1,t}$ the consumer price index of year $t + 1$ compared to year t.

The standard errors are quite large. To obtain standard errors for the differences between two means for different values of k, the corresponding variance estimates can be added, due to independence (means for different values of k are based upon disjoint sets of observations). In many cases, the means for consecutive values of k are not significantly different.

For a large sample size, (9) implies that the sample means should increase with k. This is usually the case. Only for extreme predictions ($k = 1$ or $k = 5$) is this violated in some years, but never significantly. More specific tests can be carried out if prior information on the threshold values $m_{k,i}$ is used. For example, it seems reasonable that $m_{1,i}$ and $m_{2,i}$ are negative, while $m_{3,i}$ and $m_{4,i}$ should be positive. This implies that the means for

$k = 1$ and $k = 2$ should be negative, and those for $k = 4$ and $k = 5$ should be positive. For $k = 2$ only, we find significant violations in Table 4, for the nominal as well as the real percentage income change. Thus, as in Section 5, the conclusion is that the group of households expecting a moderate decrease is overly pessimistic, on average.

6.2   Median category assumption

Using the quantitative data on income changes we can also (nonparametrically) estimate the cumulative distribution function (cdf) of the realized income change conditional on the expected income change category. From now on, we assume that the threshold values are constant across time and individuals, and we use the pooled data set.

Figure 2 presents smoothed empirical distribution functions of the realized percentage real income change (y_i) for given expected income change category (p_i). The function is smoothed with an integrated Epanechnikov kernel (see, for example, Härdle and Linton, 1994). The cdf's for higher p_i are to the right of those with lower p_i, confirming that those who are more optimistic have a higher probability of a change exceeding $a\%$, for each a. We find similar patterns when disaggregating by year, with some exceptions for the extreme categories with few observations. All figures are available upon request.

Figure 2 : Distribution functions of the realized real income change given the expected change category (data pooled across years).

INSERT FIGURE 2 HERE

Let us assume that the best-case scenario holds. From Section 3.2 we know that the α-quantile assumption then implies

$$P\{y_i \leq m_{k-1}|p_i = k\} \leq \alpha \leq P\{y_i \leq m_k|p_i = k\}. \tag{10}$$

If $\xi_{\alpha,k}$ denotes the α-quantile of y_i conditional on $p_i = k$, this can be written as

$$m_{k-1} \leq \xi_{\alpha,k} \leq m_k.$$

For $\alpha = 0.5$, Figure 2 shows that $\xi_{\alpha,2}$ is about zero, suggesting that m_2 is nonnegative. This seems unreasonable, since it would lead to the implausible asymmetry that the no change category $(m_2, m_3]$ contains nonnegative changes only. Working with nominal instead of real changes makes the asymmetry even stronger. Thus the best-case hypothesis is rejected for the group of households expecting a moderate decrease. This means that the results based upon the quantitative measures confirm the test outcomes based upon the qualitative measures. While the latter could be due to inaccurate qualitative reports, the former shows that this is not the explanation. Whether misreporting plays a role can also be checked more directly by comparing quantitative and qualitative measures of realized income changes. Table 5 presents 90% confidence intervals for the median real income change conditional on the qualitative report of income change. For those who report no change, for example, the median real income change was about 1%. More precisely, the median real income change for those who report no change and expected a decrease was 0.20% (a 90% confidence interval is given by [-0.12%;0.52%]). Thus there is no evidence that the test results based upon qualitative outcomes are due to recall errors.

Table 5 : 90% confidence intervals for the median real income change categorized by qualitatively reported income change.

	90% confidence interval	
qualitative data	lowerbound	upperbound
strong decrease	-11.01	-6.58
decrease	-1.43	-0.64
no change	0.82	1.17
increase	4.57	5.23
strong increase	12.49	16.83

7   Conclusions

Manski (1990) has compared realizations with predictions for the case of two possible outcomes. We have generalized his framework to the case of more than two outcomes. We discuss which assumptions are necessary to derive bounds on the theoretical relationship between expectations and realizations under the best-case scenario that respondents form rational expectations and report best predictions. We have focused on the case of ordered outcomes that can be interpreted as categories of an underlying continuous variable. Unlike in Manski's case, the inequalities to be tested appear to depend on the assumption on which location measure of the subjective distribution of the variable of interest is reflected by the subjective prediction. We discuss three possibilities: the modal category, the median (or α-quantile), and the mean. The former two can be applied to conduct

rational expectations tests if categorical data on predictions and outcomes are available, while the latter can only be applied if quantitative realizations data are available. The three assumptions lead to different bounds, none uniformly sharper than any other.

The tests are applied to Dutch household data on predicted and actual income changes, using panel data for 1984 to 1989. On the basis of the categorical realizations data, we find the same results for the modal and median category assumption: the best-case hypothesis is rejected for the group of households expecting a moderate income decrease. For too many of these, the realization is "no change". This result has various interpretations. One is that observations are not independER, due to common shocks. That this result obtains for a number of years reduces the plausibility of this explanation. Some insights may be gained by considering macro-economic trends in Dutch incomes. Real disposable household income decreased, on average, during the years preceding the survey (about 5% per year from 1982 to 1984). Incomes stabilized in 1985 and then increased gradually (about 3% in 1986 and 1987 and about 2% in 1988 and 1989). Thus, the experience of the early-80's may have led to persistent pessimism, despite the income growth during the latter half of the decade. A second interpretation could be that people use asymmetric loss functions, leading to the α-quantile assumption with $\alpha \neq 0.5$. Using the categorical realizations, we found that there is no single value of α which can explain the data for all years under the best-case scenario. A third explanation is that substantial groups of households do not form rational expectations.

To make a definitive choice among these interpretations of our findings, more research seems necessary, for example based upon data with more detailed information on individuals' subjective income distributions. Such data now exist in the Dutch VSB-panel (Das and Donkers, 1997), the American Survey of Economic Expectations (Dominitz and Manski, 1997a), and the Italian Survey of Household Income and Wealth (Guiso et al., 1992).

References

Browning, M., and Lusardi, A. (1996), "Household Saving: Micro Theory and Micro Facts," *Journal of Economic Literature*, 34, 1797-1855.

Carlson, J.A., and Parkin, M. (1975), "Inflation Expectations," *Economica*, 42, 123-138.

Carroll, C., Fuhrer, J., and Wilcox, D. (1994), "Does Consumer Sentiment Forecast Household Spending? If So, Why?," *American Economic Review*, 84, 1397-1408.

Das, M., and Donkers, B. (1997), "How Certain are Dutch Households about Future Income? An Empirical Analysis," *CentER Discussion Paper 9738*, Tilburg University.

Das, M., and Van Soest, A. (1996), "A Panel Data Model for Subjective Information on Household Income Growth," *CentER Discussion Paper 9675*, Tilburg University.

(1997), "Expected and Realized Income Changes: Evidence from the Dutch Socio-Economic Panel," *Journal of Economic Behavior and Organization*, 32, 137-154.

Dominitz, J., and Manski, C. (1996), "Eliciting Student Expectations of the Returns to Schooling," *Journal of Human Resources*, 31, 1-26.

(1997a), "Using Expectations Data to Study Subjective Income Expectations," *Journal of the American Statistical Association*, 92, 855-867.

(1997b), "Perceptions of Economic Insecurity: Evidence from the Survey of Economic Expectations," *Public Opinion Quarterly*, 61, 261-287.

Fuhrer, J. (1988), "On the Information Content of Consumer Survey Expectations," *Review of Economics and Statistics*, 140-144.

Garner, C. A. (1991), "Forecasting Consumer Spending: Should Economists Pay Attention to Consumer Confidence Surveys?," *Federal Reserve Bank of Kansas City Economic Review*, May/June, 57-71.

Guiso, L., Jappelli, T., and Terlizzese, D. (1992), "Earnings Uncertainty and Precautionary Saving," *Journal of Monetary Economics*, 30, 307-338.

(1996), "Income Risk, Borrowing Constraints and Portfolio Choice," *American Economic Review*, 86, 158-172.

Härdle, W., and Linton, O. (1994), "Applied nonparametric methods," in *Handbook of Econometrics*, Vol. IV, eds. R. Engle and D. McFadden, North-Holland, Amsterdam.

Hurd, M., and McGarry, K.(1995), "Evaluation of Subjective Probabilities of Mortality in the HRS," *Journal of Human Resources*, 30, S268-S292.

Ilmakunnas, S., and Pudney, S. (1990), "A Model of Female Labour Supply in the Presence of Hours Restrictions," *Journal of Public Economics*, 41, 183-210.

Juster, F.T. (1966), "Consumer Buying Intentions and Purchase Probability: An Experiment in Survey Design," *Journal of the American Statistical Association*, 61, 658-696.

Keane, M.P., and Runkle, D.E. (1990), "Testing the Rationality of Price Forecasts: New Evidence from Panel Data," *American Economic Review*, 80, 714-735.

Kiewiet, D. R. (1983), *Macroeconomics and Micropolitics*, Chicago: University of Chicago Press.

Leonard, J.S. (1982), "Wage Expectations in the Labor Market: Survey Evidence on Rationality," *Review of Economics and Statistics*, 64, 157-161.

Linden, F. (1982), "The Consumer as Forecaster," *Public Opinion Quarterly*, 46, 353-360.

Lovell, M.C. (1986), "Tests of the Rational Expectations Hypothesis," *American Economic Review*, 76, 110-124.

MacKuen, M., Erikson, R., and Stimson, J. (1992), "Peasants or Bankers? The American Electorate and the U.S. Economy," *American Political Science Review*, 86, 597-611.

Maddala, G.S. (1994), "Survey Data on Expectations: What Have We Learnt?," in *Econometric Methods and Applications*, vol. 1, ed. G.S. Maddala, Brookfield, VT: E. Elgar.

Manski, C. (1990), "The Use of Intentions Data to Predict Behavior: A Best-Case Analysis," *Journal of the American Statistical Association*, 85, 934-940.

Nerlove, M. (1983), "Expectations, Plans, and Realizations in Theory and Practice," *Econometrica*, 51, 1251-1279.

Norpoth, H. (1996), "Presidents and the Prospective Voter," *The Journal of Politics*, 58, 776-92.

Otoo, M. (1997), "The Sources of Worker Anxiety: Evidence from the Michigan

Survey," *Finance and Economics Discussion Series*, 1997-48, Federal Reserve Board, Washington, D.C.

Tobin, J. (1959), "On the Predictive Value of Consumer Intentions and Attitudes," *Review of Economics and Statistics*, 41, 1-11.

Throop, A. (1992), "Consumer Sentiment: Its Causes and Effects," *Federal Reserve Bank of San Francisco Economic Review*, Winter.