Modeling Wealth with Multiple Observations of Income: Redesign of the Sample for the 2001 Survey of Consumer Finances

Arthur B. Kennickell
Senior Economist and Project Director, Survey of Consumer Finances
Federal Reserve Board
Washington, DC 20551
v: 202-452-2247
f: 202-452-5295
e: arthur.kennickell@frb.gov
SCF web site: http://www.federalreserve.gov/pubs/oss/oss2/scfindex.html

August 2001
Joint Statistical Meeting, Atlanta, Georgia

Revised October 2001

The author is grateful to Ryan Bledsoe for research assistance. The work reported in this paper would not have been possible without the cooperation and efforts of a number of people at the Statistics of Income Division of the IRS, including Nick Greenia, Tom Petska, Michele Rhone, Mike Strudler, and particularly Barry Johnson, whose role was far too important to summarize adequately here. James Nunns at the Office of Tax Analysis at the Department of the Treasury provided essential support. The opinions expressed in this paper are those of the author alone and they do not necessarily reflect the views of the Federal Reserve Board, the Internal Revenue Service, or the Department of the Treasury.

This paper reports on the research that underlies the redesign of a key part of the sample for the Survey of Consumer Finances (SCF) for the 2001 wave of the survey. The paper also documents the implementation of the new design. As discussed in more detail below, this work builds on the history of sampling research for the survey dating back to the 1983 wave. In particular, the paper follows a line of research laid out in Kennickell (1999a) on the problems of modeling wealth in terms of income.

The sample for the SCF includes an oversample of relatively wealthy families. The stratified design that underlies this oversample is based on a mapping from observations of components of income to an estimate of wealth. The mapping is imperfect, and three factors seem particularly important: First, at any given time, rates of return that connect assets with capital income may vary widely across individuals depending on luck, information, and local economic conditions. Second, some assets, such as 401(k) accounts and residences, do not generate returns that are observable on a regular basis. Finally, temporally transitory factors, such as the timing of income receipts or unusually good luck, may cause the income that is observed in a given period to have a noisy relationship to the underlying assets that generate those returns. It is very difficult to address the first two concerns directly. However, the third can be addressed by using multiple observations of income to model wealth; an effort to do so is the principal focus here.

The first section of this paper gives a brief overview of the SCF and a general summary of the sample design. The next section provides motivation for a reexamination the sample design. The third section describes the implementation of the redesign of the sample for the 2001 SCF. A final section discusses areas for future research.

I. Overview of the SCF and Sample design

The SCF is undertaken every three years by the Federal Reserve Board (FRB) in cooperation with the Statistics of Income Division (SOI) of the Internal Revenue Service, primarily as a study of household wealth and use of financial services.[1] To this end, the subject matter of the questionnaire focuses on detailed components of assets and liabilities along with

[1]See Kennickell *et al.* (2000) for an overview of the data and Kennickell (2000a) for a detailed description of the survey and its methodology.

items that help in understanding the financial decisions households make and that describe the institutional structure that underlies each household's balance sheet.

Because household wealth in the U.S. is highly concentrated (Kennickell, 2000b), a very large sample would be required to make reliable wealth estimates without some type of oversampling by wealth. At the same time, the problem of adequate representation of wealth is amplified by the nonrandom nature of nonresponse to the survey. Although many factors enter into nonresponse in the SCF (Kennickell, 1999b), it is very clear that wealthier respondents are less likely to participate. Their lower participation rate may be driven by greater sensitivity about privacy issues, greater perceived value of the time required for the interview, greater difficulty interviewers face in getting beyond gatekeepers in order to request participation, or other factors. The net result is that without some means of addressing the nonrandom nonresponse, many of the survey estimates would be severely biased.

The SCF sample is a dual-frame design. One part is a standard multi-stage national area-probability (AP) design (Tourangeau et al., 1993), which provides good coverage of the general population. This sample is selected by the National Opinion Research Center at the University of Chicago, the contractor for data collection since 1992.

The second part is a list sample, which is selected by FRB staff from statistical records derived from tax returns. These records are provided by SOI to the FRB under strict controls on the use of the data.[2] This file is a sample from the full set of tax returns filed in a given year, and it is specially edited using procedures designed to yield data to support research at the Office of Tax Analysis and the Joint Economic Committee of the Congress (see IRS, 2001). Beginning with 1989, the SOI data used for the SCF sample have come from the tax year two years before

[2]Proper treatment of data to minimize the intrusion on the privacy of respondents is a very high priority. This treatment is regulated formally by contracts between the FRB and SOI, FRB and NORC, and NORC and SOI, and by the provisions of IRC 6103. FRB staff select the list sample from a file containing no names, and they pass the SSNs of the cases selected to SOI along with a temporary identification number. SOI retrieves the names and addresses corresponding to the SSNs, and passes that information to NORC. NORC provides the FRB with the answers to the survey questions along with the original temporary ID number for each case. In no case is the FRB allowed to know the name of any survey participants. SOI is never given the link between the survey answers and the final ID numbers; indeed, SOI encourages FRB staff to take rigorous measures to minimize the likelihood that anyone at SOI or elsewhere might be able to use the survey answers to guess the identity of the participants. NORC never receives any financial data from tax returns.

the date of the survey.[3] As discussed in more detail below, the list sample is stratified using a "wealth index" computed using income data in order to predict a rank ordering of people by wealth. This stratification allows both oversampling of wealthy families and targeted nonresponse adjustment. The AP and list samples are pooled using weights that are designed to respect the relative strengths of each part of the sample (Kennickell and Woodburn, 1999 and Kennickell 1999c).

Because the wealth index used in the list sample stratification has such a powerful effect on the estimation efficiency of the SCF, refinement of the index has always been an important part of the methodological research supporting the survey.[4] The list sample for the 1983 survey, the first of the series, was stratified by income categories created using a set of rules that have not been cleared for public release. The 1989 SCF was the first to use an explicit wealth index approach, and the construction of that index is discussed in Heeringa *et al.* (1994). In its pure form, this general type of index for a given time might be expressed as

$$WINDEX0_i = \sum_{\mathbf{j}} \frac{1}{r_{ij}} \mathbf{Y}_{ij}$$

where "i" indexes individuals, "j" indexes components of capital income, and r_{ij} is the rate of return associated with capital income component Y_{ij}. Thus, if a person had interest income of $100 and the rate of return were 5 percent, then the contribution of this item to the wealth index would be $2,000.

One problem in implementing such an index is that no direct information is available *a priori* that would allow the use of person-specific rates of return. Thus, as WINDEX0 has been implemented, a common rate of return is assumed for all individuals for each income type. Another problem is that not all assets generate regularly observable returns.[5] For example, income associated with owned personal residences usually only appears when such properties are sold; returns on 401(k) accounts, IRAs, and Keogh accounts only appear as income when

[3]The "tax year" is taken to indicate the year in which the income reported on a return was earned.
[4]This history is summarized in Kennickell (1999a) and more detail on the construction of the wealth index is also given there.
[5]Kennickell and McManus (1993) explore both of these issues in more detail.

funds are withdrawn from such accounts. Some items may only appear as entries on estate tax returns.

An index of this broad form was used in selecting the list samples for the 1989 and 1992 surveys, taking an average rate of return for all cases for each of the different sources of capital income. As implemented, the index differed further from the ideal by the inclusion of a crude estimate of housing wealth, the use of the absolute values of negative income components, and the inclusion of the amount of realized capital gains with a capitalization rate of one. The latter two adjustments were made to reduce the possibility of serious misclassifications.

Beginning with the 1995 SCF, the wealth index used for stratification was altered to include a component calculated from a model estimated from a regression of actual wealth observed in the previous survey on its corresponding original frame data.[6] This model, referred to here as "WINDEX1", offers a flexible way of accounting for systematic patterns in the structure of various types of income–including non-capital income–on wealth. To the extent that the structure of asset returns can be taken to vary systematically with income, deductions, region, filing characteristics etc., there is the prospect that such modeling might ameliorate problems associated with differential rates of return and unobserved asset returns.

Of course, it is almost certain that the WINDEX1 model does not fully capture idiosyncratic variations in returns and in holdings of assets with unobserved returns. Morever, there is a risk that the estimated coefficients may be too dependent on the distributions of rates of return in the period for which the model is estimated. For example, suppose the model were the same as the WINDEX0 model, so that (ignoring the effects of omitted structural variables, which would cause the included variables to proxy for the omitted variables) the regression coefficients were actually estimates of the average rates of return. In this case, if the average dividend yield rose temporarily by half while other returns remained the same, then the model would over estimate the wealth of people who own stocks and underestimate that of others. If the model components were actually the same as the WINDEX0 model, then it might be possible to make a correction, perhaps by rescaling all income flows to the same level as in the base period. However since what power there is in this approach lies in a the use of a broader set of variables, and since it is not known just what parts of the model proxy for omitted variables, it is not

[6]The model is proposed in Frankel and Kennickell (1995).

feasible to adjust the model in a simple way for changes in rates of return since the period used for estimation.

To hedge against the possibility of misclassification from the WINDEX1 model, the 1995 and later surveys have pooled estimates of WINDEX0 and WINDEX1; that pooled estimate is referred to here as "WINDEXM".[7] As discussed in Kennickell and Woodburn (1999), the combined approach offers an improvement over what would have been obtained from either model alone.

Figure 1 shows the distribution of the pooled wealth index by the unweighted deciles of the net worth of the portion of the SCF list sample that was interviewed in 1998.[8] In the ideal, the plots would show strong clustering of the mass of each distribution along the diagonal from the lower left-hand corner to the upper right-hand corner. Although it is clear that WINDEXM does do a reasonably good job of discriminating between very high wealth and very low wealth, across the middle of the distribution the power of the model is lower. It is clear that there is room for improvement.[9]

[7]To put the two indices on a common scale, they are adjusted to have the same population median and interquartile range.

[8]The list sample is strongly weighted toward the upper end of the wealth distribution, but the exact proportions are withheld for reasons related to disclosure limitation.

[9]To some extent, the imprecision in the relationship may reflect misreporting: respondents may intentionally or inadvertently misstate their conditions, and interviewers may make errors in administering the interview and recording the respondent's answers. Despite extensive examination of the data and monitoring of interviewers, it is likely that some such defects remain. Some errors may even be the result of interviewing the wrong person.

Figure 1: Distribution of WINDEXM by unweighted deciles of net worth, 1998 SCF.

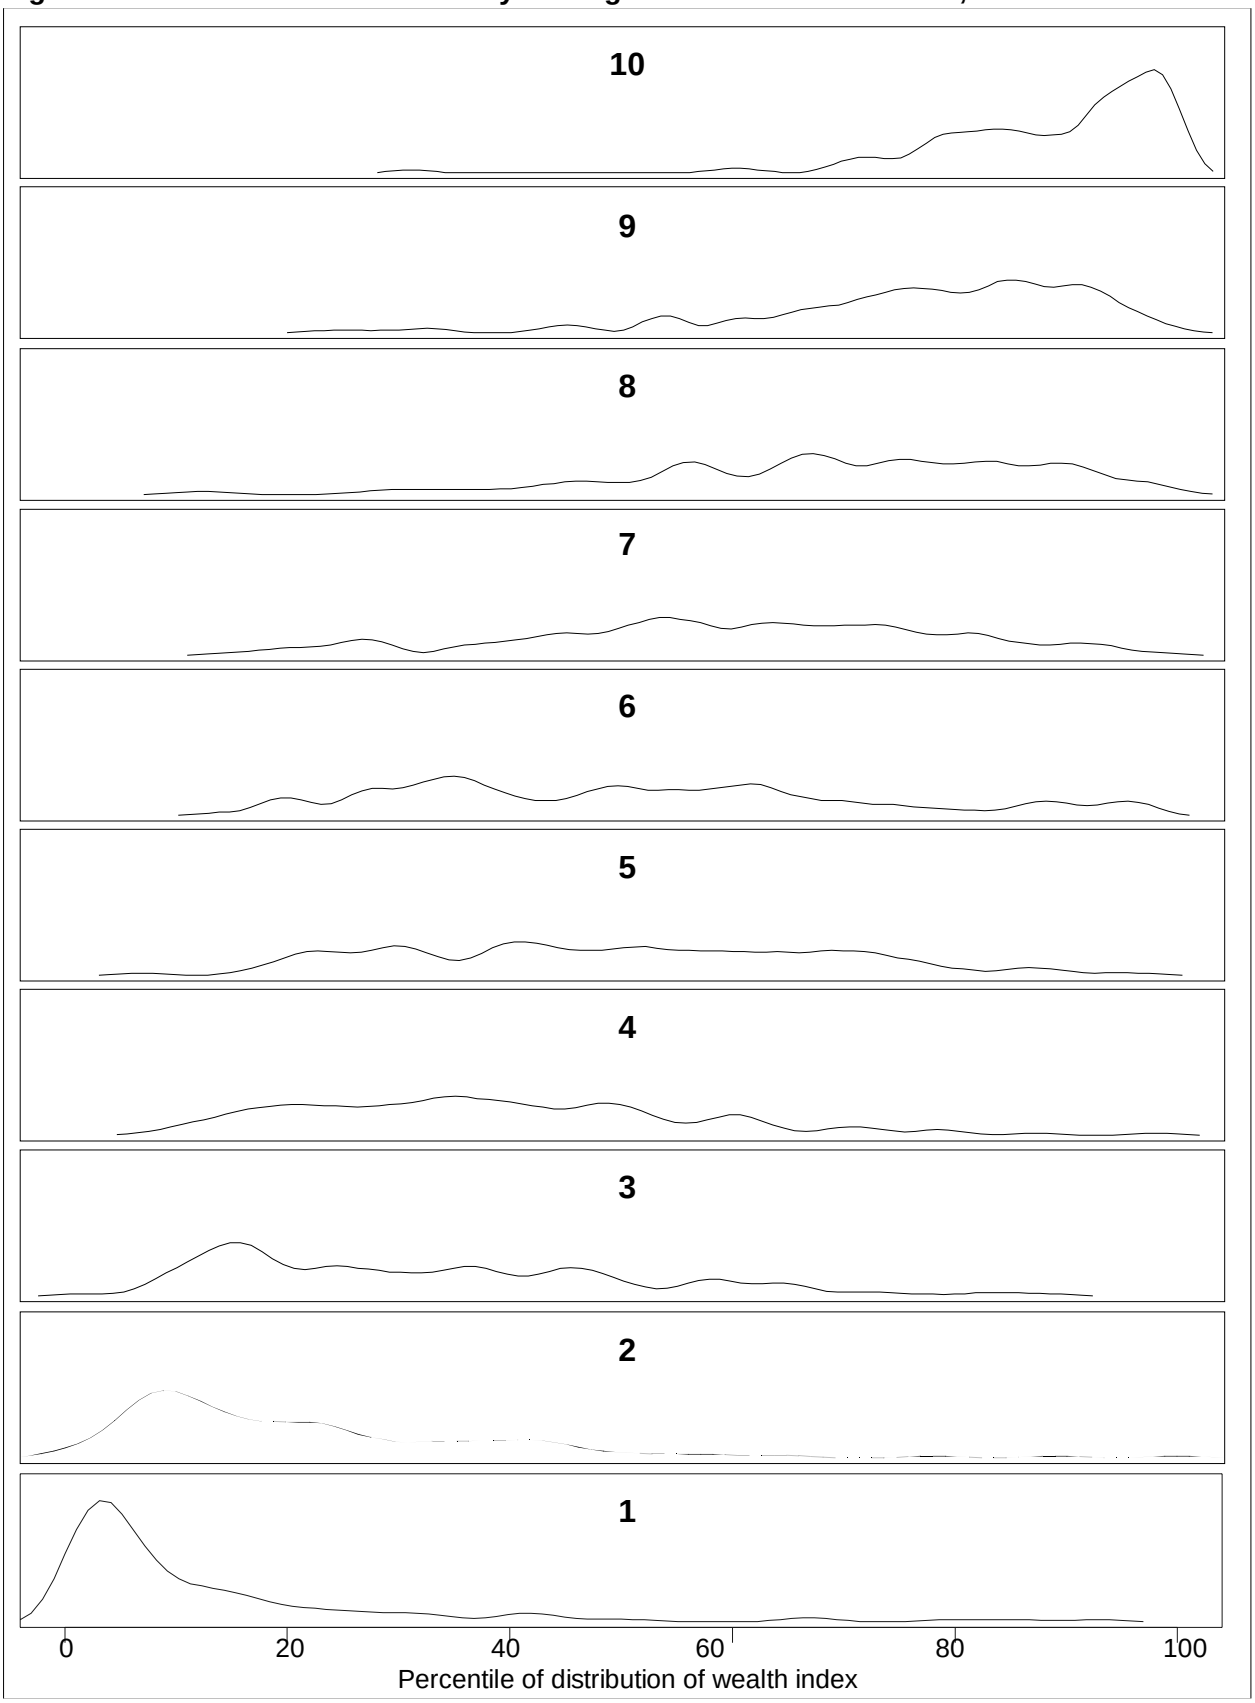

II. Motivation for redesign of the list sample

It is unlikely that any other source of data currently available for a broad range of the population could be used to index wealth better than is possible using the SOI data. Unfortunately, despite the efforts described above, the wealth model based on SOI data remains notably noisy. Of course, the model is only a sketch of the deeper structural model that would be needed to explain all the connections between wealth and income, which should include modeling of portfolio decisions, employment choices, other planning decisions, effects of the larger economy, etc. Idiosyncratic factors are also very likely to be important. Although the SOI data offer only limited scope for identifying additional structural and idiosyncratic factors directly, there is still a prospect that better proxies might yield a more reliable model for indexing wealth.

One dimension of the data that has previously not been exploited for the SCF sample is time. Income may vary over time for a number of reasons. Some income variations may parallel changes in the valuation of the underlying assets that generate the incomes. Using cross-sectional data for the sample from two tax years before the survey poses a risk of over- or under-stating wealth at the time of the survey. But if positive and negative shifts are equally likely, then pooling multiple years of income would damp such shifts without introducing net bias in estimated wealth rankings.

Other changes in income may have more limited connections to wealth. Some people with substantial resources may be able to time the receipt of their incomes to correspond to particular needs or for an advantage in their taxes. Others may have assets that have a relatively spikey payout profile–for example, a start-up investment. Many people are affected by changes in the larger economy, through their employment and wages, their interest earnings, and other such factors, and these events may affect both asset holdings and returns of households differently over time. Averaging over observations of the same unit over time would have the effect of smoothing out such transitory fluctuations in income.

The idea of income pooling for the sample has been raised many times (e.g., Frankel and Kennickell, 1995). Although extensive time series of data are available in principle, the practical obstacles to obtaining the necessary data for sampling are large. First, access to tax-based information is very limited. Although the FRB has a contract with SOI that allows the use of

IRS data, every instance of use must be separately justified. Given the abstract potential for abuse, it would be difficult to argue that access to tax-based data should be substantially easier. Arranging for access to even the types of records used previously for the SCF requires substantial time on the part of many people who are already busy with regular tasks. Second, it also requires time to process the data needed for the sample. Individuals mostly file tax returns by April following the tax year, but some people, most often those with particularly complex tax returns, may request an extension and file later. The SOI file that serves as the basis of the SCF list sample requires substantial editing. Some of this editing corrects errors and reclassifies the reported data in a way that is more appropriate for tax research. Although "advance" data for a given tax year are typically available by around December of the year following the tax year, the final data have never been available for the SCF sooner than April two years after the tax year. Thus, the SCF sample can be based on no more recent data than such a file. Third, the SOI file is itself a sample, and observations are not necessarily retained from year to year. Thus, linking individual records would require access to IRS master file data, an even more restricted source of information than the SOI file.[10]

For the design of the 2001 SCF list sample, extraordinary efforts were made by SOI staff to make it possible to use both the type of SOI data file used in the past, as well as IRS master file data. Four sets of data were provided: master file data for tax years 1996 and 1997 for cases corresponding to the full original list sample in the 1998 SCF, the edited SOI file for tax year 1999, and a set of master file data for tax year 1997 for observations in the full 1999 SOI file. The 1999 SOI file and the 1997 master file data are the most recent files of these types that are available. In order to gauge the costs and benefits of using these data sources together, it is important to assess the nature of the SOI editing process as well as the variability of income measures over time. To this end, the master file and SOI data for 1997 may be compared to assess the effects of mixing data from the two files, and the 1996 SOI data may be compared with the 1997 master file data to assess the variability of income over time.

Figures 2a-2f show various statistics on the variation in measurements induced by the editing of the 1996 SOI data, the basis of the 1998 SCF list sample, conditioned on WINDEX0 computed from the SOI data. In each plot, the solid line indicates the percent of tax filers who

[10]The data set referred to here as the "master file" is known within the IRS as the "returns transactions file" (RTF).

had a change in the value of the item as a result of editing, the dashed line indicates the mean percent change relative to the SOI value, and the dotted line shows the corresponding standard deviation of the percentage change.[11]

As shown in figure 2a, the distribution of adjusted gross income (AGI) is different for only a small fraction of observations away from the very top of the WINDEX0 distribution, and the mean and standard deviation of the change is small across the distribution. At the same time, variations in the percent of observations with changes in the components of AGI that are shown (figures 2b-2e) are substantially larger. With increasing values of WINDEX0, wage and salary income shows an increasingly greater value in the SOI data, largely reflecting what appears to be a systematic reclassification of S-corporation and farm income and Schedule C income to wages and salaries.

Figure 2f shows that the net effect of editing changes on WINDEX0 is to induce a large number of relatively small changes across the distribution of WINDEX0 computed with the SOI data. The standard deviation of the percent change is roughly 10 percent across the range of WINDEX0. However, the critical question is how large these changes are relative to the changes in the data over time.

Figures 3a-3f address the magnitude of the changes over time in the same variables shown in figures 2a-2f.[12] The dotted lines in figures 3a-3f indicate the percent of cases for which the value of the item was not zero in both 1996 and 1997; the solid lines show the mean percent change from 1996 to 1997; and the dashed lines gives the standard deviations of the percentage changes. As in the earlier figures, these values are given conditional on WINDEX0 as calculated

[11]The sample used as the basis for this plot was a subset of the 5,642 observations selected for the 1998 SCF. Of those original cases, 5,452 were successfully matched to the master file data. By constructions, every observation in the SOI data must have a corresponding observation in the master file data. However, there may be errors in the key variables used for matching the data, and some observations that were actually present may have been missed by the complex sequence of sorting and comparisons needed to combine the two sources of data. The statistics shown are unweighted, but the conditioning on WINDEX0 should produce results that are approximately equivalent to the proper weighting. To avoid extreme distortions as a result of outliers and to deal with cases where the base value was zero, the absolute value of the percent changes was bounded at 200 percent. Truncations at a lower point have an increasingly greater effect, but a somewhat larger limit has little effect. The results shown are the result of a type of kernel estimate.

[12]This comparison excludes cases where the filing status (e.g., “married filing separately”) of the underlying tax return was substantively different in the two years.

from the 1996 SOI data, and all values were computed using the same data and assumptions with an additional adjustment to put the changes into constant-dollar terms.[13]

The difference in AGI over the period is substantially greater across the entire range of WINDEX0 than was the case for the corresponding comparison of SOI and master file data. By income components, the largest mean percent differences overall are in interest and dividends; the mean difference there is roughly 10 percent on average, and the standard deviation is about 70 percent. Mean differences in wage and salaries, S-corporation and farm income, and Schedule C income are fairly small, but the standard deviations are quite large. The net effect on WINDEX0 is relatively small changes on average, but a standard deviation of about 70 percent overall. The variation in WINDEX0 is clearly much larger than what one would expect purely from the reclassification error seen in figure 2f. Thus, the data suggest that there may be more to be gained in sampling efficiency by accounting for intertemporal variability in income than would be lost from the noise introduced by using edited and unedited files together.

For completeness, appendix figures A1a-4f show results comparable to those in figures 3a-3f, for changes between the 1997 master file data and the 1999 SOI data, the files available for the 2001 SCF list sample design. As one might expect from the longer time period, the indications of change are greater than in the comparisons of the 1997 SOI data and the 1999 master file data.

[13]The rate of return assumptions used in computing WINDEX0 were also updated to appropriate values for 1997. However, these changes were fairly small.

Figure 2: Differences between 1996 SOI and master file data; percent having any difference, mean percent change, standard deviation of percent change; by WINDEX0; 1998 SCF sample.

Figure 2a: Adjusted gross income.

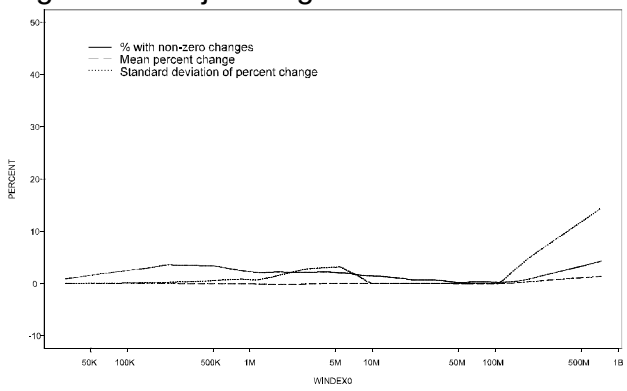

Figure 2b: Wages and salaries.

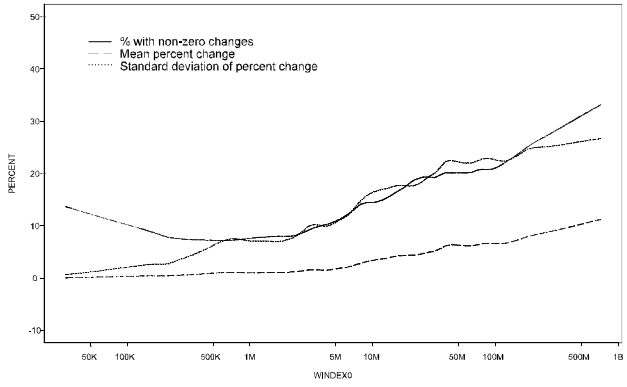

Figure 2c: Interest and dividends.

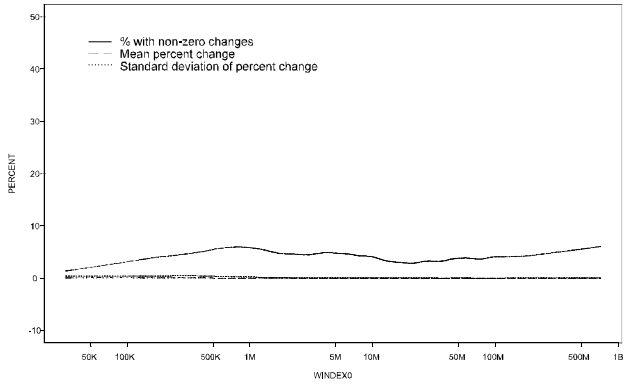

Figure 2d: S-corp. and farm income.

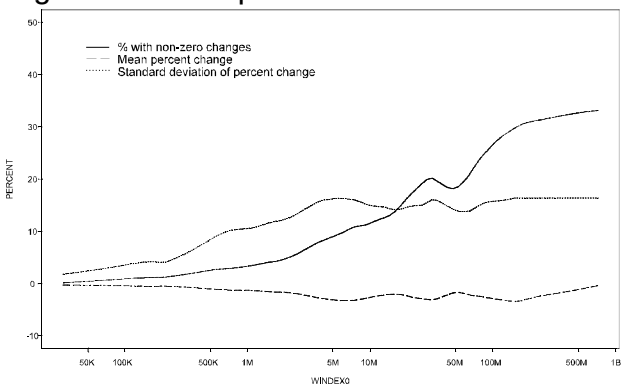

Figure 2e: Schedule C income.

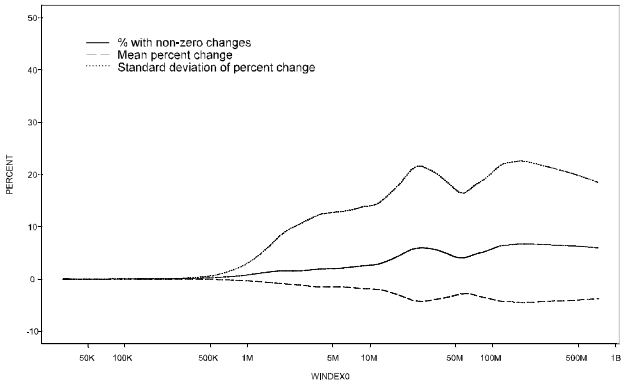

Figure 2f: WINDEX0.

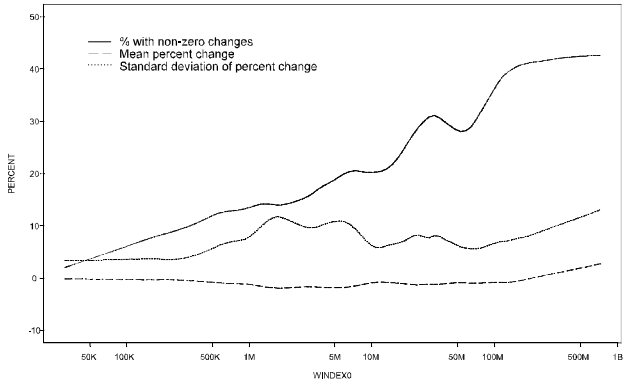

Figure 3: Differences between 1997 master file data and 1996 SOI data; percent having any item in either data, mean percent change (in constant dollars), standard deviation of percent change; by WINDEX0; 1998 SCF sample.

Figure 3a: Adjusted gross income.

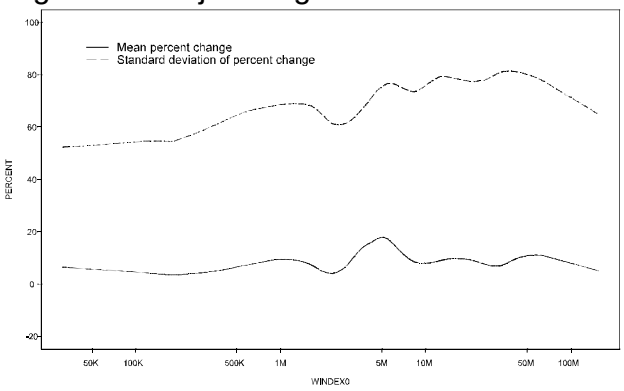

Figure 3b: Wages and salaries.

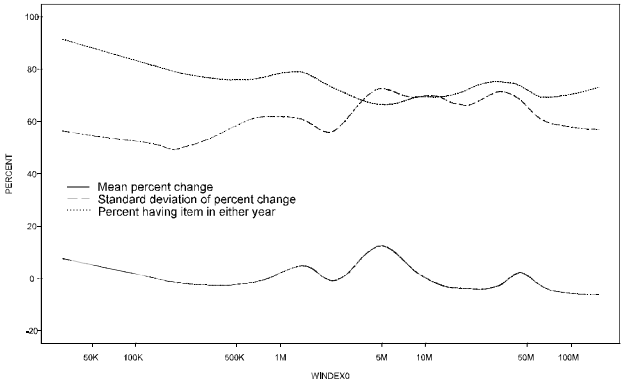

Figure 3c: Interest and dividends.

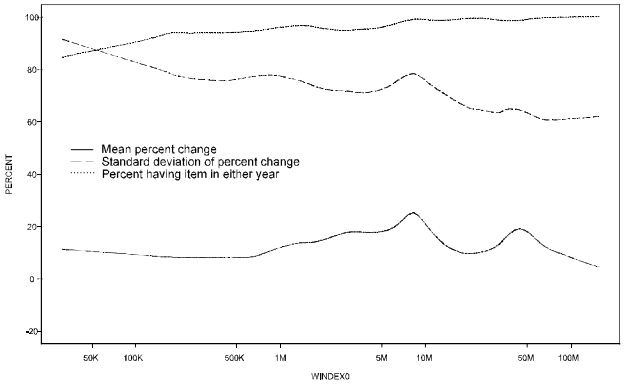

Figure 3d: S-corp. and farm income.

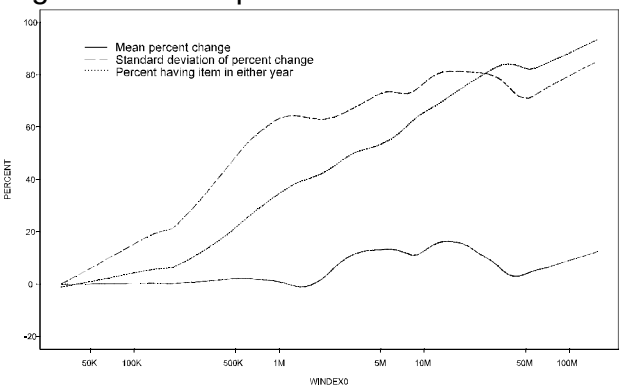

Figure 3e: Schedule C income.

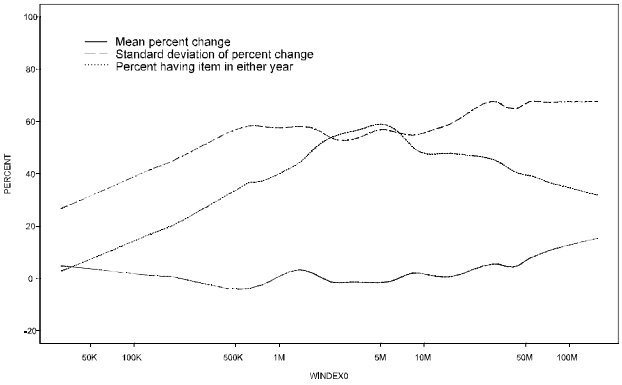

Figure 3f: WINDEX0.

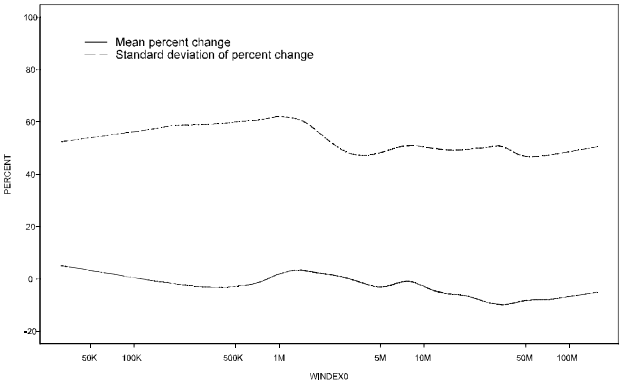

III. Implementation of the redesign for the 2001 SCF

The data examined suggest that there is a potential gain in sampling efficiency to be had by using the 1997 master file data together with the 1999 SOI data for the design of the 2001 SCF sample. There are two obvious classes of strategies for using these data.

First, one might use all of the variables in the two observations separately to model wealth. If there were a sufficient basis for modeling the mechanism that underlies the observed income variability, the increased number of conditioning variables might yield an improvement in the ability to classify observations in terms of their wealth. But variability in income reported on tax returns is probably mainly a function of two factors: real events ("luck," unemployment, etc.) and tax considerations. The available tax-based data offer little systematic help in understanding the real factors, and aggregate data from other sources are not useful in understanding individual variations. There is a rich literature that attempts to understand how people manage their finances in response to tax laws, but application of the insights of that literature to the current problem would involve a high degree of speculation. Still, even a non-structural model, if sufficiently agnostic in its construction, might yield improvement from the increased degrees of freedom. However, the practical constraint imposed by the very limited time available for the actual selection of the 2001 sample after the data became available made it impossible to consider this option seriously. Moreover, there is a good argument to be made for altering so critical an element of the SCF in well-defined steps.

The second type of strategy is to use the information from the two periods solely for the purpose of smoothing estimates. While this is a more conservative approach in some ways, it is also probably more immune to modeling errors. Motivated both by caution and the force of time pressures, a version of this strategy was applied to create a smoothed estimate of WINDEXM for the 2001 list sample, following as closely as possible the construction of the previous samples (Kennickell, 1998).

The population structure implied by the 1999 SOI file was taken as the point of reference for the sample design. Thus, it was necessary to create a file of matched data from the 1997 master file and for each observation in the 1999 SOI file.[14] However, some people did not file a

[14]The SOI file was reduced in several ways. First, observations filed from other than an address in the 50 U.S. states were removed. Second, where there were duplicate observations for the same taxpayer (mostly amended returns), only the most recent was retained. Third, observations

tax return in both years.[15] In addition, to avoid spurious variability in income measures, it was necessary to exclude cases from the match observations where there was a difference in tax filing status in the two years. Of the 172,852 observations in the 1999 SOI file, 149,148 could be matched with 1997 master file data.[16]

For every observation in the final data file–both matched and unmatched 1999 records–it was necessary to produce an estimate of WINDEXM. Using appropriately updated coefficients, estimates of WINDEX0 were made for every observation for each available year.[17] For the set of observations where both years of data were available, the 1997 estimate of WINDEX0 was adjusted to have the same mean as the 1999 estimate of WINDEX0 for the same set of observations. The smoothed estimate of WINDEX0 was computed as the simple average of the two estimates where both were available, and the 1999 estimate for the other cases. For WINDEX1, the procedure was similar.[18] Following past practice, the blended estimate of WINDEXM is taken as the average of WINDEX0 and WINDEX1 where these estimates were adjusted to have the same median and inter-quartile range.

The resulting index was used to array the sample across strata defined in terms of weighted percentiles of the distribution of WINDEXM. Although the exact breaks in this distribution cannot be revealed publicly, the lowest stratum (1) accounts for about 75 percent of

---

where the principal filer was younger than 19 were excluded; experience from earlier surveys suggests that including taxpayers as young as 18 is most likely to select individuals who were dependent members of a larger household. Finally, in cases where there were taxpayers who lived in the same 9-digit ZIP code area who also had the first four letters of their last names in common, only the cases with the oldest principal filer was retained. The point of this last deletion was to remove duplicate tax payers who are very likely to be members of the same household; an adjustment was made in the estimate of the household wealth index and weight to account for the deleted observations.

[15]In some cases, married people may change the order of their identities on a tax return. The match attempts to account for such reversals in the two years. The failure of some matches may be a result of filing extensions, clerical errors, and other errors.

[16]It might also be desirable to implement any mechanical rules used in creating the SOI files from master file data. Unfortunately, such rules were not available when this work was done. For future work, it is hoped that such adjustments can be made.

[17]Following past practice, for observations based on returns where the taxpayer filed separately from a spouse, the level of WINDEXM was doubled.

[18]The model coefficients were re-estimated using a match of 1998 SCF wealth data with the 1997 SOI data used for the original selection of the 1998 list sample.

the population of tax filers, and the five highest strata together account for about two percent of the population considered.[19]

The results of this stage of the sample construction were compared extensively under the new and old designs. This comparison is motivated principally by two concerns. First, radical differences would be likely to be a result of an error. Second, a central motivation for the redesign is to improve the classification of units across the middle of the distribution into wealth groups, so there should be notable differences in that region to justify the difficulty and risk of implementing a new procedure.

When the strata were computed using both methods, fewer than 0.15 percent of all observations were assigned to a stratum computed under the old method that differed by more than one from that implied by the new method. For each stratum, figure 5 shows the proportion of observations assigned values under the new method above and below the value under the old method. The tall central bars represent the percent assigned to the same stratum under both methods; the bars to the left of those represent cases assigned to a lower stratum under the old method, and the bars to the right represent cases assigned to a higher stratum under the old method. The figure shows some differences in classification, but relatively small ones. The largest difference occurs in stratum 3, where about 15 percent of cases are assigned a different stratum under the old method. For stratum 1, only 2 percent of cases have a different classification under the old method, while for the top stratum the corresponding figure is about 5 percent. The classification based on the blended estimate of WINDEXM uses more information than that based on the unblended estimate, and the underlying model framework is the same. Consequently, there is no reason to expect the blended estimate to be worse, and there is some *a priori* reason to think that it might be better.

After defining the stratifying variable in terms of the whole population, the data file was reduced for the actual selection to include only cases that filed returns from a county included in the 100 PSUs underlying the AP sample. An implicit substratification was imposed by sorting the resulting file by age and a measure of financial income. Within each stratum cases were selected using systematic selection from a random starting point. Cases were oversampled by an a progressively larger proportion with increasing strata.

[19]An appendix which is not available with the public version of this paper specifies the exact stratum definitions and the number of cases selected in each stratum.

The social security numbers of the selected respondents were passed to SOI for retrieval of names and addresses. SOI staff reviewed the resulting list in order to remove members of the Forbes 400 and other people who might have such high national prominence that it would be very difficult to protect their identity in the final data set without overly compromising the data they would provide. The final list of names and addresses was passed to NORC, which then provided the FRB project staff with an ID number to link back to the original sample file for the narrowly circumscribed purposes related to the statistical processing of the survey. The FRB does not receive any name or address information, NORC does not receive any tax data, and SOI does not receive the information that would allow it to match the survey participants directly with tax data. By agreement with SOI, the members of the list sample are given an opportunity to refuse participation in the survey before they are contacted by an interviewer.

Figure 5: Distribution of unblended stratum assignment around blended stratum assignment.

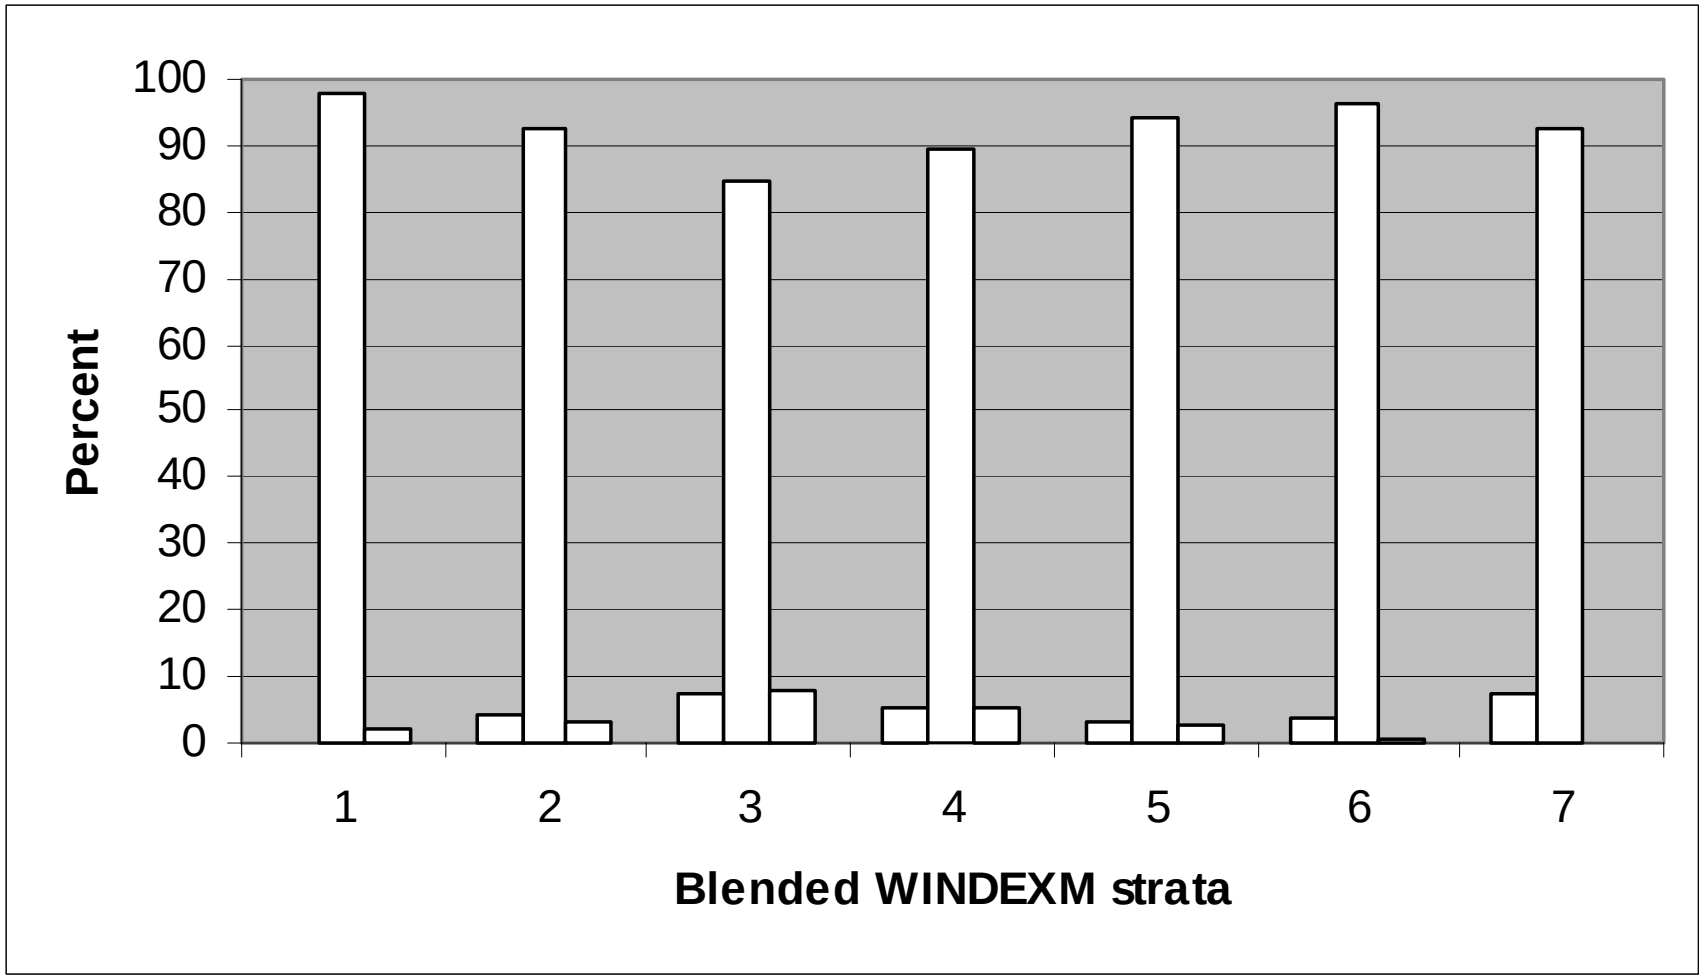

IV. Future research

Obviously, the hope is that the design used for the 2001 SCF list sample will improve the efficiency of the list sample in classifying households in terms of their actual wealth. An immediate priority after the completion of the processing of the data will be an evaluation of the classification of cases under the old and new regimes. If a significant improvement can be detected, it may justify the large effort to gain permission to incorporate a longer series of income data in the estimation of the principal stratifier for future surveys.

Had the time for section of the 2001 sample permitted, it would have been possible to explore more fully the possibility of using the 1999 and 1997 data as additional elements in the model of the wealth index, rather than simply using the data to compute a smoothed version of WINDEXM. Such modeling will also be explored using the 2001 data.

From looking in detail at the past list sample data by stratum and the 2001 data as they come in from the field, it is very clear that some of the *ex post* misclassifications are due to defects in the design of the wealth index. However, not all appearances of misclassification originate from this source. First, although strenuous efforts are made to ensure that the person actually interviewed is the same as the one selected, there have been failures in the past. Because it appears that a significant fraction of tax return are filed from addresses other than the taxpayers' home addresses–for example, the address of an attorney or accountant who may work from the selected respondent, among others–it is frequently necessary to use other means to locate the respondent. Despite strong quality control procedures, it is clear from the administrative records of past surveys that such efforts have occasionally failed to identify the correct person; when detected, such cases have been deleted. For the 2001 survey, more systematic advance verification of home addresses should reduce the scope of such locating errors. Second, sometimes interviewers have become confused under the pressure of their work, and they have interviewed someone other than the designated respondent–most often, it seems, under a misapprehension of the correct use of proxies; proxies are allowed on the SCF, but they must provide data as if they were speaking for the designated respondent, not for themselves. Despite additional training of interviewers on this subject for the 2001 survey, already it is clear that such failures continue. Clearly, some more salient means must be found to get the message across during the interviewer training. Third, mistakes in understanding, reporting, and

recording responses can generate data strongly at variance with reality. A strenuous program of data review and editing on the SCF is aimed at detecting and correcting such errors. The survey instrument design and a large part of the interviewer training are focused on eliminating such problems during the interview; much more needs to be done on both of these fronts to make it easier to do the right thing and harder to do the wrong thing. Finally, some respondents have given patently insincere interviews–possibly out of a misguided sense of being helpful or out of concerns about privacy. Future effort must be made to convince respondents that false data are worse than useless and that privacy concerns can be better addressed on other ways.

Bibliography

Frankel, Martin and Arthur B. Kennickell [1995] "Toward the Development of an Optimal Stratification Paradigm for the Survey of Consumer Finances," *Proceedings of the Section on Survey Research Methods*, 1995 Annual Meetings of the American Statistical Association, Orlando, FL.

Heeringa, Steven G., Judith H. Conner and R. Louise Woodburn [1994] "The 1989 Surveys of Consumer Finances Sample Design and Weighting Documentation," working paper, Survey Research Center, University of Michigan, Ann Arbor, MI.

Internal Revenue Service [2001] "Individual Income Tax Returns, 1998" Internal Revenue Service, Statistics of Income, Washington, DC (2001), Publication 1304(Rev. 04-2001).

Kennickell, Arthur B. [2000a] "Wealth Measurement in the Survey of Consumer Finances: Methodology and Directions for Future Research," paper presented at the 2000 annual meetings of the American Association for Public Opinion Research, Portland, May.

_____ [2000b] "An Examination of Changes in the Distribution of Wealth from 1989 to 1998: Evidence from the Survey of Consumer Finances," paper presented at the Jerome Levy Economics Institute Conference on Saving, Intergenerational Transfers, and the Distribution of Wealth, Bard College, May 2000. Levy working paper 307.

_____ [2000c] "Revisions to the Variance Estimation Procedure for the SCF" working paper, Board of Governors of the Federal Reserve System.

_____ [1999a] "Using Income Data to Predict Wealth," paper presented at the Annual Meetings of the Allied Social Science Associations, New York, 1999b.

_____ [1999b] "Analysis of Nonresponse Effects in the 1995 Survey of Consumer Finances,"*Journal of Official Statistics*, v. 15 no. 2, 1999a, pp. 283-304.

_____ [1999c] "Revisions to the SCF Weighting Methodology: Accounting for Race/Ethnicity and Homeownership," working paper, Board of Governors of the Federal Reserve System.

_____ [1998] "List Sample Design for the 1998 Survey of Consumer Finances," April 1998.

_____, Martha Starr-McCluer and Brian J. Surette [2000] "Recent Changes in Family Finances: Results from the 1998 Survey of Consumer Finances," *Federal Reserve Bulletin*, (January 2000) pp. 1-29.

_____ and R. Louise Woodburn [1999] “Consistent Weight Design for the 1989, 1992, and 1995 SCFs, and the Distribution of Wealth,” *Review of Income and Wealth* (Series 45, number 2), June 1999, pp. 193-215.

_____ and Douglas McManus [1993] “Sampling for Household Financial Characteristics Using Frame Information on Past Income," *Proceedings of the Section on Survey Research Methods*, 1993 Annual Meetings of the American Statistical Association, San Francisco, CA.

Tourangeau, Roger, Robert A. Johnson, Jiahe Qian, Hee-Choon Shin, and Martin R. Frankel [1993] “Selection of NORC’s 1990 National Sample,” working paper, National Opinion Research Center at the University of Chicago, Chicago, IL.

APPENDIX

Figure A1: Differences between 1997 master file data and 1999 SOI data; percent having any item in either data, mean percent change (in constant dollars), standard deviation of percent change; by WINDEX0; full 1999 SOI sample.

Figure A1a: Adjusted gross income.

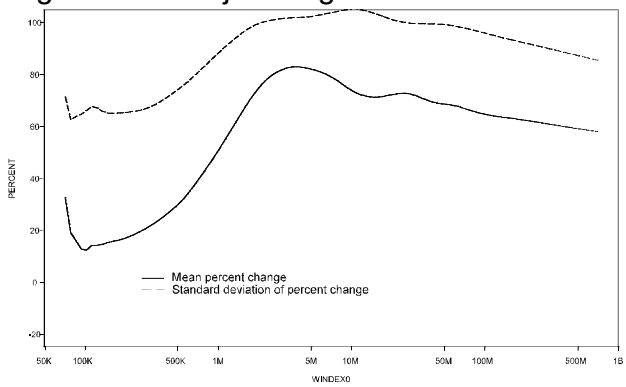

Figure A1d: S-corp. and farm income.

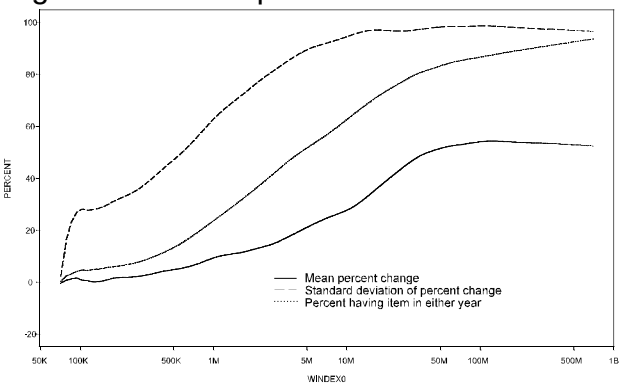

Figure A1b: Wages and salaries.

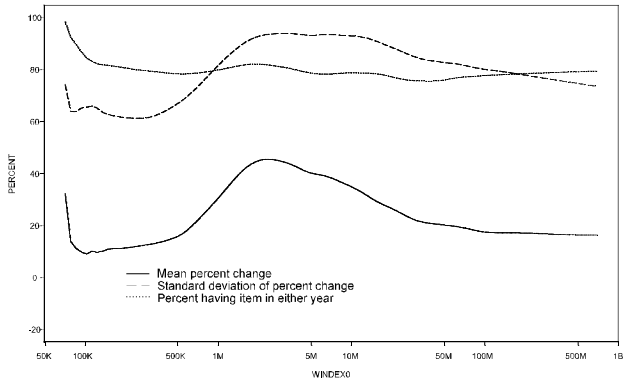

Figure A1e: Schedule C income.

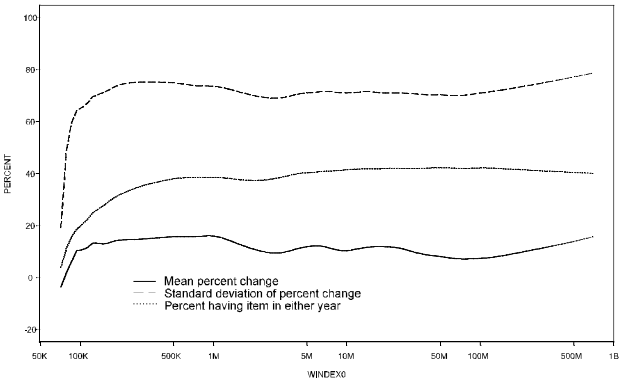

Figure A1c: Interest and dividends.

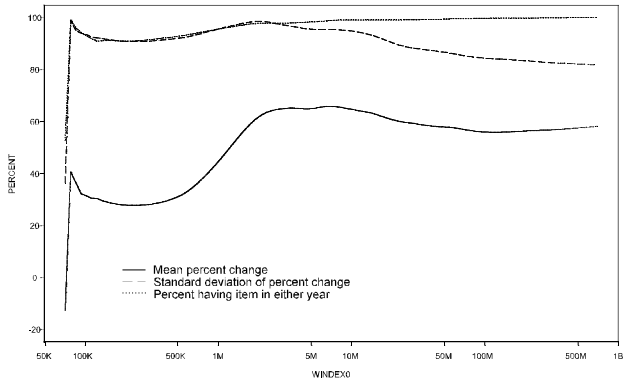

Figure A1f: WINDEX0.