The Accuracy of Economic Measurement in the Health and Retirement Study

Eric French, John Bailey Jones,
and Jeremy McCauley*

May 7, 2016

Abstract

We assess the quality of the HRS's measures of out-of-pocket medical spending and health insurance premia, both in the "core interviews" and in the "exit interview" data. We provide detailed evidence on the quality of the HRS insurance premia data, and we compare the HRS exit data to exit data in the MCBS. We document how changes in survey questions, including the introduction of "unfolding brackets", affect the HRS measures. We document what we believe are errors in the HRS imputations and provide some suggestions for improving the accuracy of some imputed variables. Overall, we find the HRS data to be of high quality. However, we believe that many interesting variables in the HRS are under-utilized because users must perform imputations themselves.

*Eric French: UCL, eric.french.econ@gmail.com. John Bailey Jones: University at Albany, SUNY, jbjones@albany.edu. Jeremy McCauley: UCL, jeremy.mccauley@gmail.com. We thank Jin Tan for excellent research assistance and Didem Bernard for helping us understand the MEPS data. The views expressed herein are those of the authors.

1 Overview

The Health and Retirement Survey (HRS) contains an unparalleled variety of measures of the well-being of the elderly. In this paper, we assess the quality of the HRS's measures of out-of-pocket medical spending and health insurance premia, both in the "core interviews" and in the "exit interview" data. We pay particular attention to issues of data imputation.

The HRS data are novel in many dimensions. One of these is the use of "unfolding brackets", where individuals who cannot give a point response to a question (such as how much they spent on medical care) are allowed to respond that the quantity in question was above some values and below others. A second dimension in which the HRS is novel is in its use of exit interviews, where data on recently deceased individuals is collected from their children or other relatives. These interviews provide information on a variety of topics, including late-in-life medical spending, the value of the deceased's estate, and bequests.

The HRS has been used extensively in the study of aging. RAND's coding, imputing, and cleaning of the HRS data has led to an easy-to-use, high-quality dataset. Our personal experience is that RAND has done a fantastic job of creating recoded HRS variables that are both well-imputed and consistent over time. RAND's recoded dataset is so accessible that it can be used by advanced undergraduates and Master's students, attracting young scholars to aging-related topics.

However, RAND recodes a relatively small share of the variables in the HRS. An informal perusal of journals suggests that very few of the variables not coded by RAND have been used in published research. This is for two reasons. First, RAND has coded the most important variables. Second, the difficulty of working with variables not coded by RAND has deterred many researchers from using them. In our opinion, there is much that can be learned from the data not coded by RAND. For the intrepid, the data not processed by RAND provides fertile ground for research.

We thus focus on the measurement of variables that have not been cleaned and recoded by RAND. Some of these variables have been imputed by HRS for some years, while others have not. We document that overall the quality of these variables is still high. This is of interest because many of the variables not coded by RAND use more complex survey methodologies (such as for the exit interviews).

However, we note several problems with the data. Perhaps most important, we document several problems with the imputations done by the HRS. The source data are of high quality, but the imputation procedures can be improved, and in some places we believe there are errors in the imputations.

This paper makes the following contributions. First, we provide detailed evidence on the quality of insurance premia data in the HRS. To the best of our knowledge these data have not been evaluated before. Second, we provide detailed evidence on the quality of the medical spending (and other) data for those recently deceased that were collected in the "exit interviews". To the best of our knowledge, we are the first to compare the HRS exit data to exit data in the Medicaid Current Beneficiary Survey (MCBS). Third, we update and expand on the estimates shown in Goldman et al. (2011). As described in Goldman et al. (2011) and De Nardi et al. (2016a), the HRS data match up well relative to other surveys such as the MCBS and the Medical Expenditure Panel Survey (MEPS); see also Hurd and Rohwedder (2009). We confirm their findings. Fourth, we provide some suggestions for data imputation. We hope that our suggestions, along with the documentation provided by the HRS (Cao, 2001) and RAND (Chien et al., 2014), alongside the work of Marshall et al. (2011) and Fahle et al. (2016), will encourage researchers to explore the HRS data beyond the RAND dataset.

2 Key findings

- The HRS has high quality data for the variables we consider. For example, the medical spending data line up well relative to the MCBS. We document lower reported Medicaid recipiency in the HRS relative to the MCBS, suggesting modest underreporting of Medicaid recipiency in the HRS. For insurance premia, the match between the HRS, the Medical Expenditure Panel Survey (MEPS), and the MCBS is not quite as good, but it is not obvious whether the problem lies with the HRS, MEPS, or with the MCBS.

- The medical spending questions in the HRS have changed over time, especially in the earliest waves. Although this can hinder comparisons across waves, for the most part the changes have improved the quality of the data. For example, prior to 2002, respondents were given flexibility in how they could respond to certain insurance questions. After 2002, however, respondents were told to report insurance payments per month. This likely improved the quality of their answers.

- Non-response rates are fairly high for many variables, including insurance premia. Issues of either partial or complete non-response (coded as "don't know" or "refused") are especially serious in the exit interviews.

- The HRS addresses these issues of non-response by using “unfolding brackets” and imputations. We assess the quality of the HRS’ imputations of health insurance premia. Overall, the HRS uses good procedures that balance the need for accuracy with the need for robustness. However, we find what we believe to be errors, some of which are serious. In particular, we document apparent errors in the 2002 and 2004 insurance premia imputations. Perhaps more seriously, the HRS no longer imputes data, and RAND only imputes a share of the data, leaving researchers to do any remaining imputations on their own. Whether individual researchers do the imputations carefully or not is unknown.

3 Key recommendations

- Overall, the HRS contains extremely high quality panel data, providing measures of key variables that are comparable over time. For this reason, we do not recommend any changes to the survey questionnaire.

- We recommend that a trusted source, such as RAND, continue to clean, impute and recode as much of the HRS data as possible.

- We document errors in the HRS imputations, and recommend that some of these specific problems with the imputations be addressed. These problems are sufficiently serious that it would be better to remove some of the imputed variables than to leave them on the HRS website as currently constructed.

- Many variables in the HRS are not imputed at all. For example, insurance premia after 2004 are not imputed. We recommend that all variables with unfolding brackets be imputed by a trusted source, such as RAND.

- The imputations should satisfy basic accounting identities. For example, the total value of the estate bequeathed to all children should equal the sum of the bequests to each child. Basic identities like this should hold in the HRS imputations, but in many cases they do not. We recommend exploiting these “adding up” conditions in the construction of the data.

- We reiterate two of Venti’s (2011) recommendations. First, the HRS should be linked to administrative Medicaid data. We document modest underreporting of Medicaid recipiency in the HRS—using administrative data should circumvent this problem. Adding the administrative data would also introduce new variables, such as the dollar value of Medicaid benefits, that should stimulate new research on Medicaid. Second, we recommend that the imputation procedures exploit the HRS's panel dimension.

4 Introduction to the HRS Medical Spending Data

We begin by briefly describing the medical spending questions in the HRS, and the way in which point responses are extracted from incomplete data.

4.1 Medical Spending Questions in the HRS

While alive, survey respondents answer a rich variety of medical spending questions in the HRS "core" interviews. Once they die, spouses, children, or other decedents are asked about their medical spending in "exit" interviews.

For researchers doing longitudinal studies it is essential that questions remain comparable across waves. In this section, we summarize how the HRS medical spending questions have evolved over time. Although the changes have generally improved the quality of the data, they have also reduced the comparability of medical spending across waves, especially in early years.

The medical expense information for wave 1 (1992) of the HRS is limited to insurance premia. Moreover, the insurance premium question only refers to insurance purchased directly from an insurance company or through a membership organization such as the American Association of Retired Persons (AARP). It does not include employee contributions to employer-provided insurance plans. Given that the information is incomplete, we do not use wave 1 data in our analyses.

The set of medical spending questions expanded significantly in wave 2:

- The insurance premia question includes employee contributions to employer-provided insurance plans and also insurance directly purchased or through membership such as the AARP.
- The respondent is asked whether the individual had any hospital stays, nursing home stays, or visits to a doctor. If the answer to any of these questions is yes, the respondent is then asked both the total cost and out-of-pocket cost for the visit or stay.

- The respondent is asked about whether the individual purchases medicines prescribed by a doctor. If the individual did purchase medicine, they are asked how much the medicines cost per year. It is not clear whether the cost measure refers to the cost paid by the individual or what the pharmacy charges the individual and the insurer. Nevertheless, we use this variable to determine drug costs.

The questions in wave 2 can be combined to compute measures of total out-of-pocket expenses and total insurance costsl. However, a number of expenditures are still excluded. The costs (both out-of-pocket and total) of outpatient surgery, dental expenses, in-home care, and special facilities and services are not included. Medicare HMO insurance costs are missing as well. We thus exclude the wave 2 data as well.

From wave 3 onwards, out-of-pocket costs are clarified to include the amount paid for doctors, hospitals, nursing homes, outpatient surgery, dental expenses, in-home medical care, special facilities and services. In addition, the procedure for determining insurance costs changes over waves 3 through 5. The insurance premium variable is now the sum of premia for all employer-provided insurance, Medicare HMO plans, supplemental plans, private/AARP/professional coverage, and long-term care plans.

Data for the oldest cohorts in the HRS were originally collected in a distinct but similar survey called the Assets and Health Dynamics of the Oldest Old (AHEAD). The first two waves of the AHEAD did not coincide with the first two waves of the HRS; in 1998 the two surveys were merged. In wave 1 (1993) of the AHEAD, we can determine out-of-pocket costs. The out-of-pocket costs include the costs of nursing home stays as well as hospital and doctor bills and any other medical or dental expenses in the last 12 months. We infer that drug costs are included in this measure, although respondents are not asked explicitly about drug costs. AHEAD wave 1 also asks the respondents if they had insurance policies in addition to Medicare, including long-term care policies, and how much they paid yearly for such policies. From this we can determine the respondent's insurance expenses. Wave 1 is not used in the this paper, however, because the sources of health insurance are incomplete.

In the 1995 wave of the AHEAD, imputation procedures for total costs, and out-of-pocket costs, insurance costs, drug costs, and medical costs are the same as in waves 3 through 5 of the HRS. This allows us to combine the AHEAD data for 1995/1996 with the HRS data for 1996.

In the analysis below we assess the accuracy of the economic measures contained in these questions, using questions that are broadly comparable across waves. All amounts, unless otherwise stated, are deflated to 2014 dollars using the PCE price index.

4.2 Unfolding Brackets and Imputation in the HRS

Non-response is a problem in almost every survey. Non-response can take multiple forms. People may refuse to participate in the survey at all, or they might participate but not answer certain questions, giving "don't know" (DK) or "refused" (RF) as an answer. The HRS has done a remarkable job in addressing both issues. The rate of participation in the HRS is much higher than in related surveys, including the aging surveys for other countries, such as ELSA and SHARE, and attrition is modest.

Many of the questions asked in the HRS are personal in nature, and many are difficult to answer. Medical spending questions belong to both categories. Many people may not know how much they paid in medical bills, as they are paid only intermittently. Many people may also be uncomfortable discussing their health care needs. For this reason a high share of all respondents respond either DK or RF when asked about their medical spending. To address this issue, the HRS uses "unfolding brackets". Respondents not providing point values for their medical spending are asked "Did it amount to less than \$ _____ , more than \$ _____, or what?" If respondents report a higher value, they will be asked a new question about whether their spending was more or less than a higher value. This procedure ultimately identifies a range of possible values.

In the early waves, the HRS replaced incomplete or missing values with imputations for many variables. See Cao (2001) for a description of these procedures. However, in more recent waves these imputations have ceased.

Many of these variables are now imputed by RAND. Chien et al. (2014) catalogue the variables available in the RAND data, and describe how the variables are imputed. Our personal experience is that RAND does an extremely good job of coding the data. Furthermore, the data are well-documented and the procedures used to construct the data are also reasonably well-explained. Our assessment is based upon the fact that we have coded many variables ourselves, and have compared many of the variables we have coded ourselves to the same variables in the RAND data. Although there are sometimes discrepancies between our own coding and RAND's coding of the data, they are minor and our procedures are not necessarily better. The only error we have ever discovered in RAND's coding was quickly corrected after contacting them. For this reason in our own work we rely on RAND's coding of the data whenever possible.

It should be noted, however, that RAND codes only a subset of the data. An example particularly relevant to this study is that RAND does not code or impute insurance data or data from the exit interviews. In general, many of the variables that are most difficult to impute are not imputed by RAND.

For earlier waves, additional imputations are available in the HRS itself. Up until 2000, the HRS released "Exit Imputations" files, which imputed all missing values in the exit data. Since that time, the exit data are no longer imputed. Likewise, HRS imputed many of the variables in the core files until 2004. Below we document some problems with the HRS imputations of insurance premia and estate dispositions. But our key message is *not* that HRS did a bad job in the imputations. On the contrary, the HRS, with its expert staff, mostly did a very good job. The most serious problem with the HRS imputations is that they are no longer being done by the expert HRS staff. Because many researchers lack the resources and/or expertise to impute the data themselves, this limits the use of the data. Perhaps even more seriously, researchers who impute the data themselves may not do it correctly. That there are issues in the coding, cleaning, and imputation of the data when done by experts shows how difficult the process is. Furthermore, there is little professional oversight of data coding and imputation. The academic refereeing process is useful for vetting the quality of ideas, but not as useful for assessing the details of data processing. The fact that the cleaning, coding, and imputation of HRS data has been delegated to the research community, with little if any oversight, is not reassuring.

We perform very little topcoding. Fahle (2016) do more in the way of topcoding. However, these differences seem fairly minor: relatively few observations have values of medical spending that are clearly implausible. For example, topcoding medical spending at $200,000 reduces mean medical spending in the elderly De Nardi et al. (2010) sample by 2%. This certainly does not mean that measurement error is unimportant: it just means that it is not at all obvious how to detect it.

5 Medical Insurance Premia

In addition to being queried about their out of pocket expenditures on deductibles and co-pays for medical care, respondents in the HRS are asked about their out-of-pocket expenditures on insurance premia. Although RAND imputes missing data and recodes out-of-pocket medical spending on co-pays and deductibles, it does not do the same for insurance payments. The HRS provides imputations only for 2002 and 2004. Below we report estimated insurance premia, using our own imputation procedures, and describe how the estimated premia have changed as the survey questions evolved. We also document problems with the HRS imputation procedures.

Individuals in the HRS are questioned about the different types of health insurance that they might have. Respondents are also asked about how much is spent on each type of insurance. The insurance types that respondents are asked about change from wave to wave, as do the reporting methods. These changes are discussed in detail below.

Prior to 2002 (wave 6) we drop observations that report having health insurance, but respond with DK or RF when asked how much was spent on insurance premia. We find that this generally amounts to 10-20% of observations. From 2002 onward, these respondents are directed to unfolding bracket questions. For those who respond to the unfolding brackets, we impute insurance premia data ourselves. Our imputation method is simple. First, we find the upper and lower bounds on the premium implied by the bracket the respondent selected. We then give the respondent the spending of a "donor" individual who gave a point response that fell in the respondent's bracketed range. Our procedure is thus a hot deck method, where the cells to be matched include everyone with premia in the desired range. We drop those who did not give a point or unfolding bracket response.

5.1 Comparing Insurance Premia in the HRS, MEPS, and MCBS

The first two columns of Table 1 shows mean insurance premia in the HRS by year for the age 65+ population. To make the results comparable to the MEPS and the MCBS, premia for Medicare HMO or other Medicare and long-term care insurance plans are not included as these do no exist in the MEPS and MCBS. The table begins in 1996 because, as discussed in Section 4.1, the medical spending questions used by the HRS prior to 1996 are not comparable to those used in later waves.

The next two columns show premia from the MEPS, and the final two show premia from the MCBS. The MCBS and MEPS both capture individual payments for employer-provided and other private insurance premia, but exclude long-term-care, Medicare HMO or other Medicare payments. The three surveys should thus capture similar types of insurance premia. Appendix A provides more background on the MEPS and MCBS, and the data that these surveys capture.

Column 1 of Table 1 shows unweighted mean premia for the HRS while column 2 shows the same premia after adjusting for population weights. Columns 3 and 4 show unweighted and population-weighted means for the MEPS, and columns 4 and 5 show the same for the MCBS. Panel A shows unconditional means, whereas panel B shows means for respondents who report non-zero premia.

Column 1 of panel A shows that the HRS measure of insurance premia is not stable over time. There is a significant jump in reported insurance premia in 2002. Column 2 shows that the jump in 2002 is not an artifact of changes in sample composition. Although the MCBS data also show a jump in insurance premia in 2002, the jump is much smaller than in the HRS. Moreover, prior to 2002 insurance premia in the HRS are well below premia in the MCBS, whereas after 2002 the insurance premia reported in the two data sets are similar.

Year	HRS data		MEPS data		MCBS data	
	unweighted	weighted	unweighted	weighted	unweighted	weighted
			Panel A: Unconditional (whole population)			
1996	734	771	–		808	827
1998	778	821	–		872	882
2000	728	774	–		937	948
2002	1,194	1,258	739	828	1,123	1,132
2004	1,286	1,352	868	1,056	1,130	1,143
2006	1,110	1,160	707	847	1,136	1,157
2008	1,161	1,241	617	691	1,202	1,198
2010	1,197	1,291	607	703	1,128	1,152
2012	1,103	1,186	766	904	–	–
			Panel B: Conditional on positive premia spending			
1996	1,725	1,768	–		1,373	1,398
1998	1,924	1,896	–		1,603	1,637
2000	2,080	2,092	–		1,643	1,667
2002	2,407	2,406	2,280	2,343	1,835	1,873
2004	2,617	2,611	2,736	2,968	1,944	1,965
2006	2,496	2,494	2,500	2,566	1,788	1,816
2008	2,801	2,848	2,376	2,438	1,871	1,919
2010	2,989	3,013	2,386	2,435	1,832	1,848
2012	2,927	2,881	2,752	2,795	–	–

Notes: Mean annual insurance premia, 2014 dollars, excluding Medicare HMO and long-term care insurance premia. MCBS estimates taken from De Nardi et al. (2016c).

Table 1: Insurance premia for 65+, comparison.

Panel B shows that once we restrict the sample to those with positive premia, the jump in the mean insurance premium between 2000 and 2002 becomes much smaller. It seems that much of the jump in insurance payments is due to the fact that we drop respondents who answer DK or RF, the majority of which can be imputed after the introduction of unfolding brackets in 2002.

The MEPS data on insurance premia information begin in 2002. Like the HRS, the MEPS uses self-reports to measure insurance premia. The MEPS also employs reporting procedures that are in many ways similar to the HRS. Panel A shows that the insurance premia reported in the MEPS data are smaller than the insurance premia reported in the HRS or MCBS. Panel B shows, however, that conditional on being positive, insurance premia in the MEPS are similar to those in the HRS. This suggests that the fraction of people with positive insurance premia is underreported in MEPS. That the HRS insurance premia are close to those in the MCBS in panel A is reassuring, although panel B shows that conditional on positive spending, spending in the HRS is higher.

Table 2 shows the same data as Table 1, but for the population aged 55-64. Because the MCBS is only representative of the age-65+ population, we omit it from this table. As in Table 1, insurance premia in the MEPS are below the insurance premia in the HRS in panel A, but very close in panel B. Also as in Table 1, in panel A there is a large jump in the HRS measure of insurance premia between 2000 and 2002, and the jump shrinks—in this case disappears—in panel B. More generally, the HRS-measured premia tend to fluctuate over time. Weighting has little effect on the estimated average insurance premia. Although the age composition of the sample changes slightly over time, with the aging of the sample and the addition of new cohorts in 1998, 2004, and 2010, these changes in sample composition are not large enough to explain the observed changes in the measured insurance premia.

Overall, the discrepancies between the surveys are not trivial. We are not sure about the sources of the discrepancies, or which survey should be trusted the most. All three surveys (MRS, MEPS, and MCBS) rely on self reported insurance premia, and it is not obvious which survey does the best job of measuring premia.

Tables 1 and 2 show that when we use the full sample (panel A) the HRS measure of insurance premia jumps in 2002, and remains high thereafter. Table 1 shows that the size of the jump decreases when we restrict the sample to people with positive premia (panel B), and in Table 2 the jump vanishes. It is thus likely that much of the difference is due to the introduction of the unfolding brackets in HRS wave 6, which significantly reduced non-response. However, panel B of Table 1 still displays a jump, and panel B of Table 2 shows that the mean premium fluctuated more in the waves

Year	HRS data		MEPS data	
	unweighted	weighted	unweighted	
	Panel A: Unconditional (whole population)			
1996	987	1,040	–	
1998	854	876	–	
2000	1,129	1,176	–	
2002	1,755	1,742	1,066	1,229
2004	2,050	2,079	1,262	1,443
2006	1,974	2,030	1,269	1,508
2008	2,193	2,313	1,261	1,423
2010	2,046	2,488	1,414	1,717
2012	1,877	2,370	1,422	1,626
	Panel B: Conditional on positive premia spending			
1996	2,399	2,452	–	
1998	2,143	2,136	–	
2000	2,776	2,803	–	
2002	2,962	2,889	2,824	2,938
2004	3,369	3,297	3,280	3,342
2006	3,264	3,268	3,193	3,341
2008	3,632	3,672	3,134	3,246
2010	3,836	4,059	3,459	3,713
2012	3,616	3,911	3,459	3,650

Notes: Mean annual insurance premia, 2014 dollars, excluding Medicare HMO and long-term care insurance premia.

Table 2: Insurance premia for age 55-64, comparison.

prior to 2002. We believe that these differences are due to differences in the wording of the HRS questions.

Although the HRS uses roughly the same insurance premia categories for wave 6 (2002) and beyond as in waves 3-5, there are differences in how the questions are asked. In earlier waves, individuals were asked whether the amount paid for insurance per period was (1) per year, (2) quarterly/every 3 months, (3) bimonthly/every 2 months, (4) per month, (5) week, (6) biweekly/every 2 weeks, (7) semi-annually/2 times per year, or (8) "other". For 2002 and beyond, respondents were asked to report medical spending on a monthly basis. Allowing a choice of reporting period in the earlier waves appears to have led to less accurate answers.

Reporting period	Observations	Mean Annual Premia	
Wave 4 (1998)	Wave 4	Wave 4	Wave 5
Overall	20,130	376	392
		(16)	(16)
Those who pay "all" or "some"	4,479	1,689	1,265
Reporting period:		(70)	(66)
1. Year	376	1,591	1,469
		(97)	(209)
2. Quarterly/Every 3 months	100	1,675	907
		(105)	(154)
3. Bimonthly/Every 2 months	49	787	1,456
		(137)	(394)
4. Month	3,225	1,740	1,208
		(29)	(80)
5. Week	427	2,627	1,562
		(692)	(249)
6. Biweekly/Every 2 weeks	295	1,581	1,296
		(74)	(153)
7. Semi-annually/2 times per year	6	1,199	517
		(302)	(319)
8. Semi-monthly/2 times per month	1	1,535	1,649
		(.)	(.)

Table 3: Mean annual employer provided insurance premia for those with employer provided insurance, 2014 dollars, by wave 4 reporting period.

Table 3 shows mean annual insurance premia for those with employer-provided insurance, by reporting period, in wave 4 (1998). We construct annual premia by multiplying the reported per-period premia by the number of reporting periods in a year: e.g., by 4 if they self-reported paying quarterly, 6 if they reported bimonthly, and so forth. The table shows that mean annual payments vary greatly across the reporting periods. In particular, the option "bimonthly" has an unusually low average annual premium, of $787. It is possible that some respondents are confused as to whether they are reporting twice a month or every 2 months. Furthermore, mean payments for those reporting weekly are extremely high, at $2,627. This is caused by the 1.4% of the sample who reported weekly insurance premia of over $500 per week ($26,000 per year). For most years, this has only a modest effect on sample means. For example, topcoding at $15,000 reduces weighted mean insurance premia in 2008 for 55-64 year olds from $2,313 to $2,238.

When we consider the same individuals in wave 5, the large differences across reporting period groups disappear. For example, both those reporting bimonthly and those reporting weekly have average insurance premia in wave 5 that are very close to the overall sample average among those who pay at least part of their premia. This suggests that some individuals were previously reporting monthly or annual premia, not weekly. As we show below, this may not be the only case of people making mistakes about the reporting period.

5.2 Problems with Imputed Insurance Premia in the HRS

Although in most years, any imputations of insurance data must be done by the individual researcher, in 2002 and 2004 the HRS imputed insurance premia. However, we recommend the researcher impute the data themselves for 2002 and 2004 as well, as the HRS imputations appear to have serious errors.

Panel A of Table 4 compares our imputed insurance premia to the HRS imputations. Our estimated mean premia are $1,397 and $1,540 for waves 6 and 7, respectively. These are very different than the HRS imputations of $437 and $1,217. In particular, there appears to be a serious problem with wave 6. As it turns out, the difference between our estimates and the HRS estimates come not from different treatments of bracketed responses, but from the HRS's recoding of exact responses.

Panel A: Our imputation and HRS imputation comparison

	Wave 6		Wave 7	
	Observations	Mean annual total premia	Observations	Mean annual total premia
Our imputation	18,167	1,397	20,129	1,540
		(14)		(18)
HRS imputation	18,167	437	20,129	1,217
		(6)		(13)

Panel B: Specific variable examples from Wave 6

	Private insurance premia		Long term care insurance premia	
Percentile	Self reported HN040_1	Imputed HN040_1X	Self reported HN079	Imputed HN079X
1%	0	0	0	7
5%	16	0	14	28
10%	30	0	44	46
25%	69	0	100	83
50%	130	0	203	139
75%	228	0	1,240	211
90%	368	0	2,760	300
95%	500	0	3,600	367
99%	750	0	6,000	592
Mean	172	1	1,038	165
Std. Dev.	151	20	3,585	142
Observations	7,842	7,842	1,300	1,300

Notes: HN040_1 is a monthly private plan insurance premia. HN079 is the monthly long-term care insurance premia. In wave 6 (2002) there were a total of 18,168 observations. When comparing the imputations with the original variables, we only compare the observations which were non-missing in the original variables. Dollar amounts are in nominal terms.

Table 4: Percentiles of Reported and Imputed Monthly Insurance Premia for those Reporting Positive Point Values, HRS Wave 6.

Panel B of Table 4 displays the percentiles of monthly private non-group and long-term care insurance premia in wave 6 (2002). It shows individual self-reported values for those who gave an exact (not bracketed) response, and the HRS-imputed values of the same premia for the same individuals. Surprisingly, the HRS-imputed values are very different from the self-reports. In the case of the private insurance premia variable, almost all of the self-reported non-zero responses have been set to 0 in the imputed version of the variable. We believe this represents an error in the HRS coding of the data. It appears that (possibly accidentally) the HRS cleans the data when it seems that there is little reason to do so.

The right-hand side of panel B shows that in the case of long-term care insurance premia, most of the self-reported non-zero responses have been revised downwards by the HRS. This could potentially be a fix to incorrect reporting, given that the premia are monthly, and the self-reported values appear too large for the 75th percentile and above. Again, this is potentially because some respondents have given annual rather than monthly values. While some of the self-reported values appear to be implausibly high, the extent to which these values have been revised downwards is dramatic. Detailed documentation explaining what has been done to the data would be especially useful in instances such as these.

6 Exit Interviews

One of the most novel aspects of the HRS data is that when an individual dies, a proxy respondent is given an exit interview, providing details of death-related expenses, the health situation of the deceased prior to death, and the disposition of the estate. To the greatest extent possible, the proxy respondents are knowledgeable about the health, family, and financial situation of the deceased: often the proxy is a widow, widower, child or some other family member. Since 1998, if the disposition of the estate is not yet settled at the time of the exit interview, the proxy respondent will be re-interviewed in the next wave, in a "post-exit" interview.

While these data are in principle a great resource to researchers, they have been used much less than data from the core interviews. The exit data have been used to estimate late-in-life medical spending (e.g., De Nardi et al., 2010, 2016a, 2016c; Marshall et al., 2011), to learn about the disposition of estates (e.g., Hurd and Smith, 2001), and to measure the decline in assets near the death of a spouse (French et al., 2006). However, there is also some skepticism about the quality of these data. Poterba et al. (2015), in their estimates of wealth dynamics, exclude the estate data because of concerns about their quality. There are fewer assessments of the quality of the exit data than there are of the core data. In this section we describe how these data are constructed, and we provide some evidence on their quality. We find the overall quality of the exit data to be high, but we also document some problems, especially with imputations.

6.1 Medical Spending Data

We describe the medical spending data first. As in the core interviews, questions are asked about components of out-of-pocket spending for items such as doctor visits, hospital stays, and nursing home care. Exit interviews tend to be shorter and less detailed than the core interviews. Furthermore, the proxy respondents may know little about the circumstances of the deceased. This lack of knowledge can be seen in the relatively high share of DK and RF responses in the exit data.

A good example is hospital stay spending. Table 5 shows the response rates for the hospital spending questions in the core and exit data for HRS wave 9 (2008). In the core data, 1,984 of individuals had a hospital stay not completely covered by insurance, 110 had a bill not yet settled, and 59 responses were either DK or RF. For those individuals, 600 respondents answered DK or RF when asked to give a point value for out-of-pocket hospital spending. In the exit data, 356 of the deceased had a hospital stay not completely covered by insurance, 11 had a bill not yet settled, and 38 responses were either DK or RF. Of these, 197 reported DK or RF when asked to give a point value for out-of-pocket hospital spending. Although most respondents who initially report DK or RF when asked for a point response later give a bracketed range, much of the data must be imputed. In the core interview, of the 600 respondents who initially reported DK or RF, 354 respondents gave a bracketed response greater than 0 and 454 respondents gave a bracketed response less than the maximum possible value. In the exit data, of the 197 DK or RF responses to "how much did you pay", 119 gave a response greater than 0 and 127 respondents gave a bracketed response less than the maximum possible value. Thus the bracketing significantly reduces the number of missing values generated by DK or RF responses. Nonetheless, in the exit data a non-trivial share of the responses provide no useful information for imputation. This raises the question of whether the exit data can be as trusted as data from the core interviews.

Table 6 compares average out-of-pocket medical spending in the last year of life in the HRS with spending the MCBS. The structures of the HRS and MCBS are not directly comparable, but we have chosen a spending measure that we believe is. In

Panel A: Core data

Hospital stays covered by insurance? (LN102)					
Completely	Mostly/Partially/Not	Not settled	DK/RF		
2,707	1,984	110	59		
	How much did you pay OOP? (LN106)				
	Point Response	DK/RF			
	1,553	600			
		Bracketed Response:			
		Min > 0	Min = 0	Max < 99,999,996	Max > 99,999,996*
		354	245	454	145

Panel B: Exit data

Hospital stays covered by insurance? (VN102)					
Completely	Mostly/Partially/Not	Not settled	DK/RF		
672	356	11	38		
	How much did you pay OOP? (VN106)				
	Point Response	DK/RF			
	208	197			
		Bracketed Response:			
		Min > 0	Min = 0	Max < 99,999,996	Max > 99,999,996*
		119	78	127	70

Table 5: Response Rates for Hospital Spending in Core and Exit Data, HRS Wave 9 (2008). *: Max > 99,999,996 means that no maximum value was given.

the HRS we use out-of-pocket medical spending for those who died within 12 months of their last interview. In the MCBS we use out-of-pocket medical spending in the year they died. The MCBS measure mixes together those who died in January (and so had only one month of spending in the year of death) with those who died in December (and so had 12 months of spending). Similarly the HRS measure mixes those who died one month after their last interview (and so had only one month of spending) with those who died 12 months after their last interview (and so had 12 months of spending).

Table 6 shows that the overall mean out-of-pocket medical expenditure in the year of death was $8,220 for the HRS sample. Once we use population weights, this decreases to $5,662. The weighted means for the 65+ and under-65 populations are $5,438 and $6,514 respectively. The under-65 population having higher spending in the year of death is similar to findings from the UK, where National Health Service spending in the year of death shows that under-65s cost more in the year of death than those 65+ (Aragon, Chalkley and Rice, 2016). Likewise, total spending on the year of death for those who die after age 65 appears lower than for other age groups in Taiwan (Chen and Chuang, 2016).

The weighted HRS and MCBS estimates are very close. This provides evidence of the quality of the HRS exit data. The MCBS data should be of good quality because the MCBS interviews individuals multiple times per year, and the MCBS uses administrative Medicare information in the construction of the data. The fact that the HRS lines up so well relative to the MCBS suggests that the longer periods between interviews does not adversely affect the HRS data's quality.

	HRS exit data		MCBS data	
	unweighted	weighted	unweighted	weighted
Out of Pocket Medical Expenditure				
Overall	8,220	5,662	–	–
	(458)	(7)		
65+ population	8,425	5,438	5,773	5,182
	(501)	(7)	(121)	(22)
Under 65 population	6,990	6,514	–	–
	(1,117)	(17)		

Notes: The HRS sample consists of individuals who died within 12 months of their last interview, and had their exit interview conducted at most 24 months after their death. We use the years 1996–2010. The population weights for the HRS are the RAND weights from the last interview. Adjusted to 2014 dollars. MCBS estimates taken from De Nardi et al. (2016c).

Table 6: Medical Spending in Year of Death

6.2 Disposition of the Estate

Up until 2000, the HRS released "Exit Imputations" files, which imputed all missing values. While our general recommendation is that the HRS (or a similarly qualified source) provide *more* imputations, not fewer, we have concerns with some of the existing imputations.

We can illustrate some of our concerns by examining questions about the disposition of the estate. Proxy respondents in 1998 are asked: “Altogether, what was the value of (his/her) total estate?” In 2004 the question was refined slightly to “Excluding any life insurance, altogether what was the value of[her/his] total estate?” From the wording of the question, it is not completely clear if the answer should include the value of the house. In later waves, responses to the follow-up question “Does that include the value of the home?” show that the large majority of respondents included the value of the house when describing the value of the estate. There are other issues in terms of measurement of the estate, such as whether it includes legal fees or other expenses. French et al. (2006) and De Nardi et al. (2016b) show that at the time of death of a spouse, the value of the estate drops significantly, more than what can be explained from medical, burial, and other death related expenses. For example, De Nardi et al. (2016b) find that assets fall by $28,000 in periods when a spouse dies (and rises a little when a spouse does not die), whereas medical spending is $13,000 (and medical spending plus death related expenses is $20,000) in periods when a spouse dies. However, given the differences between the structure of the asset questions asked in the core interview and the corresponding questions in the exit interview, the quality of the match is actually quite impressive: the reported estate values seem to have content.

Proxy respondents are asked not only about the value of the estate, but are also asked about how the estate was divided between the spouse, children, and other people. Proxy respondents can provide information on the percentage bequeathed to all children in at least three different ways. First, the proxy respondent can report the percentage of the total estate left to all children. Second, proxy respondents can report the total dollar amount left to all children. Researchers can then divide this amount by the total value of the estate to find the percentage left to the children. Third, the proxy respondent can report the percentage of the total estate left to each child, which can be summed up to provide the total percentage for all children. Each of the three approaches should result the same percentage. The equivalence of the three approaches should in principle be valuable for constructing imputed values. However, this information is ignored in the imputations.

Table 7 uses six observations from the wave 4 (1998) exit data to highlight this issue. The first two columns present the houshold ID and person number (HHID and PN) for each observation. The third column presents the answer to the question for “percent to all children”. The fourth column gives the cumulative response to “percent to child 1”, “percent to child 2”, etc., summing over all children. Respondents only ever answered one of these two questions. The fifth column gives the imputed

		Self reported		Imputed	
HHID	PN	Q2438	$\sum_{i=1}^{11} Q2449\_i$	Q2438X	$\sum_{i=1}^{11} Q2449X\_i$
22816	10	100	–	100	0
36504	10	–	100	0	100
45805	10	–	100	0	100
82260	20	100	–	100	0
83756	10	33	–	33	94
200403	10	–	100	0	100

Notes: Q2438 is the question "percent given to children"
Q2449_i is the question "percent given to child i"
Q2438_X is the imputed "percent given to children"
Q2449X_i is the imputed "percent given to child i"

Table 7: Example of imputed percentage of inheritance given to children, wave 4.

value of the variable shown in the third column, and the sixth column gives the imputed value of the variable in the fourth. In principle columns 5 and 6 should be identical, but in practice they are not—the data in column 3 was not used to impute missing values in column 4, and vice versa. For the first observation (HHID 22816) we know that 100% of the estate was given to the children. This is useful information on the share given to each child. However, it was not used in the imputations, since the imputed share given to all children is 0%.

Turning to the entire sample, if we regress the imputed value of the percentage left to all children on the sum of the imputed percentages left to each child, we would hope to get a coefficient close to 1. Equation (1) shows this is not the case. Instead, the slope coefficient is negative when regressing the former (Q2438X) upon the latter (Q2449X). The regression coefficients (and standard errors) are:

$$\begin{aligned} Q2438X &= \underset{(1.3)}{25.2} - \underset{(.02)}{.23} * Q2449X \end{aligned} \tag{1}$$

The case just described is an example of a basic identity (the percentage of an inheritance going to each child, summing over all children, should equal the percentage left to the children in total) that should hold, but is not forced to hold in the HRS imputations. We recommend exploiting these "adding up" conditions in the construction of the data. The idea of using adding-up identities in data construction is hardly new: take for example the National Income and Product Accounts, which forces adding-up (or accounting) identities to hold. It should be easy enough to use these identities to improve the data.

7 Comparing MCBS and AHEAD data

	HRS data (AHEAD cohort)				MCBS data		
Income Quintile	Total income	Annuity income	Medicaid recipiency	Out-of-pocket expenses	Total income	Medicaid recipiency	Out-of-pocket expenses
1	9,124	5,682	60.9%	3,006	7,957	69.9%	4,774
2	12,131	9,749	28.1%	5,034	11,812	41.8%	6,295
3	18,273	12,850	11.0%	5,953	16,198	15.5%	7,627
4	22,740	16,964	5.6%	7,498	23,236	8.0%	8,606
5	39,587	31,004	3.0%	8,252	52,047	5.4%	9,455

Notes: We only use the years 1996–2010, and those who were single and age 72 and older in 1996. Adjusted to 2014 dollars.

Table 8: Income, out-of-pocket spending, and Medicaid recipiency rates, AHEAD cohort versus MCBS.

Our final set of exercises is to compare the out-of-pocket medical spending, Medicaid recipiency and income measures in the AHEAD cohort of the HRS to the corresponding measures in the MCBS. Comparisons to the MCBS are of particular interest because of the quality of the MCBS data. The MCBS uses administrative records to identify Medicaid recipients. Furthermore, Medicare claims data are used to help construct data on out-of-pocket medical spending. Appendix A provides more background on the MCBS data.

Our analysis relies heavily on the work in De Nardi et al. (2016a). Whereas for the preceding analyses we consider all age groups, here we restrict the sample to singles (over the sample period) who meet the AHEAD age criteria (at least 70 in 1994, 72 in 1996, and so forth) and who are not working over the sample period. We combine data from the core and exit (and post-exit) interviews.

In addition to measures of Medicaid recipiency and out-of-pocket spending, we construct a measure of permanent income (PI) that can be used to stratify the data. As in De Nardi et al. (2016a), our PI measure for the AHEAD is the percentile rank of the individuals' average (while observed) annuity income. Because the MCBS asks only about total income (including asset and other non-annuitized income), our measure of PI for the MCBS is the percentile rank for total income. The first four columns of Table 8 show sample statistics from the full AHEAD sample in the HRS while the final three columns of the table shows sample statistics from the MCBS sample. The first set of statistics we compare are for total income. Total income in the AHEAD data lines up well with total income in the MCBS data, although income in the top quintile of the MCBS is higher than in the AHEAD. In De Nardi et al. (2016a) we compare the annuity income quintiles of the AHEAD to the annuity income quintiles of the PSID, and find that the match is good, including at the top of the income distribution. Thus if anything the problems with income measurement probably lie with the MCBS, which uses a less detailed methodology to elicit income information.

Next we compare Medicaid recipiency rates in the AHEAD and MCBS. 61% and 70% of those in the bottom PI quintile are on Medicaid in the AHEAD and in the MCBS, respectively. In the top quintile, 3% of people are on Medicaid in the AHEAD, whereas 5% are in the MCBS. The MCBS may have higher Medicaid recipiency rates because it has administrative information on whether individuals are on Medicaid, which eliminates underreporting problems. Although Table 8 suggests underreporting in the AHEAD data, it should be noted that for many programs, under-reporting in survey data is often severe. In our opinion, the underreporting of Medicaid participation in the AHEAD cohort is relatively modest.

Columns 4 and 7 of Table 8 show out-of-pocket medical spending. This measure includes insurance premia (including imputed Medicare Part B premia) in both datasets. Out-of-pocket medical expenditures (including insurance payments) average \$3,010 in the bottom PI quintile and \$8,250 in the top quintile of the AHEAD. In comparison, the same numbers in the MCBS data are \$4,770 and \$9,460. While out-of-pocket medical spending in the MCBS is higher than in the AHEAD, the two spending measures are fairly similar. This may be surprising, given the differences in survey methodology behind the two measures. The HRS and the MCBS each have their own advantages relative to the other dataset. There are more detailed questions underlying the out-of-pocket medical expense measures in the AHEAD than in the MCBS. HRS Respondents can use the unfolding brackets to give ranges for their medical spending, while in the MCBS they must provide a point estimate or DK.

The MCBS has the advantage that unreported out-of-pocket medical expenses will be imputed if Medicare had to pay a share of the total cost.[1]

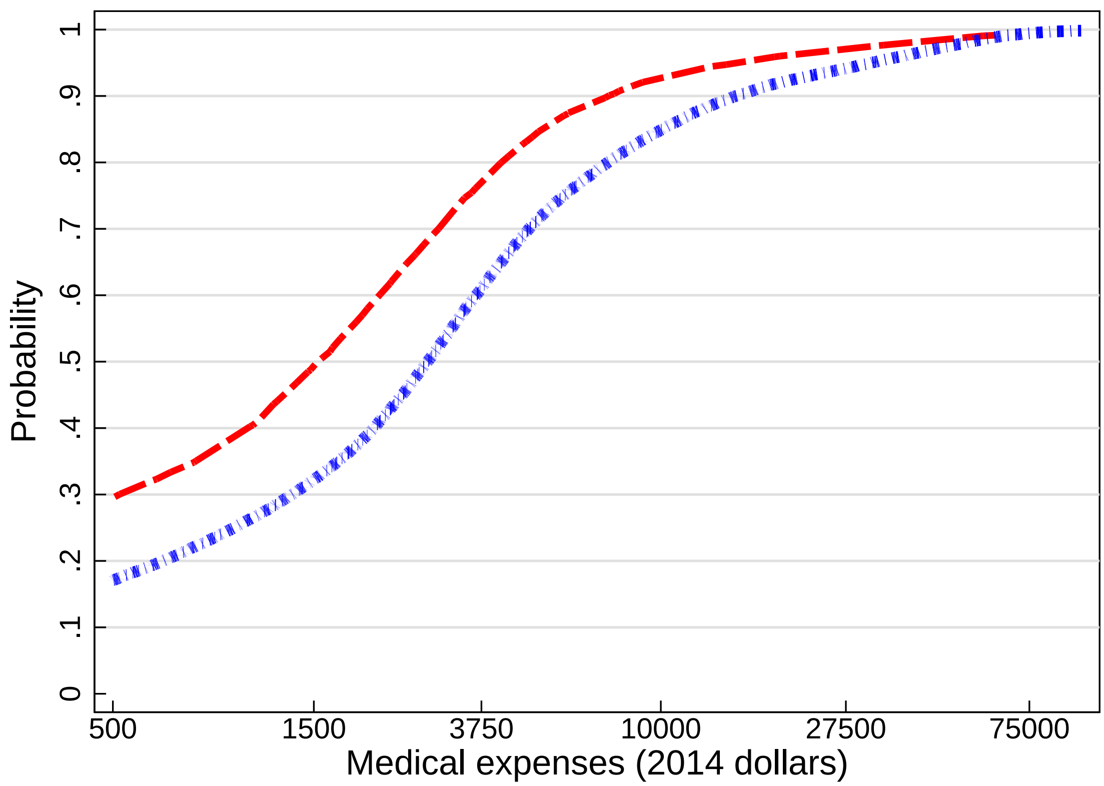

Figure 1: Distribution of Out-of-pocket Medical Expenses, AHEAD data (long dashed line) and MCBS data (short dashed line).

Turning to cross-sectional distributions, Figure 1 presents the cumulative distribution functions (CDFs) of out-of-pocket medical expenditures found in the AHEAD and MCBS data: the dashed line is the AHEAD CDF, and the dotted line is the MCBS CDF. The figure shows that, consistent with Table 8, out-of-pocket medical spending in the MCBS is higher than in AHEAD.

[1]Fahle et al. (2016) finds average out-of-pocket medical spending, excluding insurance premia, is \$2,626 for the age 65+ population, whereas De Nardi et al. (2016c) find that out-of-pocket spending for the same group is \$2,740. The estimates in Table 8 are higher because the sample are an older group of individuals, with higher medical spending.

8 Conclusion

The HRS does an excellent job constructing economic data related to the elderly population. In this paper we assess the quality of HRS variables related to out-of-pocket medical spending and health insurance premia. Comparing the HRS measures with measures from other high quality surveys show that the quality of the HRS data is high. The data from the "exit interviews" also appear to be accurate, and we encourage researchers to make more use of them. While we make a number of recommendations to improve the data, we stress that we do not recommend any changes to the survey questionnaire. Our recommendations all pertain to post-processing of the data. For example, while the source data are of high quality, the imputation procedures can, in some cases, be improved. Furthermore, much of the data are not imputed at all, which limits the usability of the data and the comparability of results across studies where the researcher is left to impute the data herself. We hope our assessments and recommendations will help make the HRS even more useful and accessible.

References

[1] Maria Jose Aragon, Martin Chalkley and Nigel Rice, "Medical Spending and Hospital Inpatient Care in England: An Analysis over Time," forthcoming in *Fiscal Studies*, 2016.

[2] Honggao Cao, *IMPUTE: A SAS Application System for Missing Value Imputations*, 2001.

[3] Stacey H. Chen and Hongwei Chuang, "Recent Trends in Taiwanese Medical Spending," forthcoming in *Fiscal Studies*, 2016.

[4] Sandy Chien, Nancy Campbell, Orla Hayden, Michael Hurd, Regan Main, Josh Mallett, Craig Martin, Erik Meijer, Michael Moldoff, Susann Rohwedder, and Patricia St. Clair, *RAND HRS Data Documentation, Version N*, 2014.

[5] Mariacristina De Nardi, Eric French and John Bailey Jones, "Why do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118(1): 37-75, 2010.

[6] Mariacristina De Nardi, Eric French and John Bailey Jones, "Medicaid Insurance in Old Age," forthcoming in *American Economic Review*, 2016a.

[7] Mariacristina De Nardi, Eric French and John Bailey Jones, “Couples and Singles’ Savings After Retirement,” mimeo, 2016b.

[8] Mariacristina De Nardi, Eric French, John Bailey Jones, and Jeremy McCauley), ”Medical Spending on the U.S. Elderly,” forthcoming in *Fiscal Studies*, 2016c.

[9] Sean Fahle, Kathleen McGarry, and Jonathan Skinner, “Out-of-Pocket Medical Expenditures in the United States: Evidence from the Health and Retirement Study,” forthcoming in *Fiscal Studies*, 2016.

[10] Eric French, Olesya Baker, Phil Doctor, Mariacristina DeNardi, and John Bailey Jones, “Right before the End: New Evidence on Asset Decumulation at the End of the Life Cycle,” *Economic Perspectives*, Third Quarter 2006.

[11] Eric French, Hans-Martin von Gaudecker and John Bailey Jones, “The Effect of the Affordable Care Act on the Labor Supply, Savings, and Social Security of Older Americans,” mimeo, 2016.

[12] Eric French and John Bailey Jones, “The Effects of Health Insurance and Self-Insurance on Retirement Behavior,” *Econometrica*, 79(3), 693-732, 2011.

[13] Dana P. Goldman, Julie Zissimopoulos and Yang Lu, “Medical Expenditure Measures in the Health and Retirement Study,” *Forum for Health Economics & Policy*, 14(3), Article 4, 2011.

[14] Michael D. Hurd and Susann Rohwedder, “The Level and Risk of Out-of-pocket Health Care Spending.” Michigan Retirement Research Center Working Paper 218 (2009).

[15] Michael D. Hurd and James P. Smith, “Anticipated and Actual Bequests,” *Themes in the Economics of Aging*, David A. Wise, editor, 2001.

[16] Samuel Marshall, Jonathan Skinner and Kathleen McGarry, “The Risk of Out-of-pocket Health Care Expenditure at the End of Life” in D. Wise (ed.), *Explorations in the Economics of Aging*, University of Chicago Press, 2011.

[17] Svetlana Pashchenko and Ponpoje Porapakkarm, “Medical Spending in the US: Facts from the Medical Expenditure Survey Dataset,” forthcoming in *Fiscal Studies*, 2016.

[18] James Poterba, Steven Venti, David A. Wise, “What Determines End-of-Life Assets? A Retrospective View,” NBER Working Paper No. 21682, 2015.

[19] Merrile Sing, Jessica S. Banthin, Thomas M. Selden, Cathy A. Cowan, and Sean P. Keehan, "Reconciling Medical Expenditure Estimates from the MEPS and NHEA, 2002," *Health Care Financing Review* 28(1), 25-40, 2006.

[20] Steven Venti, "Economic Measurement in the Health and Retirement Study," *Forum for Health Economics & Policy*, 14(3), Article 2, 2011.

Appendix A: The MCBS and MEPS Datasets

This appendix describes the MEPS and MCBS datasets.

MCBS

(This description draws heavily from De Nardi et al. (2016c), who analyze the MCBS data in some detail.)

We use the 1996 to 2010 waves of the Medicare Current Beneficiary Survey (MCBS). The MCBS is a nationally representative survey of disabled and elderly (age-65+) Medicare beneficiaries. Although the sample misses elderly individuals who are not Medicare beneficiaries, virtually everyone aged 65+ is a beneficiary. The survey contains an over-sample of beneficiaries older than 80 and disabled individuals younger than 65. We exclude disabled individuals younger than 65, and use population weights throughout.

MCBS respondents are interviewed up to 12 times over a 4-year period, and are asked about (and matched to administrative Medicare claims data on) health care utilization over 3 of the 4 years, forming panels on medical spending for up to 3 years. We aggregate the data to an annual level. These sample selection procedures leave us 66,790 different individuals who contribute 152,193 person-year observations.

The MCBS's unit of analysis is an individual. Respondents are asked about health status, income, health insurance, and health care expenditures paid out-of-pocket, by Medicaid, by Medicare, by private insurers, and by other sources. The MCBS survey data are then matched to Medicare records.

A key variable of interest is medical spending. This includes the cost of hospital stays, doctor visits, pharmaceutical, nursing home care, and other long term care. The MCBS's medical expenditure measures are created through a reconciliation process that combines survey information with administrative Medicare claims data and Medicaid recipiency data. As a result, the MCBS contains accurate data on Medicare payments and fairly accurate data on out-of-pocket, Medicaid, and other insurance payments. Out-of-pocket expenses include hospital, doctor and other bills paid out of pocket, but does not include insurance premia paid out of pocket. Because the MCBS includes information on people who enter a nursing home or die, its medical spending data are very comprehensive.

As shown in De Nardi et al. (2016c), the MCBS accurately measures the share of the population receiving Medicaid payments: the MCBS Medicaid recipiency rates of age 65+ "dual eligibles" (i.e., those who receive both Medicare and Medicaid) line up well with the aggregate statistics. The MCBS data matches aggregate Medicaid recipiency rates for the age 65+ population extremely well, which should not be surprising since MCBS uses administrative Medicaid data to create recipiency rates. However, because the MCBS does not capture those drawing Medicaid but not Medicare, the MCBS will likely understate the aggregate Medicaid recipiency rate by several percent.

MEPS

(This summary draws heavily from French et al. (2016), who analyze MEPS data in some detail.)

We use data from the 1996-2012 waves of the Medical Expenditure Panel Survey (MEPS). The MEPS is a nationally representative survey. Respondents are asked about health status, health insurance, and health care expenditures paid out-of-pocket, by Medicaid, by Medicare, private insurers and by other sources. The MEPS data are matched to information provided by providers. Although it does not capture certain types of medical expenditures, such as nursing home expenditures, it captures most of the sources of medical spending that are faced by individuals in their 50s and 60s. See Sing et al. (2006) and Pashchenko and Porapakkarm (2016) for more on comparisons between the MEPS data and the aggregate statistics.

MEPS respondents are interviewed up to 5 times over a 2 year period, forming short panels. We aggregate the data to an annual level. We use the same sample selection rules for the MEPS that we use for the HRS.