Lessons from the Spanish survey of household finances

Cristina Barceló and Olympia Bover[1]

1. Introduction

To assess the financial behaviour and the situation of different types of families aggregate levels are not enough and we need to know the distribution of real and financial assets of households, their debts, and their relationship with other variables. Tax records, even if available, do not contain information on many of the relevant variables and the only way to analyze many of these issues is to get information from surveys to households.

In Spain, the first of such surveys was carried out in 2002 by the Banco de España. In this paper we review the main challenges and features of the Spanish Survey of Household Finances (EFF) 2002. More details can be found in Bover (2004) and Barceló (2006).

The paper is organized as follows. Section 2 describes briefly the making of the questionnaire. In Section 3 the sample design is discussed, particularly how oversampling by wealth is achieved while preserving stringent tax confidentiality requirements. Section 4 presents for the EFF 2002 the problems of unit and item non response, usually faced in wealth surveys. Section 5 explains the need to provide imputations and motivates the imputation methods used. Section 6 describes the actual imputation work involved in the EFF. Finally, Section 7 concludes.

2. Questionnaire

Information is collected on: demographics, real assets and their associated debts, other debts, financial assets, pension plans and insurances, labour market situation and labour income (for all household members), non-labour income in previous calendar year, means of payments, and consumption and savings.

When designing the EFF questionnaire, the examples of wealth surveys questionnaires from other countries were important inputs that were adapted to suit the Spanish situation. One important consideration all along the making of the questionnaire was to try and keep the total length of the interview to an hour on average.

This survey is the only statistical source in Spain by which it is possible to relate incomes, assets, debts and consumption at the household level. Linking the EFF data with other data sources (e.g. register data) is ruled out by our pledge of anonymity to households which we feel is important in order to convince them to participate in such survey.

1 DG Economics, Research and Statistics, Bank of Spain.

3. Designing the sample

One distinctive characteristic of the EFF, following the example of the SCF in the US, is that there is oversampling of wealthy households. The distribution of wealth is heavily skewed and moreover some types of assets are held only by a small fraction of the population. Therefore, it was judged important to have a sample that would be not only representative of the population but also of aggregate wealth and that would also facilitate the study of financial behaviour at the top of the wealth distribution. This oversampling was achieved thanks to the collaboration of the Tax Office and the Statistics Office.

Basis for oversampling of the wealthy

In Spain there is a wealth tax ('Impuesto sobre el Patrimonio') and it is on the individual wealth tax files information that the EFF oversampling is based. People liable to the wealth tax in Spain were, in 1999 (which was the tax year used in selecting our sample), those with taxable wealth over 104,000 €. In 1999 around 980,000 individuals filed a wealth tax return. This corresponds approximately to 700.000 households, i.e. around 5% of the household population. We defined eight wealth strata which were oversampled progressively at higher rates.

Confidentiality guaranties

The Tax Office is subject to very stringent confidentiality requirements and cannot release, even to the Statistics Office, any personal tax information (not even in the form of intervals). To overcome the problem and enable wealth tax oversampling while preserving confidentiality, the National Tax Office volunteered to actually do the random sample selection herself following the sample design requirements, as instructed by the Bank of Spain and the National Statistics Office.

Thanks to the collaboration of both the Statistics Office and the Tax Office there is a unique population frame for the sampling. The population frame for the sample was the Continuous Municipal Census dated mid-2001, where the units are the households as defined by their address. With this information sent by the Statistics Office to the Tax Office, the Tax Office constructed for each address three variables based on information from both the wealth and the income tax. These data were the starting point for the sampling.

The first variable, the wealth stratum indicator, is based on total declared taxable wealth for the household, which was obtained by adding up the returns of all its members when applicable. The second one, for those filing income tax but not wealth tax, is a variable indicating to which quartile in the national taxable income distribution the household belongs. Finally, information on the per capita income of the household was also added. The income variables were helpful in the selection of sample replacements (as we shall see below), and to ensure that households from all income levels were selected into the sample. The latter was obtained by using systematic sampling with random start in a properly ordered data frame. Furthermore, the income quartile indicator was used to correct for non-response in large cities. The tax information available at the time was dated 1999. This entailed some limited mismatch between the two sources.

Sampling

The sampling design was different for the following three cases:

1. municipalities with more than 100,000 inhabitants. For large towns, the sampling was random within the eight wealth strata.

2. municipalities with 100,000 inhabitants or less. For small municipalities, the sampling was a two stage cluster design, with the primary sampling units (PSU or 'secciones censales') being selected first with probability proportional to their population. Further, within PSU the selection of households was different according to the number of wealth tax filers in the PSU.

3. Finally, in Navarre and the Basque Country where no oversampling of the wealthy was possible because the national Tax Office does not hold the personal tax file information for those regions, the sample was selected according to a two stage stratified cluster design with six strata defined according to municipality size.

Due to confidentiality reasons, stratum and cluster indicators cannot be provided. However, to calculate appropriate variance formulas replicate weights are provided instead.

Replacements

Another relevant aspect of the EFF sample design was the replacement scheme chosen. To try and preserve the oversampling scheme as much as possible, tightly controlled replacements were chosen. The use of controlled replacements is similar to post-stratification and weight adjustments done within cells when data collection is finished. An important advantage in our case for having controlled replacements was the fact that we do not have any indication of the wealth stratum to which the sample households belong so no 'directed' effort could be applied during the field work were we to discover that the response rate of certain strata was being particularly low.

In particular, up to four replacements were provided for each household originally in the sample that would serve as replacements of that household only. Those replacements were selected to be the two households immediately before and the two immediately after the household in a file ranked by income quartile (for non wealth tax filers), wealth stratum, and per capita household income. Replacements had to belong to the same income quartile (for non wealth tax payers) or the same wealth stratum as the sample household. This was done within municipalities in the case of large cities and within PSU in the case of small ones to keep replacements geographically not too distant from the original sample household. These implied that in some cases less than four replacements were available (and in a few instances, none at all). In the case of Navarre and the Basque country a more standard scheme of a pool of eight replacement households being potential substitutes for eight sample households (within the same PSU) was adopted.

Correcting for unit non-response and weights

To compensate for differential unit non-response, the sample weights are adjusted within the cells defined by the various sampling frame variables, including in particular wealth strata and income quartiles.

4. The fieldwork

Outsourcing the fieldwork

As it is usually done when Central Banks are responsible for wealth surveys to households, the Banco de España outsources the fieldwork for the EFF. The quality of potential fieldwork companies is a crucial factor for the good development of the survey. Unfortunately, in countries where major household surveys are conducted by the Central Statistics Offices, private fieldwork companies are mostly oriented towards opinion polls and marketing research.

Non-response

One of the characteristics of wealth and income surveys is high unit non-response due to the nature or the difficulty of the questions asked. The Banco de España was intensively involved in the efforts to reduce non-response, providing information to sample households and preparing written material.

Not possible to establish contact (never at home)

The number of households for which the interviewer was unable to find anybody at home (having confirmed with neighbours etc… that the address corresponds to the household) is very high despite at least five attempted visits (see Table 1). The number of these failed contacts as a proportion of the total number of attempted contacts by wealth strata has some non-random component as we can see in Table 2. Multiple residences was perceived as a potential reason for failing to establish contact with high wealth people during the field work.

Refusal

As we can see in Table 2, there is a clear non-random component in cooperation rates [defined as completed/(completed+refused)], decreasing as we move up the wealth strata, ranging from 53.6% to 29.4%. It is clear from this pattern that overall cooperation or response rates are not very informative in case of oversampling since they are dependant on the degree of such oversampling. For some meaningful comparison, we constructed cooperation rates by strata for the 1992 SCF. These cooperation rates for the list sample ranged from 52.6% for stratum 1 to 20.1% for stratum 7.

Table 1

Number of attempted contacts, by type of response

Completed	5143
Refused	5722
Never at home	6670
Out of scope	
(wrong address, not a housing unit, empty dwelling, deceased, others out of scope)	1797
Discarded after supervision	569
Total	19901

Supervision and discarded interviews

All the completed interviews were first revised by the field work agency supervisors. A large proportion of the completed interviews were re-contacted (mostly by phone but some personally). There were various reasons for re-contacting: (i) check potential inconsistencies, (ii) confirm all extreme values, and (iii) reduce item non-response.

A program was developed to detect logical inconsistencies between questions. Households sometimes provided a plausible explanation for them. For example in some cases the reference person in the household appeared as born after their main residence was bought because of having inherited that residence. However, in many cases this was useful to detect errors.

Table 2

Some measures of non-participation, by wealth strata

	Never at home[1]	Cooperation rate[2]
Total	33.5%	47.3%
Stratum 1	31.0%	53.6%
Stratum 2	38.9%	45.3%
Stratum 3	32.9%	44.7%
Stratum 4	35.5%	46.5%
Stratum 5	37.0%	38.5%
Stratum 6	38.0%	36.1%
Stratum 7	40.1%	37.8%
Stratum 8	39.8%	29.4%
Navarre and Basque Country	26.0%	46.0%

[1] Defined as (Never at home/Total attempted contacts) [2] Defined as [Completed/(Completed+Refused)]

Aside from the previous reasons, there was also extensive random re-contact to further control the work of the interviewers.

The EFF team at the Banco de España also examined the completed interviews for overall individual coherency. The process of validating the interviews is considered to be highly necessary to achieve a reliable dataset.

The degree of oversampling in the final sample

Finally, in what follows we give some figures about the degree of oversampling in our final sample. These were kindly provided by the Tax Office due to the confidentiality restrictions. Overall, slightly over 40% of the households that completed the interview correspond to wealth tax filers. Furthermore, aggregate tax returns information indicates that four per thousand of the population of households hold 40% of total taxable wealth. We would therefore expect to have at most 20 of such households in a 5,000 random sample, an upper bound since it assumes non-differential rate of response. In contrast, our sample contains over 500 of them.

5. Rationale behind imputation[2]

Item non-response

Item non-response occurs when a household agrees to participate in the survey but fails to respond to one or more questions. Together with high unit non-response, item non-response

[2] The references for this section (except the last part) are Little and Rubin (1987), Rubin (1987), and Schafer (1997).

is an inherent characteristic of wealth surveys. Moreover, they are closely related. Indeed, item non-response will partly depend on the stringency of the conditions imposed (in terms of the amount of important questions having to be completed) to declare valid an interview which in turn affects unit non-response rates. This is an issue one has often to address at early stages since it may affect the terms of the contract with the field agency. In particular, there is a trade-off because stringent conditions would give the right incentives to the interviewers but would produce self-selection into the sample in addition to the one created by overall refusals to participate. Moreover, faced with too stringent conditions the interviewers are more likely to cheat or to induce answers from the household.

Answers to the questions on whether the household holds a particular asset are usually readily provided. In contrast, households may have experienced more difficulties in providing information about the value of the asset held or about the amount of a particular income source. In Table 3 we present non-response rates to some key questions.

Why impute

Given the item non-response rates reported above, working with only the available cases ignoring item non-response would not be sensible. First, this would assume that the complete cases are a random subsample of the original sample. This is most probably not valid (as we have seen, for example, in the case of unit non-response), and therefore such an analysis could induce severe biases in the results. Second in multivariate analyses, working with only the observations for which all the variables of interest are completed would lead to far too small samples.

- *Imputation for enabling the analysis of the EFF with complete-data methods.* Correct inferences from an incomplete data set can be made using for example model based maximum likelihood methods. However, this is not technically available to all potential users of the data. Therefore, it is beneficial to provide users of the data with some imputation of (i.e. ‘filling in’) the missing data, which of course analysts are free to ignore[3]. Imputation is not meant to create artificial information or to give the impression that the data set contains more information than it actually has, but to exploit exhaustively the existing one in a way to enable the various possible analyses of the data using complete data tools.
- *Imputation as a responsibility of the data provider.* Imputation is a resources consuming process which is not at the disposal of most users and is sensibly thought to be the data provider’s responsibility [see Rubin (1996)]. An additional reason, very relevant in the case of the EFF, for the Banco de España to provide imputation is that we have access to some information (like some stratifying and location variables) relevant for imputing sensible values which will not be available in the public data file for confidentiality reasons.

[3] All imputed values are flagged accordingly.

Table 3

Reporting rates (%) of various items, unweighted sample

	Have item		Value for those having the item		
	Yes	**Unknown**	**Value**	**DK**	**NA**
Own main residence	84.5	0.0	86.5	13.0	0.5
Amount owed, 1st loan, main residence	15.0	0.0	88.6	11.2	0.3
Monthly payment, 1st loan, main residence	15.0	0.0	96.2	3.5	0.1
Rent main residence	9.9	0.0	97.4	1.0	1.6
Other real estate, 1st property	41.7	0.0	82.0	16.4	1.0
Amount owed, 1st loan, 1st other real estate	5.0	0.0	91.1	6.6	0.8
Accounts usable for payments	96.9	1.5	74.3	11.7	14.0
Accounts not usable for payments	20.8	2.2	81.8	6.5	11.8
Listed shares	20.7	0.3	76.6	15.9	7.4
Unlisted shares	6.9	0.2	51.3	34.6	14.2
Mutual funds, 1st fund	14.7	0.2	76.6	12.8	7.5
Fixed income securities	3.3	0.2	81.4	11.0	7.6
Pension plans, 1st plan	25.8	0.0	62.3	34.6	3.0
Life insurance (1st policy) coverage	8.9	0.0	63.9	33.5	2.6
Business market value (reference person)	13.1	0.0	64.3	32.3	3.4
Wage income (reference person, 2001)	36.9	0.0	97.6	1.2	1.3
Self-employment income (ref. person, 2001)	13.4	0.0	89.6	5.2	5.2
Unemployment benefits (ref. person, 2001)	1.5	0.0	94.7	5.3	0.0
Pensions (reference person, 2001)	31.8	0.0	99.2	0.2	0.6
Income from real assets (2001)	11.1	0.1	92.0	3.3	4.7
Income from dividends, coupons, etc (2001)	9.3	0.9	60.7	33.4	5.9
Bank accounts interest income (2001)	65.1	3.6	34.1	60.5	5.4
Food expenditure	100.0	0.0	93.8	5.8	0.4
Non-durable expenditure	100.0	0.0	95.9	3.6	0.5

Choice of imputation method

Before explaining our choice of imputation method we should say that they all rely on the *missing at random* (MAR) assumption [as defined in Rubin (1976) and in Little and Rubin (1987)]. This requires that the missing values behave like a random sample of all values but within groups defined by observed data. The goodness of this assumption will depend on the availability of observed variables which could plausibly explain missingness and conditional upon which the analysis can be conducted.

One of the central motivations for launching the EFF was to learn about the distribution of the real and financial assets of households, their debts and their relationship with other variables. To preserve the observed distribution of variables and the covariances between them, stochastic imputation methods should be used. Indeed, simple methods like mean imputation (conditional or unconditional) tend to produce peaked distributions of the variables and underestimation of the variances.

A very popular method of stochastic imputation is *hot deck*, with some variations. In general, with a hot deck procedure the missing item for a given household would be replaced by the value of the item reported by some similar-in-characteristics household. However, in the case of the EFF the number of characteristics/variables upon which one would like to condition, before being sensible to assume that the missing information is missing at random, is too large to produce reasonably sized cells from which to draw the hot deck imputation. Therefore, most of the EFF imputations, as we will see later, are based on random regression type models.

Finally, to take into account uncertainty about the imputation under the considered model and additional potential uncertainty when more than one model could be chosen for imputation, we provide multiple imputations (MI), as proposed by Rubin (1987).

Software used for imputation

We have been very fortunate to be allowed to use the programs written at the Board of Governors of the Federal Reserve System by Arthur Kennickell [see Kennickell (1991 and 1998)] for the SCF multiple imputation, as well as to benefit from his advice.

The multiple imputation procedure is based on the data augmentation algorithm and Markov chain Monte Carlo method and has an iterative and sequential structure [see Tanner and Wong (1987) and Schafer (1997)].

- *Iterative process:* The iterations of the imputation process are split into two steps. In the first step, missing data are imputed using the previous-iteration estimates of the parameters of the complete data distribution according to our imputation models. In the second step, we estimate the parameters of the imputation models using the observed data and the missing values previously imputed in the first step, in order to use these updated estimates for imputing missing data in the next iteration of the imputation process. Once the imputation process ends one iteration, another iteration starts repeating both steps until the convergence of this process.
- *Sequential process:* Within the same iteration of the imputation process, these two steps are repeated sequentially for imputing each one of the survey variables having missing information. The order in which the variables are imputed sequentially is not innocuous, mainly when we have missing information in the covariates of the imputation models. The imputed values are sequentially used to impute the subsequent variables with missing information. For this reason, in the EFF data we start imputing those variables not having a high fraction of missing information and those variables we consider to be very good predictors of the remaining variables to be imputed.

The programs for the SCF multiple imputation impute continuous variables stochastically using linear regression models. The imputation is not based on more complex models, since the linear regression models allow us to accommodate very easily a huge number of different patterns of item missingness across households. When we do not have available previously imputed values of the covariates with missing information, we need to impute values as if we implement different linear imputation models for each household depending on its non-missing covariates in the imputation model. Linear regression allows us to take advantage of reshaping easily and rapidly the variance and covariance matrices among the non-missing covariates and the variable to be imputed, depending on the particular pattern of item missingness in the covariates for each household with missing information on the variable to be imputed. In this way, we save enormous costs in terms of time and effort, since we must take into account a large number of covariates in the imputation models (from 100 to 200 depending on the sample size) and we must handle a huge number of different patterns of item missingness.

Binary variables are imputed using linear probability models in order to take advantage of the linear regression models for the reasons explained above. The imputation of multinomial variables is made using hot deck procedures.

Finally, the SCF imputation programs allow us to restrict the imputed values of missing data to one upper and one lower bound specific to each observation. These upper and lower bounds are constructed using the information provided by the EFF survey or previously imputed, whereby we can maintain consistency between the observed data and the imputed values of the missing information in the survey.

6. Imputation work in the EFF

This section describes how the actual imputation was carried out in the EFF, a process that has to be adapted to a large extent to each specific questionnaire and survey implementation.

Logical trees and shadow values of the EFF data

Before starting to impute the data, we have to create the flags of all variables and observations of the EFF data. These flags give information about whether or not the values provided have been answered by the households and also show the reason why the values are missing. In addition, these flags also indicate whether the existing missing values in variables are really *true missing values* or whether they have been imputed as “true missing” during the imputation process. The different values that these flags take for indicating the data origin and the reason for item missingness are called *shadow values*.

The task of constructing the flags and assigning the shadow values is carried out in two stages: in the first stage, we convert the different codes of “don’t know” and “no answer” responses (DK and NA responses) into missing values and assign their shadow values. In the second stage, we need to specify and program all the potential and logical relationships among the variables of the EFF questionnaire, so that we can assign the shadow values correctly to all observations and variables, mainly to those having either true missing values or item missing values derived from the household non-response to a previous related question.

The logical relationships that exist among the EFF variables are grouped in *logical trees* of variables; in each tree, one variable is the *head-variable* and the remaining ones are *branch-variables*. The household response (or non-response) to one head-variable affects both the values and the shadow values of the branch-variables, since the value of the head-variable may involve true or item missing values and may restrict the values of the branch-variables according to the design of the EFF questionnaire.

The flags are constructed before the imputation stage, since the shadow values are continuously used to impute the missing data mainly due to two reasons. First, we only impute the variables whose shadow value exceeds a certain threshold; and second, the imputation stage relies greatly on all the logical trees established among the variables of the survey, since the order in which the variables are imputed and the way in which the head-variable determines the values imputed subsequently to its branch-variables are based on the existing logical trees.

The meaning and the total number of different shadow values are specific to the EFF survey and depend greatly on both the survey characteristics and its implementation. The SAS programming for assigning the shadow values is facilitated by the fact that the interviews were made by CAPI and the original data were previously inspected. A small list of variables has not been imputed, due to the fact that the fraction of missing information exceeds the 60%, the number of respondents is very small to impute suitable values or due to the fact the households have not generally understood the question very well.

Covariates of the imputation models in the EFF

The goal of imputation is not to replace the missing data by their best predicted values, but to preserve the characteristics of the data distribution and the relationships among the different variables of the survey, in order not to bias the potential analyses made using different statistics (means, variances, correlations, percentiles, etc.). Thus, we need to include a large number of covariates in the imputation models that may be classified into four groups: the first group of covariates is formed by the determinants of the non-response, in order to satisfy the assumptions of "missing at random" and "ignorable missing data mechanism". Some key covariates in this first group are the following: total household income; random wealth strata indicators; regional indicators; age and education of the family head and the partner; and information provided by interviewers, such as the indicators of the type of building, neighbourhood, social status, house quality, the respondent's degree of understanding and sense of responsibility in answering, etc.

The second group of covariates is formed by the variables that are very good at predicting and explaining the imputed values. We usually include total household income; non-durable consumption; indicators of the different types of assets owned by the households; and the amounts of wealth held in the most common assets in which households generally invest, such as the owner-occupied house, other real estate properties, stocks, mutual funds and pension schemes.

The third group of covariates is formed by the variables that we expect to affect or explain the variable to be imputed according to different economic theories, in order to preserve the existing relationships between these variables and not to condition or bias the estimates made by the potential users of the data.

Finally, the fourth group of covariates is formed by determinants or good predictors of the covariates included in the rest of groups of variables. The role of this group is very important, since variables are imputed sequentially using the observed values and the values of the previously imputed variables within the iteration. Thus, we may have missing information on some key covariates and we need other variables that predict or capture the explanatory power of the missing covariates appropriately. Depending on the sample size available to impute the variable of interest, we try to include a set of key variables as large as suitable. Some of the covariates in this group are the following: characteristics of the household composition and structure (number of children, the children's age, the household head's civil status, number of adults, number of adults broken down by their labour market situation, etc.) and personal characteristics of the family head and the spouse or partner (age, education, labour history, current labour status, type of work done, economic activity, etc.).

Specifications of the imputation models

Continuous variables

To take non-linearities into account, regressors may be either formed by interactions between variables or introduced in logarithms or as polynomials. To impute euro questions that allow zero values as a response, we first impute a binary variable indicating whether the imputed value is zero or not. The imputation models of continuous variables are usually based on their logarithm.

When we impute percentages, we first impute a categorical variable indicating whether the percentage value corresponds to one of the probability mass points observed in a histogram or which is the range of values defined by these probability mass points in which the percentage lies. Afterwards, we impute the logarithm of the percentage restricting its value to the range previously imputed.

Finally, we sometimes specify the imputation model for other continuous variable highly related to the variable of interest; in this way, the model makes more economic sense and has a greater explanatory power.

Multinomial variables

The imputation is done by hot deck procedures using two discrete covariates or one discrete and other continuous variable. We usually include income or age as a continuous variable or two covariates being the result of interactions among some variables, such as indicators of the random wealth strata, the total household income quartiles, the family head's age bands and education, and other characteristics specific to the variable of interest.

Questions asked separately to each household member over 16, each particular asset within an asset type, each job, etc.

The way of imputing these variables is to construct a pooling of subsamples defined for each household member, for each job, etc. First, we generate the covariates of the imputation model separately for each household member, job, etc., and then we pool all these subsamples to estimate the parameters of the imputation model over the pooled sample. The imputation of these types of question is a very time-consuming process in SAS.

Constructed total household income variables

As total household income is a key covariate in the imputation models for all variables of the EFF survey, total income is one of the first imputed variables. We construct two total household income variables: one corresponds to the earnings obtained in 2001 and the other to the income received in the month in which the interview took place during 2002 or 2003. The total income variables are constructed as the sum of all sources of income, the property income from all household's asset holdings and the labour and non-labour income earned by all household members. The imputation models impute higher total household income values when we impute total income variables than those obtained when total income is generated as the sum of the different income sources imputed individually and separately. This may be due to the fact that we have richer information for imputing the constructed variables of total household income and due to the fact that we have no information about the ranges in which each income source may lie when the respondents neither know nor answer the exact amount of income.

Multiple responses to one question posed

The EFF survey contains questions that allow the households to make multiple responses. The way of imputing these questions are also to pool subsamples defined for each one of the household's multiple responses and to impute by hot deck procedures.

Evaluating the imputation of the EFF data

We implement two procedures specific to the imputation of the EFF data to evaluate the imputed values of continuous variables to ensure reasonable starting values in the first iteration of the imputation process and to evaluate the convergence of the imputed data across iterations, mainly due to the fact that a part of the imputed value comes from the randomisation. The implementation of these two evaluation procedures are explained in Bover (2004) and Barceló (2006).

7. Conclusions

In this paper we have described the main issues faced in conducting the first wave of the Spanish Survey of Household Finances. In particular, (i) we have reviewed the considerations that prevailed when drafting the questionnaire, (ii) we have described the design of the sample with a special emphasis on how the oversampling of the rich was achieved, and (iii) we have reported on survey and item non-response and explained the imputation work involved.

References

Barceló, C. (2006): Imputation of the 2002 Wave of the Spanish Survey of Household Finances (EFF), Occasional Paper No 0603, Banco de España.

Bover, O. (2004): The Spanish Survey of Household Finances (EFF): Description and Methods of the 2002 Wave, Occasional Paper No 0409, Banco de España.

Kennickell, A. (1991): 'Imputation of the 1989 Survey of Consumer Finances: Stochastic Relaxation and Multiple Imputation'.

——— (1998): 'Multiple Imputation in the Survey of Consumer Finances'.

Little, R. J. A. and D. B. Rubin (1987): *Statistical Analysis with Missing Data*, Wiley, New York.

Rubin, D. B. (1976): 'Inference and Missing Data', *Biometrika,* 63, pp. 581-592.

——— (1987): *Multiple Imputation for Nonresponse in Surveys,* Wiley.

——— (1996): 'Multiple Imputation After 18+ Years', *Journal of the American Statistical Association,* 91, pp. 473-489.

Schafer, J. L. (1997): *Analysis of Incomplete Multivariate Data,* Chapman & Hall, London.

Tanner, M. A. and W. H. Wong (1987): 'The Calculation of Posterior Distributions by Data Augmentation', *Journal of the American Statistical Association,* 82, pp. 528-540.