

The Effect of Inheritance Receipt on Retirement

Jeffrey R. Brown, University of Illinois at Urbana-Champaign and NBER

Courtney C. Coile, Wellesley College and NBER

Scott J. Weisbenner, Michigan State University and NBER

September 2008

JEL: J14, J26

Abstract:

This paper provides new evidence on how wealth shocks influence retirement behavior. Economic theory generally posits that leisure is a normal good, yet it is difficult to obtain reliable empirical estimates of the wealth effect because wealth is correlated with numerous unobservable characteristics that affect labor supply. We use inheritance receipt as a wealth shock, and find that it is associated with a significant increase in the probability of retirement, especially when the inheritance is unexpected. This evidence has important implications for how public policies, such as pension or tax reform, may influence retirement behavior through the wealth effect.

We thank an anonymous referee, Michael Hurd, Erzo Luttmer, John Karl Scholz, Jim Poterba, and seminar participants at the NBER Summer Institute Workshop on Aging, University of Virginia, RAND, and the University of Wisconsin for very helpful comments.

1. INTRODUCTION

A widely accepted principle of economic theory is that leisure is a normal good, and thus that positive wealth shocks can be expected to lead to reductions in lifetime labor supply. Yet it is difficult to obtain reliable empirical estimates of the wealth effect because wealth is often correlated with unobservable characteristics that affect labor supply, such as a taste for leisure or risk aversion.¹ In this paper, we leverage the fact that an inheritance, particularly if it is unexpected, generates a shock to wealth that allows one to identify the effect of wealth on retirement behavior.

Understanding the effect of wealth shocks (and inheritances in particular) on retirement behavior is of substantial interest to researchers and policy makers. For example, one of the pathways through which reforms of the tax system, Social Security, or private pension regulations will influence retirement behavior is through wealth effects. Thus, it is important for policy analysis to develop an empirical foundation for understanding wealth effects as distinct from other pathways. Further, demographic trends suggest that many households receive inheritances when they are nearing retirement age. Among older households in the 2004 Survey of Consumer Finances who report ever having received an inheritance, more than half received one at ages 50-65.² Some analysts believe that strong returns in housing and equity markets over the past several decades will make inherited wealth even more important in the coming decades.³

To our knowledge, only two prior studies (Holtz-Eakin et. al., 1993; Joulfaian and Wilhelm, 1994) examined the effect of inheritance receipt on labor supply. Unfortunately, these papers shed little light on the effect of inheritance receipt on the retirement decisions of older workers; the former examines only younger workers while the latter obtains inconsistent results for older workers, with some specifications even suggesting that inheritance receipt *reduces* the

likelihood of retirement. Further, neither study is able to make the important distinction between expected and unexpected inheritances, which means they will underestimate the effect of wealth shocks if households adjust their behavior prior to receiving an expected inheritance.

This paper explores the effect of inheritance receipt on retirement using data from the 1992 – 2002 Health and Retirement Study (HRS). We examine whether the receipt of an inheritance – particularly an unexpected one – causes the individual to retire earlier than expected. By comparing actual inheritance receipt and retirement behavior to what the individual expected at the beginning of the sample period, we are able to control for most factors that might otherwise lead to a spurious correlation.

We have several findings. First, we document that inheritances are empirically important: over an eight-year period, approximately one in five households receives an inheritance, with a median value of about \$30,000, and many of these are unexpected. Second, we find that inheritance receipt increases the probability of retirement and that this effect increases with the size of the inheritance. Third, we find that unexpected inheritances have a larger effect than expected ones, suggesting that prior estimates of the effect of inheritances on labor supply may underestimate the effect of a wealth shock. We also confirm that this effect is driven by the receipt of the inheritance, and not by the direct effect of a death of a parent (which might affect one's preferences for leisure.)

2. HOW DOES WEALTH AFFECT LABOR SUPPLY?

2.1 Ambiguity in the Previous Literature

While there is a vast literature on how various financial considerations influence labor supply decisions, the few studies that have attempted to isolate the effect of a wealth shock on labor

supply have found collectively ambiguous results. For example, Imbens et. al. (2001) find that lottery winners consume some of their winnings in the form of reduced labor earnings and that the effect is larger for older winners. In contrast, Krueger and Pischke (1992) find little evidence of an increase in labor supply for workers in the Social Security “notch” cohort, who experienced a dramatic reduction in Social Security wealth due to a law change. Even when authors use similar sources of wealth variation, the prior literature finds contrasting evidence. For example, Coronado and Perozek (2003) and Sevak (2001) find that unanticipated stock market gains led workers to retire earlier, while Coile and Levine (2006) find no effect of stock market fluctuations on retirement.

Only two earlier studies have examined the effect of inheritance receipt on labor supply.⁴ Holtz-Eakin et. al. (1993) use a sample of estate tax returns from the early 1980s matched to the income tax returns of recipients before and after inheritance receipt. They find that recipients are more likely to exit the labor force when they receive a larger inheritance. However, their sample is limited to recipients age 19 to 58, and thus, by excluding most of the retirement age population, they are unable to observe how the inheritance influences subsequent retirement behavior. Using the Panel Survey of Income Dynamics, Joulfaian and Wilhelm (1994) find small effects of inheritance receipt on hours worked for prime-age workers. They find inconsistent effects on labor force exit by older workers, possibly due to a small sample size.

2.2 Overcoming the Limitations of the Prior Literature

Our analysis is the first to use the HRS to estimate the effect of inheritance receipt on retirement, and doing so offers many advantages over the prior literature. First, there are a large number of inheritance recipients among older workers in the HRS, which thus provides sufficient

variation and power to improve upon the inconclusive results of the past literature. Second, we are able to include a richer set of covariates than was used in prior studies. Importantly, this includes the ability to control for the direct effect, if any, that the death of a parent might have on retirement behavior (for example, the sudden realization that one should “stop and smell the roses.”) Third, the HRS provides data on *ex ante* inheritance expectations, which we use to test whether the effect of inheritance receipt depends on whether the inheritance is anticipated. As authors of previous studies have noted, the inability to distinguish between inheritances that are expected and unexpected introduces a potential downward bias in estimates of wealth effects, as the adjustment to an expected inheritance may have already occurred prior to receipt. Finally, the HRS has data on retirement expectations, allowing us to control for unobserved characteristics that may be correlated with both inheritance receipt and retirement. The next section provides more detail on the data and our empirical strategy.

3. USING INHERITANCES TO TEST FOR WEALTH EFFECTS

3.1 How Common are Expected and Unexpected Inheritances?

In order to use inheritances as a source of variation for measuring the effect of wealth shocks on retirement, inheritances must be common enough and large enough to make a meaningful difference to recipients. Further, in order to distinguish between expected and unexpected inheritances, it is important to know that the expectations questions contain some, although perhaps not perfect, information about the likelihood and size of future inheritance receipt. In this section, we provide summary data to support these two points.

Specifically, we use data from waves 2 through 6 (1994-2002) of the HRS, a panel data set that provides richly detailed data on respondents’ labor supply, health and finances.⁵ We

begin with the second wave because the data on inheritance expectations was not available until wave 2. For our purposes, a critical feature of the HRS is that it provides information on inheritances at each wave, including whether any inheritances were received since the last wave, the value of any inheritances received, the respondent's self-reported probability of receiving an inheritance over the next ten years, and the conditional value of the expected inheritance. Individuals who give a positive probability of inheritance receipt but are unable to provide a conditional value are asked a series of questions that allow for the value of the expected inheritance to be put into one of the following brackets: \$0 to \$10,000; \$10,000 to \$50,000; \$50,000 to \$250,000; \$250,000 to \$1,000,000; or over \$1,000,000.

Inheritance receipt is quite common in the HRS population: 5.4 percent of workers are in households that receive an inheritance over a two-year period, as are 19.3 percent of workers over an eight-year period.⁶ When weighted by dollars received, the majority of inheritances come from parents (72 percent) or aunts and uncles (7 percent); source missing (15 percent) is the other significant category. Many of these inheritances are quite substantial. The mean and median values for inheritance received by workers over a two-year period are about \$67,000 and \$28,000, respectively, indicating that the distribution of inheritance values is highly skewed; indeed, one-quarter of inheritance recipients receive less than \$10,000 while the top 5 percent receive inheritances in excess of \$280,000. The median inheritance is equivalent to nearly 11 percent of net worth and four months of household income; the top 5 percent of inheritances are more than 1.5 times net worth or three years of household income. This analysis suggests that inheritances are sufficiently common and large that they may well affect labor supply decisions of recipients.

Next, we examine inheritance expectations among HRS workers and the accuracy of these expectations. Figure 1 is a histogram showing the respondent's self-reported probability of receiving an inheritance over the next 10 years, as reported in 1994. Nearly 60% of respondents report zero probability of receiving an inheritance. The focal responses of a 50% chance and a 100% chance are the next most common answers, with about 10% of the sample selecting each of these responses; the rest of the distribution is spread out fairly evenly.⁷

Table 1 provides information related to the accuracy of inheritance expectations. We compare 1994 expectations of inheritance receipt during the next 10 years with actual inheritance receipt between 1994 and 2002, grouping respondents by expected probability of inheritance receipt. The first two columns of Table 1 report the expected probability of receipt and the fraction of the sample falling in that range. The third and fourth columns show the median conditional value of the expected inheritance and the median value relative to household income.⁸ Column 5 shows the fraction of each group who actually received an inheritance over the next eight-year period, while the final columns show the median value of received inheritances and the median value relative to household income.

There are several key findings from Table 1. First, the inheritance expectations questions contain some useful information about future inheritance receipt. The share of respondents receiving an inheritance rises monotonically with the expected probability of inheritance receipt, while the median value of received inheritances by probability group is fairly similar to the median conditional value of expected inheritances. However, individuals are often surprised by receiving an inheritance when none was expected or by not receiving one (during our eight-year sample period) when one was expected. Among the 60 percent of the sample who said that there was no chance of getting an inheritance, more than 10 percent did receive one, while among

those who said they were certain to receive an inheritance in the next ten years, fewer than 40 percent actually did. Furthermore, individuals often receive inheritances that are larger or smaller than expected, as we document more fully in Brown et. al. (2006). For example, among inheritance recipients who expected an inheritance of \$10,000 to \$50,000, fewer than half received an inheritance value in this range; results are similar for other values of the expected inheritance.

Overall, the evidence suggests that individual inheritance expectations are correlated with inheritance receipt, but that many inheritances are unexpected and most inheritances are either larger or smaller than expected. As we discuss in more detail below, unexpected inheritances are likely to represent a true wealth shock that can be used to identify the effect of wealth on retirement.

3.2 Empirical Strategy

The core of our empirical strategy is to test for whether the receipt of an inheritance increases the probability of retirement. Because there are multiple ways of defining retirement, inheritance, and the relevant time horizon, we implement a number of alternative specifications to explore the robustness of our core finding. As alternative definitions of retirement, we consider “labor force exit” as well as “retire earlier than expected.” We measure inheritances as a dummy variable for any inheritance receipt, as a dollar value of inheritances, and as a dollar value normalized by household income. For the time horizon, we consider both retirements over a two-year period as well as over the entire 8-year window of our panel. While each of these measures has its own advantages and disadvantages, we find a robust pattern of positive

correlations between inheritance receipt and retirement behavior across these alternative specifications.

Two novel aspects of our econometric strategy are worth underscoring: to our knowledge, we are the first to (i) use individual expectations about retirement and inheritances to control for unobservables and (ii) divide inheritances into expected and unexpected components. Taken as a whole, our results are strongly consistent with the hypothesis that inheritance receipt increases the probability of retirement, and that this effect is strongest when the inheritance was unexpected and thus most likely to represent a true shock to wealth. We now turn to a more detailed discussion of our specifications.

Our first specification defines retirement as “exit the labor force” and uses the following Probit specification:

$$\Pr(\text{retire}_{it} = 1) = \Phi(\beta_0 + \beta_1 \text{Inheritance}_{it} + \beta_2 \text{ParentDeath}_{it} + \beta_3 X_{it} + \gamma_t + \varepsilon_{it}) \quad (1)$$

where $\Phi(\cdot)$ is the CDF of the normal distribution. In this initial specification, retire_{it} is a dummy equal to 1 if the individual has exited the labor force since the previous wave, Inheritance_{it} is a dummy variable equal to 1 if the household received an inheritance since the previous wave, and the time horizon for this sample is a two-year wave of the HRS. Because retirement is the behavior of interest, the sample is limited to individuals who were working at the previous wave. If both spouses are in the labor force at the previous wave, both will be included in our regression. Because our measure of inheritance receipt is at the household level, we will cluster the standard errors on the household identifier to reflect any within household correlation in the response to the inheritance receipt. When using person-wave observations, we have 17,801 observations.⁹

$ParentDeath_{it}$ is an indicator variable for whether the respondent or spouse experienced the death of a parent since the last wave, which is important for ensuring that it is the inheritance, rather than the death of a parent, that impelled the person to retire.¹⁰ X_{it} is a vector of demographic characteristics including age dummies, race/ethnicity, gender, marital status, a dummy for poor health status at the previous wave and the change in this dummy, a cubic in own earnings from the previous wave, a cubic in household income from the previous wave, and a cubic in lifetime wages¹¹, net worth at previous wave, education level dummies, pension type dummies, industry and occupation dummies, and region dummies. γ_t is a wave fixed effect.

As a second specification, we replace the inherit dummy with the dollar value of the inheritance to test whether the magnitude of the response is increasing with the size of the inheritance. A third specification scales the value of the inheritance by the household's income at the wave prior to receipt.

We will then alter our time horizon, and analyze the effect of inheritance receipt (using each of our three measures) during the entire sample period (1994-2002) on the labor force exit over this period. We refer to this as the “long-difference” estimation. In this approach, each individual contributes one observation to the sample, conditional on being in the labor force at wave 2.¹² The long-difference approach can potentially include longer-term responses to inheritance receipt, while the original approach will focus on more immediate responses. In addition, the long-difference approach matches up somewhat better with the questions about inheritance expectations, which ask about receipt over the next ten years. In our long-difference sample, we have 4,508 observations.

We will then turn to the specification that makes use of the expectations data in the HRS. As noted above, we will introduce two important innovations. First, we use as our new

dependent variable whether the individual retired earlier than expected, where expectations about retirement are measured as of wave 1.¹³ This approach allows us to control for a wide range of unobservable factors that might be correlated with both retirement behavior and inheritance receipt. For example, if individuals who receive inheritances also happen to have a stronger taste for leisure, then this information would already be incorporated into their expectations about retirement. By comparing actual to expected retirement dates, we can determine whether the receipt of an inheritance influenced this individual's behavior, while controlling for these other factors.

Second, we can regress this difference in actual and expected retirement date on measures of inheritance receipt that differentiate between expected and unexpected inheritances. This combination essentially allows us to examine the *change* in retirement (actual minus expected) on the *change* in inheritance. By comparing changes in retirement and inheritance relative to expectations, we are controlling for numerous unobservable characteristics of the individual in much the same way that individual fixed effects would do.

After presenting our core results in section 4, section 5 provides further discussion about the interpretation of our results, including whether the *source* of the uncertainty over inheritance receipt matters. We will also discuss how potential measurement error, the presence of liquidity constraints, and possible risk aversion over the size of the inheritance might lead to even greater differences than we document between the effect of unexpected and expected inheritance receipt.

4. REGRESSION RESULTS

4.1 Effects of Inheritance Receipt on Retirement

Our first regression results are shown in Table 2, with results for the person-wave sample in the first three columns and those for the long-difference sample in the last three columns.¹⁴ For ease of interpretation, we report the marginal effects of the Probit specification evaluated at the mean.¹⁵ We first discuss the person-wave results. In specification 1, we find that receiving an inheritance is associated with a 2.3 percentage point increase in the probability of retirement over a two-year period, or 12 percent of the baseline retirement rate over a two-year period; the effect is statistically significant at the 10 percent level. In the next column, we test whether the response is increasing in the size of the inheritance by instead using the continuous value of the inheritance received.¹⁶ This coefficient is positive and statistically significant. Increasing the value of the inheritance by \$100,000 is found to increase the probability of retirement by 2.0 percentage points, or 10 percent of the baseline retirement rate. In the final column, we use the inheritance value scaled by household income at the previous wave, as it may be that it is the size of the inheritance relative to household finances that matters. This variable also has a positive effect on retirement. The effect is not statistically significant, but the magnitude is quite similar to the other coefficients on the table: an increase in the inheritance value equal to household income (about \$55,000 for the median household) raises the probability of retirement by 1.0 percentage points. The *ParentDeath* coefficients (not reported on table) are positive but small and statistically insignificant. Thus, these results suggest that it is the increase in wealth associated with the receipt of an inheritance and not the psychological blow resulting from the death of a parent that accelerates the recipient's retirement.

The second half of the table shows results for the long-difference sample, where the dependent variable is labor force exit over the full eight-year sample period. The results are quite similar to those already discussed. The inheritance dummy and the inheritance value

coefficients are positive and statistically significant at the 10 percent level or better. The magnitudes of the coefficients are generally similar to those in the two-year change sample, though they are smaller relative to the baseline retirement rate in the long-difference sample. For example, receiving an inheritance raises the probability of retirement over the eight-year period by 4.0 percentage points, or 7 percent relative to baseline retirement over an eight-year period, and increasing the value of the inheritance by \$100,000 raises the probability of retirement by 3.8 percentage points. As before, the *ParentDeath* coefficients are positive but insignificant.

As detailed in Brown et. al. (2006), we have assessed the robustness of the estimates in Table 2 to several alternative measures of labor supply. These include retirement defined based on the respondent's self-report of transitioning from being not retired to being partly or completely retired, the change in the self-reported probability of working past 62 (which may pick up changes in expected retirement behavior that have not yet been realized), the change in hours worked (which may pick up responses on the intensive as well as extensive margin), and labor force re-entry. In all cases the coefficients are of the expected sign, with inheritance receipt consistently reducing labor supply, and the effects are frequently statistically significant. We conclude that the estimated effects of inheritance receipt on labor supply are quite robust.

4.2 Results Using Expectations Data

Our results indicate that when a household experiences a wealth shock in the form of an inheritance, the household members reduce their labor supply by retiring earlier. There are, however, two limitations to these initial results. First, inheritance receipt is not random in the population. If, for example, individuals with wealthy parents are more likely to receive an inheritance and are also more likely to retire early even in the absence of an inheritance due to

differences in both observable (e.g., education, income) and unobservable (e.g., financial knowledge, risk aversion) characteristics, then this would cause a spurious correlation between inheritance receipt and retirement. Second, the receipt of an inheritance may not actually constitute a wealth shock for many households, because inheritances are often expected. Some of the households that expect an inheritance may have adjusted their labor supply prior to inheritance receipt (for example, by having one spouse retire early) and thus there may be no change in their behavior when the inheritance actually arrives. In this case, treating all inheritances as unexpected will tend to understate the true effect of wealth shocks on behavior.

To address both of these concerns, we make use of the rich data on expectations in the HRS. Specifically, we now define our dependent variable to be equal to one if the household retires earlier than expected, and zero otherwise. To create this variable, we make use of the individual's planned retirement year as reported in wave 1. Because some respondents did not answer these questions in the survey, the sample size is reduced from 4,508 to 2,499 observations.¹⁷

We also make use of expectations questions about inheritances in order to distinguish expected from unexpected inheritances. Even with the rich information in the HRS, it is not necessarily clear how to delineate between expected and unexpected inheritances. For example, an inheritance may be unexpected because the recipient did not expect to receive one at all, or it may be partially unexpected because the amount received was greater than expected.

Therefore, we first divide inheritance recipients by whether they expected any inheritance (expected probability greater than zero) or not, a simple but appealing way to identify a group of individuals for whom the inheritance was truly unexpected. Our calculations indicate that more than one-third of inheritance recipients had said there was no chance they would receive an

inheritance over the next ten years. Our first specification estimates separate inheritance dummy coefficients for those with expected and unexpected inheritances, while our second specification estimates separate inheritance-dollar-value coefficients for the two groups.

Finally, to incorporate the idea that an inheritance may constitute a surprise by its size rather than its arrival, we classify recipients by whether they received more than they expected, less than expected, or about what they expected. Since many of those who reported a positive probability of inheritance receipt could not give a conditional value of the inheritance but could answer questions that allow us to determine whether their expected inheritance falls in a particular range (e.g., \$10,000 to \$50,000), we define the more than expected dummy as receiving an inheritance that fell in a higher bracket than expected, and similarly for inheritances that were less than expected and about expected.¹⁸

The results of this analysis are shown on Table 3. All of the regressions are estimated over the long-difference sample. For ease of comparison with earlier results, we first estimate models with our new dependent variable and the same inheritance variables used in Table 2. The coefficients on the prior inheritance variables are all significant at the 10 percent level or better, including the coefficient on the inheritance value scaled by household income, and are somewhat larger relative to the mean of the dependent variable than those for the long-difference sample on Table 2. For example, receiving an inheritance increases the probability of retiring earlier than expected by 4.9 percentage points, or 13 percent relative to the baseline. Increasing the value of the inheritance by \$100,000 increases the probability of retirement by 5.1 percentage points.

The final three columns on the table display the results including our new right-hand side variables. In interpreting the results, it is important to note that the median expected inheritance (i.e., the actual inheritances received from 1994-2002 by households that reported some positive

probability of receiving an inheritance back in 1994) is \$40,135 while the median unexpected inheritance (i.e., the actual inheritances received from 1994-2002 by households that reported zero chance of receiving an inheritance back in 1994) is \$17,554. As a result of the smaller size of unexpected inheritances, one does not find a difference between expected and unexpected inheritances when one uses a simple indicator variable for inheritance receipt, whereas we do find a significant difference when the value of the inheritance is considered. Specifically, the effect of a given dollar amount of inheritance on the probability of retiring early is more than twice as large if the inheritance is unexpected – the effect of raising the inheritance value by \$100,000 is to increase the probability of retiring early by 10.3 percentage points if the inheritance is unexpected and by 4.3 percentage points if it is expected. Each coefficient is individually significantly different from zero, and a test of significance of the difference between the two coefficients has a p-value of 0.116.

In the final column of Table 3, we allow for the possibility that even an expected inheritance may be unexpected in its size by including dummy variables for whether the inheritance was more than expected, less than expected, or about what was expected (where expected means within the same “bracket,” e.g. \$10,000-\$50,000). Here, too, the results suggest a stronger effect of unexpected inheritances. Receiving an inheritance that is larger than expected raises the probability of retiring early by 6.0 percentage points, and this effect is significant at the 10 percent level. The coefficient on receiving an inheritance that is less than expected is negative (but not different from zero) and the two coefficients are statistically different from each other at the 10 percent level. Overall, our results indicate that inheritances that are either entirely unexpected or unexpectedly large have larger effects on retirement than expected inheritances.

5. FURTHER DISCUSSION

The results presented in the previous section indicated that the receipt of an unexpected inheritance matters more for retirement behavior than does an expected inheritance. It is worth noting that inheritances can be unexpected for at least three distinct reasons. First, an individual may not know how much money his parents have. Second, he may not know when his parents are going to die. Third, he may not know how the parents intend to distribute their wealth across their children, other individuals, or charities. To the extent that each of these three sources of uncertainty operate on retirement behavior solely through the creation of uncertainty about the timing and size of the bequest, it is not important to distinguish among them. We acknowledge the possibility, however, that one or more of these sources of uncertainty may influence labor supply decisions through other channels, such as if the resolution of uncertainty about a parent's health (e.g., witnessing a parent's health decline) affects one's preferences about labor supply. While the HRS allows us to make the first distinction between expected and unexpected inheritances, the data simply does not allow us to go further and distinguish among the three sources of uncertainty, other than by directing controlling for the death of a parent, which we have done. The fact that the parent's death had no discernable effect on retirement and no effect on the relation between inheritance receipt and retirement increases our confidence that such concerns are not first-order, but it does suggest an interesting avenue for further research.

While we have found support for our hypothesis that inheritances have larger effects when they are unanticipated, we pause to consider additional factors that might make it difficult to find a stronger effect of unexpected inheritances in the data. One possible issue is that that there may be measurement error in respondents' self-reported probability of inheritance receipt

and conditional value of the expected inheritance, so that our measures of expected and unexpected inheritances are both quite noisy. While it is undoubtedly the case that our delineation between expected and unexpected bequests is noisy, Table 1 shows that the subjective probability questions on inheritances contain useful information, as the self-reported probability of inheritance receipt is correlated with actual receipt and the received value is correlated with the conditional value. Furthermore, our measures are constructed so as to not rely very heavily on the specific probability of inheritance receipt or the exact value of the expected inheritance. However, given our need to classify inheritances based on the bracket values, it is possible that inheritances are labeled “about as expected” when they are either far more or less than expected (e.g., if a respondent expects \$100,000 and receives \$225,000), or conversely classified as “more than expected” when in fact the unexpected component is not very large (e.g., if a respondent expects \$45,000 and receives \$55,000). Thus, it is possible that the difference between the effect of expected and unexpected inheritances on retirement is larger than what we have found here.

Another potential factor that may affect the interpretation of our results is the role played by liquidity constraints. Workers who expect inheritances might wish to consume some of their inheritance prior to receipt and retire earlier, but be unable to do so because they cannot borrow against the inheritance and hold few assets or mostly illiquid assets.¹⁹ If so, this will tend to make the effect of expected and unexpected inheritances more similar.

To explore this, we experimented with three measures of liquidity constraints, including “has financial assets < \$5,000,” “has financial assets < \$10,000,” and “has financial assets < 20% of income.” By interacting these with our measures of inheritance receipt, we can test whether there is a differential response to the inheritance based on liquidity. These results did not

produce a consistent pattern of there being a stronger response to the receipt of an inheritance (whether expected or unexpected) by liquidity-constrained households.

These results, however, should not be interpreted as evidence that liquidity constraints are not important, for at least two reasons. First, our three measures are, admittedly, imperfect proxies for liquidity constraints. Indeed, the literature in this area suggests that good proxies for liquidity constraints are extremely difficult to find. In this context, the problem is made even more intractable by the fact that most proxies of liquidity constraints (such as measures of financial assets, levels of debt, etc.) may be endogenously determined with inheritance expectations. In other words, a household that is expecting an inheritance might save less or borrow more in an attempt to smooth out the consumption effects of the future bequest.

A final theory is that recipients may be reluctant to act upon as-yet-unrealized inheritances because inheritance receipt is uncertain. To use a common expression, “a bird in the hand is worth more than a bird in the bush.” Conceptually, it would be more appropriate to use the certainty equivalent of the expected inheritance as the portion that is expected. Unfortunately, calculating the certainty equivalent is not possible with the data available to us, as we would need to know the individual’s full probability distribution of inheritance receipt, as well as parameters of the individual’s utility function.²⁰ In sum, measurement error, liquidity constraints, and risk aversion over the size of the inheritance suggest that the difference between the effect of unexpected and expected inheritance receipt on behavior may well be even greater than what we find.

6. CONCLUSIONS

Inheritances represent a shock to wealth that may provide a useful way to estimate the effect of wealth on labor supply. Our paper provides new evidence on the effect of inheritance receipt on retirement using the HRS, which has a large number of inheritance recipients among its sample of older workers and includes data on ex-ante inheritance and retirement expectations.

We find that inheritance receipt is associated with a significant increase in the probability of retirement and that the effect is increasing in the size of the inheritance. These findings contrast with those of the previous literature, which failed to find large and consistent effects of inheritance receipt on retirement. We find that the effect is more than twice as large when the inheritance is unexpected, suggesting that earlier studies may have underestimated the wealth effect due to an inability to distinguish between expected and unexpected inheritances. Our findings may be of use to economists and policy makers seeking to project the effect of other wealth changes on retirement, such as those that might result from changes to Social Security.

A second contribution of our work is that we document that inheritance receipt is an important phenomenon for households nearing retirement age. About 20 percent of HRS respondents receive an inheritance over an eight-year period and these inheritances can be quite substantial, with a median value of about \$30,000. When a household receives an inheritance, it can spend it in a variety of ways – by reducing labor supply and increasing the consumption of leisure, by increasing its consumption of goods and services, or by increasing transfers to family and friends via bequests or *inter vivos* gifts. Studying some of these other behavioral responses to inheritance receipt may be a fruitful area for future research.

REFERENCES

- Barro, Robert J.. “Are Government Bonds Net Wealth?” *Journal of Political Economy* 82 (November-December 1974), 1095-1117.
- Bernheim, B. Douglas, Andrei Shleifer, and Lawrence H. Summers. “The Strategic Bequest Motive,” *Journal of Political Economy* 93 (December 1985), 1045-1076.
- Bernheim, B. Douglas. “How Strong Are Bequest Motives? Evidence Based on Estimates of the Demand for Life Insurance and Annuities,” *Journal of Political Economy* 99 (October 1991), 899-927.
- Blanchflower, David and Andrew J. Oswald. “What Makes an Entrepreneur? Evidence on Inheritance and Capital Constrains,” *Journal of Labor Economics* 16 (January 1998), 26-60.
- Blundell, Richard and Thomas MaCurdy. “Labor Supply,” in Orley Ashenfelter and David Card (eds.), *Handbook of Labor Economics* (Amsterdam: North-Holland, 2000).
- Brown, Jeffrey R. and Scott J. Weisbenner. “Intergenerational Transfers and Savings Behavior,” in Perspectives on the Economics of Aging, David A. Wise, ed. (Chicago: University of Chicago Press, 2004).
- Brown, Jeffrey R., Courtney C. Coile, and Scott J. Weisbenner. “The Effect of Inheritance Receipt on Retirement,” National Bureau of Economic Research Working Paper #12386 (2006).
- Coile, Courtney C. and Phillip B. Levine. “Bulls, Bears, and Retirement Behavior.” *Industrial and Labor Relations Review* 59 (April 2006), 408-429.
- Coronado, Julia and Maria Perozek. “Wealth Effects and the Consumption of Leisure: Retirement Decisions During the Stock Market Boom of the 1990s,” Board of Governors

- of the Federal Reserve System, Finance and Economics Discussion Series #2003-20 (2003).
- Gale, William, and John Karl Scholz. "Intergenerational Transfers and the Accumulation of Wealth," *Journal of Economic Perspectives* 8 (Autumn 1994), 145-160.
- Gale, William and Joel Slemrod. "Rethinking the Estate and Gift Tax: Overview," in W. Gale and J. Slemrod, eds., Rethinking Estate and Gift Taxation (Washington, D.C.: Brookings Institution, 2000).
- Gokhale, Jagadeesh and Lawrence Kotlikoff. "The Impact of Social Security and Other Factors on the Distribution of Wealth," in Martin Feldstein (ed.), *The Distributions Aspects of Social Security and Social Security Reform* (Chicago: University of Chicago Press, 2002).
- Holtz-Eakin, Douglas, David Joulfaian, and Harvey S. Rosen. "The Carnegie Conjecture: Some Empirical Evidence," *Quarterly Journal of Economics* 108 (May 1993), 413-435.
- Holtz-Eakin, Douglas, David Joulfaian, and Harvey S. Rosen. "Sticking It Out: Entrepreneurial Survival and Liquidity Constraints," *Journal of Political Economy* 102 (February 1994), 53-75.
- Holtz-Eakin, Douglas, David Joulfaian, and Harvey S. Rosen. "Entrepreneurial Decisions and Liquidity Constraints," *RAND Journal of Economics* 25 (Summer 1994), 334-347.
- Hurd, Michael D. and Kathleen McGarry. "Evaluation of the Subjective Probabilities of Survival in the Health and Retirement Study," *Journal of Human Resources* 30 (Supplement 1995), s268-s292.
- Imbens, Guido W., Donald B. Rubin, and Bruce I. Sacerdote. "Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption:

- Evidence from a Survey of Lottery Players," *American Economic Review* 91 (September 2001), 778–794.
- Joulfaian, David and Mark O. Wilhelm. "Inheritance and Labor Supply," *Journal of Human Resources* 29 (Autumn 1994), 1205-1234.
- Kessler , Denis, and Andre Masson. "Bequest and Wealth Accumulation: Are Some Pieces of the Puzzle Missing?," *Journal of Economic Perspectives* 3 (Summer 1989), 141-152.
- Killingsworth, Mark and James Heckman. "Female Labor Supply: A Survey," in Orley Ashenfelter and Richard Layard (eds.), *Handbook of Labor Economics* (Amsterdam: North-Holland, 1986).
- Kotlikoff, Laurence, and Lawrence Summers. "The Role of Intergenerational Transfers in Aggregate Capital Accumulation," *Journal of Political Economy* 89 (August 1981), 706-732.
- Kotlikoff, Laurence. "Intergenerational Transfers and Savings," *Journal of Economic Perspectives* 2 (August 1988), 41-58.
- Krueger, Alan B. and Jorn-Steffan Pischke. "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation," *Journal of Labor Economics* 10 (October 1992): 412-437.
- Light, Audrey and Kathleen McGarry. "Why Parents Play Favorites: Explanations for Unequal Bequests," *American Economic Review* 94 (December 2004), 1669-1681.
- McGarry, Kathleen M.. "Inter Vivos Transfers and Intended Bequests," *Journal of Public Economics* 73 (September 1999), 321-351.
- McGranaham, Leslie. "Charity and the Bequest Motive: Evidence from Seventeenth-Century Wills," *Journal of Political Economy* 108 (December 2000), 1270-1291.

- Modigliani, Franco. "The Role of Intergenerational Transfers and Life Cycle Saving in the Accumulation of Wealth," *Journal of Economic Perspectives* 2 (Spring 1988), 15-40.
- Pencavel, John. "Labor Supply of Men: A Survey," in Orley Ashenfelter and Richard Layard (eds.), *Handbook of Labor Economics*. (Amsterdam: North-Holland, 1986).
- Perozek, Maria G.. "A Reexamination of the Strategic Bequest Motive," *Journal of Political Economy* 106 (April 1998), 423-445.
- Saint Clair, Patricia, Delia Bugliari, Sandi Chien, Steven Haider, Orla Hayden, Michael Hurd, Serhii Ilchuk, Gabriela Lopez, David Loughran, Constantijn Panis, Philip Pantoja, Monika Reti, and Julie Zissimopolous. "RAND HRS Data Documentation, Version D," mimeo, RAND Corporation (2004).
- Schervish, Paul G. and John J. Havens. "Gifts and Bequests: Family or Philanthropic Organizations?" in Alicia H. Munnell and Annika Sunden (eds.), *Death and Dollars: The Role of Gifts and Bequests in America*. (Washington, D.C.: Brookings Institution Press, 2003).
- Sevak, Purvi. "Wealth Shocks and Retirement Timing: Evidence from the Nineties." Michigan Retirement Research Center, Working Paper WP00D1 (2001).
- Wilcox, David W. "Social Security Benefits, Consumption Expenditure, and the Life Cycle Hypothesis," *Journal of Political Economy* 97 (April 1989), 288-304.
- Wilhelm, Mark O. "Bequest Behavior and the Effect of Heirs' Earnings: Testing the Altruistic Model of Bequests," *American Economic Review* 86 (September 1996), 874-892.
- Zeldes, Stephen P. "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy* 97 (April 1989), 305-346.

Table 1: Expected vs. Received Inheritances

| Prob. of Inheritance Receipt during 1994-2002 Period | % of Sample | Median Cond. Value of Expected Inheritance | Median of Cond. Value of Expected Inheritance/ HH Income | % Who Received Inheritance by 2002 | Median Value of Inheritance Received | Median of Inheritance/ 1994 HH Income |
|--|-------------|--|--|------------------------------------|--------------------------------------|---------------------------------------|
| 0 | 0.598 | 0 | 0.000 | 0.109 | 17,554 | 0.350 |
| .01-.25 | 0.113 | 12,139 | 0.226 | 0.181 | 18,021 | 0.339 |
| .26-.49 | 0.020 | 24,278 | 0.420 | 0.190 | 17,199 | 0.486 |
| 0.50 | 0.083 | 36,417 | 0.617 | 0.245 | 39,231 | 0.597 |
| .51-.75 | 0.042 | 36,417 | 0.667 | 0.329 | 35,728 | 0.618 |
| .76-.99 | 0.054 | 60,695 | 0.872 | 0.367 | 63,219 | 0.741 |
| 1 | 0.088 | 60,695 | 0.893 | 0.388 | 48,592 | 0.708 |
| All | 1.000 | 0 | 0.000 | 0.178 | 30,474 | 0.500 |

Note: Values are reported in \$2002. The two final columns are conditional on receipt of an inheritance.

Table 2: Effect of Inheritance Receipt on Retirement

| Variable | Dependent Variable: Labor Force Exit | | | | | |
|-----------------------|--------------------------------------|-----------|---------|------------------------|-----------|---------|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| | Person-Wave Sample | | | Long-Difference Sample | | |
| Inheritance Flag | 0.0911 * | | | 0.1020 * | | |
| | (.0520) | | | (.0535) | | |
| | [.0233] | | | [.0402] | | |
| Inheritance Value | | 0.0791 ** | | | 0.0964 ** | |
| | | (.0377) | | | (.0292) | |
| | | [.0195] | | | [.0382] | |
| Inh Value / HH Income | | | 0.0394 | | | 0.0292 |
| | | | (.0256) | | | (.0214) |
| | | | [.0097] | | | [.0116] |
| # of Obs | 17,801 | 17,801 | 17,733 | 4,508 | 4,508 | 4,485 |
| Mean of Depend Var | 0.192 | 0.192 | 0.192 | 0.541 | 0.541 | 0.541 |

Note:

- 1) Coefficient estimates from a Probit model are reported with standard errors shown in parentheses. Marginal effects (evaluated at the sample means) are shown in brackets.
- 2) The sample is limited to individuals who were working at the previous wave. The person-wave and long-difference samples are described in more detail in the text.
- 3) Inheritance value is measured in 100,000s of \$2002.
- 4) All regressions include controls for the death of a parent, age, gender, marital status, race, education, current and lifetime income, net worth, health status, pension type, industry, occupation, region, and wave; see text for details.
- 5) * indicates significance at the 10% level, ** indicates significance at the 5% level.

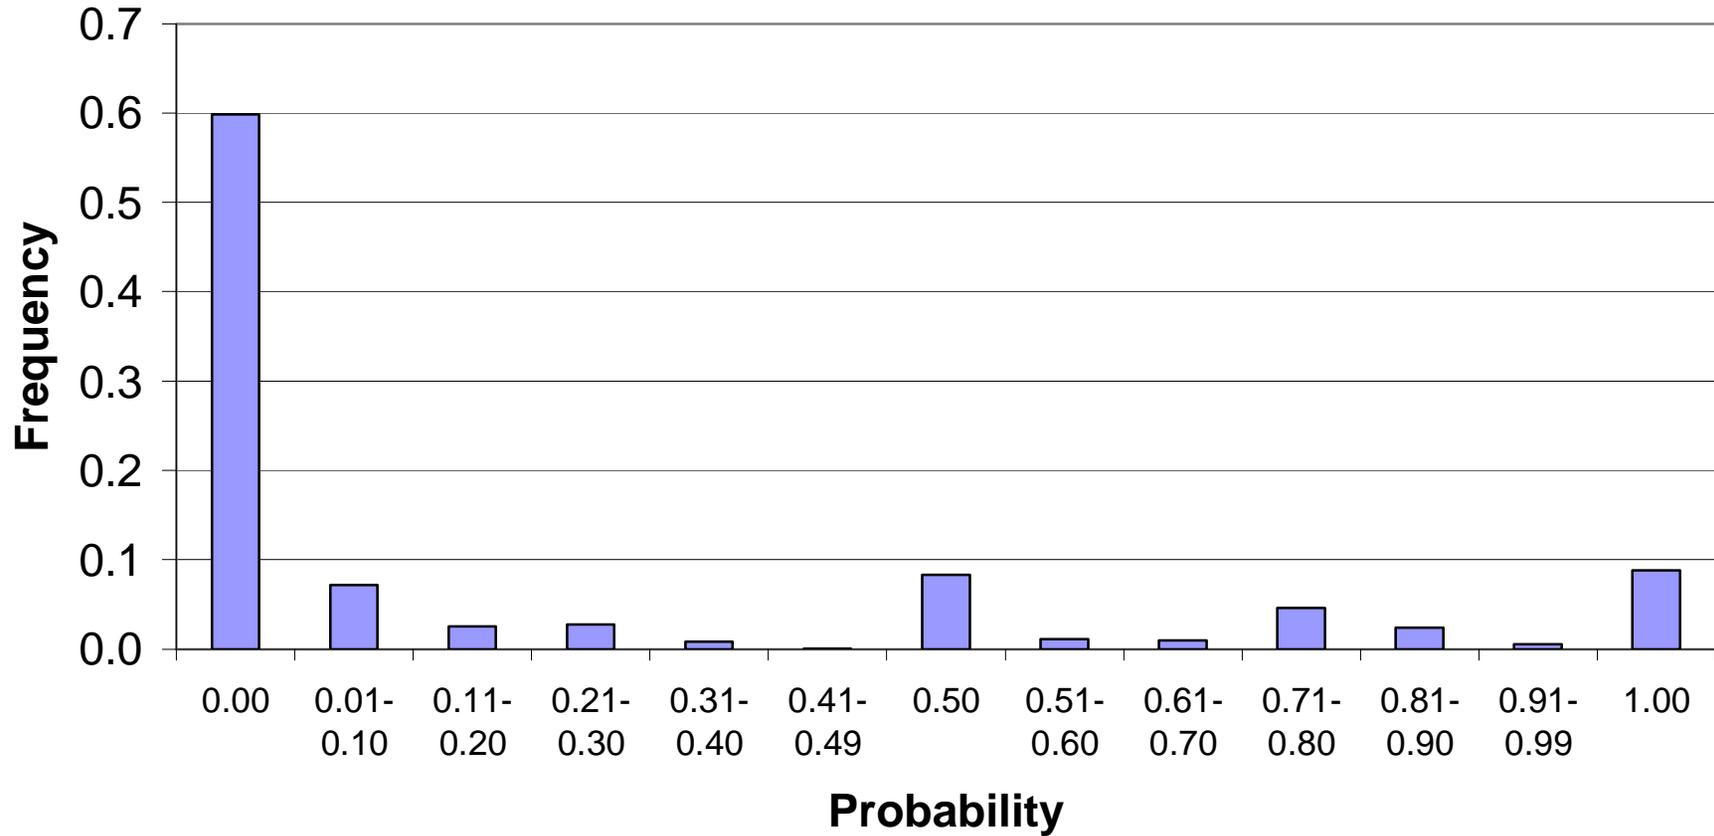
**Table 3: Effect of Expected vs. Unexpected Inheritance Receipt on Retirement
Long-Difference Sample**

| Variable | Dependent Variable: Retire Earlier Than Expected | | | | | |
|---|--|-----------|-----------|---------|-----------|----------|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Inh Flag | 0.1280 * | | | | | |
| | (.0688) | | | | | |
| | [.0491] | | | | | |
| Inh Value | | 0.1333 ** | | | | |
| | | (.0414) | | | | |
| | | [.0506] | | | | |
| Inh Value / HH Income | | | 0.1047 ** | | | |
| | | | (.0258) | | | |
| | | | [.0398] | | | |
| Inh Flag - Expected | | | | 0.1141 | | |
| | | | | (.0862) | | |
| | | | | [.0438] | | |
| Inh Flag - Unexpected | | | | 0.1477 | | |
| | | | | (.0989) | | |
| | | | | [.0570] | | |
| Inh Value - Expected | | | | | 0.1126 ** | |
| | | | | | (.0413) | |
| | | | | | [.0427] | |
| Inh Value - Unexpected | | | | | 0.2720 ** | |
| | | | | | (.0943) | |
| | | | | | [.1033] | |
| Inh Flag - Amt Expected | | | | | | 0.1970 |
| | | | | | | (.1243) |
| | | | | | | [.0765] |
| Inh Flag - Amt More | | | | | | 0.1546 * |
| | | | | | | (.0818) |
| | | | | | | [.0596] |
| Inh Flag - Amt Less | | | | | | -0.1692 |
| | | | | | | (.1828) |
| | | | | | | [-.0623] |
| # of obs | 2,499 | 2,499 | 2,485 | 2,499 | 2,499 | 2,499 |
| Mean of Depend Var | 0.382 | 0.382 | 0.382 | 0.382 | 0.382 | 0.382 |
| F-test: Expected=Unexpected (or Amt More=Amt Less) | | | | 0.785 | 0.116 | 0.096 |

Note:

- 1) Coefficient estimates from a Probit model are reported with standard errors shown in parentheses. Marginal effects (evaluated at the sample means) are shown in brackets.
- 2) The sample is limited to individuals who were working at wave 2 and provided an expected retirement date at wave 1. Long-difference samples are described in more detail in the text.
- 3) Inheritance value is measured in 100,000s of \$2002.
- 4) All regressions include controls for the death of a parent, age, gender, marital status, race, education, current and lifetime income, net worth, health status, pension type, industry, occupation, and region; see text for details.
- 5) * indicates significance at the 10% level, ** indicates significance at the 5% level.

Figure 1: Distribution of Expected Probability of Inheritance Receipt Over the Next Ten Years (as reported in 1994)



¹ For surveys of the large literature estimating the effect of unearned income on labor supply, see Pencavel (1986), Blundell and MaCurdy (2000), and Killingsworth and Heckman (1986).

² Authors' calculations.

³ For example, Schervish and Havens (2003) predict that from 1998 to 2052, over \$45 trillion of wealth (in 2002 dollars) will be transferred from estates. Not all analysts agree with the idea that inheritances are likely to grow in importance as a source of wealth for future retirees. Gokhale and Kotlikoff (2002), for example, point out that for most households, inheritances are small relative to lifetime labor earnings.

⁴ A small but growing literature examines the effect of inheritance receipt on other outcomes. Blanchflower and Oswald (1998) and Holtz-Eakin et. al. (1994a, 1994b) find that inheritance receipt increases the probability of becoming an entrepreneur. Joulfaian and Wilhelm (1994) look at the effect of receipt on consumption. There are numerous other papers focusing on inheritances, but most of it has focused on implications for aggregate wealth accumulation (e.g., Kotlikoff & Summers (1981), Modigliani (1988), Kotlikoff (1988), Kessler and Masson (1989), Gale & Scholz (1994), Gale & Slemrod (2000), and Brown and Weisbenner (2004)) or on the reasons that households make bequests (e.g., Barro (1974), Bernheim, Shleifer, and Summers (1985), Bernheim (1991), Wilhelm (1996), Perozek (1998), McGarry (1999) McGranahan (2000), Light and McGarry (2004)).

⁵ The HRS began in 1992 as a survey of people who were ages 51-61 and their spouses, with re-interviews of these individuals every two years. Although the HRS has subsequently expanded to include persons born in years up through 1953, we use only the original HRS cohort (born 1931-1941) in our analysis. For some labor force and demographic variables, we make use of the RAND version of the HRS, a user-friendly subset of the data that offers cleaned and consistent variables.

⁶ We defer detailed discussion of the samples used for these calculations until the next section of the paper. Calculations that refer to inheritance receipt over a two-year period are based on our “person-wave” sample, while calculations that refer to receipt over the full eight-year sample period are based on our “long-difference” sample.

⁷ The tendency of respondents to give focal answers such as 50% in response to subjective probability questions has been documented in other cases – see, for example, Hurd and McGarry (1995) for the case of subjective survival probability questions.

⁸ For this table only, in cases where respondents reported a positive probability of inheritance receipt and could not provide a conditional value of the expected inheritance but did answer the bracket questions, we assign them a conditional value equal to the midpoint of the bracket (or \$2,000,000 for the very few observations in the over \$1,000,000 bracket). As discussed more below, we do not use these values in our later empirical analyses.

⁹ As detailed below, in some specifications we make use of data on expectations of inheritance receipt, which are first asked at wave 2. To use a consistent sample throughout the analysis, we start our person-wave sample with persons working at wave 2 and observe whether they retire by wave 3. As our data extends through wave 6, each of the HRS' 12,652 respondents may provide up to 4 observations to the person-wave sample. Starting with a potential sample of 50,608 person-wave observations, the sample is selected as follows: we lose 12,426 observations because the individual died, left the sample, or was divorced or separated before wave 6; 4,603 observations because the individual did not report a probability of inheritance receipt; 1,018 observations because the individual failed to report a conditional value for the expected inheritance or answer the bracket questions, and 14,760 observations because the individual was not working at the previous wave.

¹⁰ Results are very similar regardless of how we define death of a parent, such as whether we distinguish between in-laws and own parents or between mothers and fathers.

¹¹ Lifetime wages are defined as the sum of both spouses' real earnings from ages 25 – 50 based on Social Security administrative records. For those observations missing earnings records, we use the median earnings for that individual's gender and education group.

¹² The same sample selection criteria discussed for the person-wave sample apply here as well. For example, if the respondent dies before wave 6 or does not report a probability of inheritance receipt, he or she is not in the sample.

¹³ We use retirement expectations from wave 1 rather than wave 2 for two reasons. First, doing so increases the likelihood that such expectations were formed prior to learning any information about possible inheritance receipt in the near-term. Second, the retirement expectations data is unavailable for wave 2. Specifically, we use the variable “r1rplnya” from the RAND version of the HRS as our measure of the expected year of retirement. As explained in Saint Clair et. al. (2004), this variable incorporates the answers to several questions about expected date of retirement in different parts of the survey in order to provide non-missing data for as many observations as possible. The HRS did not ask some of these questions in the wave 2 survey and thus the RAND data does not include a variable equivalent to r1rplnya for wave 2. For reference, we use the variable r6retyr from the RAND data as our measure of the actual date of retirement.

¹⁴ In all regressions, standard errors are clustered at the household level to correct for serial correlation in the error term for members of the same household across survey waves.

¹⁵ Estimating our specifications as linear probability models yields results very similar to those reported below.

¹⁶ To allow for the possibility of a non-linear effect of inheritances on retirement, we try including squared and cubic terms as well; however, these are never statistically different from zero and so the results are not shown here.

¹⁷ Our retired earlier than expected dummy is 1 for people who are retired by wave 6 and did so earlier than their expected retirement date, 0 for those who are retired by wave 6 and retired on time or later than their expected retirement date, and 0 for those who are still working at wave 6 and have passed their planned retirement date. This variable is missing for those who do not report a planned retirement year (1,321 observations) or who have neither yet retired nor reached their planned retirement date (677 observations). There is no significant difference in the probability of inheritance receipt among those observations with missing values of the retired earlier than expected dummy and those with non-missing values.

¹⁸ While it is tempting to try to decompose the *value* of the inheritance into its expected and unexpected components using the difference between the received value and the conditional value or expected value, such attempts are stymied by the large number of people who expect an inheritance but do not report its conditional value (about 40% of those who give a positive probability of inheritance receipt). Thus, it is not clear how much of their inheritance should be considered expected vs. unexpected. The approach we have adopted is, in our opinion, the cleanest way to test for differences between expected and unexpected inheritances given the available data.

¹⁹ There is a very large literature exploring how liquidity constraints influence a range of behaviors, including consumption of the elderly (e.g., Wilcox 1989, Zeldes 1989, and references therein).

²⁰ We have conducted weaker tests of whether uncertainty matters based on the measure of risk aversion available in the data (based on answers to questions about income gambles), on the theory that those who are relatively less risk averse will be more willing to spend expected inheritances prior to receipt, so they will respond less to inheritance receipt than will the most risk averse individuals. However, we do not find any evidence that response to inheritance receipt depends on risk aversion.