Retirement and Consumption Insurance*

Frank N. Caliendo[†] Erin Cottle Hunt[‡] Aspen Gorry[§] Jonathan M. Leganza[¶]

October 20, 2025

Abstract

We use a life-cycle model and causal empirical methods to study how individuals respond to changes in wealth and income. We show that endogenous retirement provides substantial consumption insurance. First, we use a simple model to show that individuals fully insure changes in wealth by adjusting their retirement date, keeping consumption constant, when wages and fixed costs of work do not vary with age. In a richer model, we show that the retirement margin still provides powerful consumption insurance to unanticipated wealth and income shocks. Finally, to assess the predictions of the theory, we estimate the effects of receiving an inheritance on retirement and spending. We find large increases in retirement and relatively modest increases in spending, supporting the idea that endogenous retirement insures consumption in practice.

Keywords: Retirement, Consumption Insurance, Permanent Income, Wealth, Inheritance

JEL codes: D15, E21, J26.

^{*}Acknowledgments: We thank Toni Braun, Maria Casanova, Paulo Lins, and participants at the 2024 WEAI Annual Meeting and 2024 Liberal Arts Macro Conference for helpful comments.

[†]Utah State University. (email: frank.caliendo@usu.edu)

[‡]Reed College. (email: erinc@reed.edu)

[§] John E. Walker Department of Economics. Clemson University. (email: aspen.gorry@gmail.com)

[¶]John E. Walker Department of Economics. Clemson University. (email: jleganz@clemson.edu)

1 Introduction

A fundamental question in economics asks how people insure themselves against uncertain realizations of wealth and income. While an extensive literature has studied consumption and savings decisions (e.g., Friedman, 1957; Deaton, 1991; Hubbard et al., 1994; Carroll, 1997) and labor supply adjustments along the intensive margin (e.g., Attanasio and Weber, 1995; Blundell et al., 1994; Low, 2005), the ability to insure against shocks by changing a key extensive-margin labor supply decision—retirement—has been comparatively underappreciated. In this paper, we provide new evidence on how endogenous retirement influences the way that people respond to wealth and permanent income shocks and find that the retirement choice provides significant insurance against consumption risk.

We proceed in three steps. First, we develop a simple life-cycle model with endogenous retirement driven by a fixed cost of working at each date to highlight how consumption responds to an increase in wealth under full information. If wages are constant over the life cycle, then the individual will not adjust consumption at all in response to a change in wealth, but will fully insure their consumption by changing their retirement date. The intuition for this result comes from comparing the marginal utility of consumption to the fixed cost of work at retirement. In response to an increase in wealth, the individual prefers to retire earlier to avoid the fixed cost of work rather than increase consumption that is subject to diminishing returns. Consumption can partially adjust if wages change with age around the date of retirement, changing the relative cost of work. This result continues to hold when individuals can adjust their labor supply on the intensive margin with a variable cost of work. Including this feature, it is still optimal to respond to a change in wealth by adjusting only the retirement age while holding consumption and hours worked per period fixed, as avoiding the fixed costs of work is also more valuable than the linear gain from working less in the years leading up to retirement.

In the second step, we extend our model to consider richer economic environments and quantitatively assess the robustness of our key result to unanticipated changes in wealth and permanent income. We continue to find a powerful self-insurance role for endogenous retirement. First, even though evidence supports the assumption of constant wages at older ages (Casanova, 2013; Rupert and Zanella, 2015), we consider the sensitivity of our result to declining wage profiles. To assess the upper bound of the consumption response, we parameterize a wage function that attributes the entire decline in earnings at older ages to declines in wages and find that individuals still respond to changes in wealth primarily by

¹For literature reviews focusing on the empirical and theoretical approaches to measuring consumption responses to income and wealth shocks, see Jappelli and Pistaferri (2010) and Meghir and Pistaferri (2011).

adjusting retirement. For example, in the baseline case with a coefficient of relative risk aversion of 1.5, consumption adjustments account for about 13% of the change in wealth, whereas changes in retirement account for the other 87%.

Second, we show that unanticipated increases in wealth at different points in the life cycle generate very similar responses in consumption and retirement to those documented in the simple model. That is, our results from the full information environment still hold when the individual has no information about potential shocks. The exception to these results is if the wealth shock is either large enough or occurs late enough in life to induce immediate retirement. In this case, any additional lifetime wealth that cannot be used to retire earlier is consumed. Hence, the timing of the shock does not have a large impact on how people respond until the individual is very close to their anticipated retirement age.

Third, we consider permanent income shocks in the form of changes to the level of wages. This extension allows us to link more closely to the literature on partial insurance of income risk, such as Blundell et al. (2008), who estimate the propensity to consume from a permanent income shock. The key finding of that literature is that consumption does not fully respond to permanent income shocks, meaning that individuals must have partial insurance against their income risk. We find that with endogenous retirement, the percent change in consumption relative to the percent change in wages ranges from 0.5 to 0.7 for values of the coefficient of relative risk aversion between 1.5 and 2. Interestingly, these values are quite similar to those in Blundell et al. (2008), which suggests that their partial insurance finding could be explained by individuals using endogenous retirement to self-insure against permanent income shocks.

In the third step, to assess the relevance of our model predictions, we study how individuals respond to changes in wealth using data from the Health and Retirement Study (HRS). Following Brown et al. (2010), we use the receipt of an inheritance as a case study for a wealth shock. We update and extend their analysis, which documents significant associations between the receipt of an inheritance and retirement in the HRS, by using a quasi-experimental research design to estimate the effects of inheritance receipt on both retirement and spending, which proxies for consumption. Using a stacked difference-in-differences approach (Wing et al., 2024), we compare the evolution of these outcomes for a treatment group of individuals who do not receive an inheritance, before and after the treatment group receives the inheritance.

The results support the qualitative model prediction of strong retirement responses that insure consumption risk. The inheritances that we study are sizable, roughly \$100,000 on average, and we show that they reflect mostly one-time lump sum payments that people receive at older ages. In response to these inheritances, we find visually clear, statistically

significant, and large retirement responses. Our event study estimates indicate an immediate retirement response after 1 year and additional increases in retirement after 3 years. Our main estimate that summarizes the effects over this three-year time horizon indicates a 5.9 percentage point increase in the probability of being retired, on average. This estimate translates to a meaningful 12% increase when compared to the baseline mean.

We also find spending increases that align with our theoretical predictions. The inheritances we study occur later in life and induce some immediate retirements, so we expect to see some increase in consumption. Consistent with this idea, we find statistically significant increases in nondurable spending. However, our main estimate indicates a relatively modest increase of \$2,700 on average, which is less than one-third of what we might have expected based on a simple version of the permanent income hypothesis that ignores retirement responses. Importantly, this finding is robust to using a measure of "core nondurable spending" like in Aguiar and Hurst (2013), which focuses on spending categories that should be more stable through retirement and thus avoids issues associated with retirement-induced declines in work-related expenditures. Overall, our empirical results thus provide validation of the key theoretical takeaway: retirement is a powerful hedging device, and adjusting retirement helps individuals smooth consumption in response to wealth shocks.

Related Literature. Our paper connects to three primary literatures: one that emphasizes the importance of the extensive margin of labor supply, one that studies consumption insurance, and one that provides empirical estimates of the effects of wealth shocks on labor supply and consumption.

First, we extend the literature highlighting the importance of the extensive margin of labor supply by showing that this margin of adjustment is particularly important in helping individuals insure wage and income risk. Papers in this literature, such as Rogerson and Wallenius (2009), Keane and Rogerson (2012), and Rogerson and Wallenius (2013), emphasize that including the extensive margin is important for accurately measuring labor supply elasticities. These studies have found that aggregate labor supply elasticities are driven by individual adjustments on the extensive margin. Rogerson (2024) argues that labor supply changes have often been overlooked in macro models, perhaps because of the view that labor supply elasticities are small. We highlight that the large elasticities on the extensive margin documented in previous work imply that the extensive margin is also powerful in insuring individuals against income and wealth risk over the life cycle.

Another strand of this literature focuses on life-cycle models with endogenous retirement choice. This literature often focuses on understanding the determinants of retirement or assessing policy changes, but does not emphasize the role of changing retirement in insuring individuals against various shocks. For instance, French (2005) and French and Jones (2011) construct life-cycle models with heterogeneous agents and endogenous retirement, but ask how uncertainty and public programs influence the retirement decisions rather than how changes in retirement insure individuals. Likewise, Fan et al. (2024) emphasize how endogenous retirement interacts with human capital accumulation over the life cycle to explain observed patterns in labor supply and the implications for social security reform. Gorry and Oberfield (2012) and vom Lehn et al. (2018) study taxes and labor supply in models with endogenous retirement. In contrast, we emphasize the implication of having an extensive labor supply choice for consumption patterns over the life cycle.

Second, our findings speak to a longstanding literature studying the extent to which consumption is insured against economic shocks. Blundell et al. (2008) and Heathcote et al. (2014) provide empirical evidence documenting that consumption only partially responds to permanent income changes. This important finding has generated a literature that seeks to understand the sources of this partial insurance. Kaplan and Violante (2010) find that a calibrated life-cycle model with incomplete markets does not generate the amount of consumption smoothing found in the data. Guvenen and Smith (2014) develop a model to study labor income risk and partial insurance, but use an exogenous partial insurance parameter to match the data. Other work proposes additional insurance mechanisms such as family labor supply, bankruptcy law, and progressive taxation (Heathcote et al., 2009; Blundell et al., 2016). Endogenous retirement has been underappreciated by this literature, as much of it abstracts from retirement completely or treats it as fixed. We find that endogenous retirement alone implies substantial smoothing of consumption in response to shocks to permanent income, as the level of consumption over the life cycle is pinned down by the first order condition for retirement. Indeed, our results highlight how the level of partial insurance that has been documented in the literature could be fully accounted for by endogenous retirement.

Finally, our study connects to an empirical literature that uses natural experiments to estimate how labor supply, retirement, and consumption respond to wealth shocks (e.g., Krueger and Pischke, 1992; Imbens et al., 2001; Coile and Levine, 2006; Gustman et al., 2010; Gelber et al., 2016; Cesarini et al., 2017; Picchio et al., 2018; Golosov et al., 2024). Papers typically focus on estimating either labor supply or consumption responses. Our broad contribution to this literature is to emphasize the importance of—and provide a theoretically-grounded reason for—including both consumption and retirement responses in these analyses, whenever possible, like in recent work that estimates the effects of winning the lottery (Golosov et al., 2024). The results from our case study also directly contribute to a strand of this literature that focuses on the effects of inheritances on retirement and consumption (Holtz-

Eakin et al., 1993; Joulfaian and Wilhelm, 1994; Brown et al., 2010; Suari-Andreu, 2023; Belloc et al., 2025). We provide updated evidence on how people respond to this important source of wealth by using a quasi-experimental framework that allows us to estimate causal effects and to produce clear graphical evidence of dynamic responses along both of these key margins.

2 Retirement in a Simple Life-Cycle Model

We first construct a simple life-cycle model with endogenous retirement to highlight the role of retirement choice in insuring consumption. To highlight the extensive margin, the model follows Rogerson and Wallenius (2009), Gorry and Oberfield (2012), and vom Lehn et al. (2018) in modeling a fixed cost of work, but with indivisible labor supply on the intensive margin. The key implication of the model is that the first order condition for the date of retirement links the marginal utility from the level of lifetime consumption with the return to work at retirement. While this basic tradeoff is still present in more complicated models, the simple model highlights its implications for consumption insurance.

2.1 Model

Consider a standard life-cycle model with labor supply in continuous time, where time is indexed by t. The individual enters the workforce at time t = 0, may choose an endogenous retirement date t_R , and passes away at time t = T.

The individual chooses consumption c(t), labor supply h(t), and the date of retirement t_R to maximize lifetime utility subject to a lifetime budget constraint:

$$\max \int_0^T \left[u(c(t)) - v(h(t)) \right] dt \tag{1}$$

s.t.
$$\int_0^T c(t)dt = \int_0^{t_R} w(t)h(t)dt + B,$$
 (2)

where w(t) is the life-cycle wage profile and B is a known wealth endowment. If the individual chooses to work, they receive a positive wage, w(t) > 0. Utility of consumption is given by u(c), where u'(c) > 0, and u''(c) < 0. Disutility of labor is given by v(h).

In order to have an active retirement choice, we assume the individual faces a fixed utility cost of working at each date, χ . If we also assume a variable cost of work α , v(h) can be

written as:

$$v(h) = \begin{cases} \chi + \alpha h & h > 0 \\ 0 & h = 0. \end{cases}$$

To highlight the power of retirement choice, we assume labor is indivisible in our baseline model. To do this, we set h=1 while the individual is working and normalize $\alpha=0$. The assumption of labor indivisibility is consistent with small labor supply elasticities measured on the intensive margin, while allowing individuals to adjust their labor supply on the extensive margin. In Appendix A, we extend the model to include an intensive margin with a linear disutility of labor and find that even with linear costs that make intensive hours adjustment attractive, the individual will first adjust along the extensive margin to insure consumption risk.

A fixed cost of work is commonly modeled as either a utility cost or a time cost, depending on which assumption is more convenient for a given model application. We model the fixed cost as a utility cost following Diamond and Köszegi (2003), Dybvig and Liu (2010), and Gorry and Oberfield (2012) among others, as it generates a convex constraint set, but our results would still hold if we modeled it as a time cost as in Rogerson and Wallenius (2009).

To simplify the analysis, we assume the discount rate and interest rate are both equal to zero and that there is no mortality risk. These assumptions, along with separable utility and the concavity of u(c), imply that a constant consumption profile c^* is optimal and is given by:

$$c^* = \frac{1}{T} \left(\int_0^{t_R} w(t)dt + B \right). \tag{3}$$

We can generate more realistic consumption profiles by relaxing these assumptions, but our results about the level of consumption still hold.²

The optimal retirement age t_R^* is given by

$$t_R^* = \arg\max \quad Tu\left(\frac{1}{T}\left(\int_0^{t_R} w(t)dt + B\right)\right) - \chi t_R.$$
 (4)

The first order condition for an interior retirement age is

$$u'(c^*) = \frac{\chi}{w(t_R^*)}. (5)$$

²In Appendix ??, we solve a more complex model that includes discounting and find the same results.

This first order condition highlights the power of the extensive margin labor supply choice for consumption.³ The intuition is the same as a standard labor supply choice: the individual equalizes the marginal rate of substitution between leisure (avoiding the fixed cost of work) and consumption with the return to work (the wage). However, on the extensive margin, this equation pins down the level of lifetime consumption, as the marginal utility from the level of lifetime consumption is pinned down by the ratio of the fixed cost of work and the wage at retirement. While the simple model with only an extensive-margin labor-supply choice emphasizes this result, richer models with an endogenous retirement choice still contain an analogous first order condition that pins down the level of lifetime consumption.

2.2 Impacts of a Change in Wealth under Full Information

To show how endogenous retirement allows the individual to insure their consumption, we study the impact of a change in wealth, B, on the optimal choices of consumption and retirement when the individual has full information about the wealth change. The effect of wealth on consumption is given by

$$\frac{\partial c^*}{\partial B} = \frac{1}{T} \left(w(t_R^*) \frac{\partial t_R^*}{\partial B} + 1 \right), \tag{6}$$

which depends indirectly on the effect of wealth on retirement age. To see the effect of wealth on retirement directly, the implicit function theorem implies that:

$$\frac{\partial t_R^*}{\partial B} = \frac{-u''(c^*)w(t_R^*)T^{-1}}{u''(c^*)w(t_R^*)^2T^{-1} + w'(t_R^*)u'(c^*)}. (7)$$

The denominator of (7) is negative if the second order condition holds, which means the fraction is negative. The optimal retirement age is decreasing in wealth. Substituting equation (7) into equation (6), the effect of wealth on consumption is

$$\frac{\partial c^*}{\partial B} = \left(\frac{1}{T}\right) \frac{w'(t_R^*)u'(c^*)}{u''(c^*)w(t_R^*)^2 T^{-1} + w'(t_R^*)u'(c^*)}.$$
 (8)

Given these responses, it is interesting to ask what fraction of a change in wealth will be absorbed in consumption and how much will be reflected in a change in labor earnings. We can compare the share of an individual's response to changes in wealth that occur along the

The second order condition for this problem is given by $u''(c^*)w(t_R^*)^2T^{-1}+w'(t_R^*)u'(c^*)<0$. This condition imposes some restrictions on the functional forms of the wage profile w(t). For example, if wages are decreasing near retirement, $w'(t_R^*)<0$, it is clearly satisfied. Also, if w(t) is constant, then once again we have concavity and a well-defined maximum.

consumption margin relative to the retirement margin. The dollar value of the change in retirement age that results from a change in wealth is given by:

$$-\frac{\partial t_R^*}{\partial B}w(t_R^*) = \frac{u''(c^*)w(t_R^*)^2T^{-1}}{u''(c^*)w(t_R^*)^2T^{-1} + w'(t_R^*)u'(c^*)},\tag{9}$$

where the left-hand side is the reduction in the retirement age due to an increase in wealth multiplied by the wages at the retirement age. This shows the wage earnings lost due to retiring earlier. The dollar value of the change in consumption that results from a change in wealth is given by:

$$\frac{\partial c^*}{\partial B}T = \frac{w'(t_R^*)u'(c^*)}{u''(c^*)w(t_R^*)^2T^{-1} + w'(t_R^*)u'(c^*)},\tag{10}$$

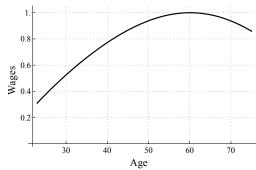
where the left-hand side shows the change in consumption due to a change in wealth multiplied by the length of the life cycle. This represents the lifetime change in consumption due to a change in wealth. Together, equations (9) and (10) can be interpreted as shares of an individual's response to a change in wealth.

These equations highlight how the individual will always adjust retirement in response to a change in wealth, but that there are conditions under which consumption may not respond. Note that the numerator of (9) cannot be zero because $u''(c^*) < 0$ and w(t) > 0, $\forall t$. However, the numerator of (10) can be zero, and if it is, then consumption remains constant. This phenomenon occurs when wages are constant at retirement or over the entire life cycle. In these cases, consumption is fully insured by a change in retirement. Next, we turn to quantitatively assessing the likely magnitudes of these responses when individuals face declining wage profiles.

2.3 Quantitative Implications of Endogenous Retirement

We have shown that if wages are constant, then an individual fully insures changes in wealth by changing their retirement date without changing their consumption. Importantly, a constant wage profile near retirement is not just theoretically possible, but it has been shown to be empirically plausible. While many papers have documented a hump shaped earnings profile over the life cycle, other evidence suggests these earnings declines in older age are driven by changes in hours and worker composition, not by declining wages. Using HRS data, Casanova (2013) shows that hump shaped profiles are an artifact of averaging over full and part-time workers and that wages only decline for those who transition into part-time work. Rupert and Zanella (2015) find that recent cohorts (born after 1937) in the Current Population Survey (CPS) do not experience reductions in wages at older ages and

Figure 1: Calibrated Wage Profile



Notes: Calibrated wage profile depicting average relative earnings using 2015 CPS data for individuals aged 23 to 75 from Caliendo et al. (2023).

replicate these findings using full life-cycle data from the Panel Study of Income Dynamics (PSID). Finally, Fan et al. (2024) show that declines in raw wages between ages 60 and 65 measured in the CPS and the Survey of Income and Program Participation (SIPP) disappear when controlling for individual fixed effects.

However, if wages do decline around retirement, our model implies that consumption will also adjust in response to a change in wealth. We can use the model to quantitatively assess the potential magnitude of such a consumption response. To do so, we make the extreme assumption that the entire decline in earnings measured around retirement is due to a change in wages. This assumption is consistent with the simple model assumption of no changes in hours worked along the intensive margin, and it allows us to calculate an upper bound on the consumption response to changes in wealth.

We parameterize the model to quantify the potential consumption response. We assume an individual ages from 23 to 80 and set wealth, B, equal to one year of peak wages. Given that hours worked at each date are normalized to one, we use the wage profile from Caliendo et al. (2023) that calculates average relative annual earnings by age in 2015 using CPS data for individuals between the ages of 23 and 70. The peak wage is normalized to 1. Figure 1 shows this wage profile.⁴

The instantaneous utility function is given by

$$u(c) = \begin{cases} \frac{c^{1-\sigma}}{1-\sigma} - 1 & \sigma \neq 1\\ \ln(c) & \sigma = 1. \end{cases}$$

Our baseline results are for a coefficient of relative risk aversion of $\sigma = 1.5$, but we also show

⁴When solving the model, we assume that the individual begins work at the start of life and chooses a single retirement date; they do not choose a start and stop date as in Rogerson and Wallenius (2009). This setup implies that the individual does not necessarily work only during their highest-wage years. One justification for this approach would be unmodeled human capital accumulation on the job.

Table 1: Consumption and Retirement Responses to Marginal Changes in Wealth

	CRRA parameter			
	$\sigma = 1$	$\sigma = 1.5$	$\sigma = 2$	$\sigma = 4$
Panel A: Initial Retirement Age 65				
Retirement response	81.9%	87.2%	90.1%	94.8%
Consumption response	18.1%	12.8%	9.9%	5.2%
Panel B: Initial Retirement Age 62				
Retirement response	93.0%	95.2%	96.4%	98.1%
Consumption response	7.0%	4.8%	3.6%	1.9%

Notes: Retirement and consumption responses to a marginal increase in wealth computed using equations 9 and 10 for different values of σ and targeted retirement dates.

how other values of σ influence our results.

Then, for a given value of σ , wage profile w(t), and value of wealth B=1, we calibrate the fixed cost of work χ to target a specific retirement age, either 62 or 65, using equations (3) and (5). For example, with $\sigma=1.5$, $\chi=2.50$ implies a retirement age of 62, while $\chi=2.15$ generates a retirement age of 65.

Using this calibration, we compute the marginal retirement and consumption responses to changes in wealth from equations (9) and (10). Table 1 shows the results for different values of σ and targeted retirement ages. In all cases, individuals respond primarily by changing their retirement date. In the baseline case with $\sigma = 1.5$ and an initial retirement age of 65, consumption only accounts for 12.8% of the response, with 87.2% from changing the retirement date. The consumption response falls to 10% for $\sigma = 2$ and to 4.8% if the targeted retirement age is 62. The smaller responses for an earlier retirement age are because wages are flatter closer to the peak of the life-cycle wage profile.

Overall, these results show that consumption responses are small even with extreme assumptions about the decline in wages near retirement. Hence, the ability to adjust retirement provides powerful consumption insurance. Another factor beyond declining wages that could mitigate the retirement response is if the fixed cost of work varies with age, $\chi(t)$. Assuming that $\chi(t)$ is increasing near retirement would have an analogous effect as declining wages, as they show up in the same term in the first order condition. While we are unaware of any direct evidence on how $\chi(t)$ evolves around retirement for individuals who choose their retirement date, allowing the fixed cost to vary with age is a flexible way to capture factors such as policy induced retirement incentives, changes in health status, and occupation specific difficulties in continuing work at older ages that may influence an individuals cost of work with age. However, the fact that many individuals move from full-time to part-time work

in their transition to retirement suggests that fixed costs are likely not rising dramatically with age, or paying the fixed cost to work part-time would not be worthwhile.

The power of this mechanism remains in richer models. For example, in Appendix A, we show that our results hold when we include an intensive margin labor supply adjustment, as it is optimal to first insure along the extensive margin. Moreover, our findings hold in models with additional features such as discounting, return on assets, and survival risk; the first order condition for retirement will still tie the level of lifetime consumption to the returns to work at the extensive margin.

3 A Model with Wealth and Permanent Income Shocks

Next, we extend the model to analyze unanticipated shocks in wealth or permanent income that can occur at different points in the life cycle. This framework allows us to study shocks commonly considered in the empirical literature, like inheritances (which we study in Section 4) or lotteries. We focus on unanticipated shocks because they provide an upper bound for the consumption response due to the lack of precautionary savings.⁵ In contrast to Section 2 where the individual has full information, unanticipated shocks constitute the other extreme case where the individual has no information. Taken together, these two extreme cases cover the information spectrum, highlighting the robustness of our results.⁶

3.1 Model

We begin with a pre-shock problem that is identical to that in Section 2, where we consider only extensive margin labor supply choices. Here we modify the notation for the pre-shock problem so that the planned date of retirement is t_{R1} and the individual expects their external wealth to be B^e . Note that in this simple setup, the timing of when they receive B^e does not matter for their optimal plan, as they can freely borrow and save to smooth consumption. The individual chooses consumption $c_1(t)$ and the date of retirement

⁵A key issue in the literature studying wealth shocks is whether the shock is anticipated or unanticipated. Our results suggest that this distinction is perhaps of secondary importance as individuals primarily absorb the wealth shock by adjusting the timing of retirement, rather than consumption. Since changing the retirement date provides substantial insurance against wealth shocks, the individual in our model does not engage in meaningful precautionary saving in anticipation of an uncertain inheritance. In fact, there is no precautionary motive if wages are constant.

⁶We have also explored an in-between dynamic-stochastic case when an individual knows ex ante the distribution of a wealth shock but not the realization. We find similar results in this case, but focus on the simpler extreme cases in the paper.

⁷The model does not require us to take a stand on how this expectation is formed. It is sufficient to assume the individual makes decisions as if they will receive B^e . The advantage of this modeling choice is that it allows us to clearly show the effect of changes in wealth on consumption and retirement.

 t_{R1} to maximize lifetime utility subject to a lifetime budget constraint:

$$\max \int_{0}^{T} \left[u(c_{1}(t)) - \chi \mathbf{1}_{t < t_{R1}} \right] dt$$
s.t.
$$\int_{0}^{T} c_{1}(t) dt = \int_{0}^{t_{R1}} w(t) dt + B^{e}.$$

The solution to the pre-shock problem is constant consumption:

$$c_1^* = \frac{1}{T} \left(\int_0^{t_{R1}} w(t)dt + B^e \right).$$

The optimal planned retirement age t_{R1}^* is given by

$$t_{R1}^* = \operatorname{arg\,max} \quad Tu\left(\frac{1}{T}\left(\int_0^{t_{R1}} w(t)dt + B^e\right)\right) - \chi t_{R1}.$$

The first order condition for an interior retirement age is

$$u'(c_1^*) = \frac{\chi}{w(t_{R1}^*)}.$$

We then assume that at a particular date before retirement, $t = t_B < t_{R1}$, the individual receives a wealth or income shock, realizing the true value of their external wealth, B, or a new wage profile, $w_2(t)$. After the realization of the shock, the individual chooses consumption and saving for the remainder of their life and a new retirement date. The individual solves the following maximization problem for date t_B forward:

$$\max \int_{t_B}^T \left[u(c_2(t)) - \chi \mathbf{1}_{t < t_{R2}} \right] dt$$
 s.t.
$$\int_{t_B}^T c_2(t) dt = \int_{t_B}^{t_{R2}} w_2(t) dt + B + \int_0^{t_B} (w(t) - c_1^*(t)) dt,$$

where the final term in the lifetime budget constraint is the value of the individual's assets at time $t = t_B$.

The solution still features constant consumption,

$$c_2^* = \frac{1}{T - t_B} \left(\int_{t_B}^{t_{R2}} w_2(t) dt + B + \int_0^{t_B} (w(t) - c_1^*(t)) dt \right),$$

and a retirement date, t_{R2}^* , that maximizes

$$(T - t_B)u\left(\frac{1}{T - t_B}\left(\int_{t_B}^{t_{R2}} w_2(t)dt + B + \int_0^{t_B} (w(t) - c_1^*(t))dt\right)\right) - \chi(t_{R2} - t_B).$$

The first order condition for an interior retirement age takes the same form as before:

$$u'(c_2^*) = \frac{\chi}{w_2(t_{R2}^*)}.$$

3.2 Impacts of Wealth and Permanent Income Shocks

The impacts of a shock to wealth in this environment are nearly identical to those discussed in Section 2. A positive wealth shock will induce changes in consumption and the retirement date guided by the first order condition. Consumption will be fully insured if wages are constant at retirement and will partially adjust if wages decline.

The impacts of a shock to permanent income are more nuanced. A positive income shock in this environment will increase consumption so that the first order condition remains satisfied. However, the retirement date could be adjusted in either direction necessary to support the new consumption level for the remainder of life. That is, if the increase in wages occurs late enough in life, the individual may end up deciding to work longer to generate the income to support the new equilibrium consumption for the remainder of their lifetime.

Additionally, this extended model allows us to consider corner solutions when the first order condition does not hold, such as wealth shocks that occur later in life that induce immediate retirement. If an individual receives a wealth shock greater than a threshold amount $B > \bar{B}(t_B)$, then they will retire immediately and increase their consumption. The value $\bar{B}(t_B)$ is the realization of the wealth shock that is *just* large enough that they choose to retire immediately $t_{R2}^* = t_B$. Considering a wealth shock only (holding permanent income fixed, $w_2(t) = w(t)$), the threshold value of wealth satisfies

$$\bar{B}(t_B) = B = (T - t_B)m\left(\frac{\chi}{w(t_B)}\right) + t_B c_1^*(t) - \int_0^{t_B} w(t)dt,$$
 (11)

where $m(.) = u'(.)^{-1}$ is the inverse marginal utility of consumption. If the shock is equal to the threshold amount $B = \bar{B}(t_B)$, the individual retires immediately, and the first order condition still holds. If it is greater than the threshold, $B > \bar{B}(t_B)$, the individual cannot reduce their retirement age any further, and the first order condition does not hold. The individual increases their consumption to satisfy the lifetime budget constraint. In this case, the level of consumption is determined by the logic of the permanent income hypothesis

Table 2: Consumption and Retirement Responses to an Unanticipated Wealth Shock

	CRRA parameter			
	$\sigma = 1$	$\sigma = 1.5$	$\sigma = 2$	$\sigma = 4$
Panel A: Initial Retirement Age 65				
Initial retirement age	65	65	65	65
New retirement age	64.15	64.10	64.08	64.03
Retirement change in years	-0.85	-0.90	-0.92	-0.97
Initial consumption	0.5769	0.5769	0.5769	0.5769
New consumption	0.5797	0.5789	0.5784	0.5777
Percent change in consumption	0.50%	0.35%	0.27%	0.14%
Retirement share	83.68%	88.56%	91.20%	95.42%
Consumption share	16.32%	11.44%	8.80%	4.58%
Panel B: Initial Retirement Age 62				
Initial retirement age	62	62	62	62
New retirement age	61.05	61.03	61.03	61.01
Retirement change in years	-0.95	-0.97	-0.97	-0.99
Initial consumption	0.5246	0.5246	0.5246	0.5246
New consumption	0.5256	0.5253	0.5251	0.5249
Percent change in consumption	0.17%	0.12%	0.09%	0.04%
Retirement share	94.80%	96.49%	97.35%	98.67%
Consumption share	5.20%	3.51%	2.65%	1.33%

Notes: Retirement and consumption responses for a wealth shock of one year of income received at t=0 for different values of σ and targeted retirement dates.

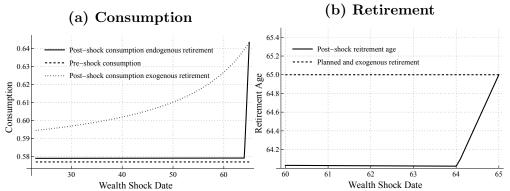
rather than the first order condition for retirement.

3.3 Quantitative Implications of Wealth Shocks

We now quantitatively assess the implications of wealth shocks. While wealth shocks have no consumption response if wages are constant around retirement, we again consider the calibrated wage profile from Section 2 to provide an upper bound of the response.

We assume the individual does not initially expect to receive a wealth shock, $B^e = 0$, and consider a discrete shock that increases wealth by one year of the maximum wages from our wage distribution, which was normalized to one. Only wealth changes at the shock date, and permanent income remains the same $w_2(t) = w(t)$. This setup means that we measure the effects of a discrete shock rather than using the model to tell us the marginal effects. We

Figure 2: Consumption and Retirement Responses to an Unanticipated Wealth Shock, by Date of the Shock



Notes: The response of receiving a wealth shock on consumption and retirement by age for $\sigma = 1.5$ and calibrated retirement age of 65. Results are shown for models with endogenous and exogenous retirement.

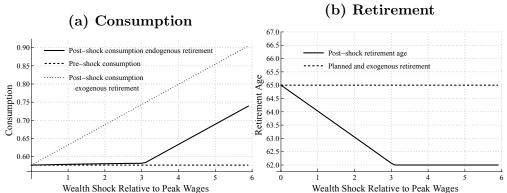
then compute the share of the wealth shock used to finance earlier retirement and increased consumption. The share of the wealth shock spent on earlier retirement is $(\int_{t_{R_2}}^{t_{R_1}} w(t)dt)/B$, and the share of the inheritance spent on increased consumption is $((c_2^* - c_1^*)(T - t_B))/B$.

We first highlight the analogous results for the effect of a change in wealth, which corresponds to a wealth shock that occurs at t=0 in this extended model. Table 2 presents the results, which are very similar to the marginal results in the simple model. If an individual receives one additional year's worth of income at t=0, they reduce their retirement age by a little less than one year, for most values of σ , and increase their consumption by less than half of one percent. That is, households choose to consume nearly the same amount and retire earlier. The share of the wealth shock spent on increased consumption is 16% or less, depending on the initial retirement age and the utility curvature parameter σ .

We next consider how the date of the shock influences changes in consumption and retirement, and compare results from our baseline model to a model with a fixed retirement date. Figure 2 illustrates these effects graphically. Panel (a) compares consumption before and after the shock based on the date of the shock. For most wealth shock dates, post-shock consumption increases only slightly. However, as the shock date approaches the individual's retirement date, post-shock consumption can grow dramatically. These large increases in consumption occur when the individual is induced to retire immediately and then consume the additional wealth over the remainder of their lifetime, as implied by the permanent income hypothesis. Panel (b) confirms this mechanism by highlighting the difference between pre- and post-shock retirement dates. For most wealth shock dates, the individual responds to the shock by reducing their retirement age by about one year, choosing to retire near age 64. However, shocks that occur after age 64 induce immediate retirement.

We emphasize that these patterns contrast with those that would arise in a model that

Figure 3: Consumption and Retirement Responses to an Unanticipated Wealth Shock, by Size of the Shock



Notes: Shock size is relative to one year of earnings at the peak of the life-cycle wage profile. The initial retirement age is calibrated to age 65, and the shock occurs at age 62. The solid lines show consumption and retirement age, assuming endogenous retirement. The dotted line shows the consumption response assuming the retirement age is fixed, exogenously, at age 65.

ignores endogenous retirement. With an exogenous retirement date, the individual would not be able to adjust retirement and would respond to wealth shocks throughout the life cycle by increasing consumption immediately and dramatically, smoothing the consumption of their increased wealth over their remaining lifespan. To highlight this contrast, the dotted lines in Figure 2 show the quantitative results for exogenous retirement.

Finally, Figure 3 illustrates how consumption and retirement respond to different sizes of realized wealth shocks at age 62. The parameters of the model are the same as the baseline calibration with $\sigma=1.5$ and an initial retirement age of 65. Age 62 is close enough to retirement that a shock larger than about 3 times peak earnings induces immediate retirement. Panel (a) shows the consumption response plotted against the size of the wealth shock, where 1 is the size of the baseline shock from the previous results. Shocks at age 62 induce small changes in consumption for inheritances up to about 3 times peak earnings. Above this threshold, the individual retires immediately and consumes additional wealth based on the permanent income hypothesis. Panel (b) shows the retirement responses. As the size of the shock increases, the post-shock retirement age decreases until it reaches the shock age. For example, shocks greater than about 3 times peak wages induce immediate retirement.

Again we contrast these patterns with those from a model with exogenous retirement. The dotted lines in the figure show that when retirement is fixed, all increases in wealth are consumed, and the increase in consumption is proportional to the size of the wealth shock.

Table 3: Consumption and Retirement Responses to an Unanticipated Permanent Income Shock, by Age of the Shock

	Age of Permanent Income Shock				
	Age 23	Age 30	Age 40	Age 50	Age 60
Panel A: $\sigma = 1.5$					
Initial retirement age	65	65	65	65	65
New retirement age	64.91	64.91	64.93	64.97	65.03
Retirement change in years	0.09	0.09	0.07	0.03	-0.03
Initial consumption	0.594	0.594	0.594	0.594	0.594
New consumption	0.599	0.599	0.599	0.598	0.598
Percent change in consumption	0.70%	0.70%	0.70%	0.68%	0.65%
Retirement share	27.44%	30.10%	29.36%	18.21%	-56.15%
Consumption share	72.56%	69.90%	70.64%	81.79%	156.15%
Panel B: $\sigma = 2$					
Initial retirement age	65	65	65	65	65
New retirement age	64.86	64.86	64.89	64.94	65.01
Retirement change in years	0.14	0.14	0.11	0.06	-0.01
Initial consumption	0.594	0.594	0.594	0.594	0.594
New consumption	0.598	0.598	0.598	0.597	0.597
Percent change in consumption	0.54%	0.54%	0.53%	0.52%	0.50%
Retirement share	43.87%	46.11%	45.82%	37.60%	-18.48%
Consumption share	56.13%	53.89%	54.18%	62.40%	118.48%

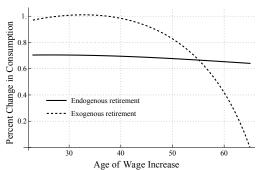
Notes: Consumption and retirement responses to a permanent 1% increase in the future wage profile by age at the time of the shock. Parameterization of the model follows the wage profile from Section 2, and the planned retirement age is calibrated to 65. Results are shown for the baseline parameterization of $\sigma = 1.5$ and for $\sigma = 2$.

3.4 Quantitative Implications of Permanent Income Shocks

We conclude our model-based exercises by assessing the quantitative implications of permanent income shocks on retirement and consumption. Specifically, we consider permanent income shocks that increase the remainder of the wage profile by 1%. For this exercise, we assume wealth is zero and does not change at the shock date $B^e = B = 0$. Following the unanticipated increase in wages at t_B , the individual re-optimizes and chooses consumption, saving, and a date of retirement. We compare these choices to the pre-shock values under the baseline calibration described in Section 2.

Table 3 shows the results for different shock ages and for two values of σ , $\sigma = 1.5$ and $\sigma = 2$. To contextualize these results, it is useful to consider the permanent income hypothesis,

Figure 4: Consumption responses to a permanent income shock



Notes: Percent change in consumption in response to a 1% increase in the future wage profile by date of the shock. Parameterization of the model follows the wage profile from Section 2 and the retirement age is calibrated to age 65 with $\sigma = 1.5$. The solid line shows consumption responses in our baseline model with endogenous retirement. The dashed line shows the consumption response in a model with an exogenous retirement age set to 65.

which would suggest that an individual should increase their consumption one-for-one with changes in permanent income. Empirical estimates have documented substantially lower consumption responses. For example, Blundell et al. (2008) estimate that the change in consumption relative to a change in permanent income is 0.64 for the US. For $\sigma = 1.5$, we find that the percentage change in consumption from a one percent increase in permanent income ranges from 0.65% if the shock occurs at age 60 to 0.71% if the shock occurs at age 23. For $\sigma = 2$, the consumption response is even smaller, ranging from 0.5 to 0.54%. These results suggest that the inclusion of endogenous retirement alone could account for the low consumption response to permanent income shocks observed in the data.

The consumption response generated by the model comes from the first order condition that equates the marginal utility of lifetime consumption with the ratio of the fixed cost of work and the wage at retirement. A higher wage implies that the individual will choose a higher level of consumption so that the first order condition holds. The results show that individuals tend to retire earlier. However, for shocks that occur late enough in life, the worker is induced to delay retirement in order to finance the higher level of consumption for the remainder of the lifetime. For example, for a shock at age 60, the individual works longer in order to finance consumption.

Finally, we emphasize that the muted consumption responses in the model are not simply due to life cycle effects. Here, consumption is determined by the first order condition that equates the marginal utility of consumption to the wage at retirement. Hence, the small changes in how consumption responds to permanent income shocks with age arise from differences in the wage at the new retirement date. The individual changes their retirement date to finance the desired level of lifetime consumption after the shock. In contrast, in a model with exogenous retirement, the individual's consumption response follows the intuition of the permanent income hypothesis. For shocks early in life, the wage change is nearly

permanent, so it is almost fully reflected in consumption. In contrast, as the individual gets closer to retirement, permanent wage shocks are more like wealth shocks as the remaining period of employment shrinks and the additional income is used for consumption both while working and in retirement. Hence, the size of the consumption response in the model with exogenous retirement is driven by two competing forces. First, the amount of additional wealth generated by the shock decreases with age, as a 1% increase in permanent income is smaller later in life when an individual has fewer working years remaining. Second, the length of remaining life over which the income will be smoothed is shorter for an older individual. The relative size of each force depends on the exogenous retirement age, the length of the life cycle, and the shape of the wage profile.

We compare the consumption response in models with and without exogenous retirement by the date of the permanent income shock in Figure 4. The solid line shows the percent change in consumption by the age of the shock in the model with endogenous retirement, while the dashed line plots the consumption response for a model with exogenous retirement at 65. The consumption response with exogenous retirement is nearly 1% early in life and decreases to zero as the shock date approaches the retirement age. In contrast, the consumption response with endogenous retirement is near 0.7% over most of the life cycle. Workers who experience an increase in their permanent income early in life respond by decreasing their retirement age to enjoy more leisure. Workers who experience the shock later in life respond optimally by increasing their retirement age to fund the higher level of consumption implied by the higher wage at retirement.

4 Empirical Analysis of Inheritances, Retirement, and Spending

In the final phase of our analysis, we study how people respond to wealth shocks empirically using inheritances as a case study. Guided by our theoretical framework, we document the effects of receiving an inheritance on both retirement and spending.

4.1 Data

We use data from the Health and Retirement Study (HRS).⁸ The HRS is a survey of older households in the U.S. Crucially, the survey asks respondents about inheritances, retirement, and spending. The data is thus well-suited for our purposes.

To access and assemble the data, we use the cleaned-and-processed HRS data products produced by the RAND Center for the Study of Aging. We use the RAND HRS Longitudinal

⁸The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan.

File 2022 (V1) to construct the basis of our analysis sample. The file contains panel data covering all people, respondents and their spouses, who appear in the HRS survey. There are eight survey cohorts corresponding to different groups of birth cohorts and 16 survey waves corresponding to the years 1992 through 2022. The Longitudinal File contains a good deal of demographic and economic information, including a variable on labor force participation that we use to define retirement.

For information on inheritances, we use the RAND HRS Detailed Imputations File 2022 (V1) and the RAND Fat Files, which are processed versions of the raw HRS data. The Detailed Imputations File contains variables that capture the dollar amounts of the three largest lump sum payments received by a household. The Fat Files contain information on the type of these lump sum payments. We merge these two data files so that we have information on the amount of received inheritances.

For information on spending, we use the RAND HRS CAMS Data File 2021 (V1). This file contains data from the Consumption and Activities Mail Survey (CAMS), which is a supplement to the core HRS survey. This supplemental survey occurs every other year, in odd-numbered years from 2001 to 2021, between the years of the core survey. These data are crucial for our analysis because they provide direct measures of spending, which proxies for consumption. A key limitation of the data is the sample size. Three factors limit the size of these data. First, the data are only available after wave 5, so we do not have spending information for the earlier years in our analysis. Second, the consumption survey is only sent to a subset of the households that take part in the core survey. Third, in the case of coupled households, there is typically only one spouse who responds to the consumption survey (which is designed to elicit household-level spending). So, while our analysis of retirement can include both members of a coupled household, our analysis of consumption is naturally limited to CAMS respondents.

To assemble the data for our analysis, we begin with the Longitudinal File and merge these data with the data on inheritances and spending. When merging the CAMS data, it is important to consider the timing of the variables. The core survey takes place in even-numbered years, but the spending survey takes place in odd-numbered years. By default, the RAND files link the CAMS data to core data from the previous survey wave, but this setup means that the timing of the spending variables is not aligned with the timing of other key variables, which typically capture behaviors that are either contemporaneous or that occurred within the last two years. For example, people observed in the core survey during

⁹We use data from all survey cohorts and all waves, except for the first three waves of the AHEAD cohort, which correspond to years for which the AHEAD survey instruments differed substantially from the HRS survey instruments.

wave 14 (which corresponds to 2018) answer questions about their current (2018) labor force status and about inheritances received in the last two years (between 2016 and 2018), but the subset of these people in the CAMS data answer survey questions that capture their spending in 2019. Because our research design involves tracking spending around the time of inheritances, we instead link the CAMS data to the subsequent core survey data. This approach ensures that our spending outcome measures are aligned with other measures. To continue the example, our observations of people in wave 14 (2018) contain information on their spending in the previous year, 2017.

After merging these data files, we have a biennial panel dataset of older individuals spanning 16 survey waves. Our sample for analysis contains 275,133 observations of 42,119 unique individuals. The subsample with information on spending contains 36,691 observations on 7,747 unique individuals.

4.2 Key Variables

To study retirement, we use a categorical and individual-level variable that captures labor force status. The variable indicates whether the respondent was working full time, working part time, unemployed, partly retired, retired, disabled, or not in the labor force at the time of the survey. We define our retirement outcome as an indicator variable that takes the value of one for people who report being retired or out of the labor force. Our goal in using this measure is to capture people exiting the labor force as in Brown et al. (2010). We also study an indicator variable that takes the value of one only for people who reported being retired as a robustness check.

To study spending, we use several of the aggregate spending variables in the CAMS data. We emphasize nondurable spending as one of our primary outcomes, as spending in this category should track consumption relatively closely. Our main outcome for nondurable spending is the aggregate measure recorded in the CAMS data, which captures spending on food, clothing, utilities, entertainment, services, and other spending.¹⁰ In the spirit of Aguiar and Hurst (2013), we also study "core nondurable spending," which we define as total nondurable spending minus spending on dining out, food and beverages, clothing, and personal care. We view this measure as an important robustness check that allows us to study spending that is less likely to be work-related.

We also study additional variables that capture spending in other categories that can

¹⁰The detailed spending categories vary across waves but the nondurable spending variable that we use generally captures spending on gifts, clothing, charitable contributions, dining out, food and beverages, utility bills, medications, health insurance and health services, medical supplies, telecommunications, vacations, personal care, furnishings, housekeeping, supplies, yard services, and hobbies and sports.

be lumpier and that therefore do not track consumption as well: housing spending (like mortgage interest, property taxes, and rent), transportation spending (like car purchases and maintenance, but also gas), and other durable spending (large household appliances, televisions, and computers). Finally, we study total household spending, which is the sum of the main category-specific measures (nondurables, housing, transportation, and other durables). We adjust spending variables (and other variables measured in nominal dollars) for inflation and report values throughout the paper in 2010 dollars. Because spending data can be noisy, we winsorize spending variables in our baseline analysis at the 1st and 99th percentiles to reduce the influence of outliers.

To study inheritances, we use two household-level variables. First, we construct an indicator variable that takes the value of one for respondents who report that they or their spouse received an inheritance since their last interview. Second, we construct a continuous variable that captures the amount of inheritance received, including zeros for observations of people who did not receive an inheritance. We also track parental mortality around the inheritances that we study by using two indicator variables, one that takes the value of one for respondents who report that their mother is alive and another that takes the value of one for respondents who report that their father is alive.

Finally, we use several other variables to carry out our analyses. To contextualize inheritances, we use a variable that captures total household wealth, which is the sum of wealth components (residences, real estate, vehicles, businesses, IRAs, stocks and mutual funds, checking and savings accounts, certificates of deposits, bonds, other savings) less debt components (mortgages, other home loans, and other debt). We also use a variable available in waves 2 through 8 of the core survey that captures the self-reported probability of receiving an inheritance in the next 10 years. For control variables that we can include in our regression analyses, we use indicator variables for being white, for being male, for having attended at least some college, and for being married.

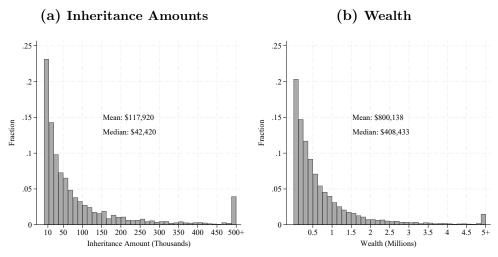
4.3 Descriptive Analysis of Inheritances

We begin by describing inheritances in the HRS data. These analyses build upon and update the descriptive statistics presented in Brown et al. (2010), who use data on only the original HRS cohort from waves 2 through 6.

First, we note that inheritances are relevant. Out of the 42,119 people in our sample, 9,349 (22%) of them report receiving an inheritance at some point in our time horizon. Moreover, 2,628 (6%) of them report receiving more than one inheritance. Inheritances are common and worth studying.

Second, we argue that inheritances are meaningfully large. Figure 5 illustrates this point.

Figure 5: Histograms of Inheritance Amounts and Wealth



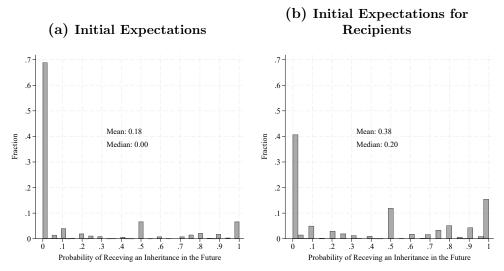
Notes: This figure presents two histograms that provide evidence on the size of inheritances. Panel (a) is a histogram of inheritance amounts. Panel (b) is a histogram of wealth. The underlying sample contains all personwave observations for which the person reports receiving an inheritance.

Panel (a) plots the distribution of inheritance amounts. The underlying sample consists of all person-wave observations for which the person reports receiving an inheritance. The average inheritance is \$117,920, and the median is \$42,420. Moreover, almost 5% of observed inheritances are over \$500,000. To put these numbers into perspective, panel (b) plots the distribution of wealth using the same person-wave observations underlying panel (a). The average inheritance is 14.7% of average wealth, and the median inheritance is 10.4% of median wealth.

Third, we leverage the breadth of the HRS data to look at inheritance expectations. We have information in waves 2 through 8 on how likely respondents believe they are to receive an inheritance in the next 10 years. Figure 6 plots histograms of the inheritance expectations variable. Panel (a) plots a histogram that we interpret as capturing initial expectations. The underlying sample consists of the first observation of each person who is surveyed in waves 2 through 8. The distribution thus captures expectations of people when they first enter the survey. Notably, almost 70% of people report no chance of receiving an inheritance in the next ten years. Panel (b) plots the distribution of initial expectations for people who eventually receive an inheritance in the data. Even for these individuals, the expectations variable suggests that many inheritances are unexpected. Over 40% of this sample of inheritance recipients reported no chance at all of receiving an inheritance, and the median expected probability of receiving an inheritance in the next 10 years is only 20%.

Overall, this descriptive evidence highlights how our earlier theoretical example, where an individual expects zero wealth and then receives a large and unexpected wealth shock, captures a relevant experience for many individuals. It also supports the idea that we can

Figure 6: Histograms of Inheritance Expectations



Notes: This figure presents two histograms that provide evidence on inheritance expectations. Panel (a) is a histogram of inheritance expectations using the initial observations of people surveyed in waves 2 through 8, when the inheritance expectations questions were asked. Panel (b) is a histogram of initial inheritance expectations of the subsample of people who we observe receiving an inheritance.

reasonably treat inheritances as shocks to wealth in the causal analysis that follows. Even people who expect to receive an inheritance at some point still face meaningful uncertainty as the timing and size of an inheritance depends on uncertain factors like late-in-life health expenses, the timing of the death, and the asset balance of the decedent.

4.4 Identification Strategy

4.4.1 Stacked Difference-in-Differences

Next, we identify the causal effects of inheritance receipt on retirement and spending using a difference-in-differences approach. We compare outcomes for a treatment group of people who receive an inheritance to a control group of people who do not, before and after the treatment group receives the inheritance. Because our setting has staggered adoption of treatment, meaning people are treated at different times, we use a stacked difference-in-differences design. The basic idea is to avoid potential problems with two-way fixed regression models (e.g., Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021) by (i) constructing several clean analysis samples for each treatment group, defined by the survey wave during which the inheritance occurs, (ii) stacking these analysis samples, and then (iii) estimating simple difference-in-differences regression models using the stacked data.

To implement this stacked difference-in-differences strategy, we follow Wing et al. (2024), who clarify the type of causal parameter estimated by stacked difference-in-differences designs

that have been used previously (e.g., Deshpande and Li, 2019) and present a weighted stacked regression that identifies the parameter of interest. They call the average treatment effect on the treated parameter of interest the "trimmed aggregate ATT." It is a weighted average of the group-time ATT parameters. We follow their approach by constructing a stacked dataset that is trimmed such that the number of pre-periods and post-periods for each group is the same, and by estimating completely standard difference-in-differences regression models, but with appropriate sample correction weights. In the robustness section, we assess the sensitivity of our estimates to the choices we make within this estimation framework and report results from a traditional two-way fixed effects estimator.

4.4.2 Constructing Stacked Samples

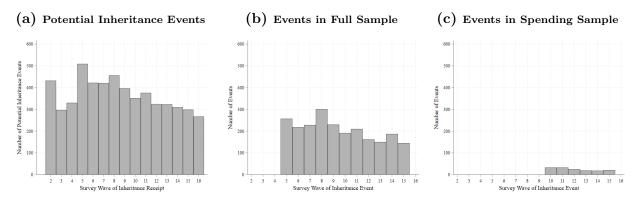
We construct two separate stacked samples in three steps. One uses as many inheritance events and as much data as possible to study the effects of an inheritance on retirement. We sometimes refer to this stacked sample as the "full sample." The other is the "spending sample," which we use to study the effects of inheritances on spending (and retirement). The spending sample is meaningfully smaller because of the consumption data availability discussed above, but it allows us to study spending.

The first step is to define inheritance events that we could potentially study. We begin with our analysis sample described above, which contains 9,349 person-wave observations of inheritances. We then implement two additional data restrictions to ensure that we can construct well-defined quasi-experimental groups. First, we drop observations of people with inconsistent records of an inheritance. Specifically, we drop observations of 799 individuals who receive an inheritance at some point but who report the amount of that inheritance to be \$0. Second, we drop observations of 3,035 individuals who report receiving an inheritance in their first survey wave, for whom we do not have any pre-inheritance observations. We are left with 7,600 observations of inheritances that are received by 5,515 unique individuals.

We define an inheritance event as the first observed inheritance a person receives. Panel (a) of Figure 7 plots the histogram of these inheritances that we can potentially study. It shows the number of inheritance events that occur in each survey wave. Most waves have between 300 and 450 events.

The second step is to create the trimmed difference-in-differences datasets for each sub-experiment, defined by the wave when the inheritance is reported. These sub-experiment-specific datasets will then be concatenated to form the stacked samples. For the full sample, we begin with everyone in our analysis data who either (i) experiences one of the potential inheritance events, who could be included in the treatment group, or (ii) never experiences an inheritance event, who are included in the control group. For the spending sample, we also

Figure 7: Histograms of Inheritance Events Across Survey Waves



Notes: This figure presents three histograms that illustrate the construction of the stacked samples used in our difference-in-differences analysis. Panel (a) is a histogram of all potential inheritance events. Panel (b) is a histogram of individuals experiencing inheritance events that we include in the full stacked sample, where the set of events is restricted to those for which we observe the recipient of the inheritance in every wave corresponding to four waves prior to the event, the wave of the event itself, and the wave after the event. Panel (c) is a histogram of individuals experiencing inheritance events that we include in the spending stacked sample, where the data is limited to observations with non-missing spending variables and the events are restricted to those for which we have a balanced panel of observations.

keep only person-wave observations with non-missing values for spending, which naturally limits the sample to the later survey waves for which the consumption survey exists and the subset of core survey households who are sent the consumption survey.

We define sub-experiments according to the wave when an inheritance is reported. For the sub-experiment corresponding to inheritances in wave w, we define the treatment group to be everyone who experiences an inheritance event in wave w, and we define the control group to be everyone who does not experience an inheritance event in any wave. For these individuals, we define event time, e, as time relative to the actual or placebo event wave.

Next, we trim the observations by balancing the panel, keeping only individuals observed in every wave of an event window. In our baseline analysis, we use an event window that includes four waves prior to the inheritance event and two post-period waves: the wave in which the inheritance is reported, and one wave after. We chose this window, which puts more emphasis on the pre-period, because we want to assess whether the outcomes of our treatment and control groups were trending in parallel before the treatment group received an inheritance. We later show the robustness of our estimates to the choice of event window.¹¹

The result for each sub-experiment, a, is a panel of treatment and control individuals that is balanced in event time $e \in \{-4, 1\}$. For example, the dataset for the sub-experiment corresponding to wave 10 consists of a treated group of people who report an inheritance in wave 10 and a control group of people who never report receiving an inheritance in our

¹¹The key tradeoff when defining an event window relates to sample stability versus sample size. A longer window allows us to observe more pre- or post-period data, but results in a smaller sample size because it requires balancing the sample over a longer horizon.

data, and we have a balanced panel of these people with observations spanning wave 6 (four waves before the events) to wave 11 (one wave after the wave of the event).

The third and final step is to stack these sub-experiment-specific difference-in-differences datasets. Panel (b) of Figure 7 displays a histogram of the inheritance events that underlie our full stacked sample. The earliest inheritance events we study occur in wave 5, since we require observations for four pre-period waves, and the latest events we study occur in wave 15, since we require two observations for two post-period waves. We have a total of 2,279 treated individuals who experience an inheritance event, and we use all of these instances when studying retirement. However, have less data for studying consumption because of the data limitations described above. Panel (c) displays a histogram of the events we can use to study consumption. The earliest events occur in wave 10, and there are fewer observations in waves 10 through 15 than in panel (b). The spending stacked sample contains data on 149 treated individuals who experience an inheritance event.

It is important to detail the timing of our outcome variables as it relates to the construction of our stacked samples. Our stacked samples are based on events defined by inheritances received since the last survey, so in the last two years. If we assume that the inheritances occur about one year before the survey interview on average, then in the wave of the reported inheritance, the retirement outcome on current labor force participation captures retirement status about one year after the inheritance and the spending outcomes on spending in the previous year capture expenditures in the year of the inheritance.

4.4.3 Summary Statistics for the Stacked Samples

Table 4 presents summary statistics for our stacked samples, using data from two survey waves before the reported inheritance event. Panel A is for the full sample, and panel B is for the spending sample. Within each panel, columns (1) and (2) report means and standard deviations for the treatment group and columns (3) and (4) present means and standard deviations for the control group. Both panels show that people in the treatment group, those who we observe receiving inheritances, are younger, more likely to be male, married, and white, and are more likely to have attended some college compared to their control group counterparts. There are non-negligible differences in these characteristics in levels between the two groups. From the perspective of our difference-in-differences framework, these differences in levels are not necessarily problematic; the key factor is whether the evolution of the control group outcomes is a good counterfactual for the evolution of the treatment group outcomes, which we assess later. However, these differences in levels also motivate a robustness check where we include covariates in our regression specifications.

Importantly, the stacked samples contain older individuals. People in our study are,

Table 4: Summary Statistics for the Stacked Samples Before Inheritance Events

	Treatment		Control		
	Mean (1)	SD (2)	Mean (3)	SD (4)	
Panel A: Full Sample					
Age	62.67	7.86	66.67	9.28	
Male	0.43	0.50	0.38	0.48	
White	0.92	0.28	0.73	0.44	
Married	0.82	0.39	0.64	0.48	
College	0.59	0.49	0.36	0.48	
Retired	0.42	0.49	0.58	0.49	
Nondurable Spending	_		_		
Total Spending	_		_		
Individuals	2,279		$13,\!825$		
Observations	2,279		63,792		
Panel B: Spending Sample					
Age	65.99	8.08	70.80	8.64	
Male	0.36	0.48	0.33	0.47	
White	0.93	0.25	0.77	0.42	
Married	0.69	0.46	0.50	0.50	
College	0.66	0.47	0.39	0.49	
Retired	0.48	0.50	0.68	0.47	
Nondurable Spending	$27,\!105$	$13,\!419$	19,605	$13,\!523$	
Total Spending	50,179	24,187	$34,\!547$	24,027	
Individuals	1	149		1,614	
Observations	149		4,684		

Notes: This table reports summary statistics for our stacked analysis samples. Panel A corresponds to the full sample, and Panel B corresponds to the spending sample. Columns (1) and (2) report means and standard deviations for the treatment group, respectively. Columns (3) and (4) report the same statistics for the control group. In both panels, the underlying data come from two survey waves before the wave of the inheritance event.

naturally, reaching typical retirement ages. One implication of this fact is that we can reasonably expect to observe any changes in retirement that occur because of an inheritance. It would be much harder to study changes in retirement timing if we instead had data on people experiencing wealth shocks earlier in life. A second implication is that we might expect to see some meaningful consumption responses when people receive an inheritance for two reasons. First, some people might respond by retiring contemporaneously, or soon after receiving an inheritance, which would correspond to the corner solution discussed in our model and would involve accompanying increases in consumption. Second, some people are already retired and can not adjust their retirement dates, and these people should respond

by increasing consumption. However, in practice, retirement is not always an absorbing state (Maestas, 2010), so inheritances could also affect retirement status by causing some people to remain retired who would have otherwise returned to work.

4.4.4 Estimating Equations

Using the stacked samples, we estimate simple, weighted event study regression models. Specifically, we estimate

$$Y_{iae} = \alpha + \beta \cdot Treat_{ia} + \sum_{k \neq -2} \gamma^k \cdot 1[e = k] + \sum_{k \neq -2} \delta^k \cdot Treat_{ia} \cdot 1[e = k] + \varepsilon_{iae}, \tag{12}$$

where Y_{iae} is an outcome (like an indicator for being retired) for person i in sub-experiment a observed during event wave e, $Treat_{ia}$ is an indicator for receiving an inheritance and thus being in the treatment group, 1[e=k] are event-time indicators that take the value of one for observations that are k survey waves away from the time of the inheritance event, and ε is an error term. Following Wing et al. (2024), we use corrective sample weights defined as

$$Q_{ia} = \begin{cases} 1 & \text{if } Treat_{ia} = 1\\ \frac{N_a^T/N^T}{N_a^T/N^C} & \text{if } Treat_{ia} = 0, \end{cases}$$

where N_a^T is the number of treated individuals in sub-experiment a, N^T is the total number of treated individuals, N_a^C is the number of control individuals in sub-experiment a, and N^C is the total number of control individuals. The coefficients of interest are the δ^k s. They are the difference-in-differences estimates that capture the dynamic causal effects of receiving an inheritance. They reflect the average difference in the outcome for the treatment group compared to the control group, relative to the omitted period.

We omit the event time indicator that corresponds to two survey waves before the inheritance. Our outcomes in this omitted wave capture retirement status approximately 3 years before the inheritance, and spending approximately 4 years before the inheritance. We omit this wave instead of the wave immediately preceding the inheritance because some parental deaths that predate inheritances begin to show up in the data one full survey wave before the inheritance, as we show below. We thus treat three of our four pre-inheritance-event waves of data as true pre-periods, one wave of pre-inheritance-event data as a period of partial treatment (when the inheritance has yet to be reported but some people report parental deaths), and two waves of post-inheritance-event data as post-inheritance data.

The key identification assumptions underlying our difference-in-differences design are the parallel trends and no-anticipation assumptions. The parallel trends assumption states that

in the absence of receiving an inheritance, the outcomes for the treated group would have evolved in parallel with the outcomes of the control group. We examine pre-period trends to assess the validity of this assumption. The no-anticipation assumption states that the treatment group is not responding to the inheritance more than two survey waves before it is received. This assumption is perhaps harder to assess, although examining pre-period trends in outcomes is still helpful. For example, if people respond to a future inheritance by retiring earlier, then we would expect to see an increasing pre-trend in retirement for the treatment group compared to the control group, which we do not find. Moreover, Brown et al. (2010) show that many inheritances are unexpected, as does our descriptive analysis above. Even for people who expect to receive an inheritance, there are likely several potential unexpected components, like the size and timing of the inheritance.¹²

Finally, the event study specification is useful for tracking dynamics, but we also want to quantify overall magnitudes and summarize the findings. Guided by the suggestion in Wing et al. (2024), we also report estimates that are the simple averages of the post-period dynamic estimates. To compute these average effects, we estimate

$$Y_{iae} = \alpha_1 + \beta_1 \cdot Treat_{ia} + \gamma_1 \cdot Post_{ae} + \delta_1 \cdot Treat_{ia} \cdot Post_{ae}$$

$$+ \sum_{k < 0, k \neq -2} \left(\gamma_2^k \cdot 1[e = k] + \delta_2^k \cdot Treat_{ia} \cdot 1[e = k] \right) + \mu_{iae}.$$

$$(13)$$

The key difference here is that we replace the two post-inheritance-event indicators with one indicator, $Post_{ae}$, that takes the value of one if the observation corresponds to either post-inheritance wave. The pre-inheritance wave-specific indicators (and their interactions with the treatment group indicator) remain in the regression as before. The coefficient of interest is δ_1 . It captures the average difference in the outcome for the treatment group compared to the control group after the receipt of an inheritance, relative to the difference two waves before the inheritance.

¹²To further examine the extent to which the inheritances we study could be anticipated, we leverage our expectations data. Specifically, we estimate our event study regression model using inheritance expectations as an outcome to examine how expectations evolve around the timing of inheritance events. Appendix Figure B.1 presents the results for two different analysis samples. There are two key findings. First, the pre-period means for people in the treatment group, who are about to receive an inheritance, indicate substantial uncertainty, as their average self-assessed probability of receiving an inheritance is just under 50%. Second, while the pre-period event study estimates do reveal some upward trend in expectations, the increase is modest (an increase of roughly 5 percentage points from 7 years before the inheritance to 1 year before). Moreover, expectations about future inheritances appear to evolve smoothly right around the time of the inheritance, before declining sharply after inheritance receipt. We conclude that the events we study are, on the whole, capturing meaningful shocks to wealth.

4.5 The Effects of Inheritances on Retirement and Consumption

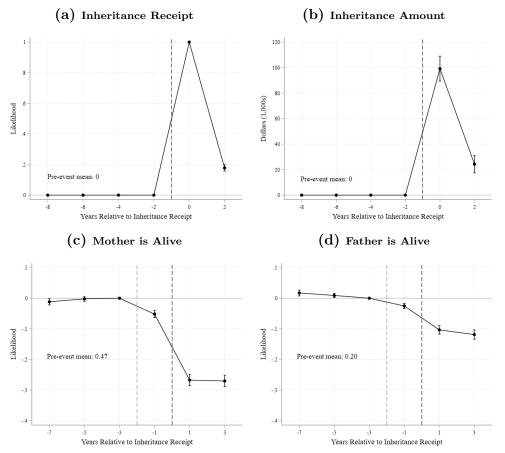
Figure 8 displays event study results for inheritances and parental mortality. These graphs are like a "first stage," in that they show how the treatment evolves for those treated compared to those not treated. The outcome in panel (a) is an indicator of receiving an inheritance. The outcome in panel (b) is the inheritance amount in dollars. Each graph plots the δ^k s from equation (12). The estimates in panel (a) are mechanically zero for the survey waves before the treatment group receives an inheritance, and the estimate for the wave of the event is mechanically 1, because the events we study correspond to the first observed inheritance for each treated person. The fact that the estimate drops to around 0.20 in the following wave indicates that the inheritances we study mostly reflect one-time lump sum payments; about 20% of people in the sample receive additional inheritances in the next wave. Panel (b) provides information on the size of the inheritances that we study. On average, the treatment group receives an average inheritance of just under \$100,000, reflecting a substantial change in wealth.

The outcomes in panels (c) and (d) are for parental mortality. Panel (c) is for an indicator for the mother of the respondent being alive, and panel (d) is for an indicator for the father of the respondent being alive. There are two notable patterns. First, the inheritances we observe are more likely to follow the deaths of mothers as opposed to fathers, who are much less likely to be alive before the inheritance event. Second, some parental deaths predate the reported inheritance by a full survey wave. For example, compared to two waves before the inheritance, the probability of having a living mother declines by about 5 percentage points in the wave before the inheritance. The decline is sharper and even more pronounced in the wave of the inheritance event, but these patterns highlight one reason we might expect to see changes in outcomes in the wave before the inheritance is received, to the extent that people respond to parental mortality by adjusting retirement or spending. For this reason, we include a second, gray dotted line on our event study graphs to highlight pre-inheritance estimates that could be influenced by these earlier parental deaths.

Figure 9 presents event study results for retirement. The probability of being retired is not trending differently in the pre-period for the treatment group compared to the control group. None of the point estimates are statistically different from zero at conventional levels, and the pattern of the estimates appears flat. Moreover, the estimate one year before the inheritance event is also small and not statistically different from zero. These patterns are consistent with a lack of meaningful anticipatory retirements in response to receiving an inheritance or to declines in parental health that might lead to parental mortality.

In stark contrast, the post-period estimates indicate large and sharp increases in retire-

Figure 8: Inheritances and Parental Mortality Around Inheritance Events

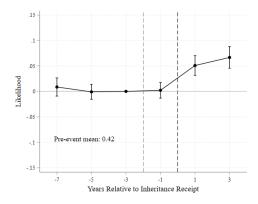


Notes: This figure presents event study graphs that illustrate the dynamic effects of inheritance events on inheritance receipt and parental mortality. The underlying sample is the full stacked sample detailed in the text. Each graph corresponds to a different outcome variable. Panel (a) is for an indicator of receiving an inheritance. Panel (b) is for inheritance amounts. Panel (c) is for an indicator of having a living mother. Panel (d) is for an indicator of having a living father. Each graph plots the δ^k s and the 95% confidence intervals from estimating equation (12).

ment after the receipt of an inheritance. Receiving an inheritance leads to an increase in the probability of being retired of 5.1 percentage points after 1 year and 6.7 percentage points after 3 years. These patterns indicate that some people respond by retiring soon after receiving an inheritance, consistent with the corner solutions discussed in our theoretical model. The fact that the point estimates are increasing over time also indicates that inheritances influence retirements several years later, consistent with some people responding by retiring not immediately, but by moving up their retirement date to stop working earlier than they otherwise would have.

Figure 10 presents event study results for spending using the spending stacked sample. Each graph in the figure corresponds to a different spending outcome. Panel (a) shows the results for nondurable spending. The pre-period estimates are small and are not statistically different from zero. The pattern of the estimates is flat. Interestingly, we see an increase in

Figure 9: The Effect of Inheritances on Retirement



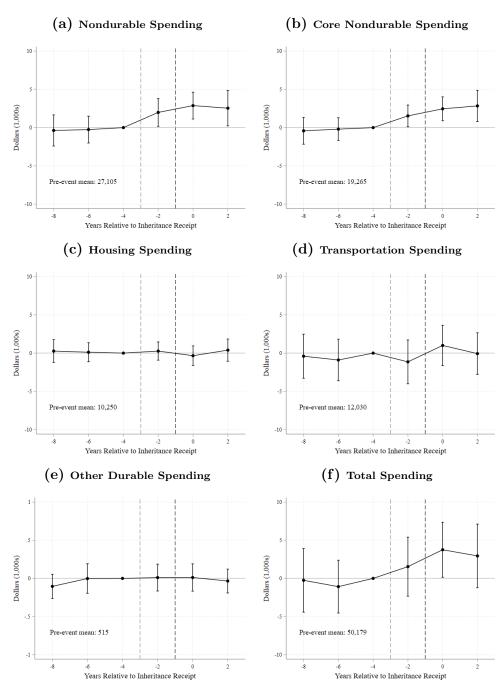
Notes: This figure presents the event study graph that illustrates the dynamic effects of inheritance events on retirement. The underlying sample is the full stacked sample detailed in the text. The outcome variable is an indicator for being retired, defined as reporting a current labor force status as either retired or out of the labor force. The graph plots the δ^k s and the 95% confidence intervals from equation (12).

nondurable spending in the wave immediately preceding the inheritance event, corresponding to about 2 years before the inheritance receipt. This increase could represent increased expenditures associated with declining parental health or parental mortality, but it is also possible that it represents a change in spending because of the impending inheritance.

The post-inheritance estimates for nondurable spending are positive, statistically different from zero, and larger than the estimate in the wave immediately before the inheritance. They indicate increases in spending that amount to roughly \$2,500 and are consistent with our predictions that spending should increase for people who were already retired or who were induced to retire soon after the inheritance event.

One potential concern related to the interpretation of these estimates is that retirement itself could have its own independent impact on spending. Many papers study whether and the extent to which consumption or spending declines in retirement (e.g., Banks et al., 1998; Hurd and Rohwedder, 2003; Aguiar and Hurst, 2005; Smith, 2006; Haider and Stephens Jr, 2007; Kolsrud et al., 2024). In our case, if retirements induced by inheritances are also accompanied by declines in spending, then the estimates in panel (a) would be the result of a mix of spending increases because of the inheritances and decreases because of retirements. To lessen this concern, we follow Aguiar and Hurst (2013), who argue for isolating core nondurable spending (nondurable spending minus food and work-related expenses) from "home production" spending (work-related expenses and food). Panel (b) shows the results for core nondurable spending, which we construct as nondurable spending (which already excludes gas and transportation) minus spending on dining out, food and beverages, personal care, and clothing. The results for core nondurable spending, which tend not to decline in retirement, are very similar to those for nondurable spending.

Figure 10: The Effect of Inheritances on Spending



Notes: This figure presents event study graphs that illustrate the dynamic effects of inheritance events on spending. The underlying sample is the spending stacked sample detailed in the text. Each panel corresponds to a different outcome variable. Panel (a) is for nondurable spending, panel (b) is for core nondurable spending, panel (c) is for housing spending, panel (d) is for transportation spending, panel (e) is for other durable spending, and panel (f) is for total spending. Each graph plots the δ^k s and the 95% confidence intervals from equation (12).

Notably, nondurable spending is the only spending category for which we find evidence of a change. Panels (c), (d), and (e) indicate no evidence of increases in spending on housing, transportation, or other durables, respectively. Panel (f) is for total spending, the sum of

these main components. The pattern of the estimates mimics panel (a), but because this measure includes the other components for which we find no statistical evidence of responses. The graph is less clear, and the estimates in the post-period years are less precise.

Table 5 displays the point estimates from equation (13) that summarize magnitudes and help to interpret these patterns. Panel A is for the full sample, which we use to document the main result for retirement. The point estimate for inheritances in column (1) corresponds to the estimate in the event study graph for the survey wave of the reported inheritance. The inheritances in this sample are, on average, \$99,189. The point estimate for retirement in column (2) is the average of the two event study post-period estimates above. It is statistically significant at the 1% level and indicates that the receipt of an inheritance causes a 5.9 percentage point increase in the likelihood of being retired. This estimate is large. It represents a 12% increase when compared to the pre-period mean.

Panel B is for the spending sample. Column (1) shows that the inheritances in this sample are on average \$96,608 and are thus similar in size to the inheritances in the full sample. Column (2) shows that the point estimate for retirement is statistically significant and larger in the spending sample compared to the full sample, although the standard error is also larger and the confidence interval contains the full-sample estimate.¹³

Columns (3) and (4) display the estimates for nondurable and total spending.¹⁴ We find an increase in nondurable spending that is highly statistically significant and that amounts to \$2,710 on average, and an increase in total spending that is marginally statistically significant and that amounts to \$3,351 on average. On the one hand, these estimates translate to sizable 9.5% and 6.7% increases when compared to the means. On the other hand, the estimates are modest (2.8% and 3.5%) when compared to the size of the inheritances.

It is also important to compare the size of these responses to what we might have predicted based on a simple version of the permanent income hypothesis. How much would spending increase in response to the inheritances if individuals could not change their retirement dates? The treatment group in the spending sample (which contains events in later waves) is roughly 70 years old on average in the wave of the inheritance event. If we ignore discounting, which will establish a lower bound for a consumption response, and assume people would live to 80 on average, then dividing the \$96,608 by 10 years of remaining life translates to a rough benchmark for spending increases of \$9,600. Our nondurable spending estimate is less than one-third, and our total spending estimate is about 35%, of this benchmark.¹⁵ Our estimates

¹³In the appendix, we present event studies for inheritances and parental mortality (Appendix Figure B.2) and retirement (Appendix Figure B.3) in the spending sample. Results are similar to the full sample.

¹⁴Appendix Table B.1 reports estimates for the other spending components (housing, transportation, and other durables), none of which are statistically distinguishable from zero.

¹⁵The upper bound of the 95-percent confidence interval for nondurable spending is \$4,502 and the upper

Table 5: The Effects of Inheritance Receipt on Retirement and Spending

	Inheritance (1)	Retirement (2)	Nondurable Spending (3)	Total Spending (4)
Panel A: Full Sample				
Diff-in-Diff Estimate	99,189*** (4,968)	0.059^{***} (0.010)	_	_
Mean	0	0.481	_	_
Households	11,201	11,201	_	_
Individuals	16,104	16,104	_	_
Observations	$396,\!426$	$396,\!426$	_	_
Panel B: Spending Sample				
Diff-in-Diff Estimate	96,608***	0.105***	2,710***	3,351*
	(16,346)	(0.035)	(913)	(1,759)
Mean	0	0.544	28,417	49,963
Households	1,750	1,750	1,750	1,750
Individuals	1,763	1,763	1,763	1,763
Observations	28,998	28,998	28,998	28,998

Notes: This table presents difference-in-differences estimates of the effects of inheritance events on inheritance amounts, retirement, and spending. Column (1) reports the estimate of δ^0 from equation (12) when the outcome is the inheritance amount. Columns (2), (3), and (4) report estimates of δ_1 from equation (13) when the outcomes are retirement, nondurable spending, and total spending, respectively. Panel A presents estimates using the full stacked sample as detailed in the text. Panel B presents estimates using the spending stacked sample as detailed in the text. In addition to the difference-in-differences estimates, the rows within each panel present the dependent variable means for the treatment group in the omitted survey wave (two waves before the inheritance event), and the number of households, individuals, and observations that contribute to each regression. Standard errors are clustered at the household level.

are thus relatively far from predictions that ignore the powerful retirement response.

Moreover, the size of the estimates are reasonably close to predictions based on our theoretical framework, which emphasizes a dominant role for retirement responses unless that response margin cannot be adjusted. Roughly half of the treatment group in the spending sample was already retired before receiving their inheritances. Our model would thus predict very strong retirement responses from the non-retired population and spending increases from the rest. Multiplying the \$9,600 benchmark by one-half would suggest spending responses of \$4,800. Our estimates are much closer to this prediction.¹⁶

Taken together, our empirical results align with the predictions from our theoretical frame-

^{*} p < 0.10, ** p < 0.05, *** p < 0.01

bound for total spending, which is less precisely estimated, is \$6,801, still below the \$9,600 benchmark.

¹⁶The sample size (149 treated individuals) is too small to draw clear conclusions from subsample analyses that attempt to identify separate responses for these two groups of people.

work. We find strong retirement responses and relatively modest spending responses when compared to what we might have predicted if we were to have ignored retirement. The evidence thus indicates that retirement insures consumption in the case of inheritances. Before concluding, we conduct several robustness checks on the empirical analysis.

4.6 Robustness Checks

First, we assess the robustness of our results to regression specification choices. Table 6 displays the results from these assessments. Within each panel, the rows of the table correspond to different robustness checks, with the baseline estimates reproduced at the top, and the columns correspond to the retirement and spending outcomes. The point estimates come from estimating different versions of equation (13).

We assess the robustness of our regression estimates along five dimensions. First, we add controls, motivated by the demographic differences in levels across the treatment and control groups described earlier. We include dummy variables for gender, race, marital status, and college attendance, as well as age fixed effects, which are important because people in our sample are navigating Social Security eligibility ages, when there are non-linear and age-specific increases in retirement. Including these control variables leads to somewhat smaller estimates in general, but the key takeaways remain. Second, we drop the stacked weights proposed by Wing et al. (2024). By doing so, we estimate the simplest possible unweighted OLS difference-in-differences models, and we obtain similar results. Third, we use respondent survey weights for population representation instead of the stacked weights. The estimates for retirement and nondurable spending are quite similar to their baseline counterparts, although the standard errors are larger, and the estimate for total spending is not statistically significant.

Fourth, we use an alternative definition of retirement that does not consider people out of the labor force as retired. Instead, we use an outcome that takes the value of one only if the respondent reports their labor force status as retired. The estimates for this definition are similar to the baseline estimates, indicating that people tend to view their inheritance-induced exits from the labor force as retirements per se. Fifth, we change how we handle outliers in the spending data. In our baseline analysis, we winsorize spending variables at the 1st and 99th percentiles. If we do not winsorize, we obtain similar point estimates. If we winsorize more aggressively, at the 5th and 95th percentiles, we obtain smaller estimates, especially for nondurable spending, which would suggest even more modest spending responses. The averages in our baseline analysis might thus reflect some larger spending responses by a relatively smaller group of people.

Next, we assess the robustness of our results to alternative constructions of our stacked

Table 6: Robustness of Estimates to Alternative Regression Specifications

	Retirement (1)	Nondurable Spending (2)	Total Spending (3)
Panel A: Full Sample			
Baseline Estimate	0.059^{***} (0.010)	_	_
Add Controls	0.041^{***} (0.009)	_	_
Drop Stacked Weights	0.062^{***} (0.010)	_	_
Use Survey Weights	0.061*** (0.013)	_	-
Alternative Retirement Definition	0.054*** (0.010)	_	_
Panel B: Spending Sample	,		
Baseline Estimate	0.105^{***} (0.035)	2,710*** (913)	$3,351^*$ $(1,759)$
Add Controls	0.084** (0.035)	2,257** (961)	2,923 (1,847)
Drop Stacked Weights	0.104^{***} (0.035)	2,614*** (913)	3,168* (1,758)
Use Survey Weights	0.103^{**} (0.050)	2,860* (1,467)	2,557 $(2,512)$
Alternative Retirement Definition	0.112*** (0.037)	_	_
Winsorize Spending More	_	1,838** (752)	2,523 $(1,560)$
Do Not Winsorize Spending	_	$3,031^{***}$ (972)	3,458** $(1,760)$

Notes: This table presents difference-in-differences estimates of δ_1 from equation (13) as we vary the regression specification. Panel A presents estimates using the full stacked sample as detailed in the text. Panel B presents estimates using the spending stacked sample as detailed in the text. The rows within each panel correspond to different regression specifications. Each column corresponds to a different outcome. Standard errors are clustered at the household level.

samples based on the choice of event windows. In our baseline analysis, we trim the samples to be balanced panels of individuals who appear in all event waves $e \in \{-4, 1\}$. This time horizon is useful because it allows us to track outcomes over several pre-period waves, one anticipatory wave during which we observe some parental deaths, and multiple post-

^{*} p < 0.10, ** p < 0.05, *** p < 0.01

period waves. However, balancing relatively long panels leads to smaller sample sizes. We therefore consider two alternative approaches to constructing our stacked samples, one that requires one fewer pre-period wave, resulting in a balanced panel of individuals in event waves $e \in \{-3, 1\}$ and one that requires one fewer post-period wave, resulting in a balanced panel of individuals in event waves $e \in \{-4, 0\}$. Appendix Figures B.4 and B.5 display the event study results for retirement and the two key spending measures, respectively. Appendix Tables B.2 and B.3 present the point estimates that summarize magnitudes for the two different approaches to constructing the samples. Overall, the results are similar to our baseline, and our takeaways do not change.

Finally, we present estimates from a traditional two-way fixed effects estimator. This estimator would identify the average treatment effect on the treated if both the parallel trends and no-anticipation assumptions hold and if the paths of treatment effects are homogeneous across treatment cohorts. This third assumption is an additional, strong assumption that our main approach does not make. We begin with the base analysis sample (not the stacked samples) detailed earlier. To produce estimates that should be compared with our main estimates, we trim the sample of people who we observe receiving an inheritance by keeping a balanced panel of treated individuals who appear in the data in each survey wave that corresponds to event times $e \in \{-4,1\}$. The control group consists of everyone in the analysis data who we do not observe receiving an inheritance, and we do not impose any balance condition on these individuals. We then estimate a simple event study equation,

$$Y_{it} = \alpha_i + \lambda_t + \sum_k \theta^k \cdot Treat_{ik} + \nu_{it},$$

where Y is an outcome for person i in survey wave t, α_i are person fixed effects, λ_t are wave fixed effects, $Treat_{ik}$ is an indicator that person i is treated and is k periods away from the inheritance event. The coefficients of interest are the θ^k s. As in our main analysis, we estimate these two-way fixed-effects models using two samples: a full sample without additional restrictions and a spending sample that is first limited to only observations with non-missing spending values.

Appendix Figure B.6 presents the results. Panels (a) and (b) show how inheritance amounts evolve around the events in both samples, panels (c) and (d) show results for retirement in both samples, and panels (e) and (f) show results for nondurable and total spending in the spending sample. Overall, the results are very similar to our main analysis.

We conclude from these sets of robustness checks that our main estimates are quite stable. There is strong and clear evidence that receiving an inheritance induces retirement. There is also clear evidence of spending responses that are relatively modest.

5 Conclusion

In this paper, we study how people respond to changes in wealth and permanent income. Using a life-cycle framework, we highlight the power of endogenous retirement. In a simple model with constant wages and constant fixed costs of work, people respond to changes in wealth by adjusting their retirement timing and holding consumption constant. Additional exercises highlight the dominant role retirement plays in richer models. We then estimate the causal effects of receiving an inheritance on retirement and spending. Consistent with our theoretical predictions, we find clear evidence that this large and late-in-life shock induces immediate increases in retirements that are accompanied by relatively modest increases in spending.

Our central point is that this substantial consumption insurance role for endogenous retirement has been underappreciated. One implication of this insurance role is that other shocks, policies, or institutions that reduce the ability to work later in life play a more prominent role in life-cycle earnings risk. While many individuals are able to choose their retirement date to insure consumption, others are less flexible. For example, health shocks that prevent individuals from being able to work longer can reduce insurance, and uncertainty about retirement timing like that studied in Caliendo et al. (2023) presents a large risk to lifetime consumption. Moreover, policies that mandate retirement or that strongly incentivize retirements at specific dates, as in Rust and Phelan (1997), can make it more costly for individuals to insure their lifetime income risk by changing when they retire. While Bronshtein et al. (2019) emphasize the power of working longer for improving late-in-life standard of living, our study adds yet another factor to consider when analyzing the importance of work capacity at older ages.

Furthermore, our results have broad implications for economic policies that aim to influence outcomes through wealth effects. Projections and analyses of the effects of these types of policies should consider retirement as a key margin that people might adjust. For example, major reforms to social security systems are currently being enacted around the world. While it is important to analyze the impacts of these reforms on labor supply and consumption-savings decisions throughout the life cycle, our findings highlight how one might expect a good deal of the responses to come in the form of changes in retirement later in life. Consistent with this general idea, Slavov et al. (2019) find little evidence of savings responses to social security reforms and García-Miralles and Leganza (2024) provide evidence of savings responses that occur primarily through delayed retirement.

References

- Aguiar, M. and Hurst, E. (2005). Consumption versus expenditure. *Journal of Political Economy*, 113(5):919–948.
- Aguiar, M. and Hurst, E. (2013). Deconstructing life cycle expenditure. *Journal of Political Economy*, 121(3):437–492.
- Attanasio, O. P. and Weber, G. (1995). Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy*, 103(6):1121–1157.
- Banks, J., Blundell, R., and Tanner, S. (1998). Is there a retirement-savings puzzle? *American Economic Review*, pages 769–788.
- Belloc, I., Molina, J. A., and Velilla, J. (2025). Consumption responses to inheritances: The role of durable goods. *Journal of Macroeconomics*, 83:103661.
- Blundell, R., Browning, M., and Meghir, C. (1994). Consumer Demand and the Life-Cycle Allocation of Household Expenditures. *The Review of Economic Studies*, 61(1):57–80.
- Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921.
- Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). Consumption Inequality and Family Labor Supply. *American Economic Review*, 106(2):387–435.
- Bronshtein, G., Scott, J., Shoven, J. B., and Slavov, S. N. (2019). The power of working longer. *Journal of Pension Economics and Finance*, 18(4):623–644.
- Brown, J. R., Coile, C. C., and Weisbenner, S. J. (2010). The effect of inheritance receipt on retirement. *The Review of Economics and Statistics*, 92(2):425–434.
- Caliendo, F. N., Casanova, M., Gorry, A., and Slavov, S. (2023). Retirement timing uncertainty: Empirical evidence and quantitative evaluation. *Review of Economic Dynamics*.
- Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.
- Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. The Quarterly Journal of Economics, 112(1):1–55.
- Casanova, M. (2013). Revisiting the hump-shaped wage profile: Implications for structural labor supply estimation. Working paper, UCLA.
- Cesarini, D., Lindqvist, E., Notowidigdo, M. J., and Ostling, R. (2017). The effect of wealth on individual and household labor supply: evidence from swedish lotteries. *American Economic Review*, 107(12):3917–3946.
- Coile, C. C. and Levine, P. B. (2006). Bulls, bears, and retirement behavior. *ILR Review*, 59(3):408–429.
- Deaton, A. (1991). Saving and Liquidity Constraints. Econometrica, 59(5):1221–1248.
- Deshpande, M. and Li, Y. (2019). Who is screened out? application costs and the targeting of disability programs. *American Economic Journal: Economic Policy*, 11(4):213–248.
- Diamond, P. and Köszegi, B. (2003). Quasi-hyperbolic discounting and retirement. *Journal of Public Economics*, 87(9):1839–1872.
- Dybvig, P. H. and Liu, H. (2010). Lifetime consumption and investment: Retirement and constrained borrowing. *Journal of Economic Theory*, 145(3):885–907.
- Fan, X., Seshadri, A., and Taber, C. (2024). Estimation of a life-cycle model with human capital, labor supply, and retirement. *Journal of Political Economy*, 132(1):48 95.

- French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. The Review of Economic Studies, 72(2):395–427.
- French, E. and Jones, J. B. (2011). The effects of health insurance and self-insurance on retirement behavior. *Econometrica*, 79(3):693–732.
- Friedman, M. (1957). A Theory of the Consumption Function. National Bureau of Economic Research, Inc.
- García-Miralles, E. and Leganza, J. M. (2024). Public pensions and private savings. *American Economic Journal: Economic Policy*, 16(2):366–405.
- Gelber, A. M., Isen, A., and Song, J. (2016). The effect of pension income on elderly earnings: Evidence from social security and full population data. *NBER Working paper*.
- Golosov, M., Graber, M., Mogstad, M., and Novgorodsky, D. (2024). How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income. *The Quarterly Journal of Economics*, 139(2):1321–1395.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. Journal of Econometrics, 225(2):254–277.
- Gorry, A. and Oberfield, E. (2012). Optimal Taxation Over the Life Cycle. Review of Economic Dynamics, 15(4):551–572.
- Gustman, A. L., Steinmeier, T. L., and Tabatabai, N. (2010). What the stock market decline means for the financial security and retirement choices of the near-retirement population. *Journal of Economic Perspectives*, 24(1):161–182.
- Guvenen, F. and Smith, A. A. (2014). Inferring Labor Income Risk and Partial Insurance From Economic Choices. *Econometrica*, 82:2085–2129.
- Haider, S. J. and Stephens Jr, M. (2007). Is there a retirement-consumption puzzle? evidence using subjective retirement expectations. *The review of economics and statistics*, 89(2):247–264.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2009). Quantitative Macroeconomics with Heterogeneous Households. *Annual Review of Economics*, 1(1):319–354.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2014). Consumption and Labor Supply with Partial Insurance: An Analytical Framework. *American Economic Review*, 104(7):2075–2126.
- Holtz-Eakin, D., Joulfaian, D., and Rosen, H. S. (1993). The carnegie conjecture: Some empirical evidence. *The Quarterly Journal of Economics*, 108(2):413–435.
- Hubbard, R., Skinner, J., and Zeldes, S. P. (1994). The importance of precautionary motives in explaining individual and aggregate saving. Carnegie-Rochester Conference Series on Public Policy, 40:59–125.
- Hurd, M. D. and Rohwedder, S. (2003). The retirement-consumption puzzle: Anticipated and actual declines in spending at retirement.
- Imbens, G. W., Rubin, D. B., and Sacerdote, B. I. (2001). Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players. *American economic review*, 91(4):778–794.
- Jappelli, T. and Pistaferri, L. (2010). The consumption response to income changes. *Annual Review of Economics*, 2(1):479–506.
- Joulfaian, D. and Wilhelm, M. O. (1994). Inheritance and labor supply. *Journal of Human Resources*, pages 1205–1234.

- Kaplan, G. and Violante, G. L. (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, 2(4):53–87.
- Keane, M. and Rogerson, R. (2012). Micro and macro labor supply elasticities: A reassessment of conventional wisdom. *Journal of Economic Literature*, 50(2):464–76.
- Kolsrud, J., Landais, C., Reck, D., and Spinnewijn, J. (2024). Retirement consumption and pension design. *American Economic Review*, 114(1):89–133.
- Krueger, A. B. and Pischke, J.-S. (1992). The effect of social security on labor supply: A cohort analysis of the notch generation. *Journal of labor economics*, 10(4):412–437.
- Low, H. (2005). Self-Insurance in a Life-Cycle Model of Labor Supply and Savings. *Review of Economic Dynamics*, 8(4):945–975.
- Maestas, N. (2010). Back to work: Expectations and realizations of work after retirement. Journal of Human Resources, 45(3):718–748.
- Meghir, C. and Pistaferri, L. (2011). Earnings, Consumption and Life Cycle Choices. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 4 of *Handbook of Labor Economics*, chapter 9, pages 773–854. Elsevier.
- Picchio, M., Suetens, S., and van Ours, J. C. (2018). Labour supply effects of winning a lottery. *The Economic Journal*, 128(611):1700–1729.
- Rogerson, R. (2024). Why labor supply matters for macroeconomics. *Journal of Economic Perspectives*, 38(2):137–58.
- Rogerson, R. and Wallenius, J. (2009). Micro and macro elasticities in a life cycle model with taxes. *Journal of Economic Theory*, 144(6):2277–2292. Dynamic General Equilibrium.
- Rogerson, R. and Wallenius, J. (2013). Nonconvexities, retirement, and the elasticity of labor supply. *American Economic Review*, 103(4):1445–62.
- Rupert, P. and Zanella, G. (2015). Revisiting Wage, Earnings, and Hours Profiles. *Journal of Monetary Economics*, 72(C):114–130.
- Rust, J. and Phelan, C. (1997). How social security and medicare affect retirement behavior in a world of incomplete markets. *Econometrica*, 65(4):781–831.
- Slavov, S., Gorry, D., Gorry, A., and Caliendo, F. N. (2019). Social security and saving: An update. *Public Finance Review*, 47(2):312–348.
- Smith, S. (2006). The retirement-consumption puzzle and involuntary early retirement: Evidence from the British household panel survey. *The Economic Journal*, 116(510):C130–C148.
- Suari-Andreu, E. (2023). Labour supply, retirement, and consumption responses of older europeans to inheritance receipt. *Empirical Economics*, 64(1):33–75.
- vom Lehn, C., Gorry, A., and Fisher, E. O. (2018). Male labor supply and generational fiscal policy. *Review of Economic Dynamics*, 28:121–149.
- Wing, C., Freedman, S. M., and Hollingsworth, A. (2024). Stacked difference-in-differences. Technical report, National Bureau of Economic Research.

Appendix

A Intensive and Extensive Labor Supply Decomposition

In our baseline model, we include only an extensive labor choice to keep the focus on the main mechanism. We now show that our main result does not depend on this simplification by considering a model with both a retirement choice with a fixed cost of work and intensive margin labor choices with a variable cost of work at each date.¹⁷ The individual chooses consumption c(t), labor supply h(t), and the date of retirement t_R to maximize lifetime utility:

$$\max_{c(t),h(t),t_R} \int_0^T [u(c(t)) - v(h(t))] dt$$

subject to a resource constraint

$$\int_0^T c(t)dt = \int_0^{t_R} w(t)h(t)dt + B.$$

Let χ be the fixed cost of work at any age t and let α be the variable cost per hour of work at any age t. Both are modeled as utility costs such that

$$v(h) = \begin{cases} \chi + \alpha h & h > 0 \\ 0 & h = 0. \end{cases}$$

Let u(c(t)) be period utility from consumption, where u' > 0 and u'' < 0 to ensure a well defined maximization problem. The individual's problem can be written as:

$$\max_{c(t),h(t),t_R} \int_0^T u(c(t))dt - \int_0^{t_R} \chi dt - \int_0^{t_R} \alpha h(t)dt$$

subject to

$$\int_0^T c(t)dt = \int_0^{t_R} w(t)h(t)dt + B.$$

Consider the case in which w(t) = w. In this setting, optimal consumption is constant over the life cycle, and optimal hours are constant up to the retirement date. The problem

¹⁷We again assume that the individual begins in the labor force and chooses a single retirement date. This avoids the common issue that many models with endogenous retirement pin down the length of the working life, but not which periods are optimal to work unless the returns to work vary over the life cycle.

simplifies to

$$\max_{c,h,t_R} Tu(c) - t_R \chi - t_R \alpha h$$

subject to

$$c = \frac{t_R w h + B}{T}.$$

We consider three cases of how the individual responds to a change in wealth B:

Case 1: The retirement date, t_R , and hours of work on the intensive margin, h, are both exogenous and fixed.

In this case, optimal consumption

$$c^* = \frac{t_R w h + B}{T}$$

absorbs the full shock, and the individual consumes the annuity value of an increase in wealth

$$\frac{\partial c^*}{\partial B} = \frac{1}{T}.$$

Case 2: The retirement date, t_R , is exogenous and fixed but the individual is allowed to adjust hours of work on the intensive margin, h.

The first order condition in hours is

$$Tu'\left(\frac{t_Rwh+B}{T}\right)\frac{t_Rw}{T}=t_R\alpha.$$

Let $m(\cdot) = [u'(\cdot)]^{-1}$ be the inverse of the marginal utility of consumption. The first order condition becomes

$$\frac{t_R w h + B}{T} = m \left(\frac{\alpha}{w}\right)$$

$$h^* = \frac{Tm\left(\frac{\alpha}{w}\right) - B}{t_R w}$$

$$c^* = \frac{t_R w}{T} \left(\frac{Tm\left(\frac{\alpha}{w}\right) - B}{t_R w} \right) + \frac{B}{T}$$
$$= m\left(\frac{\alpha}{w}\right).$$

Here, a change in B is absorbed entirely along the intensive labor margin with no change in consumption.

Case 3: Both the retirement date, t_R , and hours of work on the intensive margin, h, are free to adjust.

The first order conditions for t_R and h are

$$u'\left(\frac{t_Rwh+B}{T}\right)wh=\chi$$

$$u'\left(\frac{t_Rwh+B}{T}\right)w=\alpha.$$

Combining the above equations gives

$$\frac{\chi}{wh} = \frac{\alpha}{w}$$

or

$$h^* = \frac{\chi}{\alpha}$$
.

And,

$$t_R^* = \frac{Tm\left(\frac{\alpha}{w}\right) - B}{w_{\alpha}^{\chi}}$$

$$c^* = \frac{t_R w h + B}{T}$$

$$= \frac{1}{T} \left(\frac{T m \left(\frac{\alpha}{w} \right) - B}{w_{\alpha}^{\underline{\chi}}} \right) w_{\alpha}^{\underline{\chi}} + \frac{B}{T}$$

$$= m \left(\frac{\alpha}{w} \right).$$

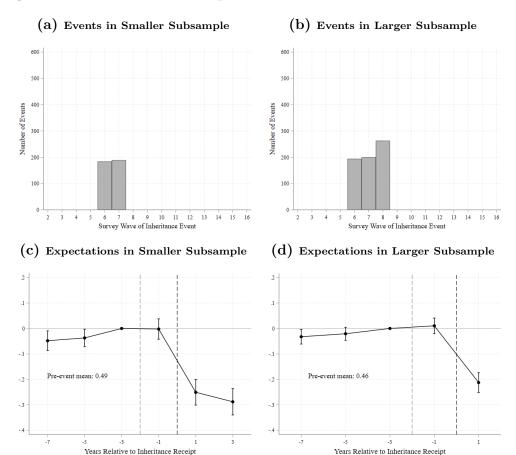
Here, a change in B is absorbed fully along the extensive labor margin, with intensive hours and consumption holding constant in the face of a change in B.

Comparing the three cases confirms the intuition and main result from the paper: when individuals are able to adjust their labor supply in response to a wealth shock, the labor margin absorbs most of the shock. In Case 2, changes in wealth are absorbed along the intensive labor supply margin. In Case 3, changes in wealth are absorbed along the extensive labor supply margin. It is only in Case 1, with exogenous intensive and extensive labor supply, that changes in wealth lead to changes in consumption.¹⁸

¹⁸Sufficiently large changes in wealth that lead to corner solutions for labor supply could also trigger changes in consumption in Cases 2 and 3, as in the main body of the paper.

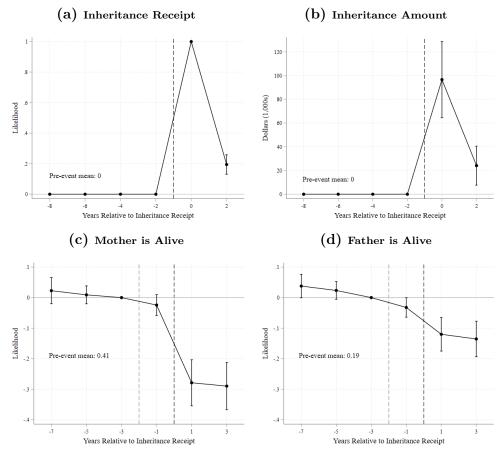
B Additional Empirical Results

Figure B.1: Inheritances Expectations Around Inheritance Events



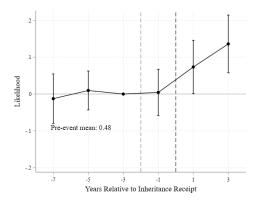
Notes: This figure presents two histograms that illustrate the construction of the stacked samples used in our difference-in-differences analysis for studying inheritance expectations, and the corresponding event study results on expectations. Because the expectations data are only available in waves 2 through 8, the number of events for which we can track expectations in a balanced, stacked sample is limited. Panel (a) shows the histogram of events that we can use if we start with non-missing expectations data and create a stacked sample like in our baseline analysis, keeping a balanced panel with 4 pre-period waves and 2 post-period waves. We can study the 373 treated individuals who experience events in waves 6 and 7. Panel (b) shows the events that we can study if we relax our sample construction requirements to instead balance the sample over 4 pre-period waves and one post-period wave, with no additional wave after the inheritance receipt. This approach allows us to study 657 unique individuals experiencing an event at the cost of not being able to track expectations several years after the inheritances, which is minor because our focus is on examining pre-period data. Panels (c) and (d) show the corresponding event study results. The outcome is the probability of receiving an inheritance in the next ten years. Each graph plots the δ^k s and the 95% confidence intervals from equation (12).

Figure B.2: Inheritances and Parental Mortality Around Inheritance Events in the Spending Sample



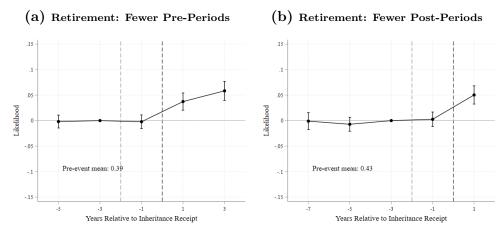
Notes: This figure presents event study graphs that illustrate the dynamic effects of inheritance events on inheritance receipt and parental mortality. The underlying sample is the spending stacked sample detailed in the text. Each panel corresponds to a different outcome variable. Panel (a) is for an indicator for receiving an inheritance. Panel (b) is for inheritance amounts. Panel (c) is for an indicator of having a living mother. Panel (d) is for an indicator of having a living father. Each graph plots the δ^k s and the 95% confidence intervals from equation (12).

Figure B.3: Event Study Estimates for Retirement in the Spending Sample



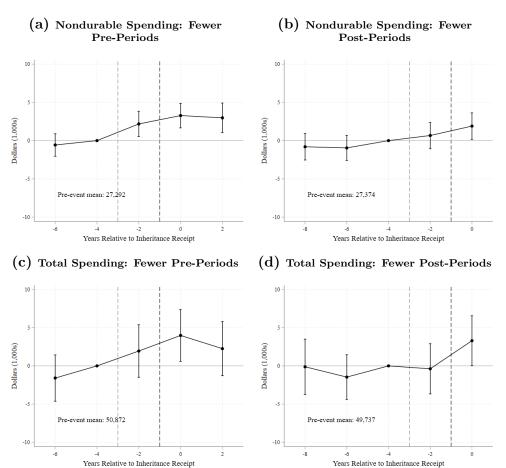
Notes: This figure presents the event study graph that illustrates the dynamic effects of inheritance events on retirement. The underlying sample is the spending stacked sample detailed in the text. The outcome variable is an indicator for being retired, defined as reporting a current labor force status as either retired or out of the labor force. The graph plots the δ^k s and the 95% confidence intervals from equation (12).

Figure B.4: Event Study Estimates for Retirement Using Alternative Approaches to Constructing the Stacked Sample



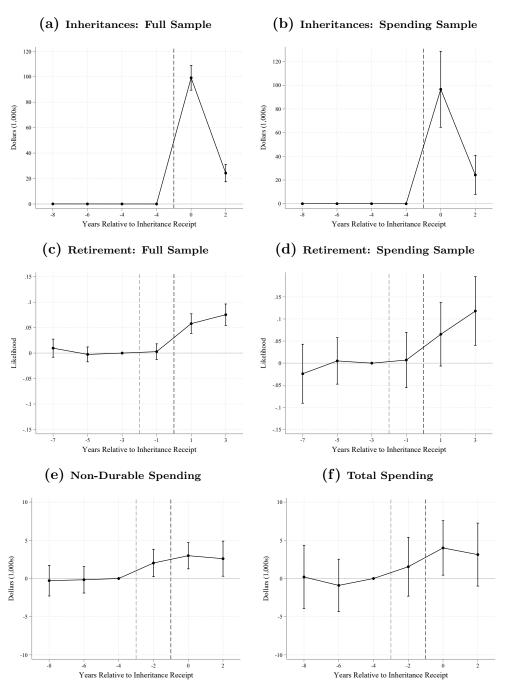
Notes: This figure presents event study graphs that illustrate the dynamic effects of inheritance events on retirement when we use alternative approaches to constructing the full stacked sample. Panel (a) shows results when we use a balanced panel in event time, e, from e=-3 to e=1. Panel (b) shows results when we use a balanced panel in event time from e=-4 to e=0. The outcome variable is an indicator for being retired, defined as reporting a current labor force status as either retired or out of the labor force. The graphs plot the δ^k s and the 95% confidence intervals from equation (12).

Figure B.5: Event Study Estimates for Key Spending Measures Using Alternative Approaches to Constructing the Stacked Sample



Notes: This figure presents event study graphs that illustrate the dynamic effects of inheritance events on non-durable spending and total spending when we use alternative approaches to constructing the spending stacked sample. Panels (a) and (c) show results for nondurable spending and total spending when we use a balanced panel in event time, e, from e=-3 to e=1. Panels (b) and (d) show results for nondurable spending and total spending when we use a balanced panel in event time from e=-4 to e=0. The graphs plot the δ^k s and the 95% confidence intervals from equation (12).

Figure B.6: Event Study Estimates Using a Two-Way Fixed Effects Estimator



Notes: This figure presents event study graphs from a two-way fixed effects estimator. Each graph corresponds to a different outcome variable. Panels (a) and (b) are for inheritances in the full stacked sample and the spending stacked sample. Panels (c) and (d) are for retirement in the full stacked sample and the spending stacked sample. Panels (e) and (f) are for nondurable spending and total spending in the spending stacked sample. The graphs plot the θ^k s and the 95% confidence intervals from estimating the equation specified in Section 4.6.

Table B.1: The Effects of Inheritance Receipt on Other Spending Components

	Housing Spending (1)	Transportation Spending (2)	Other Durable Spending (3)
Diff-in-Diff Estimate	27 (606)	469 (1,187)	-11 (74)
Mean	0	10,135	505
Households	1,750	1,750	1,750
Individuals	1,763	1,763	1,763
Observations	28,998	28,998	28,998

Notes: This table presents difference-in-differences estimates of the effects of inheritance events on housing spending, transportation spending, and other durable spending. The underlying sample is the spending stacked sample as detailed in the text. In addition to the difference-in-differences estimates, the rows within each panel present the dependent variable means for the treatment group in the omitted survey wave (two waves before the inheritance event), and the number of households, individuals, and observations that contribute to each regression. Standard errors are clustered at the household level.

^{*} p < 0.10, ** p < 0.05, *** p < 0.01

Table B.2: The Effects of Inheritance Receipt on Retirement and Spending Using Alternative Approaches to Constructing the Stacked Samples: One Fewer Pre-Period

	Inheritance (1)	Retirement (2)	Nondurable Spending (3)	Total Spending (4)
Panel A: Full Sample				
Diff-in-Diff Estimate	99,634*** (4,486)	0.049*** (0.008)	_	_
Mean	0	0.441	_	_
Households	12,944	12,944	_	_
Individuals	18,931	18,931	_	_
Observations	$414,\!645$	414,645	_	_
Panel B: Spending Sample				
Diff-in-Diff Estimate	100,479***	0.094***	3,411***	3,921***
	(12,626)	(0.029)	(739)	(1,396)
Mean	0	0.509	28,607	50,848
Households	2,169	2,169	2,169	2,169
Individuals	2,193	2,193	$2,\!193$	2,193
Observations	34,400	34,400	34,400	$34,\!400$

Notes: This table presents difference-in-differences estimates of the effects of inheritance events on inheritance amounts, retirement, and spending when we use an alternative approach to constructing the stacked samples that balances the panel in event time, e, from e=-3 to e=1. Column (1) reports the estimate of δ^0 from equation (12) when the outcome is the inheritance amount. Columns (2), (3), and (4) report estimates of δ_1 from estimating equation (13) when the outcomes are retirement, nondurable spending, and total spending, respectively. Panel A presents estimates using the full stacked sample as detailed in the text. Panel B presents estimates using the spending stacked sample as detailed in the text. In addition to the difference-in-differences estimates, the rows within each panel present the dependent variable means for the treatment group in the omitted survey wave (two waves before the inheritance event), and the number of households, individuals, and observations that contribute to each regression. Standard errors are clustered at the household level.

^{*} p < 0.10, ** p < 0.05, *** p < 0.01

Table B.3: The Effects of Inheritance Receipt on Retirement and Spending Using Alternative Approaches to Constructing the Stacked Samples: One Fewer Post-Period

	Inheritance (1)	Retirement (2)	Nondurable Spending (3)	Total Spending (4)
Panel A: Full Sample				
Diff-in-Diff Estimate	$106,918^{***} $ $(5,226)$	0.053^{***} (0.009)	_	_
Mean	0	0.491	_	_
Households	12,742	12,742	_	_
Individuals	18,666	18,666	_	_
Observations	$413,\!320$	$413,\!320$	_	_
Panel B: Spending Sample				
Diff-in-Diff Estimate	127,246***	0.070**	2,485***	3,828**
	(22,800)	(0.031)	(864)	(1,537)
Mean	0	0.544	27,423	47,791
Households	$2,\!151$	$2,\!151$	2,151	$2,\!151$
Individuals	$2,\!174$	$2,\!174$	$2,\!174$	$2,\!174$
Observations	$34,\!305$	$34,\!305$	$34,\!305$	34,305

Notes: This table presents difference-in-differences estimates of the effects of inheritance events on inheritance amounts, retirement, and spending when we use an alternative approach to constructing the stacked samples that balances the panel in event time, e, from e=-4 to e=0. Column (1) reports the estimate of δ^0 from equation (12) when the outcome is the inheritance amount. Columns (2), (3), and (4) report estimates of δ_1 from estimating equation (13) when the outcomes are retirement, nondurable spending, and total spending, respectively. Panel A presents estimates using the full stacked sample as detailed in the text. Panel B presents estimates using the spending stacked sample as detailed in the text. In addition to the difference-in-differences estimates, the rows within each panel present the dependent variable means for the treatment group in the omitted survey wave (two waves before the inheritance event), and the number of households, individuals, and observations that contribute to each regression. Standard errors are clustered at the household level.

^{*} p < 0.10, ** p < 0.05, *** p < 0.01