Comparing Predictions and Outcomes

Das, J.W.M.; Dominitz, J.; van Soest, A.H.O.

Publication date:
1997

Comparing Predictions and Outcomes: Theory and Application to Income Changes †

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

* Tilburg University, Department of Econometrics and CentER

** California Institute of Technology, Division of Humanities and Social Sciences

May 1997

JEL-classification: C52, C53, D12

Keywords: predictions, categorical data, loss function, income growth

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 have 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 realized income changes illustrates how alternative response models may be used to test the best-case hypothesis.

---

†Data were provided by Statistics Netherlands. The authors are grateful to Ben van der Genugten for valuable comments. Financial support by the Royal Netherlands Academy of Arts and Sciences (KNAW) is gratefully acknowledged by the third author.

Please address correspondence to Jeff Dominitz, California Institute of Technology, Division of Humanities and Social Sciences, 1200 East California Boulevard, Pasadena, California 91125 (e-mail: dominitz@hss.caltech.edu).

1 Introduction

Subjective data on respondents' intentions or predictions are commonly used for many purposes. For instance, in periods prior to elections, voter polls are held almost continuously and are taken seriously by politicians, journalists, and voters. Economists, however, are quite skeptical of subjective data. It has been claimed, for example, that expectations data need not match up to future outcomes, because there is no incentive for respondents to report expectations accurately (see, for example, Keane and Runkle, 1990).

Some examples in the recent literature, however, suggest that this attitude is changing. Dominitz and Manski (1996, 1997) 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 can 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 the expectations data bound but do not identify the probability of each possible outcome.[1]

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 future variables that will affect their decisions (income, prices, etc.), they

[1]For more on identification of probabilities, see Manski (1995).

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". In general, however, part of the group will actually buy a car. If the subjective distributions of the future variables are correct, and if the realizations of the future variables are independent, about 40% will buy a car. Thus there is no reason to expect that the distribution of the intention variable across the population is the same as the distribution of the actual variable. The reason is that the intention reflects some location measure of the household's subjective distribution (for instance, the mode), while the outcome is based upon one draw from the actual distribution. Even in the best-case scenario that subjective and actual distributions coincide, the two 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 which 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 which takes on more than two values. Our empirical analysis, for example, focuses on expectations of the change in household income, which respondents report by choosing among five ordered categories.

We consider three models generating best predictions of the prospective realization. Each model is based on a different expected loss function which respondents minimize. These behavioral models yield responses of (1) the modal category, (2) the category containing some quantile 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 "yes/no" expectations, different symmetric loss functions may yield different multiple-ordered-category survey responses and therefore imply different best-case bounds. Using panel data or repeated cross-sections, the best-case hypothesis may therefore be tested only under stronger homogeneity assumptions on the expected loss function respondents attempt to minimize. Our analysis illustrates how these tests may be conducted when categorical and/or continuous realizations data are available.

We apply our tests to data on income change expectations and outcomes reported in the 1984 – 1989 waves of the Dutch Socio-Economic Panel (SEP). Heads of households 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 ordered-category question is asked about the change in income over the past twelve months. In addition, quantitative reports of the actual income level are given in each interview.

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 have a tendency to underestimate their income change. This means that either households' expectations are not rational, or macro-economic shocks take place in a number of consecutive years, or both. An alternative explanation can be given using an asymmetric loss function: respondents tend to place more weight on negative than positive forecast errors. This will lead to underestimation, on average. Though we find some support for this in our data, our results do not support the best-case hypothesis combined with a uniform asymmetric loss function across the population.

The outline of the paper is as follows. Section 2 discusses the modelling of responses to questions eliciting ordered-category expectations. We consider the expected loss functions that respondents may be minimizing, and we discuss the implications of previous findings in empirical research on expectations data. Section 3 derives bounds for conditional probabilities of outcomes given predictions that should be valid under the best-case hypothesis. These bounds are derived under each of the three response models presented in Section 2. Section 4 compares the expectations and realizations of income changes across Dutch households. Both categorical and quantitative realizations data are used to test the best-case hypothesis. Section 5 concludes.

2 Modelling Responses to Expectations Questions

We consider responses to qualitative survey questions eliciting expectations of some outcome y (e.g., the income growth of a household), where respondents must choose among ordered categories. While the number of categories may vary, this type of question is quite common. The questions used to generate both the University of Michigan's Index of Consumer Sentiment and the Conference Board's Consumer Confidence Index include a series of such questions with three ordered categories. See Curtin (1982) and Linden (1982), respectively. Responses of *don't know* are typically accepted, but are often discarded in empirical analyses, as is the case with each of the aforementioned indices.

In the next subsection we present several loss functions respondents may minimize when they answer questions eliciting expectations. Section 2.2 gives some examples of expectations research.

2.1 Loss Functions

As the starting point for 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 not typically defined by the survey question, but are instead subjectively determined (and not reported) by the respondent. The answer to this question is denoted by p. We propose a model in which p is based upon minimizing some loss function. This interpretation of ordered-category responses follows directly from Manski (1990), who restricts attention to the case of two categories, but the framework is implicit in work dating back at least to Tobin (1959).

Influenced by, for example, the phrasing of the question, the respondent can use various loss functions. If the respondent interprets the K-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: $p = \text{argmax}_k \, P\{y \in C_k|s\}$. Choosing the *modal category* corresponds to minimizing subjective expected loss $\text{E}\{1(y \notin C_k)|s\}$ with respect to k.

This modal category response model appears sensible, but it is not typically adopted in analyses of ordered-category expectations data.[2] Instead, researchers typically adopt a model in which the respondent forms some point expectation p^* and then chooses the category p that contains p^*. We discuss such models in a framework where the respondent finds the value p^* that minimizes subjective expected loss for some loss function L:

$$\begin{aligned} p^* &= \underset{\pi}{\text{argmin}} \int_{-\infty}^{\infty} L(y-\pi) f(y|s) dy; \\ p &= k \text{ iff } p^* \in C_k. \end{aligned}$$

Researchers often assert that respondents interpret questions of what they "think" or "expect" to happen as questions eliciting the (conditional) mean of y. Such a respondent would choose the category that contains $\text{E}\{y|s\}$. This "category-containing-the-mean" model can be explained if that the squared loss function is adopted: $L(u) = u^2$.

It also seems reasonable to assume that respondents may interpret the question as eliciting $\text{Med}\{y|s\}$, the median of $f(y|s)$. This will be the case if the absolute loss function is adopted: $L(u) = |u|$. One may generalize this approach to allow for asymmetric loss functions. In particular, consider respondents who minimize the absolute loss function:

$$L(u) = \alpha|u|1(u \geq 0) + (1-\alpha)|u|1(u < 0), \quad 0 < \alpha < 1.$$

The value p^* which minimizes subjective expected loss is then the α-quantile of $f(y|s)$.

[2] Van der Klaauw (1996) uses this model to interpret responses to unordered-category expectations of occupational choice.

2.2 Examples in Expectations Research

Responses to a binary ("yes/no") intentions question, which may be thought of as a special case of 2-ordered-category expectations questions, have been interpreted in this way, either implicitly or explicitly, by Tobin (1959), Juster (1966), and Manski (1990). The analysis is easier in this case, where y may only take on two values, 0 ("no") or 1 ("yes"). As noted by Manski (1990), this framework requires that responses chosen to minimize subjective expected loss are invariant to the choice of loss function as long as it is symmetric. Therefore, the modal, mean, and median models described in the previous subsection generate identical responses. In particular, any symmetric loss function dictates that the respondent simply chooses category $k = 1$ ("yes") if $P\{y = 1|s\} > 0.5$ and category $k = 0$ ("no") otherwise.

The response model changes if the respondent weighs prospective losses asymmetrically. Suppose the asymmetric absolute loss function is chosen. Then the respondent chooses $k = 1$ if $P\{y = 1|s\} > \alpha$ and $k = 0$ otherwise. Thus, the asymmetry simply changes the relevant "yes/no" threshold probability.

Carlson and Parkin (1975) study 3-ordered-category inflation expectations data. In our notation their model rests upon the following two assumptions:

(a) choose category k if $P\{m_{k-1} < y \leq m_k|s\} \geq 0.5 \quad (k \in \{1, 2, 3\})$,

(b) choose *don't know* if $m_1 < \text{Med}\{y|s\} \leq m_2$ and $P\{m_{k-1} < y \leq m_k|s\} < 0.5$.

That is, 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.

The study by Carlson and Parkin represents a rare instance in which *don't know* responses are modeled, and, as such, it does not fall strictly within the framework given above. The model, however, can be seen as a modification of both the modal and median response models. In any K-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 median. If no category satisfies this restriction, then some other response rule must be followed, such as (b) choose *don't know*.

Expectations data have often been used to test predictions of models of rational expectations formation. 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 $E\{y|s\}$ and then attempts to quantify the qualitative responses (Maddala, 1994). The expectations data are then confronted with subsequent realizations and tests of unbiasedness are conducted. The framework for such an analysis is not always coherently specified in terms of stating (1) the feature of the subjective probability distribution which respondents are assumed to report and (2) the rational expectations implications of the assumed response model. Nerlove (1983), for example, confronts 3-ordered-category expectations data with realizations data provided 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).

Studies of single-valued quantitative expectations of continuous outcomes typically assume that respondents report the subjective mean (i.e., minimize squared loss). When the frequentist mean of realizations conditional on the value of the subjective expectation differs from that value, it is taken as evidence that respondents form biased expectations. Some researchers have attempted to rationalize such findings by arguing that respondents do not minimize squared loss but instead minimize an asymmetric expected loss function. Leonard (1982), for example, argues that the prospective costs of raising wages (and hiring additional workers) are less than the prospective costs of lowering wages (and firing workers), so firms' wage forecasts appear to be downward-biased.

In the remaining sections of this paper, we state the implications of rational expectations models for the relationship between K-ordered-category expectations and subsequent realizations, both categorical and continuous. These implications are sensitive to the assumed loss functions and assumptions on variation in the threshold values m_k across individuals and over time.

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. We start from the three different assumptions about the respondents' strategy for answering the subjective questions discussed in Section 2.1. The three assumptions refer to which feature of the subjective distribution is reflected by p_i, the prediction of respondent i ($p_i \in \{1, ..., K\}$).[3] Section 3.1 presents the *modal category* assumption. Section 3.2 discusses the *α-quantile* assumption which for $\alpha = 0.5$ reduces to the *median category* assumption. Section 3.3 presents the *mean* assumption.

The observed prediction p_i is always a categorical variable. We distinguish, however, two cases for the realization. We either observe the exact realization y_i, or the 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, the c_i may be more informative than the y_i, since they refer to the same categories as the predictions p_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 test on rational expectations, we need independent realizations across respondents and therefore have to exclude common shocks. Thus if we say we test the best-case scenario, we test the joint null hypothesis of

[3] We shall work with a random sample of respondents from some (sub)population. The index i refers to the i-th observation in the sample.

rational expectations and independence of y_i or c_i across respondents.

3.1 Modal category assumption

The modal category assumption can be formalized as

$$\mathrm{P}\{c_i = k | s_i, p_i = k\} \geq \mathrm{P}\{c_i = j | s_i, p_i = k\}, \quad j = 1, \ldots, K \tag{3.1}$$

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

$$\mathrm{P}\{c_i = k | x_i, p_i = k\} \geq \mathrm{P}\{c_i = j | x_i, p_i = k\}, \quad j = 1, \ldots, K. \tag{3.2}$$

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. Realizations are based upon drawings from the same distribution leading to the probabilities in (3.1) and (3.2). We can then use observations of c_i to check whether (3.2) holds. Consider the case that x_i is discrete. For notational convenience, assume that x_i and p_i are fixed, and define $P_j \equiv \mathrm{P}\{c_i = j | x_i, p_i = k\}$. Let $\hat{P}_j$ be the sample equivalent of P_j, i.e. the number of observations with $c_i = j$ and $p_i = k$ and the given value of x_i, divided by n, the number of observations with $p_i = k$ and the given value of x_i. Finally, define

$$P \equiv \begin{pmatrix} P_1 \\ \vdots \\ P_K \end{pmatrix}, \quad \hat{P} \equiv \begin{pmatrix} \hat{P}_1 \\ \vdots \\ \hat{P}_K \end{pmatrix},$$

If there are no macro-economic shocks, the c_i are independent (conditional on x_i and p_i) and the limiting distribution of $\sqrt{n}(\hat{P} - P)$ is $\mathrm{N}(0, \Sigma)$, with the j-th diagonal element of Σ given by $P_j(1 - P_j)$ and the (j,k)-th off-diagonal element given by $-P_j P_k$ $(k \neq j)$.

To test the inequality (3.2), 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

Now consider the case where the survey response corresponds to the category that contains a point prediction that minimizes some expected loss function. 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. Since the categories are ordered, this means that p_i is the median category.

Assume, for convenience, that the subjective distribution of y_i is such that the α-quantile is uniquely defined and corresponds to cumulative probability α, exactly. 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

$$\mathrm{P}\{y_i - p_i^* < 0|s_i\} = \alpha \tag{3.3}$$

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

$$m_{k-1,i} < p_i^* \leq m_{k,i}, \tag{3.4}$$

This implies

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

With (3.3), it follows directly that

$$\mathrm{P}\{y_i - m_{k-1,i} < 0|s_i, p_i = k\} \leq \alpha \leq \mathrm{P}\{y_i - m_{k,i} < 0|s_i, p_i = k\}. \tag{3.5}$$

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 4.2. 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.

The inequalities in (3.5) can be written as

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

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

$$\mathrm{P}\{c_i > k|x_i, p_i = k\} \leq 1-\alpha \tag{3.6}$$

$$\mathrm{P}\{c_i < k|x_i, p_i = k\} \leq \alpha. \tag{3.7}$$

Under this model, the best-case scenario 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 α of realized values are in lower categories and no more than $1-\alpha$ are in higher categories.

Whether the inequalities in (3.6) and (3.7) are satisfied for given k and α can be tested straightforwardly. For example, with P_j and $\hat{P}_j$ defined as in Section 3.1, a test of (3.6) can be based upon

$$\sqrt{n}(\sum_{j=k+1}^{K} \hat{P}_j - \sum_{j=k+1}^{K} P_j) \xrightarrow{\mathcal{L}} \mathrm{N}(0, (1-\sum_{j=k+1}^{K} P_j)\sum_{j=k+1}^{K} P_j). \tag{3.8}$$

Unlike the test in the previous subsection, 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 (3.6) and (3.7) for all k do not imply that (3.2) holds for all k and j, and vice versa. It is true, however, that for $k = 1$ (i.e., the lowest category) (3.6) implies (3.2) and for $k = K$ (i.e., the highest category) (3.7) implies (3.2). 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.[4]

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. As in the previous subsection, $p_i = k$ implies equation (3.4). Thus

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

and also

$$\mathrm{E}\{y_i|x_i, p_i = k\} \in (m_{k-1,i}, m_{k,i}]. \tag{3.9}$$

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

A drawback of the mean assumption is that categorical information on y_i is not sufficient 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 i.i.d. observations y_i are available, (3.9) can be used to construct a test, 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 4.2.

[4]If $K = 3$, we also have that, for $k = 2$, (3.2) implies (3.6) and (3.7). In that case therefore, the modal category assumption imposes a stronger restriction for intermediate predictions than the median category assumption.

4 Application to income change predictions

We apply the tests for the best-case scenario developed in the previous section to data on income change predictions and realizations. The data are taken from the 1984 – 1989 waves of the Dutch Socio-Economic Panel (SEP), an unbalanced panel of households in the Netherlands. The same data are used in Das and Van Soest (1996, 1997). Heads of households 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 thus compare p_i in wave t with c_i observed in wave $t+1$ (t= '84, '85, '86, '87, '88). In the next subsection, we discuss the tests based upon the qualitative data. Apart from that, we have quantitative information on household income from various sources, from which we construct a continuous measure of realized income change. These quantitative data will be used in Section 4.2.

4.1 Qualitative data on realized income

4.1.1 Modal category assumption

Under the best-case hypothesis of rational expectations and statistically independent realizations, frequencies of the income growth categories can be used to estimate the probabilities in (3.2) for the modal category assumption. Table 1 displays the frequencies for five combinations of adjacent years (ranging from 1984 through 1989). We present frequencies that do not condition on other covariates, so x_i is "year of observation." Since the SEP is an unbalanced sample, the numbers of observations per wave are different (see the final column of Table 1).

Table 1 shows that, for $k = 1$ (strong decrease predicted), the inequality (3.2) 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 significant (nor is it the case for the data pooled across years). For $k = 2$, however, the findings are stronger, possibly due to the larger numbers of observations. 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 significant in four of the five years. The result is also significant in case of the pooled data. The systematic violation of inequality (3.2) suggests that either the modal category assumption is not relevant or the best-case scenario is not realistic. For $k = 3$, $k = 4$, and $k = 5$, we find no violations of (3.2).

We also calculated the estimates in Table 1 conditional on several covariates x_i, such as the level of net household income, dummies for realized income changes in the past twelve months (lagged values of c_i), sex of the head of household, and dummies for the labor market state of head and spouse.[5] For a continuous x_i it is possible to summarize

[5] Exact results are available from the authors upon request.

the continuous variable into groups, for instance, low and high income groups. It is also possible to use nonparametric estimates (see, e.g., Härdle and Linton, 1994).

Table 1 : Estimates of $P\{c_i = c|p_i = k\}$ (in percentages), where k stands for *predicted* category and c for *realized* category of income change

		$c=1$	$c=2$	$c=3$	$c=4$	$c=5$	$n^{*)}$
$k=1$: strong decrease	'84 - '85	29.7	26.7	31.7	10.9	1.0	101
	'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
$k=2$: decrease	'84 - '85	10.6	24.6	53.2	10.0	1.6	549
	'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
$k=3$: no change	'84 - '85	3.0	10.4	68.8	15.0	2.8	808
	'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
$k=4$: increase	'84 - '85	3.9	7.7	28.7	48.1	11.6	181
	'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
$k=5$: strong increase	'84 - '85	0.0	0.0	25.0	12.5	62.5	8
	'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

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

The overall conclusion of the conditional analysis is that the pattern in Table 1 basically remains the same if subsamples with given values of x_i are used. For almost all x_i and combinations of adjacent years, the estimate of P$\{c_i = 3|x_i, p_i = 2\}$ is higher than that of P$\{c_i = 2|x_i, p_i = 2\}$. The results are not always significant, but this may be due to the small numbers of observations in some of the subsamples. Thus the violation of (3.2) 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.

4.1.2 Median and other quantile category assumptions

In this subsection we first test the inequalities (3.6) and (3.7) for the median: $\alpha = 0.5$. 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 1. By adding up the relevant probabilities and replacing the unknown variance in (3.8) with a consistent estimate, we can construct confidence intervals for the probabilities in (3.6) and (3.7). Table 2 displays (two-sided) 90% confidence intervals.[6]

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 inequality (3.6) is violated significantly. This also holds for the data pooled across years. For $k = 2$, four of the five probabilities are significantly larger than 0.5. For $k = 5$, (3.6) is violated twice, but in neither of the two cases this is significant. The conclusions are therefore similar to those 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 (3.7) for those who predict their income to fall 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

[6] Note that we perform one-sided tests, with significance level equal to 5%.

violations of (3.6) 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.

Table 2 : 90% confidence intervals for the (cumulative) probabilities (in percentages)

		$P\{c_i < k \mid p_i = k\}$		$P\{c_i > k \mid p_i = k\}$		$n^{*)}$
		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

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

For $k = 3$, the category with the highest number of observations, the data respect both inequalities, indicating that for the groups who predict their income to be stable, the best-case hypothesis cannot be rejected.

Under the median category assumption, the best-case scenario is rejected for several groups. Table 2 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 the inequalities (3.6) and (3.7) allow us to determine ranges of α for which the best-case scenario is not rejected. For example, the third row implies that the best-case scenario is rejected for $\alpha > 0.346$.

In some years, the ranges of α for which the best-case scenario is not rejected do not overlap. In '86-'87, for example, $\alpha \leq 0.346$ is obtained from $P\{c_i > 1|p_i = 1\}$, while for $P\{c_i < 4|p_i = 4\}$, we get $\alpha \geq 0.416$. A similar result is found for '87-'88 and '88-'89. This means that 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. For '84-'85 and '85-'86 the bounds do not conflict with each other and the results support the best-case scenario with a value of α less than 0.5. The interpretation of this is that respondents tend to place higher weight on negative forecast errors $(y_i - p_i < 0)$. This leads to underestimation, on average.

4.2 Quantitative data on realized income

4.2.1 Mean assumption

The categorical information on y_i is not enough to test the best-case hypothesis under the assumption that p_i reflects the category containing the mean. Instead of c_i, we need y_i itself. 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 by comparing household income in two consecutive waves.[7]

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. In the previous subsections, additional assumptions on all this were not needed. The only necessary assumption was that heads of households use the same concept for predicted and realized income changes. To use the quantitative measure of household income, however, additional assumptions cannot be avoided.[8]

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 given income change prediction. Many of them are estimated inaccurately, and carrying out the tests based upon (3.9) does not lead to meaningful results (details are available upon request).

A practical solution to this problem, is to remove the observations in the upper and lower tails of the distribution of the income change variable. In Tables 3 and 4, the 5% lowest and 5% highest observations are deleted.[9]

In Table 3, we assume that households consider absolute changes, and look at nominal as well as real changes. In Table 4, we do the same for percentage income changes. Both tables present estimates of the mean and their standard errors for all values of p_i and all years.[10] As in Tables 1 and 2, the only covariate we condition on is the year of observation.

---

[7] See Das and Van Soest (1996, 1997) for details and descriptive statistics.

[8] Moreover, we also have to assume that the subjective income on which p_i is based corresponds to our objective income variable. Results of Kapteyn et al. (1988) cast doubt upon this assumption: they compare subjectively and objectively measured income levels and find that certain income components, such as spouse's income or child benefits, are often not taken into account in the subjective measure.

[9] This is done for each income change variable and each year separately, without partitioning according to p_i.

[10] Standard errors are not corrected for the trimming procedure.

Table 3 : 5%-Trimmed sample means of the (actual) absolute change in income per prediction category k (standard errors of sample means in parentheses)

ABSOLUTE CHANGE						
		nominal		real		$\#\{i : p_i = k\}^{*)}$
'84 - '85	$k = 1$	-97.9	(632.5)	-902.9	(625.7)	88
	2	1177.2	(216.8)	352.1	(215.7)	503
	3	1417.0	(197.1)	575.3	(194.1)	727
	4	1967.0	(467.2)	1112.2	(449.4)	156
	5	1801.7	(2611.4)	928.1	(2717.9)	7
'85 - '86	$k = 1$	-3348.4	(1264.0)	-3409.9	(1265.3)	35
	2	-685.3	(319.5)	-746.6	(319.1)	338
	3	456.6	(154.5)	389.2	(154.2)	1189
	4	2338.1	(327.3)	2258.0	(327.0)	302
	5	5598.2	(1749.1)	5525.4	(1749.2)	13
'86 - '87	$k = 1$	-192.1	(1137.3)	-125.4	(1139.6)	41
	2	695.3	(379.5)	704.8	(382.5)	326
	3	1148.7	(152.6)	1228.1	(152.6)	1743
	4	1519.4	(313.2)	1599.7	(313.6)	430
	5	9100.4	(1694.1)	9198.7	(1703.6)	11
'87 - '88	$k = 1$	-2794.0	(782.5)	-3037.9	(782.7)	61
	2	-136.7	(272.3)	-376.5	(271.0)	452
	3	536.6	(146.2)	219.6	(145.0)	1745
	4	1645.1	(260.0)	1307.0	(259.4)	453
	5	1187.7	(2211.1)	786.4	(2225.1)	12
'88 - '89	$k = 1$	-3618.2	(1372.3)	-4076.3	(1372.7)	35
	2	-181.7	(335.5)	-636.6	(334.5)	325
	3	1236.1	(133.7)	692.3	(132.6)	2025
	4	2404.1	(274.9)	1716.9	(272.0)	490
	5	3734.4	(1809.5)	3004.8	(1791.5)	22

*) The outliers are determined for the nominal and real change separately. Since the difference in number of observations in a specific category k is at most one observation, we only present the number of observations for the real change.

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.

Table 4 : 5%-Trimmed sample means of the (actual) change in income, in terms of percentages, 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

The inequalities (3.9) imply that, for n large enough, we would expect that the sample means increase with k.[11] This is usually the case. Only for the extreme predictions ($k = 5$ in Table 3 and $k = 1$ or $k = 5$ in Table 4) this is violated in various 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 to assume that $m_{1,i}$ and $m_{2,i}$ are negative, while $m_{3,i}$ and $m_{4,i}$ should be positive, implying that the means for $k = 1$ and

[11] This will certainly be the case if the threshold values are constant across individuals, but may not be the case if there exists a negative correlation between the thresholds and p_i.

$k = 2$ should be negative, and those for $k = 4$ and $k = 5$ should be positive. For $k = 1$, there are some positive values, but they are never significantly different from 0. For $k = 2$ however, we find significant violations, particularly in Table 4, for the nominal as well as the real percentage income change. For $k = 3$, $k = 4$, and $k = 5$, the means are always positive. Thus, as in the previous subsections, the conclusion can be drawn that the group of households expecting a moderate decrease is overly pessimistic, on average.

4.2.2 α-Quantile 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 use the pooled data set. Figure 1 presents the cdf's of the realized percentage real income change (y_i) for given expected income change category (p_i). 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. The same pattern tends to be found for the years separately.[12]

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

$$P\{y_i \leq m_{k-1} | p_i = k\} \leq \alpha \leq P\{y_i \leq m_k | p_i = k\}. \tag{4.10}$$

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 1 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

[12]In a few cases the monotonicity is violated by the extreme categories, probably due to the low number of observations. All figures are available from the authors upon request.

change category $(m_2, m_3]$ contains nonnegative changes only.[13] An explanation could be that α is less than 0.5.

Figure 1 : Estimated cumulative distribution functions of the realized percentage change in real income conditional on the expected income change category for the data pooled across years.

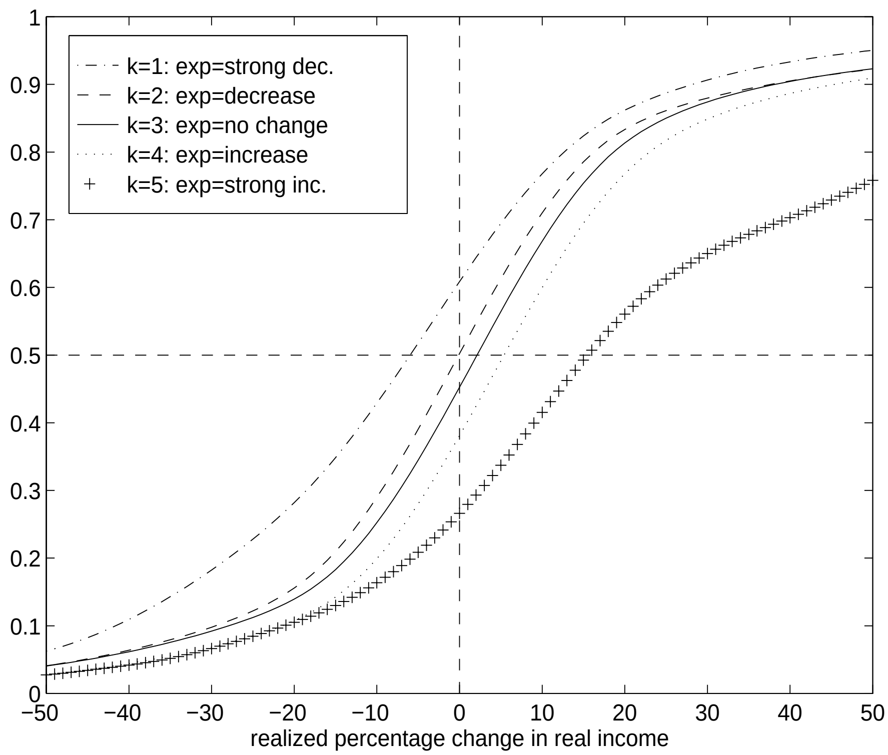

To make this more precise, we calculated [14] confidence intervals for $\xi_{\alpha,k}$ $(k = 1, ..., 5)$ for $\alpha = 0.5$ and $\alpha = 0.425$. Table 5 displays the results. Combining Table 5 with (4.10) leads to 95% one-sided confidence bands for m_k, under the best-case scenario and the α-quantile assumption. For example, $\alpha = 0.5$ implies $m_2 \geq -0.33$, and $m_3 \geq 1.31$. Thus $\alpha = 0.5$ does not allow that the no change interval $(m_2, m_3]$ is symmetric around zero. On the other hand, for $\alpha = 0.425$ we find $m_2 \geq -1.82$ and $m_3 \geq -0.18$, and symmetry is

[13] Working with nominal instead of real changes makes the asymmetry even stronger.

[14] We used the quantile regression in STATA and regressed the realized percentage change in real income on a constant. Standard errors are calculated by bootstrapping, with 1000 replicated bootstrap samples. See Gould (1992) for details.

possible. This suggests that respondents might use an asymmetric loss function. Unlike with the qualitative data, we now find that there are values of α that do not lead to evidence against the best-case scenario.

Table 5 : 90% confidence intervals for $\xi_{\alpha,k}$; pooled data set

	$\alpha = 0.50$		$\alpha = 0.425$	
	lower	upper	lower	upper
strong decrease	-6.38	-2.69	-10.7	-4.57
decrease	-0.33	0.14	-1.82	-1.11
no change	1.31	1.71	-0.18	0.14
increase	4.25	5.28	2.35	3.22
strong increase	9.17	18.5	5.89	14.8

5 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 of rational expectations and statistically independent realizations across individuals. 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, it appears that the inequalities to be tested are sensitive to the assumption on the location measure of the subjective distribution of the variable of interest reflected by the subjective prediction. We discussed three possibilities: the modal category, the median or α-quantile, and the mean assumption. The former two can be applied if comparable categorical data on predictions and outcomes are available, while the latter can only be applied if the actual outcome is measured as a continuous variable. The three assumptions lead to different bounds, none of them uniformly sharper than any of the others.

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 independent, due to common shocks. That this result obtains for a number of years reduces the plausibility of this explanation. A second interpretation is that people have asymmetric loss functions. We investigated this with more general α-quantile assumptions. 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. Using an alternative continuous measure of household income change, however, we concluded that values of α lower than 0.5 could be plausible. A third explanation is that substantial groups of households do not have rational expectations.

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

References

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

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

Curtin, R. (1982), Indicators of Consumer Behavior: The University of Michigan Surveys of Consumers, *Public Opinion Quarterly*, 46, pp. 340-352.

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

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

Das, M. and A. van Soest (1997), Expected and Realized Income Changes: Evidence from the Dutch Socio-Economic Panel, *Journal of Economic Behavior and Organization*, 32, pp. 137-154.

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

Dominitz, J. and C. Manski (1997), Using Expectations Data to Study Subjective Income Expectations, *Journal of the American Statistical Association*, forthcoming.

Gould, W.W. (1992), Quantile regression with bootstrapped standard errors, *Stata Technical Bulletin*, 9, pp. 19-21.

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

Guiso, L., T. Jappelli, and D. Terlizzese (1996), Income Risk, Borrowing Constraints and Portfolio Choice, *American Economic Review*, 86, pp. 158-172.

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

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

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

Kapteyn, A., P. Kooreman, and R. Willemse (1988), Some methodological issues in the implementation of subjective poverty definitions, *Journal of Human Resources*, 23, pp. 222-242.

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

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

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

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

Maddala, G.S. (1994), Survey Data on Expectations: What Have We Learnt?, in G.S. Maddala, ed., *Econometric Methods and Applications*, vol. 1, 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, pp. 934-940.

Manski, C. (1995), *Identification Problems in the Social Sciences*, Cambridge, MA: Harvard University Press.

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

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

Van der Klaauw, W. (1996), Expectations and Career Decisions: An Analysis of Teaching Careers Using Expectations Data, *Working Paper*, Department of Economics, New York University.