University of Michigan Retirement Research Center

Working Paper

WP 2008-174

Adequacy of Economic Resources in Retirement and Returns-to-scale in Consumption

Michael D. Hurd and Susann Rohwedder

Project #: UM07-07

Adequacy of Economic Resources in Retirement and Returns-to-scale in Consumption

Michael D. Hurd
RAND and NBER

Susann Rohwedder
RAND

May 2008

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104
http://www.mrrc.isr.umich.edu/
(734) 615-0422

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98362-5-04). The findings and conclusions expressed are solely those of the author and do not represent the views of the Social Security Administration, any agency of the Federal government, or the Michigan Retirement Research Center.

Adequacy of Economic Resources in Retirement and Returns-to-scale in Consumption

Michael Hurd and Susann Rohwedder

Abstract

Most assessments of the adequacy of retirement resources are expressed as a comparison of pre-retirement income to immediate post-retirement income. Yet, among couples a substantial fraction of retirement years is eventually spent by the surviving spouse living alone. To the extent that singles need less than couples to maintain the same standard of living, assessments of the adequacy of economic resources that make no adjustment for widowing will systematically misstate economic preparation. We estimate returns-to-scale parameters in spending by older households, using data from the Consumption and Activities Mail Survey and apply these to assessments of adequacy of retirement resources.

Authors' Acknowledgements

We are grateful for financial support from the Social Security Administration via a grant from the Michigan Retirement Research Center (UM07-07). We thank the National Institute on Aging for additional financial support for data development (1P01AG08291). The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA or any agency of the Federal Government.

Introduction

In considering whether prospective retirees will have the resources they need for retirement, a commonly used indicator is post-retirement income expressed as a proportion of pre-retirement income, that is, the replacement rate. Complete replacement of income or of some fixed fraction such as 80 percent are often put forward as a yard stick to ensuring a comfortable retirement. This kind of thinking is simplistic in that it makes no systematic accounting of a number of things: the differing role of taxes for households at different points in the income distribution; work-related expenses; financing consumption out of savings; the time horizon or survival curve of the household; returns-to-scale in consumption: couples' need to assess the risk of increased per capita spending once one of the partners dies; the changing consumption profile with age; a household's use of its increased leisure in retirement in ways that may either increase or decrease spending. For example some households may want to use their increased leisure-time to engage in activities that are associated with elevated expenses such as travel, while some may engage in home production or more efficient shopping to reduce spending. The overall goal of this paper is to define replacement rates that take into account many of these aspects.

In prior work we defined an alternative concept, the wealth replacement rate. We asked whether observed bequeathable wealth, pension income, Social Security benefits and other income sources would be able to support the observed life-cycle path of consumption from the beginning of retirement to the end of life (Hurd and Rohwedder, 2006). While we believe that paper made a useful contribution by moving away from income replacement rate toward a more accurate and comprehensive measure, it was incomplete along a number of dimensions. First, its measure of adequacy was in terms of actual bequeathable wealth holdings compared with "necessary" wealth, where necessary wealth was defined to be the minimum level of wealth necessary to carry out a life-cycle consumption plan. For example, lack of financial preparation for retirement would occur if the household's actual wealth holdings fell short of the "necessary" wealth holdings. We believe a more informative measure of adequacy would be based on the necessary adjustment to consumption because it reflects directly the required adjustment to a household's living standards in order not to run out of wealth late in life. Second, and what is the focus of this paper, is returns to scale in consumption. Among couples a substantial fraction of the total retirement years will be spent by the surviving spouse living as a single person. To the extent that a single person needs less than a couple to maintain the same standard of living, assessments of the adequacy of economic resources that make no adjustment for widowing will systematically misstate economic preparation. In our previous work we simply used the returns-to-scale parameter that is implicit in the poverty line: a couple with income that is 26% greater than the income of a single person who is just at the poverty line will also be just at the poverty line. That is, according to this scale, a couple needs 26% more income than a single person to achieve the same level of well-being.

In this paper, rather than assuming a value for the returns-to-scale parameter, we estimate it in several ways using data from the Consumption and Activities Mail Survey (CAMS). We apply these estimates to the assessment of adequacy of retirement resources developed in Hurd and Rohwedder (2006) and make several other advances over the methods in that earlier paper.

Background and Contribution to the Literature

Returns-to-scale in consumption refers to the ability of a couple to spend less than twice what a single person spends to achieve the same level of well-being. Returns-to-scale arise because the couple may share some goods such as housing services or an automobile or consumer durables. There may be returns-to-scale in household production. For example, meal production requires approximately the same amount of work for two people as for one, and the couple is likely to experience less wastage per person than the single person. Theoretically, returns-to-scale could vary between complete and none. If the couple spends the same amount as the single person yet each spouse achieves the same level of well-being as the single person, returns-to-scale are complete: two can live as cheaply as one. If the couple requires twice the spending of a single person for each spouse to achieve the same well being as a single person, there are no returns-to-scale. The truth is somewhere in between.

A common way to characterize returns-to-scale is by a returns-to-scale parameter, which is used to adjust consumption by a couple to consumption per person. If there are no returns-to-scale, the returns-to-scale parameter is 2.0: a couple requires twice the spending to achieve the level of well-being of a single person, and consumption per person is half the total consumption by the couples. If there are complete returns-to-scale, the returns-to-scale parameter is 1.0, and consumption per person is the same as consumption by the couple.

Given the trends towards less generous retirement benefits, assessments of the well being of the elderly population become an even more important ingredient to informing the policy debate as further reforms are considered, e.g., in Social Security. The estimates of returns-to-scale in consumption among the elderly that we obtain taken together with the further development of the methodology in Hurd and Rohwedder (2006) achieve a significant improvement over existing ways of assessing adequacy of retirement resources. Our estimates of returns-to-scale also have direct implications for government programs that have returns-to-scale parameters embedded in their program rules, such as Social Security or other programs that define their benefits in relation to the poverty line. For example, Social Security benefits of a widow amount to two-thirds of the benefit that the couple was receiving prior to the husband's death. The implicit returns-to-scale parameter is 1.5; that is, to achieve the same level of well-being as a single person, a couple is deemed to require 1.5 times that person's income. However, the poverty line implicitly defines the returns-to-scale parameter to be 1.26. As a consequence, a couple whose only income is Social Security benefits will experience a decline in income at widowing that is greater than the decline in the poverty line, with the result that widowing will be associated with an increase in the poverty rate. Indeed, Hurd and Wise (1997) found that aligning the reduction in Social Security benefits to the decline in poverty-level income would reduce the poverty rate of widows by eight percentage points. Our results provide estimates of realistic returns-to-scale parameters to use in the design of government programs.

As far as we know, the returns-to-scale parameter has been estimated only from systems of demand equations, by studying the way households with differing characteristics vary their spending in response to changes in prices (Donaldson and Pendakur, 2004). In these estimations, the data are cross-sections such as the CEX in the U.S., the Canadian Family Expenditure Surveys, and the Family Expenditure Survey in the U.K. This type of estimation depends on the price variation of goods to be a reflection of actual prices, not variation in quality, and on prices being determined by supply, not by demand. Both of these requirements are subject to dispute. These estimates do not allow for the returns-to-scale for two older persons to be different from the returns-to-scale among younger couples, possibly with children.

A related literature estimates the determinants of consumption in a life-cycle model and almost universally includes variables that indicate household composition.[1] Returns-to-scale are implicitly modeled because the demographic variables indicate the number and age of children and the number of adults rather than simply studying consumption per capita. Surely household composition is an important determinant of spending. In fact, variation in consumption by age is closely matched by variation in household size by age (Attanasio and Weber, 1995). However, adding demographic variables to a life-cycle model in this manner is purely descriptive, and does not reveal anything about returns-to-scale. To see this consider a one-good model of a household. Suppose that all household members are permanent household members, and that all have a subjective time rate of discount equal to the interest rate. There is no uncertainty. The desired consumption path of the household would be flat regardless of how household preferences are aggregated to produce household consumption rules. In this setting two households with the same income but with different demographics would have the same consumption path and the same saving rate; so demographics could not reveal anything about returns-to-scale.

In fact we do observe that, holding income constant, consumption levels vary with household composition. The reason is that some household members are transitory, children for example. Then the household should have high spending while transitory members are present, but the level of spending will depend on returns-to-scale. An implication is that the level will depend on both the duration of stay by family members in the household and on returns-to-scale. Without knowledge of the duration of stay, we

---

[1] For example, Gourinchas and Parker (2002); and Blundell, Browning and Meghir (1994).

cannot separate them. However, in principle we can study returns-to-scale by finding how spending varies when household members enter or leave. In the context of the older population, transitions in household composition are mostly due to widowing, which is observed with fairly high frequency in the HRS.

We adopt a very different approach, which places the problem in the dynamic framework of the life cycle model. It corresponds to the thought experiment of adding or removing a person from the household and observing how spending changes. Our study focuses on returns-to-scale among older couples, comparing their total spending as a couple with the spending of the widowed spouse. The basic idea is that the couple will choose its level of consumption so that its marginal utility of consumption will match the marginal utility of consumption of a surviving spouse. The level of consumption will depend on returns-to-scale in consumption and on the economic resources of the surviving spouse.

Widowing takes place with fairly high frequency. For example, if both spouses are 65, the probability that one of the spouses will die before both reach age 76 is about 50%, and the surviving spouse can live for a number of further years: life expectancy for a woman at age 76 is 11.4 years and for a man is 9.2 years. About 71% of 65 year-olds are married, so any assessment of the adequacy of economic resources should take into account that many of the life-years spent by married people after retirement will be as singles. The assessment depends on returns-to-scale in consumption, as well as the actual economic resources.

Returns-to-scale in a life-cycle model

Before we pass on to the formal exposition, we give an illustration with the simplest two-period model. Assume that in the first period the couple chooses consumption. At the beginning of the first period, the couple knows that at its end, the husband will die; then the widow consumes during the second period. The objective of the couple is to choose first period consumption to maximize the sum of the utility of the couple and of the widow. The interest rate and the subjective time rate of discount are both zero. Let $u(c)$ be the utility from consumption by a single person and $U(C)$ be the utility from consumption by the couple. In line with the prior discussion we specify that

$$U(C) = 2u(\tfrac{C}{k}),$$

which assumes that each spouse has the same utility function as a single person and that resources are shared equally between them; k is the equivalence scale: if it is 2.0, the utility of the couple depends on per-capita consumption and there are no returns-to-scale; if it is 1.0, the couple has twice the utility that a single person would have were that single person to consume the same total amount as the couple household.

In that there is no uncertainty, the lifetime utility maximization will equate marginal utilities as

$$U'(C) = u'(c)$$

where C is consumption by the couple in the first period, and c is consumption by the widow in the second period. Thus the couple and the widow will choose consumption such that

$$2u'(\tfrac{C}{k})\tfrac{1}{k} = u'(c).$$

The relationship between C and c will depend on k and on the marginal utility schedule. As an example suppose that the utility function is constant relative risk aversion (CRRA). Then

$$\ln(\tfrac{c}{C}) = \tfrac{1}{\gamma}(\ln k - \ln 2) - \ln k$$

In that $1 \le k \le 2$, $C > c$. Increasing γ will increase consumption by the widow relative to consumption by the couple because the marginal utility of consumption by the couple at the initial level is reduced. When $\gamma > 1$, which we take to be the base case, increasing k will reduce consumption by the widow: at the smaller returns-to-scale, more spending is needed by the couple.

If we knew γ, we could simply calculate k from the change in consumption at the death of the husband. Our first approach will be to take γ from previous empirical research. For example, Hurd (1989) estimated γ to be 1.12. Then we will calculate k directly.

Mortality Risk

When mortality risk is an important determinant of consumption, we need a model that accounts for it. We will use the structure of the life-cycle model to estimate returns-to-scale in consumption by couples.[2] When the utility function is constant relative risk aversion, the first-order condition for utility maximization is

$$\frac{d\ln C_t}{dt} = -\frac{\theta_t}{\gamma} + \frac{k^{1-\gamma} C_t^{\gamma} \Omega_t}{2\gamma}$$

where $\theta_t = h_t + \rho - r$, h_t is mortality risk, ρ and r are the subjective time-rate of discount and interest rate, and Ω_t is the expected marginal utility of wealth to survivors. The level of wealth, survivor annuities, and a bequest motive influence Ω_t, which depends on the level of consumption of the survivor. Using the definition of $\Omega_t = \mu_t c_f^{-\gamma} + f_t c_m^{-\gamma}$ where μ_t is mortality risk of the husband, $c_f^{-\gamma}$ is consumption by the surviving widow, ϕ_t is mortality risk of the wife and $c_m^{-\gamma}$ is consumption by the widower, we write

$$\frac{d\ln C_t}{dt} = -\frac{1}{\gamma} h_t + \frac{1}{\gamma}(r-\rho) + \frac{1}{\gamma}\frac{k^{1-\gamma}}{2\gamma} C_t^{\gamma} (\mu_t c_f^{-\gamma} + f_t c_m^{-\gamma})$$

Assume $c_m = c_f$ and rewrite as $\frac{d\ln C_t}{dt} = -\frac{1}{\gamma} h_t + \frac{1}{\gamma}(r-\rho) + \frac{1}{\gamma}\frac{k^{1-\gamma}}{2\gamma} C_t^{\gamma} c_t^{-\gamma} (\mu_t + f_t)$

or $\frac{d\ln C_t}{dt} = -\frac{1}{\gamma} h_t + \frac{k^{1-\gamma}}{2\gamma} h_t C_t^{\gamma} c_t^{-\gamma} + \frac{1}{\gamma}(r-\rho)$

or $\frac{d\ln C_t}{dt} = \alpha_1 + \alpha_2 h_t + \alpha_3 (h_t C_t^{\gamma} c_t^{-\gamma})$

where the interpretation of $\alpha_2 = -\frac{1}{\gamma}$ and the interpretation of $\alpha_3 = \frac{k^{1-\gamma}}{2\gamma}$

[2] See Hurd (1999) for an explication of the model and derivation of these results.

The slope of the consumption path depends on α_2 and α_3 which in turn depend on the two parameters γ and k. There are three right-hand variables, the constant, h_t and $h_t C_t^{\gamma} c_t^{-\gamma}$. From observations on couples prior to widowing and on the surviving spouse we have observations on $C_t^{\gamma} c_t^{-\gamma}$ and we take data on h_t from life tables. We estimate γ and k directly from panel data on consumption. We use three waves of the Consumption and Activities Mail Survey (CAMS) and multiple waves of the HRS.

Data

Our analyses are based on data from the Health and Retirement Study (HRS) and from the Consumption and Activities Mail Survey (CAMS). The HRS is a biennial panel. Its first wave was conducted in 1992. The target population was the cohorts born in 1931-1941 (Juster and Suzman, 1995). Additional cohorts were added in 1993 and 1998 so that in 2000 it represented the population from the cohorts of 1947 or earlier. In 2004 more new cohorts were added, making the HRS representative of the population 51 or older.

In September 2001, CAMS wave 1 was mailed to 5,000 households selected at random from those participating in HRS 2000. In households with couples it was sent to one of the two spouses at random. The fact that the sample was drawn from the HRS 2000 population allows linking the CAMS data to the vast amount of information collected in prior waves in the core survey on the same individuals and households. In September 2003 and in October 2005, CAMS wave 2 and wave 3 were sent to the same households. Wave 3 was also sent to an additional 850 households representing the new cohort of 51-56 year-olds who were inducted into HRS in 2004.[3] To facilitate panel analysis, the structure of the questionnaire was almost the same across waves. In this paper we will use data from all three waves. Descriptive statistics of data quality are similar across waves. We will therefore restrict their discussion to the first wave of CAMS.

[3] We do not use these new households in the analysis of this paper because we do not yet have panel data on them: their second wave is CAMS 2007, which is not yet available.

CAMS wave 1 consists of three parts. In Part A, the respondent is asked about the amount of time spent in each of 32 activities such as watching TV or preparing meals. Part B collects information on monetary expenditures in each of 32 categories, as well as anticipated or recollected spending change at retirement (Hurd and Rohwedder, 2005). Part C asks about prescription drugs, current marital status, and current labor force status.

The instructions requested that for Part B, the person most knowledgeable about the topics be involved in answering the questions. The addressee answered Part B in 88% of households, possibly with the assistance of the spouse; 5% of the cases report explicitly that the spouse answered the questions; 2% had the children or children-in-law of the addressee help out in answering the questions, and the remaining 5% was a mix of miscellaneous responses including nonresponse.

Of the 5,000 questionnaires mailed in 2001, there were 3,866 returned questionnaires for a unit response rate of 77.3 percent. (The second wave of CAMS had a unit response rate of 78.3 percent, not adjusted for mortality and undeliverable questionnaires).[4]

The Consumer Expenditure Survey (CEX) is the U.S. survey that collects the most detailed and comprehensive information on total spending. But CAMS could not ask about spending in as many categories as the CEX – 260 in the survey's recall component. Given that limitation, CAMS would then ideally have chosen spending categories starting from the CEX aggregate categories that are produced in CEX publications, so as to have direct comparability with the CEX. However, to reduce the burden to respondents, the categories had to be aggregated further. The final questionnaire collected information on 6 big-ticket items (automobiles; refrigerators; washers and dryers; dishwashers; televisions; computers) and on 26 non-durable spending categories.

The reference period for the big-ticket items is "last 12 months." For the non-durables it varied: the respondent could choose the reference period between "amount spent monthly" and "amount spent yearly" for regularly occurring, relatively invariant expenditures like mortgage, rent, utilities, insurance, and property taxes, and between

[4] A total of 4,156 questionnaires were mailed out for the second wave of CAMS in 2003, resulting in 3,254 returned questionnaires. The remainder of the original sample was lost due to death (n=372), due to loss to follow-up (n=173), and some respondents (n=298) participated in another HRS supplemental study and were therefore excluded from CAMS wave 2.

"amount spent last week," " amount spent last month," and "amount spent in last 12 months" for all other categories.[5] For all non-durable categories there was a box to tick if "no money spent on this in last 12 months." To refrain from inviting item nonresponse, the questionnaire had no explicit provision for "don't know" or "refuse."

The rate of item nonresponse was very low, in the single digits for most categories.[6] The maximum rate was 13.8%. For some of the categories, we imputed item nonresponse from HRS core data. For example, spending for rent had a relatively high rate of item nonresponse (13.2%), but almost all was by households who, according to HRS, were home owners. Thus with considerable confidence we imputed zero rent to such households. Because item nonresponse was so low, total imputed spending was a small fraction of total estimated spending, just 6.0 percent.[7]

Wave 2 of CAMS had the same spending categories as wave 1, augmented by personal care products and services, and gardening and housekeeping services. These amounted to 3.1% of total spending for households age 55 and above according to the CEX. Wave 3 of CAMS had the same categories as wave 2 with the addition of household furnishings and equipment, which accounts for 3.7 percent of total spending of households age 55 and above according to the CEX. Our panel comparisons are spending change found by comparing spending prior to widowing with spending after widowing, and comparing the spending change to that observed among singles who did not experience any widowing (or remarriage) or to couples who did not experience any widowing. For this comparison we always use just those categories that are measured in both pre- and post-retirement waves. For example, in calculating spending change for those who experienced widowing between waves 1 and 2, we exclude personal care products and services, as well as gardening and housekeeping services, from the wave 2 measure. Following this method, we construct panel measures of non-durable spending.

---

[5] In CAMS wave 2 and 3 the "last week" option was eliminated to reduce the risk of observing outliers that arise from unusually high values reported as "last week" that are subsequently multiplied by 52 to arrive at annualized values.

[6] The rates of item nonresponse are similar on other waves of CAMS.

[7] In CAMS wave 2, the fraction of total spending that was imputed amounted to 5.0 percent and to 5.5 percent in CAMS wave 3.

Widowing

Using information on marital status from section C in CAMS we find those households that were coupled in one wave and widowed in a subsequent wave of CAMS.

We observe that there are a total of 144 households where widowing occurs over the period covered by the CAMS surveys:

1) 57 widowed between CAMS 2001 and CAMS 2003
2) 19 widowed between CAMS 2001 and CAMS 2005
3) 68 widowed between CAMS 2003 and CAMS 2005

Cross-checking information on household size in the HRS core data, we find that a number of couples were not living by themselves or that the surviving spouse was not living alone. Because we are interested in returns-to-scale and our identification strategy focuses on the case where household size drops from two to one due to widowing we restrict our sample accordingly. That leaves us with 83 households; for 72 of these we also observe the date of the widowing in the HRS core survey.

For the model-based estimations we limit the sample further to people who were observed in CAMS 2001 and widowed between CAMS 2003 and CAMS 2005, resulting in a total of 62 observations of which 40 also meet the household size criterion experiencing a household size reduction from two to one.

Consumption Measure

In our analysis we approximate consumption by household spending on non-durable goods. We exclude spending on durables to avoid noise or biases resulting from the lumpy nature of durable purchases which would introduce apparent jumps in our consumption measure that are not reflected in actual consumption as the services from durables are enjoyed over several years (e.g., a car, fridge, TV, etc.).

Timing of consumption measure and widowing

Our consumption measure covers the household's spending over the last twelve months and it is observed every two years (see green lines in time line in Figure 1).

Figure 1:

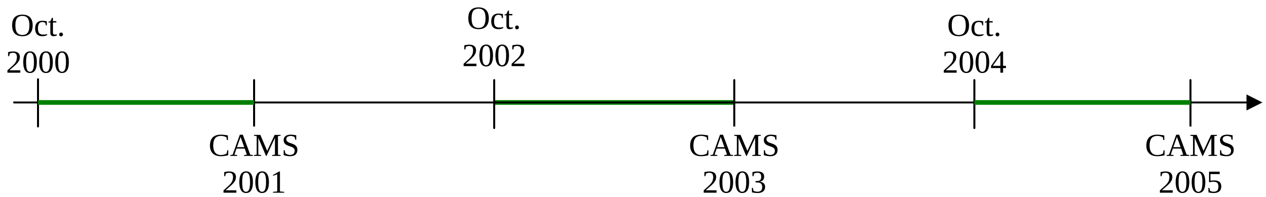

We obtain information on the timing of the widowing from the HRS core data. For constructing the observations of declines in consumption due to widowing we distinguish four different cases:

Widowing occurred	**N**	**Consumption measure of couple**	**Consumption measure of survivor**
1) after CAMS 2001 interview, but before October 2002	14	CAMS 2001	CAMS 2003
2) in or after October 2002, but before CAMS 2003	18	CAMS 2001	CAMS 2005
3) after CAMS 2003, but before October 2004	20	CAMS 2003	CAMS 2005
4) in or after October 2004, but before CAMS 2005	20	CAMS 2003	CAMS 2005

When comparing the timing of the spending measurements for each of the four cases with the time line above, one can see easily that for cases (1) and (3) we observe spending of the couple in the CAMS wave immediately preceding the widowing and the spending of the survivor in the next wave of CAMS. Cases (2) and (4) need some additional consideration because the widowing occurs during the reference period of the next observed spending measure of the household, that is, during the 12 months preceding the next CAMS survey. In case (2) we therefore use the after-next CAMS observation for measuring the spending of the survivor, as that is the next measurement of household spending that reflects the spending only of the survivor. For case (4) this is not an option, because CAMS 2007 is not available yet. So we use CAMS 2005, but note that this will bias our results slightly towards finding a smaller drop or larger returns-to-scale.[8]

Estimates of Returns-to-Scale

Raw data

Table 2 shows the drops in spending observed in the raw data for households upon widowing. We follow the scheme laid out above for the four groups based on the timing of the widowing with respect to the observations on spending in CAMS. Time *t* refers to household spending immediately before widowing takes place; time *t+1* is spending of the survivor, using the first observation on spending that fully reflects that of the survivor (i.e., neither period for reported spending includes the time of widowing, except for group 4). Whether we look at the change in the population median of spending before and after widowing or at the household-level change at the median, both imply a drop in spending of about 25 percent.[9] We take this to be our first estimate of returns-to-scale from the raw data. Alternative estimates based on the raw data come from the differential drop in spending observed among widowed households and simultaneous drops in spending among a control group of households where marital status and household size did not change. We consider two such control groups: (a) single households living alone; and (b) couple households living alone. The rationale for taking the differential is that spending might have dropped even in the absence of widowing and we want to

[8] The bias depends on (1) the fraction of the reference period of the post-widowing spending measure that the household was still a couple; and (2) the relative weight of those spending categories for which the survivor in the household used the long reference periods for reporting spending. If s/he reported most spending in terms of "last month" rather than "last 12 months" then the spending would usually only refer to the survivor and the bias may not actually be all that big. In future research we will verify this.

[9] The change in the population mean is susceptible to measurement error due to outliers, so we do not discuss it here.

identify that portion of the drop in spending that is due to widowing. The second panel in Table 2 shows spending changes for singles of about 8 percent, both at the population median and household-level median. Differencing with the drop observed among widowed households, we find that widowing appears to lead to a drop in spending of about 16 percent. This differential is subject to the caveat that singles might not be a suitable control group. Having lived alone all along, singles tend to be poorer than those who only just recently widowed. For couples living alone we find declines in spending of 3.5 percent (third panel in Table 2) so that the differential drop for widowed households would amount to 21 percent.

Results from estimations

Based on the raw data the observed ratio of spending by a surviving spouse to that by a couple is about 0.75, which corresponds to a *k* of 1.33. If the ratio is based on the differential with respect to the control group of singles it is 0.84 and k = 1.19; and with respect to the control group of couples it is 0.79 and k=1.27.

Using $\ln(\frac{c}{C}) = \frac{1}{\gamma}(\ln k - \ln 2) - \ln k$ and assuming $\gamma = 1.12$ we calculate $k = 0.025$. This value is not credible as k should be between 1.0 and 2.0. The problem is not in our particular values of the spending ratio or in our assumed values of γ or k: there are no "reasonable" combinations of γ and k that will produce the observed (and reasonable) value of c/C. This is illustrated in Table 3 where we have used $\ln(\frac{c}{C}) = \frac{1}{\gamma}(\ln k - \ln 2) - \ln k$ to calculate hypothetical values of c/C for a number of combinations of γ and k. No combinations of γ and k in the table produce a consumption ratio of 0.80. For $\gamma = 1.3$ and $k = 1.2$ which seem *a priori* to be reasonable values, the consumption ratio would be just 0.563. That is, the surviving spouse would consume just 56% of the consumption by the couple. The conclusion we draw is that the problem needs to be put in the dynamic situation with uncertainty about survival.

Thus we estimate $\frac{d\ln C_t}{dt} = \alpha_1 + \alpha_2 h_t + \alpha_3 (h_t C_t^{\gamma} c_t^{-\gamma})$ or rather the discrete time version

$$\frac{C_t - C_{t-1}}{C_{t-1}} = \alpha_1 + \alpha_2 h_t + \alpha_3 (h_t C_t^{\gamma} c_{t+1}^{-\gamma})$$

The data requirements are that consumption by a couple, C, is observed in two successive panels and that widowing takes place between the second and third waves when we observe c. This means that our sample size is very small. To reduce the influence of outliers we estimate by median regression. Conditional on γ our estimating equation is linear in regressors. So, our method is to choose various values of γ, estimate the α conditional on the value chosen, and then to choose the optimal γ based on a global criterion such as R^2. Table 4 has the results of those conditional regressions based on 62 observations.

The right-hand variables are significant, but even so, they are estimated very imprecisely. For example, the interpretation of the coefficient on h_t is $-1/\gamma$. Take, for example, the estimate conditional on $\gamma = 1.12$. The 95% confidence interval for α_2 is (-4.81, -0.11), which implies a 95% confidence interval for γ of (8.78, 0.21). The assumed value of γ is contained in this interval, but so are many unreasonable values. The coefficient on $h_t C_t^{\gamma} c_{t+1}^{-\gamma}$ has the interpretation of $\frac{k^{1-\gamma}}{2\gamma}$. The last column of the table has the implied estimate of k using as γ the value implied by α_2. For conventional values of γ (values greater than 1) the value of k is not plausible. Of course, "conventional" values of γ are not necessarily correct.

Among the values considered here, the pseudo R^2 takes its maximum value at $\gamma = 0.90$. It is possible that even smaller values of γ would lead to a greater R^2. But given the imprecision of the estimates, we will await further data before further exploring the optimal values of γ and k.

Consumption-based assessments of adequacy of retirement resources

In Hurd and Rohwedder (2006) we derived a method of resource-based assessments of adequacy of retirement resources. But our estimate of the returns-to-scale parameter was not based on an estimated model. Having now estimated k, we use it to re-estimate the adequacy of resources in retirement, but only after first implementing several extensions of our prior work.

Much of the basic methodology remains the same, but the metric of adequacy assessment will be different: our approach relies on simulating consumption paths over the remaining life cycle for a sample of households observed shortly after retirement. For this purpose we need the initial level of consumption, which we observe directly in the CAMS data, and the slope of the consumption path, which we estimate from observed panel transitions based on CAMS waves 1 to 2 and 2 to 3.

For the simulations we construct life-cycle consumption paths for each household: we begin with the observed consumption level at retirement age and then apply the observed rates of change to trace out a life-cycle path whose slope is given by the estimated rates of change. Whereas a model would specify that the slope of the consumption path depends on the interest rate, the subjective time-rate of discount, mortality risk, and utility function parameters, we estimate these slopes directly from the data. Practically all model estimation uses the constant-relative-risk-aversion utility, which specifies that the slope of log consumption is independent of the level. The observed paths do not necessarily have that shape.

In this paper we propose a consumption-based replacement rate as a measure of the adequacy of retirement resources. We will start out by illustrating the approach for singles.

We observe the resources at retirement of a single person. We ask: can the resources support the projected consumption path? The consumption path is anchored at the initial post-retirement consumption level and follows the path given by the slopes of consumption paths that we have estimated from the CAMS panel. If the consumption path cannot be supported by the economic resources, we find the initial level of consumption that would permit the person to follow the life-cycle path. The consumption replacement rate is the ratio of the feasible of affordable consumption to the actual consumption. If the replacement rate is greater than one, economic resources are more than sufficient to finance the actual consumption path. If it is less than one, there is a shortfall in resources.

Because lifetime is uncertain, and wealth is not typically annuitized, we also ask whether the observed initial consumption level permits the person to follow the life-cycle path with a high degree of probability. Here the uncertainty is length of life, so the question is equivalent to finding whether the resources will sustain the path until advanced old age where the probability of survival is very small. Someone with a moderate level of pre-retirement consumption could sustain post-retirement consumption with a moderate level of Social Security benefits, some pension income, and a moderate amount of wealth. Someone with low pre-retirement consumption may need only Social Security and a small amount of savings. These requirements are likely to differ substantially from what would be required to consume at the pre-retirement *income* level.

We do these calculations for each single person in our CAMS sample who is in his or her early retirement years.

For couples the basic method is similar. However, the consumption path followed while both spouses survive will differ from the consumption path of single persons, so it is separately estimated from the CAMS data. The couple will follow that consumption path as long as both spouses survive, and then the surviving spouse will switch to the consumption path of a single person. The shape of the single's path is estimated as described above, but the level will depend on returns-to-scale in consumption by the couple. At the death of the first spouse, the surviving spouse reduces consumption to the level specified by the returns-to-scale parameter. In our prior work, we assumed a returns-to-scale parameter that is consistent with the poverty line, which for a couple at the poverty line implies an income 1.26 times the income of a single person at the poverty line. This implies that consumption by the surviving spouse should be 79% of consumption by the couple to equate effective consumption.

Knowing the consumption path of the surviving spouse, we find the expected present value of consumption for the lifetime of the couple and the surviving spouse. We also determine the fraction of households that can finance their expected consumption path with, say, 95% probability, and by how much a household would have to adjust consumption to keep the chances of running out of wealth towards the end of the life cycle reasonably small.

Results

Because we want to observe Social Security and pension income, we select a sample shortly after retirement and of a sufficient age that they should be receiving Social Security if they are eligible. For each of the three waves of CAMS we select couples in which one spouse is 66 to 69 and the other is 62 or older. We make the age restriction on the younger spouse because spouses younger than 62 would not yet be receiving Social Security benefits and so we would miss a significant fraction of retirement resources. In addition we require that the couple was also stably married in the adjacent HRS core interviews so that other information of the couple that we take from the HRS core survey pertains to the same persons in the couple. If the same couple household fits the sample criteria for more than one wave of CAMS we use the latest observation as information on pension and Social Security income will be more complete. For singles, we proceed accordingly and select those who were 66-69 in any wave of CAMS; impose the same restriction of stable marital status in adjacent HRS waves and retain the latest wave if the person meets the criteria in multiple waves of CAMS. Our analytical sample has 478 single households and 757 couple households.

We perform 100 simulations of the consumption and wealth paths of each single person who is in the age range 66-69.

Tables 5 and onward show the results of the simulations. Because we are interested in the fraction of individuals that runs out of resources at the end of the lifecycle, we have arranged all subsequent tables at the individual level. They show the results for 65-69-year-old singles and for 65-69-year-olds living in couple households at baseline.

Our individual-level metric for adequate preparation is based on the following concept.

By how much does the household have to adjust initial consumption compared with current initial consumption to keep the probability of running out of wealth at the end of life below a desired threshold?

We set our adjustment threshold to -15%. That is, in a particular simulation, someone is adequately prepared if he can afford initial consumption that is at least 85% of actual initial consumption. Overall we say that the individual is adequately prepared if the chances are 95% or greater that he can afford this initial adjusted level of consumption.

For couples we mean the consumption by the couple as long as both spouses survive and the subsequent consumption by the survivor. Although we begin with 757 households as shown in Table 8, we have only 924 married persons who are age-eligible (66-69), the other spouses being outside the given age range. The economic circumstances of the 924 age-eligible persons will enter the tables. In the first set of simulations we use the poverty line returns-to-scale and thus assume that the annuity of the survivor is 0.79 times the annuity of the couple.

Table 5 shows that among singles, consumption could be increased on average at all education levels. For example, among those with less than a high school education, average actual initial consumption is $19.1 thousand, whereas average affordable consumption is $23.0 thousand. At the individual level, the average of the ratio of affordable to initial consumption is 1.87; however, this ratio is influenced by a few outliers who, because of observation error, have small initial consumption. A better measure is the median of the affordable consumption ratio. In the lowest education band this is just 1.05, showing that the typical person can just afford his or her initial consumption level. An implication is that the typical single person would have to reduce initial consumption in about half of the simulations. At the other education levels the ratios are much higher indicating that actual initial consumption is more consistent with the available resources.

As indicated above, we put consumption shortfalls or excesses in a probabilistic framework. Table 6 shows that for singles, 74 percent are adequately prepared according to our definition. But as would be suggested by Table 5, just a little more than half of those in the lowest education band are adequately prepared.

Our definition of adequate preparation makes some ad hoc choices regarding the cut off points. It is not clear how small the chances of running out of wealth should be kept. We have presented results for a cut off of 5 percent or less, but some might argue that this could also be higher or possibly smaller. Similarly we have chosen a required reduction of initial consumption by 15 percent or more to signal inadequate preparedness. We have tested the sensitivity of our results with respect to these cut off points. Table 7 shows a matrix with different cut off points. The results are surprisingly insensitive to these definitions, especially with respect to the cut-off for the *probability* of having to reduce consumption. The reason is that most households either fall substantially short of the thresholds of adequacy or they exceed them by a large margin, resulting in floor and ceiling effects in the statistics for preparedness.

Our baseline simulations for couples use the returns-to-scale parameter implicit in the poverty line; that is, a couple needs 26 percent more than a single person to achieve the same level of well-being, which implies that widows or widowers will consume 79.4 percent as much as the couple did prior to widowing. We actually observed in the data a median drop of 21.1 percent which is not materially different from the drop implied by the poverty line.   As shown in Table 8, the average affordable consumption is $98.5 thousand, yet average initial consumption is just $43.0 thousand. Thus on average couples could increase their consumption substantially. Even the median of the individual ratios of affordable consumption to actual consumption is 1.84. Although there is a gradient by education level, among those lacking high school graduation, the median of the ratio of affordable to actual consumption is 1.55

Table 9 shows that for married persons about 87 percent are adequately prepared, and that females and males are about equally likely to be prepared. As with singles our overall results are not sensitive to the cut-off points we have used in our definitions of adequate preparation (Table 10).

The remaining tables repeat the information in Tables 8, 9 and 10, but with different assumptions about the returns-to-scale parameter. Thus, Tables 11, 12, and 13 are based on the assumption that a single person needs just half of the consumption by a couple to achieve the same utility; that is, there are no returns to scale. Tables 14, 15, and 16 are based on the assumption that a single person needs the same amount of consumption as a couple to achieve the same utility; that is, there are complete returns-to-scale, for a parameter of 1.0. These results bound the outcomes.

If singles need just 50 percent of the consumption by couples, 88.7 percent of couples are adequately prepared (Table 12). If they need the same amount, 84.0 percent are adequately prepared (Table 15). The difference is in the adjustment to consumption following the death of a spouse: if singles need the same amount of couples, there is no change at the death of the spouse; if they need just 50 percent, consumption drops by that amount, making the overall path more affordable.

We conclude that the fraction of couples adequately prepared for retirement is not very sensitive to the amount of returns-to-scale.

Conclusions

We have taken a novel approach to estimating returns-to-scale in household spending, by using the event of widowing for identification. From raw data, we found drops in spending ranging between 16 and 25 percent at widowing, which is fairly close to the returns-to-scale implied by the difference in the poverty lines for couples and singles. From model-based estimations we found the importance of accounting for uncertainty of survival. Overall we found that the number of households observed becoming widowed during the period covered by CAMS surveys is so far too small to obtain informative model-based estimates on key parameters, including the returns-to-scale parameter.

Using consumption-based replacement rates, we assessed the adequacy of financial preparation for retirement. This expanded on and refined the approach in Hurd and Rohwedder (2006) along a number of important dimensions. Most notably, the metric of assessment was rescaled to reflect necessary changes to current living standards of the household, i.e. adjustments to household spending shortly after retirement.

We conducted simulations for several different values of the returns-to-scale parameter. Results were not very sensitive to this parameter, because most couples are very well prepared financially for retirement in that their retirement resources exceed substantially what they need to maintain their consumption path into advanced old age.

References

Donaldson, David and Krishna Pendakur, 2004, “Equivalent-Expenditure Functions and Expenditure-Dependent Equivalence Scales,” *Journal of Public Economics*, 88 (January), 175-208.

Gan, Li, Guan Gong, Michael Hurd and Daniel McFadden “Subjective Mortality Risk and Bequests,” with , NBER Working Paper 10789, 2004.

Hurd, Michael D. and David Wise, 1997, "Changing Social Security Survivorship Benefits and the Poverty of Widows," in Hurd and Yashiro, eds., *The Economic Effects of Aging in the United States and Japan*, University of Chicago Press, pp. 319-332.

Hurd, Michael D. 1995, "Mortality Risk and Consumption by Couples," presented at the IFS-Bank of Portugal Conference on The Microeconomics of Saving and Consumption Growth, Lisbon, November, 1995, revised, January, 1997, and NBER Working paper 7048, March, 1999.

Hurd, Michael D. and Susann Rohwedder, “Alternative Measures of Replacement Rates,” MRRC Working Paper No. 2006-132.

Table 1: Age distribution of households experiencing widowing between 2001 and 2005, restricting household size*

Age of surviving spouse	***Frequency***	***Percent***
50-54	1	1.2
55-59	4	4.8
60-64	8	9.6
65-69	12	14.5
70-74	14	16.9
75-79	17	20.5
80-84	20	24.1
85-89	7	8.4
All	83	100.0

Author's calculations.
* Restricting the sample to those who were living alone, i.e., household size changes from two to one due to widowing; no other household members.

Table 2: Total spending change of widowed households and of singe households. Amounts in thousands of 2003 dollars. N = 72.

		Mean	Median	Household-level median 2-year change
Widowed	t	29.6	26.2	
	t+1	28.0	19.8	
	% change	*-5.5*	*-24.6*	*-24.6*
Singles	t	23.9	19.3	
	t+1	21.2	17.7	
	% change	*-11.3*	*-8.5*	*-8.0*
Couples	t	23.9	31.2	
	t+1	21.2	30.2	
	% change	*-11.3*	*-3.2*	*-3.5*

Table 3
Hypothetical ratio of consumption by surviving spouse to consumption by couple for various values of k and γ. No mortality uncertainty.

	γ						
k	1.1	1.3	1.5	1.7	1.9	2.1	2.3
1.0	0.533	0.587	0.630	0.665	0.694	0.719	0.740
1.2	0.524	0.563	0.593	0.617	0.637	0.653	0.667
1.4	0.516	0.543	0.563	0.579	0.592	0.603	0.612
1.6	0.510	0.526	0.539	0.548	0.556	0.562	0.567
1.8	0.505	0.512	0.518	0.522	0.526	0.528	0.531
2.0	0.500	0.500	0.500	0.500	0.500	0.500	0.500

Table 4
Median regression estimates of $\Delta C_t / C_t$ conditional on γ

	Right-hand variable						
	mortality risk, h_t		$h_t C_t^{\gamma} c_{t+1}^{-\gamma}$		Constant		
gamma	estimate	SE	estimate	SE	estimate	SE	Implied k
0.9	-2.86	1.34	1.62	0.80	0.07	0.13	1.21
0.95	-2.75	1.29	1.51	0.75	0.07	0.13	1.16
1.05	-2.57	1.22	1.33	0.65	0.07	0.13	1.06
1.1	-2.49	1.19	1.25	0.61	0.07	0.13	1.00
1.12	-2.46	1.17	1.22	0.60	0.07	0.14	0.98
1.15	-2.42	1.16	1.18	0.58	0.07	0.14	0.95
1.2	-2.36	1.13	1.11	0.54	0.07	0.14	0.90
1.25	-2.30	1.11	1.05	0.51	0.07	0.14	0.85
1.3	-2.04	1.00	0.80	0.38	0.07	0.14	0.62

Table 5: Baseline Estimates for singles; 100 simulations.

	N	mean initial wealth	mean PV of future earnings	mean PV annuities	Mean initial consumption	Mean affordable consumption	Mean Ratio: affordable/ initial consumption	Median ratio: affordable/ initial consumption
Less than high-school	129	51.3	8.6	101.4	19.1	23.0	1.87	1.05
High-school	188	202.7	15.3	183.5	24.0	55.0	2.31	1.58
Some college	98	274.4	26.8	190.9	28.1	62.7	2.37	1.78
College and above	63	422.6	33.7	276.8	36.7	80.6	2.25	1.58
All	478	205.5	18.3	175.2	25.2	51.3	2.20	1.47

Table 6

Percent of single persons adequately prepared.

Chances are 5 percent or less that household would need to reduce consumption

By more than 15 percent

	N	All	Males	Females
Less than high-school	129	57.36	74.07	52.94
High-school	188	80.85	78.00	81.88
Some college	98	79.59	86.36	77.63
College and above	63	79.37	75.00	80.85
All	478	74.06	78.26	72.73

Table 7

Percent of single persons adequately prepared.

Chances are x percent or less that household would need to reduce consumption

by more than y percent

	Drop in consumption		
Chances	< 5 percent	< 10 percent	< 15 percent
<=5 percent	66.9	70.1	74.1
<=10 percent	66.9	70.5	74.5
<=15 percent	67.6	71.1	74.9
<=20 percent	68.0	71.5	75.1

Table 8: Baseline Estimates for couples, singles consumption 79.4% of consumption by couples ; 100 simulations.

	N	mean initial wealth	mean PV of future earnings	mean PV annuities	Mean initial consumption	Mean affordable consumption	Mean Ratio: affordable/ initial consumption	Median ratio: affordable/ initial consumption
Less than high-school	165	220.8	19.7	245.8	30.7	50.5	1.92	1.55
High-school	406	437.4	23.4	407.6	37.6	81.1	2.33	1.80
Some college	175	1,021.3	35.9	493.3	50.3	134.9	2.56	1.97
College and above	178	1,253.1	89.0	631.7	59.6	146.7	2.46	2.19
All	924	666.5	37.7	438.1	43.0	98.5	2.32	1.84

Table 9

Percent of married persons adequately prepared.

Chances are 5 percent or less that household would need to reduce consumption

By more than 15 percent

	N	All	Males	Females
Less than high-school	165	78.8	79.2	78.4
High-school	406	86.5	86.5	86.4
Some college	175	93.1	88.2	96.3
College and above	178	91.6	92.4	90.7
All	924	87.3	86.8	87.8

Table 10

Percent of married persons adequately prepared.

Chances are x percent or less that household would need to reduce consumption

by more than y percent

	Drop in consumption		
Chances	< 5 percent	< 10 percent	< 15 percent
<=5 percent	82.5	85.0	87.3
<=10 percent	83.0	85.7	88.0
<=15 percent	83.3	86.1	88.3
<=20 percent	83.8	86.8	88.6

Table 11: Baseline Estimates for couples, singles consumption 50% of consumption by couples; 100 simulations.

	N	mean initial wealth	mean PV of future earnings	mean PV annuities	Mean initial consumption	Mean affordable consumption	Mean Ratio: affordable/ initial consumption	Median ratio: affordable/ initial consumption
Less than high-school	165	220.8	19.7	245.7	30.7	57.2	2.17	1.69
High-school	406	437.4	23.4	407.3	37.6	91.6	2.63	2.00
Some college	175	1,021.3	35.9	492.6	50.3	151.6	2.87	2.16
College and above	178	1,253.1	89.0	631.0	59.6	161.4	2.70	2.39
All	924	666.5	37.7	437.7	43.0	110.3	2.60	2.05

Table 12

Percent of married persons adequately prepared.
Chances are 5 percent or less that household would need to reduce consumption
by more than 15 percent

	N	All	Males	Females
Less than high-school	165	80.6	80.5	80.7
High-school	406	87.7	87.2	88.0
Some college	175	94.3	91.2	96.3
College and above	178	93.3	93.5	93.0
All	924	88.7	88.1	89.2

Table 13

Percent of married persons adequately prepared.
Chances are x percent or less that household would need to reduce consumption
by more than y percent

	Drop in consumption		
Chances	< 5 percent	< 10 percent	< 15 percent
<=5 percent	83.5	86.5	88.7
<=10 percent	84.0	87.3	89.2
<=15 percent	84.3	87.7	89.6
<=20 percent	85.2	87.9	89.8

Table 14: Baseline Estimates for couples, singles consumption 100% of consumption by couples; 100 simulations.

	N	mean initial wealth	mean PV of future earnings	mean PV annuities	Mean initial consumption	Mean affordable consumption	Mean Ratio: affordable/ initial consumption	Median ratio: affordable/ initial consumption
Less than high-school	165	220.8	19.7	245.9	30.7	46.9	1.78	1.44
High-school	406	437.4	23.4	407.7	37.6	75.3	2.16	1.67
Some college	175	1,021.3	35.9	495.1	50.3	125.9	2.38	1.83
College and above	178	1,253.1	89.0	631.5	59.6	138.8	2.33	2.09
All	924	666.5	37.7	438.5	43.0	92.0	2.17	1.72

Table 15

Percent of married persons adequately prepared.

Chances are 5 percent or less that household would need to reduce consumption by more than 15 percent

	N	All	Males	Females
Less than high-school	165	75.8	75.3	76.1
High-school	406	83.7	85.1	82.9
Some college	175	86.9	80.9	90.7
College and above	178	89.3	91.3	87.2
All	924	84.0	83.9	84.0

Table 16

Percent of married persons adequately prepared.

Chances are x percent or less that household would need to reduce consumption by more than y percent

	Drop in consumption		
Chances	< 5 percent	< 10 percent	< 15 percent
<=5 percent	78.9	81.8	84.0
<=10 percent	80.0	82.3	85.3
<=15 percent	81.3	82.6	85.6
<=20 percent	82.0	83.1	86.6