

NBER WORKING PAPER SERIES

A STRUCTURAL ANALYSIS OF THE EFFECTS OF THE GREAT RECESSION
ON RETIREMENT AND WORKING LONGER BY MEMBERS OF TWO-EARNER
HOUSEHOLDS

Alan L. Gustman
Thomas L. Steinmeier
Nahid Tabatabai

Working Paper 22984
<http://www.nber.org/papers/w22984>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2016

We gratefully acknowledge support for this research from a grant provided by the Sloan Foundation to Dartmouth College. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Alan L. Gustman, Thomas L. Steinmeier, and Nahid Tabatabai. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Structural Analysis of the Effects of the Great Recession on Retirement and Working Longer
by Members of Two-Earner Households

Alan L. Gustman, Thomas L. Steinmeier, and Nahid Tabatabai

NBER Working Paper No. 22984

December 2016

JEL No. C61,D31,D91,E21,E24,E32,H55,I3,J11,J14,J16,J32,J63,J64,J82

ABSTRACT

This paper uses data from the Health and Retirement Study to estimate a structural model of household retirement and saving. It applies that model to analyze the effects of the Great Recession on the work and retirement of older couples who were both employed full-time at the beginning of the recession. We analyze the effects of job loss, changes in wealth and changes in expectations.

The largest overall effects of the Great Recession are observed for 2009 and 2010. In 2009, an additional 2.5 percent of all 55 to 59 year old husbands were not working full-time as result of the Great Recession, amounting to a reduction of 3.2 percent in full-time work. In 2010, 2.8 percent of 55 to 59 year old husbands were not working full-time as a result of the Great Recession, amounting to a 3.8 percent reduction in full-time work. For wives the reductions in full-time work due to the Great Recession were 1.7 percent and 2.2 percent of those who initially held a job, or reductions of full-time work of 2.3 and 3.0 percent respectively. For those 60 to 64, the reductions were 1.2 percent of men and 0.9 percent of women. Having been laid off in the last three years reduces full-time work by 30 percent. There also are lingering effects of layoff on the probability of working longer. Having been laid off three or more years in the past reduces full-time employment in the current year by about 12 percent. This reflects the reduced work incentives for full-time work arising from lower earnings due to the loss of job tenure with a layoff as well as the additional earnings penalty from a layoff.

The effect on own work of a spouse having been laid off is much smaller. The reason is that, as found in the estimation of our structural model, having one spouse not working increases the value of leisure for the other. In contrast, when one member of the household loses their job, the value of consumption increases relative to leisure. For recent layoffs, these effects are roughly offsetting.

All told, the effects of the Great Recession on retirement seem relatively modest. These findings are consistent with our earlier descriptive analyses.

Alan L. Gustman
Department of Economics
Dartmouth College
Hanover, NH 03755-3514
and NBER
ALAN.L.GUSTMAN@DARTMOUTH.EDU

Nahid Tabatabai
Department of Economics
Dartmouth College
Hanover, N.H. 03755
Nahid.Tabatabai@dartmouth.edu

Thomas L. Steinmeier
Department of Economics
Texas Tech University
Lubbock, TX 79409
thomas.steinmeier@ttu.edu

This paper analyzes the effects of the Great Recession on the work and retirement of older working couples who were both employed full-time at the beginning of the recession. We focus first on older workers who lost their jobs. To understand the effects of layoff on retirement, we analyze how the loss of earnings from current employment, the reduction in the likelihood of reemployment, and the impact of job separation on the wage offer in future employment, each affect the likelihood of retirement. Our analysis distinguishes between those who lost their jobs due to the Great Recession and those who would have experienced a layoff in a more stable time

We then turn to the many older workers who did not experience a layoff. Induced changes in their expectations of the likelihood of job and earnings loss in the future may have led members of this group to adjust their retirement and saving. Our analysis indicates the importance of this effect on working longer and retirement.

Then for all groups of older workers who were approaching retirement at the time of the Great Recession, whether they lost their job or not, we estimate the extent to which the reduction in the value of their assets induced by the recession affected their retirement outcomes. We also ask how expectations as to lower future returns to assets affected their retirement.

We begin by briefly reviewing the effects of the Great Recession on the probability of being able to keep one's current job, and on key outcomes for those who lost their jobs, including duration of subsequent unemployment and the future wage. Next we review key equations describing the effects of the Great Recession on liquid assets. We use these results to project the likely effects of the Great Recession on employment expectations and on expectations of future returns to assets.

Next, we specify a model of retirement and saving for two-earner, couple households. The model is estimated using panel data from two cohorts from the Health and Retirement Study,

War Babies born from 1942 to 1947, and Early Boomers born from 1948 to 1953.¹ Retirement outcomes examined include the probabilities of full-time work, and of full and partial retirement for each spouse (including reversals from states of lesser to greater work), and the degree of coordination of retirement for husbands and wives.

We fully model the behavior of each spouse. In contrast, as seen in a recent survey by Blundell et al. (2016), most models of retirement consider the labor market behavior of the single individual. Our model also differs from household models of retirement surveyed by Blundell et al., which typically use a unitary model of household utility. We allow each spouse to maximize own utility, but do not constrain the spouses to agree on the optimal course of retirement and consumption for each.

State variables in the model include the level of assets at the end of the previous year, whether in the previous year the individual was still in a career job, had been laid off or was disabled, Social Security and pension entitlements in the previous year, and the state of leisure preferences. Stochastic transitions from the previous states include mortality of each spouse, the rate of return to assets, independent layoff probabilities for husbands and wives, the likelihood of disability, and a change in the assessment of the desirability of leisure after retiring. Given the preferences and the opportunity set for each spouse in the household, work, consumption and saving decisions are made, generating transitions for the state variables that are not stochastic, including being on a career job, working in partial retirement, pension entitlements, and the levels of Social Security benefits for each spouse.

¹ This model is a modified version of a model we constructed to explain retirement in two-earner households, the latest version of which is Gustman and Steinmeier (2014). Although our previous models have incorporated the effects of layoffs on wages for subsequent employment, they did not treat layoffs as uncertain. The major modification from our previous work involves the introduction of stochastic states related to layoffs and future employment prospects. For related analyses, see Gustman and Steinmeier (2000 and 2004).

Incorporating the equations describing how the Great Recession affected employment, assets, and related expectations, we use the structural model to simulate the effects of Great Recession on retirement and saving. The effects on retirement of the various changes induced by the Great Recession, in layoffs, future wage and employment prospects, the value of assets, and returns to assets in the future, are examined separately.

In the literature, the work closest to this paper is by Merkurieva (2016), who uses the HRS to explore the retirement effects of the Great Recession. There are a number of differences between her model and the model in this paper, the most significant being her assumption of homogeneity in time preference. As a result, her simulations fail to capture any of the observed retirement spike at the Social Security early entitlement age. Other differences include her omission of defined benefit pensions from the budget constraint and her use of individuals rather than couples as the unit of observation. The use of individuals as the unit of observation creates issues as to what part of household assets should be attributable to the individual, and it also negates the ability of a spouse to adjust labor supply to compensate for a layoff.

II. Layoffs

Consider first how layoffs induced by the Great Recession affected the incidence and duration of unemployment of near retirees. Table 1 gives the probability of HRS individuals being laid off by a given wave of the study, given that they had been working full-time in the previous wave. In each wave of the HRS, the study asks individuals who were employed in the previous wave if they are still employed by that employer. If they are not still employed by the previous wave's employer, the study asks the individual the reasons why. Multiple responses are permitted, and the individual is considered to have been laid off if the responses include "business closed," "laid off/let go," and/or "would have been laid off."

Table 1
Probability of Layoff from Full-Time Job by Wave

Wave	Year	Probability of Layoff	Standard Deviation	Number of Observations
2	1994	7.0%	0.4%	4823
3	1996	6.1	0.4	4177
4	1998	5.5	0.4	3181
5	2000	5.0	0.4	3702
6	2002	6.8	0.5	3026
7	2004	5.6	0.5	2124
8	2006	4.9	0.4	3126
9	2008	5.1	0.4	2587
10	2010	11.5	0.7	2190
11	2012	5.2	0.4	3221
All		6.2	0.1	32157

The table looks only at the first time the individual is laid off during the study; if the individual is laid off more than once during the study, waves after the first layoff are not counted in either the numerator or denominator of the probability. The third column gives the probability of the layoff by year, and the fourth column gives the approximate standard deviation. For instance, the first line says that between the first wave in 1992 and the second wave in 1994, about 7 percent of the respondents who were working full-time in the 1992 survey had been laid off.

In general, the layoff percentages are within a percentage point of the overall mean, which is 6.2 percent over a two-year interval. The one outlier, as might be expected, is the interval between the 2008 and 2010 waves, when the layoff percentage almost doubled and is clearly significantly above the others. This spike in layoff probability does not seem to have persisted, however. By the interval between the 2010 and 2012 waves, the percentage had dropped to the overall mean. Nor does this pattern seem to be an anomaly; the statistics from

BLS on major layoffs follow a similar pattern where the layoffs spiked shortly after the Great Recession started and then fell rapidly to roughly pre-recession levels.

Table 2
Probit for Layoff from Full-Time Job

	Marginal Effect	Absolute t-statistic
Tenure at employer	-3.94%	10.98
Tenure squared	0.62	6.53
Full-time experience	0.60	1.14
Experience squared	-0.063	0.69
Education: < High school	0.69	1.77
Some college	-0.92	1.78
College graduate	-0.61	1.54
Advanced degree	-2.51	5.61
Race: Black	-0.51	1.45
Hispanic	0.53	1.16
Married	-0.56	1.93
Fair or poor health	0.91	2.43
Industry: Manufacturing	3.54	10.21
Government	-2.98	5.68
Occupation: Mgmt/Prof	0.88	2.36
White Collar	1.14	3.34
Pension at employer	-2.81	8.39
Log of wage	0.32	1.33
Wave 10	5.65	10.43

Notes: Tenure and experience are both measured as years divided by 10. Number of observations: 32,198

Table 2 gives the results of a probit estimation of the probability of a layoff between the waves on a set of personal and job characteristics, including a binary variable for the 2008-2010 interval. The entries in this table are the marginal effects, that is, the change in the probability from a one unit increase in the corresponding explanatory variable. Thus, the first entry says that the effect of each additional year of tenure is to reduce the probability of a layoff by 0.39 percentage points. This effect of tenure is modified by the entry in the second row, which gives the effect of tenure squared. Comparing the effects of tenure and experience, it is clear that

tenure is by far the more important determinant of layoff probabilities. Having 10 years of tenure reduces the probability by around 3 percentage points, and having 20 years reduces the probability by around 5 percentage points. This compares with an overall average of 6 percentage points. The effect of experience, on the other hand, is not only insignificant and relatively small numerically, but the direction of the effect is in the opposite direction of what would be expected.

The remaining effects are more or less as expected. More education reduces the probability of layoffs, with the effect being particularly strong for those with advanced degrees. Race does not seem to have a significant effect, while marriage has a marginally negative effect. Those in manufacturing industries have a 3-4 percentage point greater risk of layoffs, and those in management, professional, and white collar occupations (as opposed to blue collar occupations) have around a 1 percentage point greater risk of layoffs. Government workers and employees with pensions in their jobs have around a 3 percentage point reduced risk of layoffs.²

The largest effect in this equation, however, is the binary variable indicating the period between wave 9 in 2008 and wave 10 in 2010, which corresponds to the initial period of the Great Recession. During this time frame, the probability of layoff rose by 5.7 percentage points on top of a base of 6.2 percentage points. Thus, the spike in layoffs evident in Table 1 carries over largely intact when covariates related to personal and job characteristics are introduced. The fitted values from this probit, less the effect from the wave 10 variable, will be used in the stochastic retirement model as the ex-ante probabilities of future layoffs.

In any stochastic model, the realized values of the stochastic variables may vary from their expected values. For instance, the draws for the realized returns on assets may be above or

² Note that because of the nonlinear form of the probit relationship, the total effect of a group of characteristics is not necessarily the sum of the marginal effects.

below the ex-ante expected values. The same is true for layoffs. The layoffs in the early years of the Great Recession were higher than expected, but the main question is whether these layoffs resulted in an upward shift in the probabilities of layoffs in future years. The transitory nature of this increase, severe though as it was, suggests that the ex-ante probabilities in future years were probably not affected. Thus, the model will treat the large but temporary increase in layoffs as positive deviations from a distribution and not as a change in the distribution applicable for the future. This is in contrast to the assumptions about asset returns in the next chapter, where the long string of negative real short-term interest rates has probably lowered permanently the perceptions of what the distribution of those interest rates looks like.

The next item in the sequence is the question of how long, if an individual is laid off from a full-time job, it takes to find another full-time job. Fortunately, the HRS has a sequence of questions which attempts to ascertain, month by month, what an individual's employment status is subsequent to a layoff. If an individual is not in the previous job, the study asks the month and year when the individual left that job. If the individual is currently working, the study asks the starting month and year of the current job and the number of hours in the current job. Unless the ending date of the previous job coincides with the starting data of the current job, the study inquires about which months, if any, the individual was working in intermediate jobs, as well as the hours in those jobs. When the information is complete, it is possible to construct a measure of the number of months between a layoff and the start of another full-time job, even if the next full-time job begins in a wave after the initial layoff.

Table 3 gives the distribution of the length of the spells before the next full-time employment by the year of the layoff. The column labeled "censored" refers to cases where no full-time employment was observed in the study after the layoff. This can occur either because the individual left the survey before any full-time employment was recorded, or because the

individual effectively retired and never re-entered full-time employment. It can also occur in cases where the individual had gaps in the data and had no full-time employment before such a gap.

Table 3
Duration of Spells Between Full-Time Jobs After Layoffs

Year of Layoff	Months Until Next Full-Time Job					Censored
	0-1	2-6	7-18	19-30	31+	
1992	19	5	11	7	2	18
1993	52	29	36	8	8	61
1994	43	38	25	13	10	59
1995	31	19	20	3	6	44
1996	35	12	13	8	5	47
1997	32	21	14	4	4	34
1998	29	12	6	2	4	45
1999	26	18	18	4	8	35
2000	35	6	7	3	1	42
2001	22	10	16	2	4	38
2002	22	13	15	5	1	45
2003	16	5	8	1	3	38
2004	19	8	6	4	4	31
2005	24	21	12	1	2	33
2006	22	17	11	3	4	24
2007	13	17	12	1	2	41
2008	32	11	10	5	3	59
2009	20	10	18	9	1	66
2010	17	13	8	3	1	57
2011	31	26	7	0	0	52
2012	17	13	2	0	0	56
All	557	324	275	86	73	925

Total number of observed layoffs: 2240

Table 4 gives the distribution of the lengths of time observed after a layoff for those who did not enter full-time employment during the time they were continuously observed. This table gives the overwhelming impression that the primary cause of censored observations is that individuals never re-entered full-time employment, not that the period of observation ended before a full-time job was observed. For instance, for layoffs in 2002, 28 of the 45 cases of layoffs that year in which subsequent full-employment was not observed had at least 31 months of observation after the layoff. From Table 3, the cases where full-time work was observed after the layoff were strongly skewed toward relatively short periods, and one would expect only a few cases where individuals with censored spells of 31 months or more would eventually be employed full-time. Only for layoffs after 2010 does it appear that the last observed wave of the study may have cut short the period of observation before individuals found other full-time jobs.

The distribution of observed spells from Table 3 suggests a distribution that is strongly declining over time, such as a Weibull distribution. Normally, an estimation of the Weibull distribution takes into account both the distribution of completed spells and the distribution of spells which are censored before completion. The number and lengths of incomplete spells in Table 4, however, appears to be much greater than would be expected from a Weibull distribution, or any distribution with a rapidly falling tail, consistent with Table 3. These data suggest that the distribution of spells may perhaps best be described by some kind of a compound distribution. In the first stage, let the probability that the individual will effectively retire after the layoff be p , with $1 - p$ being the probability that the individual will eventually return to full-time work. For those who return to full-time work, let $f(x)$ be the distribution of times until the return. A censored observation, that is, a layoff for which no subsequent return to full-time work is observed, can occur either because the individual has effectively retired or because the subsequent period of observation is too short to capture the return to work.

Table 4
Duration of Observed Spells Without Full-Time Work
After a Layoff for Censored Observations

Year of Layoff	Duration of Observed Spell					Total
	0-1	2-6	7-18	19-30	31+	
1992	9	1	0	3	5	18
1993	14	0	8	2	37	61
1994	6	8	3	8	34	59
1995	5	0	6	4	29	44
1996	4	6	2	6	29	47
1997	10	0	6	2	16	34
1998	7	1	2	3	32	45
1999	4	1	4	1	25	35
2000	7	3	3	2	27	42
2001	6	0	4	1	27	38
2002	2	3	4	8	28	45
2003	2	2	9	0	25	38
2004	6	1	2	0	22	31
2005	6	1	1	1	24	33
2006	1	1	1	5	16	24
2007	6	0	4	1	30	41
2008	6	2	4	6	41	59
2009	3	0	9	2	52	66
2010	5	3	4	35	10	57
2011	8	1	36	7	0	52
2012	4	35	17	0	0	56

To make things a little more specific, Let $f(x)$ be a Weibull distribution, with a cumulative distribution $F(x) = 1 - e^{-(x/\lambda)^k}$, where x is the length of the time between the layoff and the next full-time job. In a Weibull distribution, the λ parameter is called the scale parameter and governs how stretched the distribution is along the horizontal axis. The k parameter is called the shape parameter and, as its name implies, governs the shape of the distribution. For values between 0 and 1, the Weibull is a continuously declining distribution.

The closer the value is to 0, the sharper the distribution falls. Both the λ and the k parameters must be positive values.

For observations for which a complete spell from the layoff to the beginning of the next full-time job is observed, the log-likelihood value is given by

$$(1 - p) f(x) = (1 - p)(k/\lambda)(x/\lambda)^{k-1} e^{-(x/\lambda)^k},$$

which is the probability that the layoff does not result in an effective retirement, times the conditional probability density for the spell, conditional on a completed spell. For observations for which the layoff does not result in subsequent full-time employment, the log-likelihood value is given by

$$p + (1 - p)F(x) = p + (1 - p)[1 - e^{-(x/\lambda)^k}].$$

This is the probability of retirement plus the probability of an incomplete spell, conditional on the individual returning to full-time work. The overall likelihood function is the sum of the likelihood values for these two types of observations.

The likelihood values are thus functions of parameters p , λ , and k . But these parameters may in turn be functions of explanatory variables. Let $p = N(X_p\beta_p)$, so that the probability of retirement is a probit relationship dependent on a group of explanatory variables X_p . This guarantees that the probability lies between 0 and 1. Similarly, let $\lambda = \exp(X_\lambda\beta_\lambda)$ and $k = \exp(X_k\beta_k)$, where the exponential functional form guarantees that both λ and k will be greater than 0, as required for the Weibull distribution. The likelihood function is then maximized with respect to β_p , β_λ , and β_k to obtain estimates.

Table 5
Estimates of Distribution Function for Spells from Layoff until Full-Time Work

	β_p	t	β_λ	t	β_k	t
Constant	0.493	2.09	2.419	6.47	-0.607	3.21
Tenure at Employer	0.494	4.51	0.233	1.58	0.032	0.39
Tenure Squared	-0.061	2.16	-0.053	1.27	-0.007	0.29
Full-Time Experience	-0.743	5.41	-0.407	1.88	0.113	1.02
Experience Squared	0.145	6.20	0.055	1.49	-0.017	0.84
Education:						
Less than High School	-0.065	0.63	-0.074	0.55	0.052	0.74
Some College	0.006	0.04	0.168	0.80	-0.088	0.82
College Graduate	0.020	0.17	0.129	0.92	0.014	0.18
Advanced Degree	0.203	1.30	-0.256	1.53	0.173	1.71
Race:						
Black	-0.041	0.37	0.369	2.54	0.058	0.74
Hispanic	0.150	1.25	0.274	1.74	0.119	1.31
Married	-0.125	1.59	0.060	0.61	-0.064	1.17
Fair or Poor Health	0.159	1.66	0.042	0.32	0.084	1.14
Industry:						
Manufacturing	0.082	0.99	0.179	1.75	0.115	1.96
Government	0.682	2.93	0.467	1.15	-0.077	0.38
Occupation						
Management/Prof	-0.146	1.42	0.050	0.39	0.036	0.54
White Collar	0.113	1.26	0.187	1.58	0.072	1.14
Pension in Job	-0.112	1.35	-0.023	0.22	0.005	0.10
Log of wage	-0.219	3.21	-0.022	0.25	0.020	0.41
Trend	0.292	4.60	-0.312	3.49	0.129	2.78
Square of Trend	-0.084	0.67	-0.299	1.83	0.207	2.29
Great Recession			0.555	2.85		

Notes: See text for an explanation of the three sets of columns. Tenure and experience are both measured as years divided by 10. Trend is (year of layoff – 2002) / 10 and ranges from -1 (1992) to 1 (2012). Great Recession is a binary variable with a value of 1 for 2009 and 2010. The columns labeled t are absolute values of asymptotic t-statistics. Number of observations: 1940.

The results of the estimation are presented in Table 5. A quadratic trend term is included in each of these linear forms to capture any trends in retirement and/or changes in the distribution of spells over time. Most of these trend terms are significant. About the only other terms which reach 95% confidence in the Weibull distribution are the coefficient for black race in the linear form for λ and the coefficient for manufacturing in the linear form for k . A term for layoffs in the 2009-2010 time frame is included in the linear form for λ and is significant at better than 99% confidence. Its coefficient indicates that spells between layoffs and successive full-time work lengthened considerably over this period. Similar coefficients for the other two linear forms were completely insignificant and were dropped from the specification.

For a typical individual with average characteristics in the sample, these coefficients indicate that, under normal circumstances, there is a 45% chance that a spell after a layoff before the next full-time job will last less than 3 months, a 16% chance that the spell will last between 3 and 6 months, and a 17% chance that it will last between 6 and 12 months.³ When the Great Recession variable is factored in, these percentages change to a 34% chance that the spell will last for less than 3 months, a 14% chance it will last between 3 and 6 months, and a 17% chance that it will last between 6 and 12 months. The median duration is normally 3.7 months but rose to 6.5 months for layoffs during the Great Recession.

The question arises whether this increased duration of nonworking spells after a layoff could be perceived as permanent, or whether it was a transitory increase. To answer this question, a binary variable indicating a layoff in the 2011-2012 period was added to the specification in Table 5. Measuring the effects on duration of layoffs during this period is a little tenuous, since the last wave was in 2012 and only relatively short durations would have been

³ This individual has 11 years of tenure and 32 years of experience, is white with a high school education, is a blue-collar worker in a non-manufacturing industry, had a pension, and had an hourly wage of \$10.18 in 1992 dollars.

completed. Nonetheless, the coefficient of the 2011-2012 layoff variable is of the wrong sign and gives no indication that the duration effects of the Great Recession were long lasting. As with the probability of a layoff, it appears that the increase in duration of spells before the next full-time job was a transitory phenomenon and not a permanent shift in the distribution of durations.

Table 6
Fixed Effects Indexed Full-Time Wage Equation Coefficients

	Coefficient	t	Coefficient	t
Constant	2.015	97.03	2.014	96.04
Tenure at Employer	0.123	13.12	0.123	12.95
Tenure Squared	-0.009	3.52	-0.009	3.41
Full-Time Experience	0.113	8.90	0.112	8.79
Experience Squared	-0.021	9.50	-0.021	9.41
Married	0.033	2.63	0.035	2.78
Fair or Poor Health	-0.022	2.43	-0.022	2.32
Self-Employed	-0.120	9.95	-0.118	9.75
Industry:				
Manufacturing	0.103	9.87	0.103	9.72
Government	0.052	3.13	0.050	2.97
Occupation:				
Management/Prof	0.120	11.78	0.118	11.51
White Collar	-0.004	0.39	-0.005	0.51
Pension in Job	0.141	19.41	0.142	19.39
Prior Layoff	-0.073	7.52	-0.075	7.33
Layoff Between Waves 9 and 10			0.019	0.37
Number of Observations	51,924		51,343	
Number of Individuals	15,619		15,568	

Notes: Dependent variable is the log of the full-time wage rate. Tenure and experience are both measured as years divided by 10. The columns labeled t are absolute values of asymptotic t-statistics.

Thus, the Great Recession both increased the probability of being laid off and increased the length of time it took laid-off workers to find other full-time employment. The next issue is the nature of the subsequent job, and in particular the wage rate in the subsequent job. Table 6 contains coefficients from fixed effects regressions of wages taken from the HRS. These wages include wages in full-time jobs held in any of the study wave years plus wages in full-time jobs in the job histories collected in the first interview. All wages are indexed to 1992 (the first year of the HRS) using the average weekly earnings series reported in various Economic Reports of the President. The use of fixed effects nets out any differences due to different personal characteristics and only uses information on the wage changes over time for the same individual.

If an individual gets laid off, the wage in the subsequent job should be less, reflecting the loss of tenure. The question of interest here is whether there is an additional effect from having been laid off, relative to leaving a job for other reasons, and furthermore whether this effect was larger for layoffs which happened during the Great Recession. In the table, the first two columns include a term related to whether the individual has ever been laid off in a prior job. The coefficient, which is highly significant, suggests that having been laid off in a prior job results in a 7.3 percent reduction on top of the reduction from the loss of tenure. The last two columns include an additional binary variable which has a value of 1 if the layoff occurred between wave 9 (2008) and wave 10 (2010) of the study, which includes the first two years of the Great Recession. This term is slightly positive, which is the opposite of what might be expected, but it is not significant. The conclusion is that while the Great Recession resulted in a greater probability of being laid off and a longer wait until the next full-time job, it does not appear to have depressed wages in subsequent jobs below the levels which would have occurred anyway with a layoff.

III. Rate of Return

In this section we consider how the Great Recession affected the assets working individuals had accumulated for retirement. Since different classes of financial assets experienced widely varying results during the Great Recession, the first step in this analysis is to determine how much of each type of asset was held by the HRS respondents.

We begin by looking at the composition of financial assets for the typical individual who had a nontrivial level of financial assets. For this exercise, we consider total financial assets above \$5,000 to be nontrivial, where financial assets are considered to be stocks, bonds, checking accounts, CD's, and IRA's. We array households with over \$5,000 in total financial assets by those assets and look at those between the 45th to 55th percentiles. These percentiles correspond to total financial assets of between \$76,575 and \$116,940, respectively. For households within this range of total financial assets, the average level by type of asset is as follows:

Stocks	\$ 16,828
Bonds	548
Checking Accounts	28,100
CD's	11,002
IRA's	38,577
Total Financial Assets	95,054

These figures are from the 2006 wave of the HRS, before any of the effects of the Great Recession were realized. They include the original HRS cohort, the War Babies, and the Early Baby Boomers, but not the original AHEAD nor the CODA cohorts.

The importance of stock holdings in total financial assets is somewhat unclear from this breakdown because IRA's are the largest component of financial assets, and stocks may figure significantly in IRA's. The HRS did inquire about the percentage of stocks in IRA's. For this group, among the individuals who had IRA's, 199 reported that they held stocks in their most

important IRA, 103 reported that they did not hold stocks in that IRA, 39 didn't know and 2 refused. Thus, among the group that gave an answer, 66 percent said that they did have stocks in the IRA. The next question is, among those who did have stocks in their IRA, what percentage of the IRA was held in stocks. Among the 199 who reported that they held stocks in the IRA, 147 were able to give an approximate percentage of stocks in the IRA, with an average percentage of 84 percent.

Taking these figures as representative for the group as a whole, if 66 percent had stocks in their IRA's and if 84 percent of those IRA's were invested in stocks, then approximately 55 percent of IRA's as a whole were invested in stocks. Since these individuals had an average of 38,577 in IRA's, 55 percent of that amount would be \$21,217. Combining this with the amount of stock held directly yields a figure of \$38,045 invested in stocks either directly or through an IRA. As a percentage of total financial assets, then, stocks represent 40 percent of the total.

It is clear from these figures that for the typical HRS individual, bond holdings are trivial. The remaining two categories are checking accounts and CD's. Both of these are money-like in the sense that their principal is very stable, but they have relatively low rates of expected returns. The HRS does not ask about the remaining 45 percent of assets in IRA's, but it seems likely that it is also invested in money-like instruments. Thus, a typical individual in the HRS can be characterized as having 40% percent of his or her financial assets invested in stocks and the remainder in money-like instruments.

Unfortunately, it appears that even among individuals in the middle decile of the asset distribution, very few have an average composition of assets. The first indication of this was from the information on IRA's, where 34 percent reported no stock holdings and the remaining 66 percent reported 84 percent of their IRA's were held in stock. In fact, among the 147 who gave a percentage of stocks in their IRA's, 102 reported that 100 percent of their IRA's were

held in stock. Thus, the large majority of IRA holders seem to fall into one of two camps: either they do not hold any stock in the IRA or they hold their IRA completely in stock. Almost no one holds the average percent in their IRA.

This same behavior seems to carry over to holdings of non-retirement assets. Arraying these same middle-decile households by the percentage of stock holdings and the percentage of money-like holdings yields the following results:

Percentile	Percentage in Stocks	Percentage in Checking & CD's
10	0	9
20	0	39
30	0	77
40	0	100
50	0	100
60	0	100
70	20	100
80	59	100
90	90	100

These percentages are percentages of non-IRA financial assets. Thus, more than 60 percent of these households do not have any stocks aside from whatever they have in their IRA's, while 10 percent have more than 90 percent of their non-IRA financial assets in stocks. On the other hand, at least 60 percent have 100 percent of their non-IRA financial assets in checking accounts and CD's. Thus, while on average 30 percent of the non-IRA financial assets of this group are held in stocks, most individuals in this group in fact do not have any stock holdings outside of their IRA's. At the other end of the spectrum, there are a nontrivial number of individuals who have most of their non-IRA holdings in stock.

The relative composition of financial assets also appears to vary according to the total value of financial assets. Table 7 looks at the composition of financial assets by asset decile among the HRS households. This table includes households regardless of the level of financial assets. The first decile in fact has no financial assets of any type. The next three deciles have

total financial assets less than \$5,000 and hence would have been excluded from the preceding discussion. The financial assets that this group has are overwhelmingly checking accounts, with very little in the remaining categories of assets. As you move up the asset distribution, there is a declining percentage in money-like accounts (checking and CD's) and an increasing percentage in stocks and IRA's. Except for the top asset decile, bonds never account for more than a trivial percentage of assets, and even in the top decile the fraction of assets in bonds is only about 6 percent.

Table 7
Composition of Assets by Asset Decile

Asset Decile	Asset Range	Stocks	Bonds	Checking	CD's	IRA's	Total Assets
1	0 - 0	0	0	0	0	0	0
2	0 - 200	0	0	27	0	0	27
3	200 - 1,200	8	1	610	5	8	632
4	1,200 - 5,000	93	0	2,627	94	125	2,939
5	5,000 - 15,100	478	10	6,432	594	1,513	9,027
6	15,100 - 41,000	3,174	179	12,966	3,003	7,948	27,271
7	41,000 - 91,000	9,137	578	21,932	7,036	24,524	63,206
8	91,000 - 185,000	24,272	1,569	33,428	17,488	55,309	132,066
9	185,000 - 411,500	62,912	4,818	51,981	28,793	127,787	276,292
10	411,500+	516,449	69,942	97,223	62,881	449,355	1,195,850

Since stocks and money-like instruments have substantially different distributions of returns, and since there is a wide dispersion of the fraction of assets held in these two forms, it seems like it would be a good idea to treat each household as having its own distribution of rates of return, depending on the composition of that household's assets. Computationally, treating each household as having its own distribution of rates of return does not impose too much of an additional burden. Imputations of the amounts of wealth in stocks, bonds, checking accounts, CD's and IRA's are available in the RAND files for each wave through 2010. Rand does not impute the percentage of stocks in IRA's, but relatively few respondents don't know whether the

IRA has stocks, and for those who do have stocks, most of the IRA appears to be invested in stocks. These factors mean that reasonable imputations can be made for the percentage of stocks in total financial assets, even if some of the answers relating to stocks in IRA's are missing.

For those who do not have any financial assets, there is no information about the composition, and hence about returns, to potential assets. This is not a problem for two reasons. First, given that almost all of the assets in the second through fourth deciles of the wealth distribution are in checking accounts, it is probably not unreasonable to suppose that on the margin, any potential assets of those who have none would be invested in checking accounts. Secondly, individuals who have no assets are treated as having a sufficiently high time preference rate that they consume all their resources, and for such individuals the rate of return on potential assets is irrelevant. A related problem is that the composition of assets, and hence the rate of return on those assets, varies with the level of assets and is probably in some sense endogenous. However, previous attempts to explain the composition of assets have not been very successful, and in any case treating the rate of asset composition of a particular individual as fixed for small deviations in assets is probably approximately correct.

Given that each individual has a specific percentage of assets in stocks, money-like accounts, and (for high asset individuals) in bonds, we can calculate both the ex-ante and ex-post returns using data from Ibbotson Associates. Ibbotson reports the annual total returns for six classes of assets: large company stocks, small company stocks, corporate bonds, government bonds, government notes, and government bills. They also report an annual inflation percentage. These figures span from 1926 through 2012. For the returns to stocks, we take a weighted average of 75 percent large company stocks and 25 percent small company stocks. Their "large companies" correspond approximately to the S&P 500 companies, and the valuation of those companies comprise approximately three-quarters of the total stock market valuation. For

money-like assets, we use the Ibbotson returns on treasury bills, since these are most like the checking account and CD accounts in the HRS. For the small percentage of assets in bonds, we will use the Ibbotson returns on government bonds.

Table 8
Total Real Returns for Stocks and Treasury Bills

Stock Returns							
1973	-27.5%	1983	23.0	1993	10.1	2003	34.8
1974	-37.0	1984	-0.9	1994	-0.9	2004	9.5
1975	34.1	1985	26.2	1995	34.3	2005	1.7
1976	27.5	1986	14.6	1996	18.3	2006	13.3
1977	-5.8	1987	-2.8	1997	29.0	2007	-1.3
1978	1.8	1988	13.8	1998	18.0	2008	-37.0
1979	11.5	1989	21.7	1999	20.5	2009	24.1
1980	21.9	1990	-13.8	2000	-11.1	2010	17.6
1981	-9.2	1991	31.0	2001	-4.8	2011	-2.2
1982	19.3	1992	8.7	2002	-22.3	2012	14.8
			1926-2012		1947-2012		
	Arithmetic Average		9.9		9.5		
	Geometric Average		7.5		7.9		
Treasury Bills							
1973	-1.9	1983	5.0	1993	0.2	2003	-0.9
1974	-4.2	1984	5.9	1994	1.2	2004	-2.1
1975	-1.2	1985	4.0	1995	3.1	2005	-0.4
1976	0.3	1986	5.0	1996	1.9	2006	2.3
1977	-1.7	1987	1.1	1997	3.6	2007	0.6
1978	-1.9	1988	1.9	1998	3.3	2008	1.5
1979	-2.9	1989	3.7	1999	2.0	2009	-2.6
1980	-1.2	1990	1.7	2000	2.5	2010	-1.4
1981	5.8	1991	2.5	2001	2.3	2011	-2.9
1982	6.7	1992	0.6	2002	-0.7	2012	-1.7
			1926-2012		1947-2012		
	Arithmetic Average		0.5		0.7		
	Geometric Average		0.4		0.7		

Table 8 gives the annual real (inflation adjusted) rates of return for stocks and treasury bills for the previous 40 years. The complete series is available as far back as 1926. Ex-post rates for any particular household can be formed by taking, year by year, a weighted average of these rates, where the weights are the percentages of stocks and money-like instruments in the household's portfolio. Ex-ante distributions can be calculated by treating these annual rates as coming from a log-normal distribution and calculating the mean and variance of that distribution. The more stocks and fewer money-like instruments in a household's portfolio, the greater will be the mean and the greater will be the variance of the expected ex-ante returns.

There is some question as to the proper period to calculate the ex-ante distributions, but the data in Table 8 suggest that it probably doesn't matter very much. For both stocks and treasury bills, the table reports the averages over the entire 1926-2012 period and for the post-war period from 1947 to 2012. The two averages are in fact relatively close, suggesting that dropping the depression and wartime periods doesn't change the results very much. The table reports both the arithmetic and geometric means, although the geometric mean is the correct measure of a long-term compounded rate, as well as the correct measure for a lognormal distribution.

The final question is whether the ex-ante expectations of returns have changed after the Great Recession. For stocks, the answer is probably no. After the initial decline in stocks in 2008, the average returns in the four year period 2009-2012 are if anything above the long-term averages. Thus, the impact of the Great Recession on stocks looks like a transitory decline in returns which didn't affect the expected future returns. This view is consistent with the long run behavior over time of stocks, for which the year to year correlation in returns is, if anything, slightly negative. That is, low returns in the current year do not seem to predict either low or high returns next year.

For treasury bills, one could make more of an argument that the Great Recession did change expectations. The Federal Reserve quickly lowered the nominal treasury bill rate to effectively zero and has left it there for the past six years. Combined with a small positive inflation rate, this has resulted in a negative real rate on treasury bills of around 2.1 percent. Unlike the effects on layoffs and durations of spells until subsequent employment, the negative real interest rates appear to be more than transitory, though perhaps not permanent. One approach would be to characterize the effects of the Great Recession on interest rate expectations as lowering them for an extended period, perhaps as long as ten years. For those whose assets were primarily money-like instruments, such a lowered expectation would diminish the prospective growth of their retirement assets and would create a need to either save more or delay retirement.

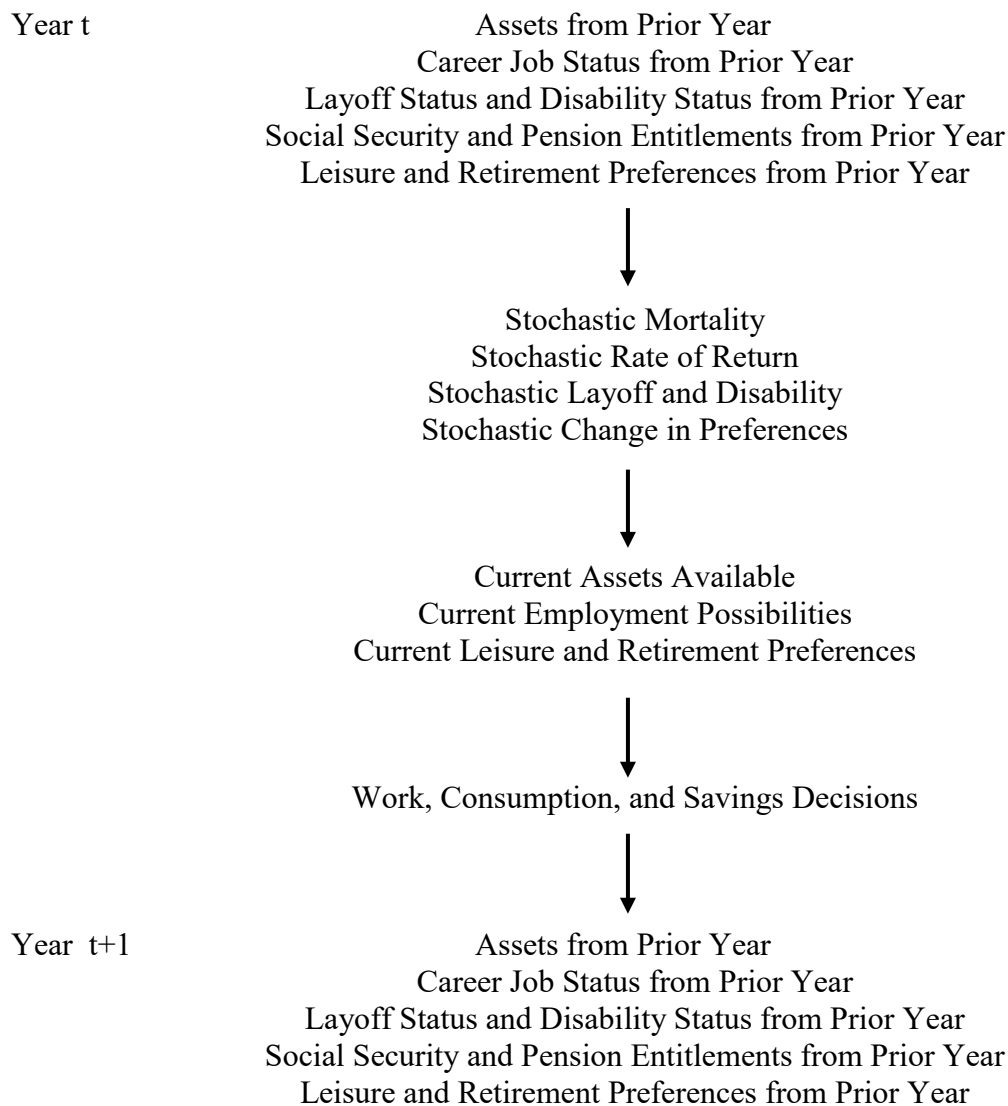
IV. Sequence of Events in the Model

In order to have a functioning model of retirement, it is necessary to specify the sequence of decisions and stochastic events. The overall sequence of these events is depicted in Figure 1. In any particular year, the sequence begins with the state variables determined by decisions and stochastic events in prior years. The state variables include the level of assets at the end of the previous year, whether or not the individual was still in the career job in the previous year, whether the individual had been laid off or disabled in the previous year or prior years, the level of Social Security and pension entitlements at the end of the previous year, and the strength of leisure preferences at the end of the previous year. The pension entitlements may include both defined benefit and/or defined contribution components.

These state variables undergo a set of stochastic changes between the previous year and the current year. Mortality may affect one or both spouses, with consequences for the Social Security and pension amounts available to the surviving spouse. A stochastic rate of return, as

estimated in the previous section, will change the level of assets available during the current period as well as the level of any defined contribution amounts. A stochastic layoff, as estimated in Section II, will affect the individual's ability to work full-time. If a layoff occurs and the individual chooses to work full-time, the individual will experience a period of non-employment equal to the mean of the layoff duration equations of Section II. Stochastic disabilities may also occur, after which the individual is unable to work. Finally, for those who have retired, the preference for leisure may change as the individual actually experiences retirement.

Figure 1
Sequence of Events



The combination of the levels of the state variables in the previous period in conjunction with the stochastic changes described in the previous paragraph defines the possibilities available to the individual in the current period. These include the assets available in the current period, the employment possibilities available in the current period, and the individual's current preferences for leisure and retirement. If the individual has previously left the career job, full-time earnings reflect the loss of tenure. If the individual has previously been laid off, earnings in full-time work are adjusted to reflect an additional penalty as estimated in the wage equations of Section II. If the individual is currently being laid off, the opportunities for full-time work reflect the fact that there will be a period of non-employment as described in the previous paragraph. And if the individual has previously been disabled or becomes disabled in the current period, the only employment choice available is retirement.

Given these possibilities defined by the state variables of the previous period and the stochastic outcomes, the individual chooses the amount of work (if the individual has not become disabled), the level of consumption, and the amount of savings. The amount of work is limited to three categories: full-time, part-time, or complete retirement. In order to reduce the computational burden, individuals are presumed to work full-time before age 50 and to retire at age 70.

These decisions in turn affect the levels of the state variables at the end of the period. End of period assets are computed from the assets of the previous period augmented by the stochastic rate of return, increased by the earnings of the work decision and any Social Security and/or pension benefits available, and reduced by consumption. Pension and Social Security benefits are assumed to be taken as soon as they are available. Defined contribution amounts are assumed to be available when the individual leaves the career job. If the individual continues to work in the career job, defined benefit entitlements are increased accordingly and defined contribution

amounts are added to the defined contribution account. If the individual works either full-time or part-time, Social Security entitlements are adjusted accordingly.

Relative to our previous family retirement model, the computational implications of this approach are significant. The previous model had one state variable to indicate whether or not the individual was still in the career job, had left the career job, or was deceased. In the current model that state variable now indicates whether the individual is still in the career job, has left the career job without a layoff, has left the career job with a layoff, or is deceased. That means that this state variable has increased from three categories to four categories, an increase of 33% in the state space. But a two-earner couple has two variables like this, and the total compounded increase in the state space is 78% ($1.33 \times 1.33 - 1 = 0.78$). The total number of calculations is approximately equal to the size of the state space, so this means a not quite doubling of the computational burden.

An additional increase in computation comes from treating SSDI as stochastic. To allow for SSDI, we expand the state variable describing leisure and retirement preferences to include one additional category. In general, the original variable is computed for seven evenly spaced values from three standard deviations below to three standard deviations above the mean value, with intermediate values being interpolated. To reflect SSDI, an eighth value is added with a value sufficiently high that the individual would always choose to be retired. In addition, for this value, an amount approximating SSDI payments is added to the income available. Thus, including SSDI raises the number of categories for this state variable from 7 to 8, a 14 percent increase. Since this applies to each spouse, the compounded effect is the raise the number of states by approximately 31 percent.

A final complication is the possibility of expectations that change due to the Great Recession. As detailed in Section II, the increased layoff probabilities and the increased time

spent between a layoff and subsequent full-time employment appear to have persisted for only a couple of years, and hence it may be best to view these as bad random draws from the corresponding distributions, but not as a permanent change in those distributions. However, the change in rates of return, and in particular the return to bonds, money market funds, savings accounts and CD's has persisted for a much longer period of time, and it may be reasonable to suppose that expectations regarding interest rates have moved lower. To incorporate a change in expectations in the model, it is necessary to do the backwards calculations twice, one for the changed expectations (which govern the decisions after the change) and one for the original expectations (which govern the decisions before the change). The first set of calculations only needs to be done back to the point where the expectations change, but this still involves a substantial increase in computational effort. One bright spot here is that this increase does not compound for two-earner households, in contrast to the increases having to do with introducing layoffs.

V. Model Specification and Estimation

In each period (a period is a year in this version of the model), there are two sets of variables. One set of variables includes the state variables defining the possibilities for the period. These contain all the information from previous periods which is relevant to the current period. In this model, the state variables include: the assets at the beginning of the period, whether or not the individual was still in the career job at the end of the previous period; whether the individual has ever been laid off or entered the Social Security disability program; the level of Social Security and pension entitlements as a result of previous work; and the level of leisure preferences. The other set of variables includes the decision variables, including the choice of the level of work effort by each spouse and the amount of savings and consumption.

Both the husband and wife have value functions that can be defined recursively as

$$V^h(S_t) = u^h(S_t, d_t) + e^{-\delta} E[V^h(S_{t+1} | S_t, d_t)]$$

for the husband and

$$V^w(S_t) = u^w(S_t, d_t) + e^{-\delta} E[V^w(S_{t+1} | S_t, d_t)]$$

for the wife. S_t is the set of state variables at time t , d_t is the set of decision variables at time t , and δ is the time preference rate. Essentially, this says that for each spouse, the current period value function is the sum of current utility plus the discounted expected value of the next period value function.

A. The Decision Process.

For the husband, current period utility is a function of current consumption C_t , his amount of leisure L_t^h , and the wife's amount of leisure L_t^w according to

$$u^h = \frac{1}{\alpha} C_t^\alpha + e^{X_t^h \beta_h + \varepsilon_h} u_L^h(L_t^h, L_t^w),$$

and similarly for the wife. The amount of leisure can take on three values corresponding to full-time work, partial retirement, and full retirement. The utility value of leisure of each spouse is allowed to depend on the amount of leisure of the other spouse in the same period. This allows for the possibility that retirement may be more valuable if the spouse is also retired. The exponential form preceding u_L^h is a multiplicative factor for the value of leisure. It consists of a standard linear form $X^h \beta^h$ plus an individual effect ε which reflects the strength of the husband's preferences for retirement over work. The elements of X^h include a constant, age, and health status, with the corresponding parameters β_0^h , β_1^h , and β_2^h . As age increases, work gradually becomes more onerous and retirement more desirable.

If the wife is working full-time, the basic value of leisure u_L^h for the husband is normalized to zero if he is working full-time, unity if he is fully retired, and u_p^h if he is partially retired. We assume the utility function is concave in leisure. If partial retirement is equated with

approximately half-time work, the value of u_p^h should fall between one-half and one. The closer the value is to one, the greater the value of partial retirement is relative to full retirement, and the more frequently and longer should be the spells of partial retirement.

If the wife is fully or partially retired, the value of retirement for the husband is increased if he prefers to spend time with his wife. More specifically, the value of his retirement is increased by a factor $\eta = e^{\beta_3^h}$, but only to the extent that her retirement overlaps with his. If the husband is partially retired, it doesn't matter whether the wife is fully retired or partially retired, since the additional leisure of the wife if she is fully retired doesn't add anything to the husband's leisure during the time he is at work. If the husband is fully retired, however, it does make a difference whether the wife is partially retired or fully retired, since only the part of his leisure that he shares with the wife is augmented. If the augmentation factor is η , the values of the function u_L^h as a function of its two arguments can be written as follows:

		Husband's Retirement Status		
		Full-Time Work	Partial Retirement	Full Retirement
Wife's Retirement Status	Full-Time Work	0	u_p^h	1
	Partial Retirement	0	ηu_p^h	$\eta u_p^h + (1 - u_p^h)$
	Full Retirement	0	ηu_p^h	η

If the husband is working full-time, the value of his leisure is zero regardless of the retirement status of his wife. If the husband is partially retired, the value of leisure is u_p^h if his wife is working full-time, and this gets multiplied by a factor of η if his wife is working no more than part-time. If the husband is fully retired, the value of leisure is unity. This again is multiplied by a factor of η if his wife is also fully retired. If his wife is partially retired, we can divide the husband's leisure into two parts. His partial retirement leisure has a value of u_p^h ,

which is multiplied by η because his wife is also partially retired. His remaining leisure has a value of $1 - u_p^h$, but this leisure does not get multiplied because the wife is not there for this leisure. The total value of the husband's leisure is the sum of these two parts. A similar situation governs the utility value of the wife's full or partial retirement.

Given a common value of α , the two spouses should agree on the amount of consumption, given initial assets and the level of work chosen by each spouse. Let V_{ij}^s be the value function of spouse s given the husband's work effort of i and the wife's work effort of j . For example, V_{PR}^h would be the value to the husband of a decision of the husband to partially retire (P) and the wife to fully retire (R). There are nine combinations of the husband's retirement and the wife's retirement, and the values of these for the husband may be arrayed as follows:

		Husband's Retirement State		
		Working	Partially Retired	Fully Retired
Wife's Retirement State	Working	V_{WW}^h	V_{PW}^h	V_{RW}^h
	Partially Retired	V_{WP}^h	V_{PP}^h	V_{RP}^h
	Fully Retired	V_{WR}^h	V_{PR}^h	V_{RR}^h

There are nine corresponding values for the value to the wife of these nine combinations. Note in particular that each valuation depends on the other spouse's choice of work effort, and these valuations are not the same for the two spouses. Also, each spouse chooses his or her own work effort, but each knows that his or her choice may influence the other spouse's choice.

The details of the choices of the two spouses are contained in Gustman and Steinmeier (2014), but the general nature of the choices can be described as follows. Each spouse initially calculates whether it would be advantageous to retire, given the other spouse's likely response.

For instance, if the wife were to retire first, and she knows that if she retired the husband would also retire, she would value her decision to retire as V_{RR}^W . If neither spouse finds it advantageous to retire, then they both will continue working. If one spouse finds that it is advantageous to retire but the other does not, then the first spouse will retire and the other spouse will optimize given the retirement decision of the first spouse. If both spouses want to retire first, then sometimes problems can arise. If the wife retires first, the husband may continue to work to provide income, and the wife would prefer that. But if the husband retires first, the wife may continue to work, also to provide income, and he would prefer that. Essentially, both spouses want to retire and have the other spouse work. If they both retire, there may not be sufficient income, however, and they would be ahead to both continue working. In this case, it is assumed that the couple coordinates enough to avoid an outcome that is undesirable to both.

This procedure is applied for years in which both spouses are between the ages of 50 and 69. For both spouses, it is assumed that they work in their career jobs at least until age 49 and are completely retired at age 70 and later. If a spouse is outside the age range of 50 to 69, there is no labor supply decision to be made by that spouse, and the other spouse makes his or her labor supply decision without having to consider the repercussions of the decision on the labor supply decision of the first spouse.

B. The Stochastic Transitions.

The stochastic transitions govern how the distribution of the state variables in year $t+1$ relate to the state variables and the decision variables in year t . The first stochastic transition is mortality, governed by the mortality tables published by the Social Security Administration. These tables give mortality rates by gender and birth cohort and are the same tables used in the Social Security projections. The households start out as a couple, and the mortality of the two spouses is taken to be independent. If one of the spouses dies, the remaining spouse makes labor

supply and consumption decisions based on his or her own value function. Any defined contribution amounts for the deceased spouse are made available to the remaining spouse, and any Social Security benefits are adjusted for the death. The consumption function is adjusted so that a given marginal utility is achieved at a consumption level that is approximately 70% as high as that of the couple. For example, if it takes \$50,000 of consumption to achieve a given level of marginal utility as a couple, it only takes \$35,000 to achieve the same level of marginal utility as a surviving spouse.

The next stochastic transition is the rate of return to assets. For each household, we observe the mix of assets between stocks and money-like assets, as discussed in Section III. This mix is allowed to be heterogeneous among households, but for any particular household it is assumed to be fixed over time at the average value observed over the waves of the survey. Using this mix, a distribution of returns is calculated for a particular household. The assets at the beginning of year $t+1$ is calculated by applying the stochastic return to the assets at the end of year t , as follows:

$$A_{t+1} = (1 + r_t) (A_t + E_{ht} + E_{wt} + P_{ht} + P_{wt} + SS_t - C_t)$$

Where A_t is assets at the beginning of year t , r_t is the stochastic rate of return, E_{ht} and E_{wt} are the labor earnings of the husband and wife, respectively, P_{ht} and P_{wt} are the pension benefits of the two spouses, SS_t is the amount of Social Security benefits received by the couple, and C_t is the level of consumption. Essentially, the amount in the last set of parentheses is the total financial resources available to the couple during year t , less the amount consumed. This is multiplied by the stochastic rate of return, and the result is the amount of assets available beginning at year $t+1$. For those in career jobs with defined contribution pensions, the same stochastic rate of return applies to accumulated defined contribution balances.

A third stochastic transition is the possibility of layoffs. The probability of these layoffs is estimated by the equations described in Section II and include the effects of both age and tenure on the probability of layoffs. Layoffs have several effects. First, if the individual is in the career job, that job ends and the tenure value of any subsequent full-time job is set equal to zero. Secondly, earnings for any subsequent full-time job is reduced by an amount estimated in the earnings equations in Section II. This results in an earnings drop of approximately 7 percent on top of the drop in earnings due to reduced tenure. Finally, if the individual has just been laid off, and if the individual chooses to continue to pursue full-time work, the individual experiences a period of non-work according to the mean duration of non-work specified by the duration equations specified in Section II. Layoff probabilities of husbands and wives are assumed to be independent of one another.

A fourth stochastic transition is the possibility of disability resulting in Social Security Disability (SSDI). An equation is estimated relating the log of the observed percentage of respondents on SSDI to age, and the probability of transitioning to SSDI is taken as the difference in the estimated percentage on SSDI in adjacent years. This is done separately for husbands and wives, and the observed percentages are the percentages from the HRS cohorts discussed above. The transition probabilities are assumed to be independent between husbands and wives. If an individual makes the transition to SSDI, he or she is assumed to be out of the labor force, and income is increased by the amount of SSDI benefits. In this model, SSDI is an absorbing state, meaning that transitions out of SSDI are not considered.

A fifth stochastic transition has to do with leisure and retirement preferences for those who are retired. Initial draws of ε_h and ε_w are made for the two spouses, indicating their initial preferences for leisure and retirement. After each spouse leaves full-time work, the values of ε_h and/or ε_w change according to the autocorrelation parameters ρ_h and ρ_w . This allows for

individuals to change their assessment of the desirability of leisure and retirement after they have started to experience it. Higher values of ρ_h and ρ_w are associated with more stability of these preferences. The potential changes in the values of ε_h and ε_w are treated as unanticipated. When a model was estimated treating these potential changes as anticipated, those with low values of these preferences would retire early so as to start the process of reversion of these effects toward the mean and so increase the value of their later retirement. This anticipation does not appear to be plausible, so we treat individuals as expecting their current assessment of the value of leisure and retirement to remain constant.

C. Other Transitions.

There are some transitions of state variables which are not stochastic, but which depend on the decisions of the previous year. The most obvious of these is the state variable associated with being in the career job. If the individual was not in the career job at the beginning of the previous year, or if the individual was in the career job at that point but chose to leave the career job during the previous year, he or she will be treated as having left the career job at the beginning of the current year.

Defined benefit pension entitlements are another set of state variables which depend on the decisions of the previous year. If the individual had left the career job before the previous year began, or if the individual was in that job when the year began but decided to leave the job in the previous year, the defined benefit pension amount remains the same as in the previous year, as does the date when the individual can start to collect the pension. If the individual was in the career job at the start of the previous year and remained in that job in the previous year, then the potential defined benefit pension amount is increased to reflect the added tenure in the career job and possibly the increased wage. The date the individual can collect the pension may also

change if the individual remained in the career job. Individuals are assumed to collect their pensions on the earliest date they are eligible to do so.

The level of Social Security benefits is another variable which operates in much the same way as defined benefit pensions. The two Social Security state variables are the level of individual benefits for the husband and the level of individual benefits for the wife. If either spouse elected to work either full-time or part-time in the previous year, these individual benefits are adjusted upward to reflect both the additional wages and any reductions in the early retirement penalty to get the level of individual benefits for the current year in accordance with the Social Security rules. To calculate the potential Social Security benefits for the couple under different retirement choices in the current year, the levels of individual benefits are combined to reflect any spouse or survivor benefits and any reductions in benefits due to current work effort. As with pensions, Social Security benefits are assumed to be collected on the earliest date that an individual or couple may do so.

D. Exogenous Variables.

In addition to the state variables and decision variables discussed above, there are a number of variables that are considered exogenous to the model. Basically, this set of variables includes any variables which are not affected by either the decision variables made by the husband or wife or the stochastic processes described above.

The most conspicuous of the exogenous variables is the level of potential wages if employed. These include wages if the individual is employed full-time or part-time and if the individual has or has not been subject to a previous layoff. For individuals who have provided permission for the HRS to access their Social Security earnings history, wages are observed for years before the individual enters the HRS. If the individual was working full-time at the initial wave of the study, that job is considered to be the career job, and potential wages are projected

forward on the basis of the tenure and experience variables estimated in Section II. If the individual was not working full-time at the start of the study, the last full-time job is considered to be the career job, and potential wages are projected forward from the point the last full-time job ended. Actual earnings are used for years before the start of the study or the end of the last full-time job, whichever occurred earlier. For individuals who did not provide permission for the Social Security records to be accessed, wages are constructed on the basis of the employment history information collected in the first year the individual entered the study. The employment history only provides wages for a limited number of years; wages in other years are projected again on the basis of the tenure and experience coefficients of the wage equations.

It is certainly true that wages are stochastic, and it would have been preferable to include stochastic wages if they did not require another set of continuous state variables which would have made the model computationally nearly impossible to solve. However, the most important source of stochastic variation in wages for older individuals is the possibility of a layoff which sharply lowers the potential wages from the point of the layoff forward, and stochastic layoffs are explicitly included in the model.

Another variable which is considered exogenous in the model is health. We have previously estimated a model in which health was stochastic, but that model took the labor force participation of the spouse as exogenous. Including both stochastic health and endogenous spouse labor force decisions runs again into problems of computational complexity, and for the current problem at hand, i.e. the Great Recession, it seems that it is probably more important to include the possibility of endogenous spouse labor force reactions to the stochastic layoffs than to include stochastic health. Note that there is still an element of stochastic health included to the extent the stochastic entry into the SSDI program is included in the model.

For career jobs with pensions, the parameters of the pensions, both defined contribution and defined benefit, are considered exogenous. For individuals who provided information and permission to collect pension information from their employers, the pension parameters come from the pension documents provided by the employers. For individuals who did not provide permission to collect the pension documents, pensions are randomly imputed from a sample of the observed pensions for individuals with similar industries, similar occupations, and similar earnings levels. Either way, the exogenous parameters of the pension plan are combined with the endogenous decision as to how long to remain in the career job to determine the level of the pension and how long it can be collected. A similar procedure is used to determine how the exogenous parameters of the Social Security system combine with the work decisions of the couple to determine the amount of Social Security benefits. For the Social Security system, exogenous parameters include the formula for determining the primary insurance amount, the reduction factors for retirement before the full retirement age, and the rules governing how the entitlements for the two spouses combine into a family amount, given the possibility of spouse and survivor benefits.

E. Model Solution.

The model is solved by the usual backwards induction method for stochastic dynamic models. In the final year (taken to be age 99 in this model), the value functions of the two spouses are evaluated for all possible combinations of the state variables. For any continuous state variable, the function is evaluated at a set of discrete points covering the range of possible values. Since all individuals are assumed to retire by age 70 in the model, there are no employment decisions at age 99, and the only remaining decision is how much to consume. Since anything left over has no utility value, the maximizing choice is to consume all the assets initially held at that age plus any other income such as Social Security and pension benefits. In

short, the value functions in the final year are the consumption value of the resources available that year.

For the immediately preceding year, the value functions are again evaluated for all possible combinations of the state variables. For each combination of the state variables, the functions are evaluated for all possible combinations of the decision variables, and the choice of the decision variables for the two spouses is governed by the process described in Part A of this section. The procedure is repeated for each prior year in turn until the husband is age 25, which is the age at which the model begins. Essentially, this procedure produces the decisions which the spouses would make for any combination of the state variables at the start of a particular year.

To simulate the model, the couple starts off when the husband is at the initial age. The couple is assumed to begin with no assets, and both spouses are in their career jobs. Thus, the set of state variables is determined at the initial age, and the decisions as to labor supply and consumption are made in accordance with the analysis described in the previous paragraph. Given these decisions, any assets at the end of the period, as well as any defined contribution balances, are adjusted according to the observed rate of return for the couple's mix of assets, and random draws are made for any other stochastic variables. This gives a set of state variables to begin the process at age 26, and the decisions at that age are made in the same way as before. This process is repeated for subsequent ages, and the result is a sequence of simulated labor supply and consumption decisions, along with a simulated trajectory of assets.

F. Specification of Heterogeneous Parameters.

In addition to the stochastic processes discussed above, the model contains three variables which are constant over time but heterogeneous among individuals. The first of these is the time preference parameter δ . This parameter is closely related to the amount of wealth that a

couple accumulates over time in preparation for retirement, with lower values of δ being associated with higher amounts of wealth. The data on wealth suggest that it is difficult to capture the distribution of δ with any sort of usual distribution, such as normal or chi-squared. Rather, there appears to be roughly two-thirds of the households who have accumulated at least a reasonable amount of wealth, consistent with reasonable low values of δ . However, about a third of the households have accumulated almost no financial or retirement wealth on their own, and for these households a fairly high value of δ would be necessary. It is tempting to model δ as a bivariate variable, being either high or low, but that would mask important but small differences for those with low values of δ which lead to substantial differences in the amount of wealth accumulated. In the end, it seems best to allow each household to have its own value of δ , not necessarily taken from any fixed distribution. We specify in the next part how this is done.

The other two parameters are the relative utility values of part-time leisure of the husband and wife, u_p^h and u_p^w , which must be between $\frac{1}{2}$ and 1 in order to satisfy diminishing marginal utility of leisure. We assume the husband's value comes from an exponential distribution $f(u_p^h) = ce^{\gamma_t^h u_p^h}$ defined over the interval $\frac{1}{2}$ to 1, where the factor c is the magnitude necessary to make the integral of the distribution over the allowable range equal to 1. The higher γ_t^h , the greater the probability the husband's value of u_p^h is close to 1, making it more likely that he will go through a period of partial retirement. Conditional on employment, the probability of partial retirement for the husband increases with age. Accordingly, we allow the value of γ_t^h to increase with age according to the equation $\gamma_t^h = \gamma_o^h + \gamma_1^h a_t^h$, where a_t^h is the husband's age at time t . We assume the husband's relative preference for partial retirement, and therefore the husband's relative position within this distribution, does not change even as the entire

distribution changes to higher values of u_p^h . The values of u_p^w for the wife are treated symmetrically and independently.

G. The Estimation Procedure.

The estimation method is the generalized method of simulated moments (MSM). In this procedure, a group of moments is gathered into a column vector \mathbf{m} . These moments are generally the difference between some observed statistic, such as the percentage retired as of a specific age, and the percentage that is simulated for the sample using specified values of the parameter. In general, these moments come from an asymptotically normal distribution with a mean value of zero. The estimation procedure seeks the parameter values which minimize $q =$

$\mathbf{m}'\mathbf{W}^{-1}\mathbf{m}$, where $\mathbf{W} = \sum_{i=1}^n \mathbf{m}_i\mathbf{m}_i'$. The \mathbf{m}_i vectors are the moments of the individual

observations, and the \mathbf{W} matrix is essentially the observed variance-covariance matrix of the moments. Variances of the estimates are calculated from $\text{var}(\Theta) = [\mathbf{G}'\mathbf{W}^{-1}\mathbf{G}]^{-1}$, where Θ is the vector of parameters and \mathbf{G} is the derivative of the moments with respect to the parameters. If the model is correctly specified, \mathbf{m} is distributed around zero, and q should have a χ^2 distribution with $\lambda - k$ degrees of freedom, where λ is the number of moments and k is the number of parameters estimated.

To construct the moments for a specific observation using a specific set of parameters, we follow a two-step procedure. The first step is to calculate the approximate value of the time preference parameter δ using the observed retirement dates and the value of wealth as of the last wave the spouses worked full-time. If one or both spouses have not retired by the last observed wave, we use the expected retirement date, and if the retirement from full-time work was into partial retirement, retirement from partial retirement work is calculated in the same way. This establishes the earnings stream of the household, and we solve the remaining consumption model

with stochastic returns and mortality. This simplified model is solved by backward induction to yield the saving decision at any age for any level of wealth. Wealth is then started at zero and accumulates using the observed returns and savings decision. The value of δ is adjusted up or down until the computed level of wealth in the last wave of full-time work just matches the observed wealth in that wave.

In general, it is not possible to establish a unique value for an individual effect such as δ in a model with unobserved stochastic elements, as is the case with the full model. However, extensive experimentation with similar models shows that the range of values of δ which yield particular retirement dates and a particular wealth level is fairly narrow. Further, the value of δ calculated from the simplified model is generally within that range. The intuition for this result is that given the observed opportunity set, the unobserved leisure preferences are fairly closely tied to the retirement dates, so the assumption of fixed retirement dates does not greatly affect the calculation of the time preference δ .

The second step is to solve the full model using the value of δ calculated from the first step. This is done using the method of backward induction discussed in the previous part of this Section. It yields the decision rules applicable to any set of the state variables. Random draws are made from the distributions of ε_h , ε_w , u_p^h , and u_p^w , given the parameters σ_h , σ_w , ρ_ε , γ_o^h , γ_1^h , γ_o^w , and γ_1^w of those distributions. These values are combined with the observed rates of returns and the decision rules to calculate the work-retirement patterns used in the moments. Then another draw is made for ε_h and ε_w , and the process is repeated 10,000 times. Each time, another set of retirement states for both spouses is calculated, and the values associated with the various moments are updated. Once the model has been estimated, the calculation of the 10,000 simulations does not add appreciably to the time required, so we use this number in order to

make the simulated moments close to the theoretical moments. In order to make the calculations feasible, new draws of u_p^h and u_p^w are not made for each of the 10,000 simulations per observation, but new draws of these values are made for each observation.

The moments used in the simulation are chosen to provide identifying information on the parameters. The moments used are as follows:

Description of moments	Number of moments
The percentage retired from full-time work at ages 54-66	
Husband	13
Wife	13
The percentage completely retired at ages 55, 58, 60, 62, and 65	
Husband	5
Wife	5
The percentage retired from full-time work at ages 55, 58, 60, 62, and 65 among families in the bottom third of potential earnings	
Husband	5
Wife	5
The percentage retired from full-time work at ages 55, 58, 60, 62, and 65 among families in the upper third of potential earnings	
Husband	5
Wife	5
The percentage retired from full-time work at ages 55, 58, 60, 62 and 65 among those in poor health	
Husband	5
Wife	5
The percentage of reversals where the respondent was working full-time after having been partially or fully retired in the previous interview	
Husband	7
Wife	7
The percentage of couples in each interview where both spouses were working full-time	8
The percentage of cases where one partner was retired in one interview and the other was working full-time three interviews later	6
Total moments	98

The first 26 moments help to establish the overall pattern of retirements governed by β_0^h , β_1^h , β_0^w , β_1^w , σ_h , and σ_w . The next 10 moments reflect the amount of partial retirement by age, and hence help to establish the values of γ_0^h , γ_1^h , γ_0^w , and γ_1^w , which determine the distribution of the value of partial retirement. The next 20 moments help to establish the value of α . The higher the value of α , the later will be the retirement of high earning individuals relative to the retirement of lower earning individuals. The next 10 moments, which involve individuals in poor health, help to determine β_2^h and β_2^w . The following 14 moments, which have to do with reversals, help to establish the values of ρ_h and ρ_w , the changes in leisure preferences after retirement which could induce a return to work. The next eight moments, which relate to joint retirement, help to determine the values of β_3^h and β_3^w , which govern how much the leisure of one spouse is augmented by the presence of the other spouse. The final six moments, which measure how often one spouse retires early and the other spouse retires much later, help to establish the value of ρ_ε , and the correlation between ε_h and ε_w . With high values of ρ_ε , cases where one spouse retires early and the other much later occur less often than when ρ_ε is low.

H. *Estimation Results.*

The model is estimated for a set of 728 couples from the War Babies and Early Baby Boom cohorts of the Health and Retirement Study. The War Babies were 51-56 years old when they were introduced into the study in 1998, and the Early Baby Boomers were 51-56 years old when they were introduced into the study in 2004. The model is estimated on data from the Health and Retirement Study through 2012. In 2008, at the beginning of the Great Recession, these cohorts were 55 through 66 years old, and in 2012 they were 59 through 70 years old.

The HRS has two important supplements, which are available on a restricted basis. First, Social Security earnings records are attached for roughly three quarters of the War Baby and

Early Boomer cohorts, allowing fairly precise estimates of Social Security earnings and benefits for this part of the sample. Secondly, for about two thirds of those respondents from these cohorts who indicated they had a defined benefit pension on their current job, the survey obtained and coded the pension's summary plan description. This enables a much more precise determination of the retirement incentives of pensions than is normally obtainable from the respondents themselves (Gustman, Steinmeier and Tabatabai, 2010, Table 4.2, p. 67.)

Our analysis uses data from wave 4 of the study, when the War Babies were first interviewed, through wave 11, conducted in 2012. By 2012, the households were four years past the beginning of the Great Recession, and the major impacts in terms of layoffs and stock market gyrations had passed. Observations are from households of married couples where both spouses appear to have career jobs from which retirement is a meaningful concept.

The results of the estimation are given in Table 9. The coefficients for own health, spouse retired, and the standard deviation of ε are expressed as multiples of the basic age coefficient. For example, the coefficient of 5.19 for the husband's own health means that the effect of poor health on utility is equivalent to an additional 5.19 years of age. Most of these coefficients are similar to those estimated in an earlier model based on data for the original HRS cohort (Gustman and Steinmeier, 2014), although that model did not allow for stochastic layoffs or disability. In particular, the coefficients of both the health variables and the variables indicating the augmentation of leisure if the spouse is retired are similar.

One notable difference between these estimates and those of the previous work is the pair of coefficients governing the reassessment of leisure after retirement, ρ_h and ρ_w . These are considerably higher in the current model, at 0.98 and 0.95, respectively, vs. 0.86 and 0.87 in the previous model, indicating that in this model there is much less change in retirement preferences after retirement. In the previous model, changes in retirement preferences were necessary in

order to generate the observed level of retirement reversals, where individuals returned to work after leaving work. In the current model, evidently, the stochastic layoffs generate enough instances where the individual exits work for a while and then returns so that changes in retirement preferences play much less of a role.

For 80 degrees of freedom (98 moments less 18 parameters), the 5% critical value of the χ^2 distribution is 101.9 and the 1% critical value is 112.3. The value calculated for the model is 126.6, which is slightly above the upper critical value and which indicates that at least some of the moments are higher in absolute value than would be observed by chance. The overestimate is only about 13 percent, though, and the simulated results described in the next Section appear to be fairly reasonable.

Table 9
Parameter Estimates

Parameter Symbol	Description	Estimated Value	Absolute t-Statistic
Husband's Parameters			
Parameters for the Value of Leisure			
β_0^h	Constant	-9.97	403.61
β_1^h	Age	0.19	14.41
β_2^h	Own Health	5.79	16.83
β_3^h	Spouse Retired	2.13	9.73
σ_h	Standard Deviation of ε_h	6.45	27.59
ρ_h	Correlation of ε_h After Retirement	0.98	91.27
Parameters for the Value of Partial Retirement			
γ_0^h	Constant	-2.18	6.79
γ_1^h	Age	0.24	2.26

Parameter Symbol	Description	Estimated Value	Absolute t-Statistic
Wife's Parameters			
Parameters for the Value of Leisure			
β_0^w	Constant	-10.04	355.45
β_1^w	Age	0.175	16.96
β_2^w	Own Health	4.45	13.41
β_3^w	Spouse Retired	0.88	4.29
σ_w	Standard Deviation of ε_w	6.24	20.32
ρ_w	Correlation of ε_w After Retirement	0.95	47.16
Parameters for the Value of Partial Retirement			
γ_0^w	Constant	-7.38	16.87
γ_1^w	Age	-0.28	1.52
ρ_ε	Correlation Between ε_h and ε_w	0.69	5.24
α	Consumption Parameter	-0.25	4.19
	Number of observations	728	
	Q Value	126.57	

Section VI. Base Simulation

The base simulation considers the counterfactual to the Great Recession, i.e., what might have happened had the Great Recession never occurred. To simulate what might have happened in ordinary circumstances, we make three changes to the budget sets and expectations. First, we assume that layoffs would have continued at more or less their normal rates, which can be calculated by setting the Great Recession term in the layoff probability and layoff duration equations to zero even during the 2009-2010 period. Secondly, we assume that the returns to stocks would have continued at their long-term averages during the 2009-2012 period rather than the sharp decline and recovery which in fact occurred. Finally, we assume that expectations of

the returns to money market-like assets would have continued at their long-term averages instead of being depressed in the wake of the Great Recession.

Table 10
Retirement States in Base Simulation, By Age and Gender

Age	Percent Pseudo-Retiring		Percent		
	From FT Work	From All Work	in FT Work	Partially Retired	Fully Retired
Husbands					
55	2.9	1.9	88.1	4.8	7.1
56	3.7	2.0	84.5	6.5	9.1
57	4.2	3.0	80.3	7.7	12.0
58	4.2	3.3	76.0	8.6	15.3
59	4.7	3.6	71.3	9.7	19.0
60	6.0	4.5	65.3	11.2	23.5
61	6.6	4.9	58.8	12.9	28.4
62	11.4	7.3	47.3	17.0	35.7
63	6.1	5.4	41.2	17.8	41.0
64	5.6	5.6	35.6	17.8	46.6
65	5.9	5.3	29.7	18.4	51.9
66	4.3	4.8	25.4	17.9	56.7
67	6.6	5.7	18.8	18.8	62.4
Wives					
55	3.0	2.1	81.4	7.7	10.8
56	3.6	2.8	77.8	8.6	13.6
57	3.8	2.7	74.0	9.7	16.3
58	4.2	3.2	69.8	10.7	19.5
59	4.4	3.7	65.4	11.4	23.2
60	5.9	5.0	59.5	12.3	28.2
61	5.5	4.7	54.0	13.1	32.9
62	11.4	7.9	42.5	16.6	40.9
63	4.3	5.3	38.2	15.6	46.1
64	5.3	6.1	32.9	14.8	52.2
65	5.4	6.2	27.6	14.0	58.4
66	2.8	5.3	24.7	11.6	63.7
67	4.7	5.5	20.1	10.7	69.2

Table 10 shows the base level retirement patterns. The last three columns give the simulated percentages of husbands and wives in full-time work, partial retirement, and full retirement, respectively for individuals in the 55-67 age range. The figures correlate fairly well with what we might expect, with declining full-time work and increased full retirement with advancing age. The partial retirement of husbands increases to around 17-18 percent by age 62 and remains high, while the partial retirement of wives increases to around 17 percent at age 62 and then begins to decline somewhat. The first column gives what might be called the amount of “pseudo-retirement” from full-time work, which is just the difference in the percent of individual in full-time work in one year subtracted from the corresponding percentage the year before. The second column gives the “pseudo-retirement” from all work, which is the difference between the percentage fully retired in one year and the corresponding percentage the year before.

The most notable feature of this table is perhaps the spike in retirement from full-time work at age 62. This means that for both husbands and wives, around 11.4 percent of individuals retire at that age. Compared to previous work using the older original HRS cohort, these figures are somewhat lower. In earlier work, the corresponding figures were around 15 percent. However, these lower percentages do correspond reasonably well to the underlying data from which the model is estimated. In the raw data, using the differences between the percentages retired at 61 and 62 are 11.5 and 10.5 percent for husbands and wives, respectively, with standard errors of around 2 percentage points.

Table 11 examines the differences between the first retirement from full-time work for the husbands and wives. The upper part of the table indicates that the husband has retired first, and the left column gives the number of years that the husband has retired before the wife. The bottom part of the table gives the corresponding figures for instances in which the wife retires first. The figures in the right column give the percentage of simulations in which the indicated

result occurs. Note that no adjustment is made for differences in the ages of the two spouses, just the number of years between retirement dates. The variation is rather large, but there is a definite spike at 0 years, indicating that the husband and wife retire in the same year. This agrees with other research that suggests some coordination of retirement dates.

Table 11
Distribution of Differences in Retirement Dates
Base Simulation

	Difference in Retirement Dates (Years)	Percent of Households
	10+	16.7
	9	2.7
	8	3.0
	7	3.3
Husband	6	3.4
Retires	5	3.8
First	4	4.2
	3	4.4
	2	4.3
	1	4.1
	0	9.6
	1	2.8
	2	3.6
	3	3.9
Wife	4	3.7
Retires	5	3.4
First	6	3.2
	7	3.0
	8	2.7
	9	2.4
	10+	11.7

Table 12 indicates the distribution of time preference rates among the households in the sample. The vast majority of households fall into two groups, one having very low time preference rates and the other having relatively high time preference rates. Most of the households in the high time preference rate group have accumulated almost no financial assets

that can be used to finance retirement, while those in the low time preference rate category have often accumulated substantial financial assets. About a quarter of the households fall into an intermediate time preference rate category and have accumulated some but not a large amount of financial assets.

Table 12
Distribution of Time Preference Rates

Time Preference Rate	Percent of Households
0-5%	42.4%
5-10	10.3
10-25	10.3
25-50	4.4
>50	32.6

Table 13 switches focus to examine the years around the Great Recession, starting in 2008 and ending in 2012. The table looks at three five-year age groups, starting at age 55 and ending at age 69, and examines the simulated percentages of the sample in each of the three retirement categories. The top half of the table is for husbands, while the bottom half is for wives. In looking at the table, recall that the sample consists of the War Babies and Early Baby Boom cohorts, who were born between 1942 and 1953. These individuals would have been between 55 and 66 years old in 2008 and 59 to 70 years old in 2012. Thus, the 55 to 59 year old age range has an increasing portion of individuals near the top of the age range as time advanced. In 2008, the entire age range of 55 to 59 year olds is represented, but by 2012 the youngest of these cohorts was 59 years old. A similar situation applies to the 65 to 69 year olds. In 2008, only the 65 and 66 year olds are included in this range, but by 2012 the entire 65 to 69 year old age

range is included. The 60 to 64 year age range does not suffer similar problems, since the entire age range is included in every year between 2008 and 2012 for these cohorts.

Table 13
Retirement Rates in Base Simulation
By Age Group, Year, and Gender

Age Group	Full-Time Work			Part-Time Work			Full Retirement		
	55-59	60-64	65-69	55-59	60-64	65-69	55-59	60-64	65-69
Year	Husbands								
2008	78.6	48.1	23.7	7.7	14.1	18.9	13.7	37.9	57.3
2009	77.1	49.8	22.9	8.0	13.3	16.6	14.9	36.9	60.5
2010	73.8	48.6	21.7	8.3	13.7	15.9	17.9	37.7	62.4
2011	71.5	47.3	19.3	8.3	14.6	16.2	20.2	38.1	64.5
2012	68.9	47.5	18.9	9.6	14.4	14.6	21.4	38.1	66.4
	Wives								
2008	74.1	44.3	20.6	8.1	13.7	11.3	17.8	42.0	68.1
2009	73.1	44.0	21.5	8.8	13.4	12.4	18.1	42.7	66.0
2010	73.4	42.9	18.6	9.3	14.1	11.2	17.3	43.0	70.2
2011	71.9	43.2	18.0	10.2	13.8	11.4	17.9	43.0	70.6
2012	70.8	44.3	18.8	10.1	13.3	9.5	19.2	42.4	71.7

In the base simulation, the percentage of individuals aged 60 to 64 years old, both husbands and wives, who are working in the various retirement categories (full-time, part-time, and fully retired) does not vary greatly over the 2008-2012 time period. The percentage of husbands working full-time in this age range holds fairly steady at just below 50 percent, with the percentage of wives being slightly lower. Not quite 15 percent of both husbands and wives are working part-time, while the percentage fully retired hovers around 40 percent. For the 55-59 year olds, the percentages working full-time are higher but declining over the time period, reflecting the fact that the composition of this group is becoming older over time as the younger ages drop out of the sample. For the 65-69 year olds, the percentages working full-time are

considerably lower, as one might expect, and are declining due to the fact that the composition of this group is also becoming older over time as older individuals in the cohorts are included in the sample. While these composition effects affect the base percentages of individuals in the 55 to 59 year old age range and the 65 to 69 year old age range who are working in the various retirement categories, they should be less important when we look at changes due to the Great Recession in the next section, since the composition of these age ranges is held constant in all of the simulations.

VII. Simulated Effects of the Great Recession

This section investigates the effects of three separate effects of the Great Recession. The first of these is the increase in layoffs which accompanied the start of the Great Recession. Although the Great Recession started in late 2008, in the simulations we take the start of the Great Recession to be the beginning of 2009, since the periods in the model are complete years. Table 1 indicates that layoffs almost doubled in the 2010 wave, but by the 2012 wave they had returned to levels which are about the same as in 2008 and earlier. As a result, the probit for layoffs reported in Table 2 includes a term indicating wave 10 (2010), and the duration estimates reported in Table 5 include a term for 2009 and 2010. In the simulation, we take the effects of the Great Recession on layoffs to be primarily in 2009 and 2010, with the increased probability of recession given by the coefficient in Table 2 and the increased duration given by the coefficient in Table 5.

Table 14a
Effects of Components of Great Recession
On Retirement and Work Effort

Age Group	Husbands								
	Full-Time Work			Part-Time Work			Full Retirement		
	55-59	60-64	65-69	55-59	60-64	65-69	55-59	60-64	65-69
Great Recession Layoffs Only									
2009	-2.8	-1.8	-0.7	2.8	1.7	0.6	0.0	0.1	0.1
2010	-2.7	-1.7	-0.7	2.7	1.6	0.6	0.0	0.0	0.0
2011	-0.1	-0.1	-0.1	0.1	0.2	0.1	0.0	0.0	-0.1
2012	-0.2	-0.2	-0.1	0.3	0.2	0.2	0.1	-0.1	0.0
Great Recession Actual Returns Only									
2009	1.0	2.2	1.1	0.0	0.3	0.5	-1.0	-2.5	-1.7
2010	0.6	1.1	1.0	0.0	0.2	0.2	-0.7	-1.3	-1.1
2011	0.7	0.9	0.7	-0.1	0.0	0.3	-0.6	-0.9	-1.0
2012	0.9	1.4	1.0	-0.2	0.1	0.4	-0.7	-1.5	-1.4
Great Recession Expected Returns Only									
2009	-0.6	-0.5	-0.2	0.1	-0.1	-0.3	0.5	0.6	0.4
2010	-0.9	-0.6	-0.3	0.2	-0.1	-0.2	0.6	0.7	0.6
2011	-0.7	-0.5	-0.6	0.1	-0.2	-0.2	0.6	0.8	0.8
2012	-0.6	-0.7	-0.3	0.0	0.1	-0.3	0.7	0.7	0.7
Total Effect of Great Recession									
2009	-2.5	-0.2	0.1	3.0	2.0	1.0	-0.5	-1.8	-1.1
2010	-2.8	-1.2	-0.1	2.9	1.7	0.7	0.0	-0.5	-0.6
2011	-0.2	0.1	0.0	0.1	0.0	0.2	0.1	-0.2	-0.2
2012	0.0	0.5	0.5	0.1	0.2	0.2	0.0	-0.7	-0.6

Figures are percentage point differences between the indicated effect and the corresponding amount in the base simulation. For instance, the -2.8 figure in the first row indicates that full-time work among 55-59 year olds in 2009 was 2.8 percentage points lower when Great Recession layoffs are introduced, e.g., 77.1 percent with base case layoffs vs. 74.3 percent with Great Recession layoffs.

Table 14b
 Percentage Point Effects of Components of Great Recession
 On Retirement and Work Effort

Age Group	Wives								
	Full-Time Work			Part-Time Work			Full Retirement		
	55-59	60-64	65-69	55-59	60-64	65-69	55-59	60-64	65-69
Great Recession Layoffs Only									
2009	-2.3	-1.4	-0.6	2.2	1.2	0.5	0.1	0.1	0.2
2010	-2.2	-1.3	-0.4	2.1	1.2	0.4	0.1	0.1	0.1
2011	-0.3	-0.2	-0.1	0.2	0.2	0.1	0.1	0.0	0.0
2012	-0.4	-0.2	-0.2	0.3	0.2	0.1	0.0	0.0	0.0
Great Recession Actual Returns Only									
2009	1.2	1.7	1.7	-0.1	0.1	0.1	-1.1	-1.9	-1.7
2010	0.7	1.1	1.3	-0.1	0.0	0.1	-0.6	-1.1	-1.4
2011	0.7	0.8	0.6	-0.1	0.2	0.2	-0.6	-1.0	-0.9
2012	0.8	1.3	1.1	0.1	0.2	0.4	-1.0	-1.4	-1.5
Great Recession Expected Returns Only									
2009	-0.6	-0.5	-0.4	-0.1	-0.1	0.1	0.7	0.5	0.4
2010	-0.5	-0.5	-0.3	0.1	-0.3	-0.2	0.4	0.8	0.5
2011	-0.8	-0.5	-0.5	-0.2	-0.2	-0.3	1.0	0.7	0.8
2012	-0.9	-0.6	-0.4	-0.1	-0.3	-0.3	0.9	0.9	0.7
Total Effect of Great Recession									
2009	-1.7	0.0	0.6	2.0	1.2	0.6	-0.3	-1.3	-1.2
2010	-2.2	-0.9	0.5	2.1	1.1	0.4	0.1	-0.3	-0.9
2011	-0.5	0.1	0.0	0.1	0.2	0.0	0.4	-0.3	0.0
2012	-0.5	0.4	0.5	0.3	0.2	0.2	0.1	-0.5	-0.7

The top panels of Table 14a and 14.b give the effects of the increased layoffs on the percentages of individuals working in the various retirement categories for both husbands and wives. The numbers represent the percentage point differences between the simulation involving Great Recession layoffs minus the base simulation. For instance, for 55-59 year olds working full-time in 2009 (the first number in the table), 74.3 percent of the simulations with Great Recession layoffs were for full-time work and 77.1 percent of the simulations in the base case scenario without Great Recession layoffs were for full-time work, for a difference of -2.8 percentage points.

The first two rows indicate a substantial reduction in full-time work for the two years in which the increased Great Recession layoffs were assumed to have occurred. The reduction is greatest for the 55-59 year age group, with the 60-64 year old age group having a somewhat smaller reduction and the 65-69 year old age group having an even smaller reduction. However, bear in mind that these figures are percentage point reductions and are occurring on the base of a smaller number of individuals working full-time at increasing ages.

The top panel of Table 15 expresses these reductions as a percentage of individuals still working full-time. The 2.8 percentage point reduction in full-time work (the first figure in Table 14a) relates to the 77.1 percent of husbands working full-time in the base simulation. If we express the reduction as a percentage of those still working full-time in the base simulation, the reduction becomes 3.6 percent of those working full-time ($2.8 / 77.1$), as reported in the first figure of Table 15.

Table 15
 Percent Change in Full-Time Work
 Due to Components of Great Recession

Age Group	Husbands			Wives		
	55-59	60-64	65-69	55-59	60-64	65-69
Great Recession Layoffs Only						
2009	-3.6	-3.6	-3.1	-3.1	-3.2	-2.8
2010	-3.7	-3.5	-3.2	-3.0	-3.0	-2.2
2011	-0.1	-0.2	-0.5	-0.4	-0.5	-0.6
2012	-0.3	-0.4	-0.5	-0.6	-0.5	-1.1
Great Recession Actual Returns Only						
2009	1.3	4.4	4.8	1.6	3.9	7.9
2010	0.8	2.3	4.6	1.0	2.6	7.0
2011	1.0	1.9	3.6	1.0	1.9	3.3
2012	1.3	2.9	5.3	1.1	2.9	5.9
Great Recession Expected Returns Only						
2009	-0.8	-1.0	-0.9	-0.8	-1.1	-1.9
2010	-1.2	-1.2	-1.4	-0.7	-1.2	-1.6
2011	-1.0	-1.1	-3.1	-1.1	-1.2	-2.8
2012	-0.9	-1.5	-1.6	-1.3	-1.4	-2.1
Total Effect of Great Recession						
2009	-3.2	-0.4	0.4	-2.3	0.0	2.8
2010	-3.8	-2.5	-0.5	-3.0	-2.1	2.7
2011	-0.3	0.2	0.0	-0.7	0.2	0.0
2012	0.0	1.1	2.6	-0.7	0.9	2.7

Table 16 gives the percent of full-time workers who are laid off by age group and year for both the simulation using Great Recession layoffs and the base simulation. For 2009 and 2010, the volume of layoffs mirrors fairly closely the reduction in full-time work associated with the layoffs. For instance, for the 55-59 year olds in 2009, the volume of layoffs in the simulation with Great Recession layoffs was 6.7 percent vs. 3.3 percent in the base simulation, for a difference of 3.4 percent. This compares with a 3.6 percent reduction in full-time work for the same group in Table 15.

Table 16
Percent Laid Off by Year and Age Group

Age Range Year	Husband			Wife		
	55-59	60-64	65-69	55-59	60-64	65-69
	Base Case					
2008	3.4	3.5	3.6	3.1	3.4	3.9
2009	3.3	3.4	3.8	3.0	3.4	3.7
2010	3.4	3.4	3.7	2.9	3.4	3.9
2011	3.1	3.5	3.7	2.8	3.4	3.6
2012	3.0	3.4	3.8	2.9	3.3	3.6
	Great Recession					
2008	3.4	3.5	3.6	3.1	3.4	3.9
2009	6.7	6.7	7.6	6.1	6.8	7.3
2010	6.7	6.9	7.4	6.0	6.9	7.5
2011	3.1	3.5	3.8	2.9	3.5	3.6
2012	3.1	3.5	3.8	3.0	3.4	3.7

For 2011 and 2012, there is a lingering effect of the layoffs in the previous two years. For those two years, the layoff rates are returned to the levels calculated in the layoff equation for both simulations, and in Table 16 the layoff rates for the two simulations are nearly identical for those two years. However, there is still a small but noticeable reduction in full-time work in those years even though the layoffs are nearly the same. This may be due to the fact that for the additional individuals who were laid off in 2009 and 2010 in the Great Recession simulation,

future earnings were depressed because of the loss of tenure and the additional earnings penalty for a layoff in the earnings equation. The loss of earnings, in turn, makes full-time work less attractive and results in less full-time work, as indicated in the tables.

Returning to Tables 14a and 14b, it appears that the reduction in full-time work is accompanied by a corresponding increase in part-time work, but this is probably an artifact of the assumptions made about part-time work. In the model, if an individual is laid off from full-time work, there is an interval determined by the duration estimates before he or she can return to full-time work, but part-time work is assumed to always be available. If part-time work were subject to the same duration of non-work as full-time work is, the reduction in full-time work as a consequence of the increased layoffs during 2009 and 2010 would be accompanied by a corresponding increase in full retirement (non-work).

Not surprisingly, with regard to layoffs, the basic story is that layoffs by themselves lead to a decrease in full-time work, with some carryover into the following years as a result of the reduction in potential full-time earnings for those laid off. The question is how much of this reduction in full-time work resulting from layoffs is offset by the other changes that occurred during the Great Recession, most notably the sharp drop in stock prices during the initial years of the Great Recession.

The second panel of Tables 14a and 14b analyzes the simulated effects of the drop in stock prices during the initial years of the Great Recession. These figures are percentage point changes between a simulation in which stock returns followed their actual path during the period 2008-2012 and a simulation in which stock returns adhered to their historical levels. The drop in stock prices during the Great Recession induced a substantial drop in the financial assets and defined contribution balances that individuals were accumulating for retirement, and one might expect that individuals would increase their work so as to make up some of the loss.

The tables show, for both husbands and wives, an increase in the amount of full-time work of around one percentage point. The increase is slightly higher in the 60-64 year old age range, when a majority of the retirements occur, and is lower for both the younger age group and older age group, when fewer retirements occur. For husbands in the base simulation, 20 percent of the sample retires in the 55-59 age range, 35.3 percent retires in the 60-64 age range, and 24.9 percent retires in the 65-69 age range. The decision to postpone retirement in order to offset asset losses is a marginal decision, and the frequency of postponed retirements is related to the density of retirements at a particular age, not necessarily to the number of individuals still working.

To illustrate, if an individual who is currently age 55 and working full-time and intends to retire at 62 suffers an asset loss, he or she will continue to work full-time whether or not the asset loss occurs, and for such an individual the asset loss will not cause any change in current employment, assuming the individual is not laid off. On the other hand, if an individual who is currently at his or her intended retirement age of 62 suffers a similar asset loss, he or she may decide to continue working at 62 rather than retiring, so that such an individual increases work effort at the current age. That is, an increase in work effort due to an asset loss occurs for individuals on the cusp of retirement, and individuals who are years away from retirement are not affected.

For both husbands and wives who increase full-time work effort in response to the asset losses, the movement appears to come from full retirement, not from part-time work. This means that individuals who had intended to retire completely in the simulation without the asset losses change behavior to work full-time in the simulation with the asset losses. Relatively few of the individuals appear to leave full retirement to part-time work in order to offset the asset losses caused by the Great Recession.

The third panel of the tables looks at another aspect of the Great Recession, namely, the apparent long term reduction in the returns to short term bonds, savings accounts, money market accounts, and similar cash-like equivalents. In the immediate aftermath of the onset of the Great Recession, the nominal returns on these rates plunged to approximately zero, and the real returns plunged to the negative of the inflation rate, approximately minus 2 percent. Moreover, unlike the value of stocks, which by 2012 were returning to levels close to pre-Great Recession values, the returns on the cash-like accounts continued to be around zero in nominal terms and negative in real terms, with little movement above these levels even by summer of 2016. It is certainly plausible to suppose that starting in 2009, individuals may have changed their long run expectations about the returns to these instruments. The third panel of the table looks at the implications regarding such expectations.

More specifically, it compares two scenarios. In the first (the base simulation), individuals continue to believe that the future returns on these short term assets will correspond to their long-run averages. In the second, individuals believe, starting in 2009, that the return on short term assets will decline to around negative 2 percent. For the simulation period through 2012, the actual returns are used in both simulations. The only difference is the expectations regarding future returns. Note that since the stock prices relatively quickly returned to pre-Great Recession levels, a similar assumption about future returns being lower is not made.

In terms of labor supply, lower future returns should have two effects in different directions, holding actual returns constant. On the one hand, the lower future returns mean that for a given level of saving, there will be fewer assets in the future to finance consumption in retirement, and individuals should save more in order to even out consumption. This is the income effect. Offsetting this is the fact that in order to gain a given amount of future consumption, more consumption will have to be given up currently, a tradeoff change which

leads to increased current consumption at the expense of future consumption. This is the substitution effect. In the simulations, the substitution effect evidently wins out. If expectations about future short term returns are lowered, individuals are simulated to work less, as indicated by the negative numbers for full-time work in Tables 14a and 14b. The value of future consumption relative to current leisure is tilted toward current leisure, and individuals retire more and work full-time less. Again, the effect on part-time work is fairly minimal. Note that the model does not try to model the choice of stocks vs. short term assets endogenously, but rather holds the proportion of stocks and short term assets constant at the observed proportions.

The final panel of these tables looks at the combined effect of Great Recession layoffs, stock returns, and changed expectations regarding future short term asset returns. In general, the overall effect seems to be approximately equal to the sum of the individual effects. The largest reductions in full-time work effort are for the 55-59 year old males in 2009 and 2010, years in which the layoffs had increased. After the period of high layoffs, the net effect is very small, with the increases due to the stock crash approximately offset by a reduction due to the lower expected returns for short term assets. The overall effect is much smaller for the 60-64 age group, even in 2009 and 2010. For those years, the reduction in full-time effort due to layoffs is largely offset by the increased work effort of individuals who were on the cusp of retirement but who suffered from the stock crash. The overall effect is also small for the older 65-69 year olds, for much the same reasons.

Table 17a
Retirement Rates by Layoff Status - Husbands

Laid Off in Last 3 Years		Full-Time Work		Part-Time Work		Fully Retired	
		No	Yes	No	Yes	No	Yes
	Year						
55-59	2008	89.2	57.7	4.7	29.3	6.2	12.9
	2009	88.9	57.3	4.5	28.1	6.6	14.6
	2010	86.8	55.4	4.7	27.8	8.5	16.7
	2011	85.3	53.8	5.5	28.7	9.2	17.5
	2012	84.2	50.8	6.5	32.7	9.3	16.5
60-64	2008	71.0	40.9	12.3	33.7	16.8	25.4
	2009	72.3	42.5	11.3	33.2	16.4	24.3
	2010	70.9	40.7	11.3	33.6	17.8	25.7
	2011	69.3	39.8	13.2	35.3	17.5	24.9
	2012	69.2	40.8	13.3	34.7	17.5	24.5
65-69	2008	59.8	29.5	18.7	42.3	21.5	28.2
	2009	60.7	30.2	17.0	39.1	22.4	30.7
	2010	59.1	28.0	19.1	40.2	21.9	31.8
	2011	52.9	26.4	21.8	39.4	25.4	34.2
	2012	52.8	26.8	19.2	36.9	28.1	36.3
Laid Off More than 3 Years Ago		No	Yes	No	Yes	No	Yes
55-59	2008	80.4	74.4	7.1	12.0	12.5	13.6
	2009	79.0	73.2	7.4	11.9	13.6	14.9
	2010	75.8	69.8	7.7	12.3	16.5	17.9
	2011	73.8	67.4	7.9	12.7	18.3	19.9
	2012	71.5	63.7	9.3	15.0	19.3	21.3
60-64	2008	50.9	45.3	13.5	18.9	35.6	35.8
	2009	52.6	46.5	12.8	18.4	34.6	35.1
	2010	51.4	45.2	13.0	18.6	35.6	36.2
	2011	50.2	43.8	13.9	19.7	35.9	36.5
	2012	50.1	44.6	14.0	19.5	35.9	36.0
65-69	2008	26.6	22.0	19.2	23.8	54.3	54.1
	2009	25.4	21.5	16.8	20.8	57.8	57.7
	2010	24.4	20.4	16.2	20.4	59.3	59.2
	2011	21.5	18.3	16.6	20.6	61.8	61.1
	2012	21.1	18.5	14.9	18.5	64.0	63.0

Table 17b
Retirement Rates by Layoff Status - Wives

Laid Off in Last 3 Years		Full-Time Work		Part-Time Work		Fully Retired	
		No	Yes	No	Yes	No	Yes
	Year						
55-59	2008	89.2	57.7	4.7	29.3	6.2	12.9
	2009	88.9	57.3	4.5	28.1	6.6	14.6
	2010	86.8	55.4	4.7	27.8	8.5	16.7
	2011	85.3	53.8	5.5	28.7	9.2	17.5
	2012	84.2	50.8	6.5	32.7	9.3	16.5
60-64	2008	71.0	40.9	12.3	33.7	16.8	25.4
	2009	72.3	42.5	11.3	33.2	16.4	24.3
	2010	70.9	40.7	11.3	33.6	17.8	25.7
	2011	69.3	39.8	13.2	35.3	17.5	24.9
	2012	69.2	40.8	13.3	34.7	17.5	24.5
65-69	2008	59.8	29.5	18.7	42.3	21.5	28.2
	2009	60.7	30.2	17.0	39.1	22.4	30.7
	2010	59.1	28.0	19.1	40.2	21.9	31.8
	2011	52.9	26.4	21.8	39.4	25.4	34.2
	2012	52.8	26.8	19.2	36.9	28.1	36.3
Laid Off More than 3 Years Ago		No	Yes	No	Yes	No	Yes
55-59	2008	80.4	74.4	7.1	12.0	12.5	13.6
	2009	79.0	73.2	7.4	11.9	13.6	14.9
	2010	75.8	69.8	7.7	12.3	16.5	17.9
	2011	73.8	67.4	7.9	12.7	18.3	19.9
	2012	71.5	63.7	9.3	15.0	19.3	21.3
60-64	2008	50.9	45.3	13.5	18.9	35.6	35.8
	2009	52.6	46.5	12.8	18.4	34.6	35.1
	2010	51.4	45.2	13.0	18.6	35.6	36.2
	2011	50.2	43.8	13.9	19.7	35.9	36.5
	2012	50.1	44.6	14.0	19.5	35.9	36.0
65-69	2008	26.6	22.0	19.2	23.8	54.3	54.1
	2009	25.4	21.5	16.8	20.8	57.8	57.7
	2010	24.4	20.4	16.2	20.4	59.3	59.2
	2011	21.5	18.3	16.6	20.6	61.8	61.1
	2012	21.1	18.5	14.9	18.5	64.0	63.0

Tables 17a and 17b move on to examine a slightly different question, namely what are the lingering effects of being laid off. Table 17a is for the husbands, and Table 17b is for the wives. The top half of the table looks at the effect of being laid off either in the current year or in one of the last two years, while the bottom half of the table looks at the effect of having been laid off 3 or more years ago. Each pair of numbers represents the percentage of respondents working, for instance, full-time if no layoffs occurred in the time frame specified and the percentage if a layoff did occur in that time frame. Only simulations in which the spouse was not laid off before the year in question are included in order to avoid confounding the effects of a layoff of the spouse.

In deriving these numbers, one cannot simply look at the percentage working full-time among those not laid off and compare them to the percentage for those who are laid off. The reason is that the percentage working full-time among those not laid off includes all those who are already retired in the denominator, and if the calculations are done this way the percentage working full-time for those not laid off is frequently lower than the percentages for those laid off. In order to avoid this effect the calculations are done to omit the influence of those already retired. This method may be illustrated for the line for the year 2009 in the top part of the table.

First, all simulations in which the individual was not working full-time in 2008 are eliminated, as are simulations in which the individual had been laid off prior to 2008. The remaining simulations are divided according to whether the individual was laid off in 2009, and the percentages are calculated. Next, all simulations in which the individual was working full-time in 2007 and not previously laid off are considered. Again, these are divided according to whether the individual was laid off in 2008, and the percentage of each group that was working full-time in 2009 is calculated. Finally, all simulations in which the individual was working full-time in 2006 and not previously laid off are considered. Again, these are divided according to

whether the individual was laid off in 2007, and the percentage of each group that was working full-time in 2009 is calculated. For each group, there are three percentages for full-time work in 2009, depending on whether we are considering potential layoffs in 2007, 2008, or 2009. In each case, the percentages only consider simulations in which the individual was working in the prior year. To get the figures in the top part of the table, the three percentages are weighted by the number of simulations in which the individual worked in the prior year.

The bottom part of the table is constructed similarly. For the row labeled 2009, the percentages of individuals working full-time in 2009 are calculated by looking at potential layoffs in 2006, 2005, 2004, and prior years. Again, for each year of potential layoff, only simulations in which the individual had been working full-time the previous year and had not been laid off prior are considered. Again, the figures in the table are obtained by weighing the percentages for each potential layoff year by the number of simulations in which the individual had been working in the previous year.

The figures in the top part of these tables are not surprising, given that the three year time period includes the year of the layoff. These tables are compiled for the base simulation, that is, the simulation without the increased layoffs generated by the Great Recession. To illustrate, for the year 2010 and for 60-64 year olds, the full-time employment rate in 2010 is simulated to be 40.7 percent if there was an initial layoff during 2008, 2009, or 2010, and 70.9 percent if the individual entered those years working full-time the year before and was not laid off. For all three age ranges and for all five years considered, there is about a 30 percentage point differential in full-time work for those laid off vs. those not laid off.

The figures in the bottom part of these tables indicate the lingering effects of layoffs. For instance, for the year 2010 and for 60-64 year olds, the full-time employment rate in 2010 is simulated to be 45.2 percent if the individual had been laid off in 2007 or before and 51.4 percent

if the individual had been employed full-time in those years but not laid off. The difference of 6.2 percentage points means that being laid off three or more years in the past reduces full-time employment in the current year by about 12 percent ($6.2 / 51.4$). This figure is fairly consistent across years, age groups, and spouses, and reflects the reduced work incentives for full-time work arising from lower earnings due to the loss of job tenure with a layoff as well as the additional earnings penalty from a layoff.

Tables 18a and 18b are similar to Tables 17a and 17b, except that these tables look at the effect of a spouse's layoff. Again, the top half of the tables considers recent layoffs (including the current year and the previous two years), and the bottom half of the tables considers layoffs further in the past. One would expect the result of a spouse's layoff to be the net result of two separate effects. On the one hand, a spouse's layoff, whether current or past, results in a lower probability of current full-time work, as indicated in the previous tables. Given the positive effect of the spouse's layoff on an individual's evaluation of leisure, as evidenced by the coefficients β_3^h and β_3^w , a reduction in the spouse's full-time work will induce an increase in the evaluation of leisure and hence should reduce the individual's full-time work effort. On the other hand, a spouse's layoff will reduce the income available to the household, and this effect should increase the individual's incentive to work full-time.

The tables suggest that if the spouse has been laid off recently, the two effects very nearly offset one another. In some instances a layoff for the spouse increases full-time work effort slightly, but in other instances the reverse is true. For layoffs further in the past, however, the effect of a layoff for the spouse is more uniformly positive for the full-time work effort of the individual. Overall, if the spouse was laid off more than three years ago, current full-time work effort is increased by around 2 to 3 percentage points. For these cases the income effect evidently wins out.

Table 18a
Retirement Rates by Spouse's Layoff Status - Husbands

Spouse Laid Off in Last 3 Years		Full-Time Work		Part-Time Work		Fully Retired	
		No	Yes	No	Yes	No	Yes
55-58 Year Olds	Year						
	2008	82.9	83.4	4.9	5.1	12.2	11.6
	2009	81.2	81.0	5.2	6.1	13.6	12.9
	2010	78.7	77.8	5.3	6.8	16.0	15.4
	2011	77.0	75.7	5.9	6.5	17.1	17.8
	2012	74.6	73.2	7.0	7.3	18.4	19.4
60-64 Year Olds	2008	55.2	53.2	11.0	12.1	33.8	34.7
	2009	56.0	54.2	10.6	12.1	33.4	33.8
	2010	54.8	53.0	11.0	13.0	34.2	34.0
	2011	54.3	52.0	11.6	13.8	34.1	34.2
	2012	54.1	51.9	11.8	13.5	34.1	34.6
65-69 Year Olds	2008	31.0	28.5	17.5	18.5	51.5	53.0
	2009	31.2	28.4	15.4	16.8	53.5	54.8
	2010	29.9	26.6	15.1	16.9	55.0	56.5
	2011	26.2	23.5	16.0	17.5	57.8	58.9
	2012	26.1	22.9	14.0	15.6	59.8	61.4
Spouse Laid Off More Than Than 3 Years Ago		No	Yes	No	Yes	No	Yes
55-59 Year Olds	2008	81.5	83.7	5.1	4.3	13.5	12.0
	2009	80.2	82.1	5.4	5.2	14.4	12.7
	2010	77.3	78.6	5.5	5.9	17.2	15.5
	2011	75.6	77.0	6.1	5.5	18.3	17.5
	2012	72.9	75.0	7.1	5.9	20.0	19.0
60-64 Year Olds	2008	51.8	53.8	11.2	11.3	37.0	35.0
	2009	52.8	54.7	10.8	11.2	36.4	34.1
	2010	51.6	53.4	11.3	11.7	37.1	35.0
	2011	50.8	52.3	11.5	12.4	37.7	35.3
	2012	50.0	51.8	11.8	12.1	38.2	36.1
65-69 Year Olds	2008	26.2	27.2	16.6	17.3	57.1	55.5
	2009	24.9	26.0	14.5	15.3	60.6	58.7
	2010	24.2	25.1	13.7	14.3	62.1	60.6
	2011	21.0	22.3	14.2	14.6	64.8	63.1
	2012	20.7	21.5	13.0	13.5	66.4	65.0

Table 18b
Retirement Rates by Spouse's Layoff Status - Wives

Spouse Laid Off in Last 3 Years		Full-Time Work		Part-Time Work		Fully Retired	
		No	Yes	No	Yes	No	Yes
	Year						
55-59	2008	79.4	80.9	4.8	5.5	15.8	13.7
	2009	78.9	80.5	5.3	6.0	15.8	13.5
	2010	79.4	80.9	6.2	6.3	14.5	12.8
	2011	78.7	80.2	7.1	6.7	14.2	13.1
	2012	78.6	78.8	6.2	7.0	15.3	14.2
60-64	2008	54.3	54.4	9.9	11.4	35.9	34.2
	2009	54.3	54.6	9.7	11.3	36.1	34.2
	2010	53.5	53.4	10.4	11.9	36.1	34.7
	2011	53.5	53.6	10.3	11.7	36.3	34.7
	2012	53.7	54.7	10.0	10.8	36.3	34.5
65-69	2008	33.0	31.9	8.0	10.4	59.0	57.7
	2009	35.8	34.6	8.8	10.5	55.4	54.9
	2010	32.1	31.0	9.5	11.0	58.4	58.0
	2011	30.8	29.5	10.5	11.9	58.8	58.7
	2012	32.9	31.3	8.0	9.1	59.1	59.6
Spouse Laid Off More Than Than 3 Years Ago		No	Yes	No	Yes	No	Yes
55-59	2008	76.5	79.6	5.8	5.6	17.8	14.8
	2009	76.3	79.5	6.1	5.7	17.6	14.8
	2010	76.8	80.2	6.8	6.0	16.4	13.8
	2011	75.2	78.7	7.9	6.7	16.9	14.6
	2012	74.2	76.8	7.1	7.4	18.7	15.9
60-64	2008	46.5	50.3	11.4	11.3	42.1	38.4
	2009	46.2	49.7	11.0	11.0	42.9	39.4
	2010	45.2	49.0	11.6	11.4	43.2	39.5
	2011	45.6	49.5	11.2	11.3	43.2	39.2
	2012	46.7	50.9	11.0	11.2	42.3	38.0
65-69	2008	22.7	23.7	9.1	10.6	68.2	65.7
	2009	23.1	25.0	10.5	10.9	66.4	64.1
	2010	20.2	23.5	9.5	9.4	70.3	67.0
	2011	19.6	22.3	9.6	10.4	70.9	67.3
	2012	20.3	22.8	7.6	8.6	72.1	68.6

IX. Conclusion

We have compared retirement outcomes between a base simulation that assumes the absence of the Great Recession, with outcomes produced in the presence of disturbances created by the Great Recession. These disturbances include additional layoffs, changes in future earnings prospects for both job losers and those who retain their jobs, and reductions in the value of financial assets and expected returns to these assets. The largest overall effects of the Great Recession on full-time work by members of two-earner households are observed for 2009 and 2010. In 2009, an additional 2.5 percent of all 55 to 59 year old husbands were not working full-time as result of the Great Recession, amounting to a reduction of 3.2 percent in full-time work. In 2010, 2.8 percent of 55 to 59 year old husbands were not working full-time as a result of the Great Recession, amounting to a 3.8 percent reduction in full-time work. For wives the reductions in full-time work due to the Great Recession were 1.7 percent and 2.2 percent of those who initially held a job, or reductions of full-time work of 2.3 and 3.0 percent respectively. For those 60 to 64, the reductions were 1.2 percent of men and 0.9 percent of women.

Most of these reductions were the direct result of additional layoffs during the Great Recession, which not only cost affected workers their jobs, but reduced their future earnings prospects should they continue in full-time work. Declines in the value of wealth due to the Great Recession increased the likelihood of full-time work, but this was substantially offset by the effect of decreases in expected returns, which reduced the likelihood of full-time work.

Having been laid off in the last three years reduces full-time work by 30 percent. There also are lingering effects of layoff on the probability of working longer. Having been laid off three or more years in the past reduces full-time employment in the current year by about 12 percent. This reflects the reduced work incentives for full-time work arising from lower earnings due to the loss of job tenure with a layoff as well as the additional earnings penalty from a layoff.

The effect of a spouse having been laid off on own work is much smaller. The reason is that, as found in the estimation of our structural model, having one spouse not working increases the value of leisure for the other. In contrast, when one member of the household loses their job, the value of consumption increases relative to leisure. For recent layoffs, these effects are roughly offsetting. If a spouse was laid off more than three years ago, current full-time work effort is increased by around 2 to 3 percentage points.

All told, these effects of the Great Recession on retirement were more modest than the press and initial research suggested. Our findings here of limited effects of the Great Recession on retirements and saving are consistent with our earlier descriptive research (Gustman, Steinmeier and Tabatabai, 2010b, 2012, 2014 and 2015). The more elaborate and detailed analysis contained in the present paper suggests that those descriptive studies did not obscure large offsetting effects from different influences. To be sure, as we show those who lost their jobs due to the Great Recession paid a significant price. But there were fewer of them than were initially expected. In addition, housing prices are recovering, erasing a major source of decline in wealth. On the whole, most of those nearing retirement at the outset of the Great Recession seem to have dodged a bullet.

References

- Blundell, Richard, Eric French and Gemma Tetlow. 2016. "Retirement Incentives and Labor Supply". *Handbook on Aging*. 458-566.
- Gustman, Alan L. and Thomas L. Steinemier. 2000. "Retirement in a Family Context: A Structural Model for Husbands and Wives". *Journal of Labor Economics* 18(3): 503-545.
- _____. 2004. "Social Security, Pensions and Retirement Behavior Within the Family". *Journal of Applied Econometrics* 19(6): 723-738.
- _____. 2014. "Integrating Retirement Models: Understanding Household Retirement Decisions". *Research in Labor Economics* 80: 81-114. (Selected as outstanding author contribution to *Research in Labor Economics* in 2014.)
- _____. 2016. "Effects of Social Security Policies on Benefit Claiming, Retirement and Saving". *Journal of Public Economics*. 129: 51-62.
- _____. Forthcoming. "The Role of Health in Retirement". *Research in Labor Economics*.
- Gustman, Alan L., Thomas L. Steinmeier and Nahid Tabatabai. 2008. "Do Workers Know About Their Pension Plan Type? Comparing Workers' and Employers' Pension Information". In Anna Lusardi, editor, *Overcoming the Saving Slump; How to Increase the Effectiveness of Financial Education and Saving Programs*. Chicago: University of Chicago Press, pp. 47-81.
- _____. 2010a. *Pensions in the Health and Retirement Study*. Cambridge: Harvard University Press, 2010.
- _____. 2010b. "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.
- _____. 2012. "How Did the Recession of 2007-2009 Affect the Wealth and Retirement of the Near Retirement Age Population in the Health and Retirement Study? "How Did the Recession of 2007-2009 Affect the Wealth and Retirement of the Near Retirement Age Population in the Health and Retirement Study?" *Social Security Bulletin*. 72(4), November: 47-66.
- _____. 2014. "The Great Recession, Decline and Rebound in Household Wealth for the Near Retirement Population". *Journal of Retirement* Fall 2014, 2(2): 27-44.
- _____. 2015. "Retirement and the Great Recession". *Journal of Retirement* 3(1), Summer 2015: 87-106.
- Merkuriva, Irena. 2016. "Late Career Job Loss and the Decision to Retire". University of Saint Andrews, School of Economics and Finance Discussion Paper No. 1606.