

Ippolito, Benedic N.

Working Paper

The selection effects of tied health insurance contracts

AEI Economics Working Paper, No. 2015-14

Provided in Cooperation with:

American Enterprise Institute (AEI), Washington, DC

Suggested Citation: Ippolito, Benedic N. (2015) : The selection effects of tied health insurance contracts, AEI Economics Working Paper, No. 2015-14, American Enterprise Institute (AEI), Washington, DC

This Version is available at:

<https://hdl.handle.net/10419/280535>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



AMERICAN ENTERPRISE INSTITUTE

The Selection Effects of Tied Health Insurance Contracts

Benedic Ippolito
American Enterprise Institute

AEI Economic Policy Working Paper 2015-14
November 2015

© 2015 by the American Enterprise Institute for Public Policy Research. All rights reserved.

The American Enterprise Institute for Public Policy Research (AEI) is a nonpartisan, nonprofit, 501(c)(3) educational organization and does not take institutional positions on any issues. The views expressed here are those of the author(s).

THE SELECTION EFFECTS OF TIED HEALTH INSURANCE CONTRACTS

BENEDIC N. IPPOLITO*

November 2015

Employers typically offer one of two types of health plans: tied contracts (coverage ends at retirement) or retiree contracts (coverage continues in retirement). In comparison to a retiree plan, a tied contract provides an obvious incentive to delay retirement, but which workers stay? Some papers suggest that the sickest workers are most likely to stay, because they presumably attach the highest value to health insurance. It seems odd, however, that a firm would voluntarily institute a contract that delivers such adverse selection. Empirically I show that, in fact, tied contracts produce advantageous selection. That is, firms that use tied contracts are successful not only in increasing their retention rate, but also in improving the health composition of their older workforce. Additionally, I investigate the fact that tied contracts are consistently found to delay retirement even after age 65, at which point workers are eligible for Medicare coverage. I show that this “excess retention” effect is driven by the fact that Medicare does not cover dependents. These results point to rational decision making by firms and workers, but also raise questions about why firms offer one contract over the other.

* Benedic Ippolito, American Enterprise Institute. Email: benedic.ippolito@aei.org

The author is grateful to Karl Scholz, Chris Taber, Jesse Gregory, John Mullahy, and Bobbi Wolfe. This work was supported by the National Institutes of Health under Ruth L. Kirschstein National Research Service Award No. T32 MH18029 from the National Institute of Mental Health.

1 Introduction

Employers typically offer one of two types of health insurance plans: tied contracts (coverage ends upon retirement), or retiree contracts (coverage continues after retirement). A tied contract provides an obvious incentive for workers to delay retirement prior to age 65, the age of Medicare eligibility.¹ Not surprisingly, studies have repeatedly shown that tied-workers retire at older ages than retiree-workers.²

These findings raise an odd outcome from the perspective of the firm. Compared to their healthier counterparts, the sickest workers presumably attach a higher value to the (subsidized) health insurance, which implies that a tied contract disproportionately encourages sicker workers to postpone retirement.³ But why would a profit-maximizing firm offer a contract that selects comparatively sickly workers at older ages?

The literature identifies another related puzzle: Not only does tied health insurance lower retirement rates before age 65, it also appears to discourage retirement after 65 as well. Why would a tied contract discourage retirement after Medicare eligibility? French and Jones (2011) suggest that, because retiree firms are more accommodating to earlier retirement, they might attract more leisure-loving workers, leaving those with higher preferences for work to populate the tied firms. But that begs the question: Why do retiree firms want to attract leisure lovers into their workforces?

In an attempt to resolve these issues, I focus my study on the selection effects of tied and retiree health contracts. Using 20 years of data from the Health and Retirement Study to examine retirement behavior in both types of firms, I find no evidence that firms are generating perverse selection effects. In fact, the results show that tied contracts disproportionately retain relatively healthy workers. Further, I show that the excess retention of tied contracts past age 65 can be fully explained by Medicare dependent eligibility rules, an idea which stems from Madrian and Beaulieu (1998).

The results of this study are consistent with rational decision making by both firms and workers, but point to a clear difference in age-tenure profile in these firms. In particular, while tenure at retirement is remarkably similar in both types of firms, particularly among the healthy populations, the age distributions are not. Retiree firms encourage workers to enter and retire earlier, while

¹Thus, if a worker with tied health insurance retires prior to Medicare eligibility, he must either buy private health insurance on the individual market, buy health insurance (without the firm contribution) under COBRA, or go uninsured (assuming the worker cannot obtain coverage through a spouse). The first two options are traditionally much more expensive to the individual than employer-sponsored insurance, while going uninsured obviously exposes an individual to heightened financial risk.

²For a review of early work see Gruber and Madrian, 2002, while more recent papers include Kapur and Rogowski (2011), Strumpf (2010), French and Jones (2011), Nyce et al. (2011), Marton and Woodbury (2013), Leiserson (2013), Fitzpatrick (2013).

³This issue is raised in Karoly and Rogowski (2000) who study the first few waves of the HRS, but has received limited attention since.

tied firms appear to use production functions that are more amenable to an older workforce.

2 A Simple Model of Retirement

In this section I introduce a simple model of retirement under tied and retiree insurance contracts that incorporates health status. While health likely impacts a number of factors relevant to this decision, for simplicity I allow health to affect only a worker's disutility of work and their value of health insurance. These two countervailing forces are sufficient to show how tied contracts can increase the retention of older workers without necessarily retaining the sickest workers. Consider a three-period model with a discount factor equal to one. In period one, all workers are employed at a fixed number of hours, enjoy perfect health, and suffer no disutility of work. Each has a value of marginal product equal to one, which equals the period-one wage. In period three all workers are retired and Medicare eligible. The worker's only decision is whether to work or retire during period two. After period one, each worker draws a random variable $S \sim \text{Unif}[0, 1]$, which determines the worker's period-two health status (with $S = 1$ being the sickest). The number of workers is normalized to one.

2.1 Retirement in a retiree contract

The cash wage in period two is w , which equals one minus the worker's contribution to his or her retiree health insurance. Utility is linear in consumption (there is no risk aversion and no value added by distributing consumption over periods). The money value of the disutility of work in period two is linear in sickness. Thus, the worker's period-two utility from working compared to retirement is given by:

$$U_W = w - (\beta + \gamma S), \quad \beta, \gamma > 0, S \in [0, 1] \quad (1)$$

where β is a parameter that captures the disutility from work in period two for a perfectly healthy worker ($S = 0$), and γS describes the increasing disutility affiliated with ever-higher sickness draws. If U_W is negative the worker retires; otherwise, the worker postpones retirement until period three:

$$\text{retire if: } \phi = (w - \beta) < \gamma S \quad (2)$$

where ϕ is the net value of work for a perfectly healthy individual ($S = 0$), and $\phi - \gamma$ is the net value of work for the sickest ($S = 1$).

Figure 1 depicts this solution. Workers with a sickness draw above S_R^* retire, leaving the portion S_R^* employed.

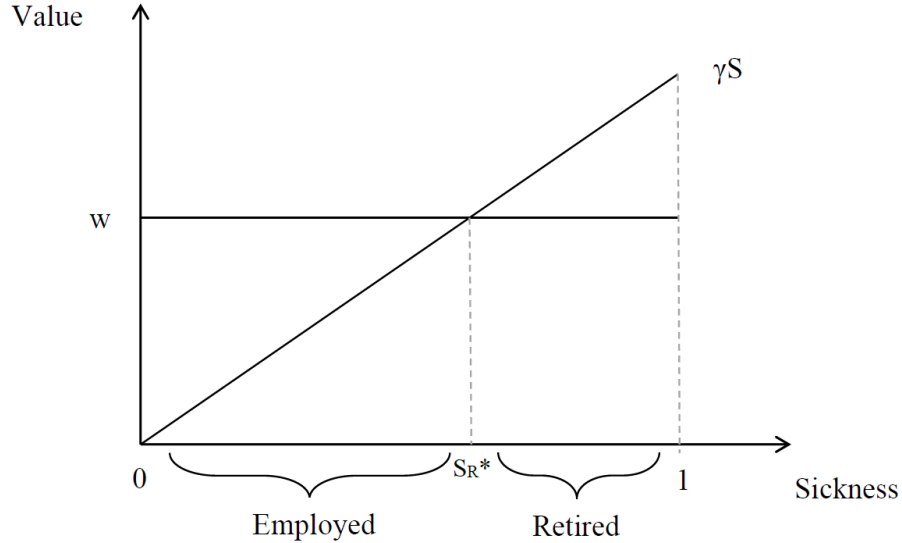


Figure 1: Retirement Outcome: Retiree Contract

Implicitly I am assuming that the firm's choice of wage w (which generates the retention rate S_R^*) corresponds to the firm's optimal retention rate of older workers. If this were not the case, the firm could restructure its compensation to raise or lower the rate.⁴ I abstract from this complication and assume that the outcome in Figure 1 conforms to the firm's optimal retirement rate.

The Cost of Insurance

In order to finance their health insurance, workers must accept reduced cash wages, which I assume occurs in period one. Suppose health expenses, $C(S)$, are proportional to a worker's level of sickness:

$$C(S) = h \cdot S, \quad h > 0. \quad (3)$$

Then the average health cost for all workers and retirees in period two is $I^R = .5h$. Each worker effectively purchases insurance in period one against ill health in period two. Recalling that the number of young worker is normalized to one, then the per-worker cost of insurance is:

$$I^R = 0.5h. \quad (4)$$

I assume that the period-two cash wage equals the value of marginal product; that is, I assume all compensating differentials show up in the period-one cash wage. In the retiree firm, period-one cash wage is $1 - I^R$.

⁴For example, the firm could use a defined benefit pension plan or could alter the period one and period two wage distribution as in Lazear (1979).

2.2 Retirement in a tied contract

Now suppose a firm wants a higher retention rate of older workers, and suppose for simplicity that the flexibility offered by a defined benefit pension contract is insufficient to generate the desired outcomes, or that the pension is otherwise uneconomic for the firm.⁵ Instead, the firm uses a tied health insurance contract.

The tied contract covers medical expenses only if the worker remains employed in period two. Consequently, a worker with a sickness draw S perceives period-two compensation, \tilde{w} , to be higher than wage rate w by the amount hS :

$$\tilde{w} = w + hS. \quad (5)$$

The worker in a tied firm retires if the net value of work, ϕ , plus the health cost is less than the disutility of work:

$$\text{retire if: } \phi + hS < \gamma S. \quad (6)$$

Figure 2 depicts the outcome. In comparison to a retiree firm, the tied firm attains a higher retention rate of older workers ($S_T^* > S_R^*$). However, in this scenario, as long as the solution is interior (if $0 < w < \beta - h$), the tied contract encourages the healthier workers to stay, but not the sickest ones. In other words, the sickest workers are least likely to react to the tied contract, even though they have the highest expected health costs.⁶

⁵Defined benefit contracts offer the firm significant flexibility to manipulate parameters to influence the workers optimum retirement age. These plans give rise to the notion of a pension delta which is described in French and Jones (2011). Also, see Gustman and Steinmeier (1994).

⁶See Appendix A for an illustration of how adding noise to the model generates a similar result, but with some healthy workers staying and some sick leaving early.

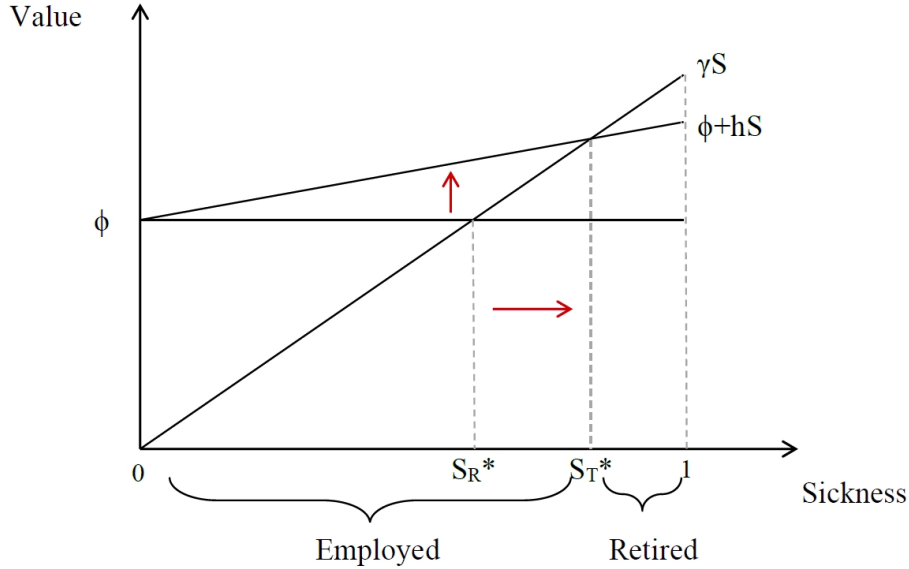


Figure 2: Retirement Outcome: Tied Contract

While highly stylized, this model illustrates how under reasonable conditions one can easily generate a scenario where tied contracts do not disproportionately encourage the sickest workers to delay retirement the most. I now turn to the data to test this question empirically.

3 Data Description

For the empirical work I use the Health and Retirement Study (HRS). The HRS is a nationally representative panel survey conducted bi-annually since 1992. I restrict the sample to individuals who, in 1992, are 50 to 55 years old, employed, covered by an employer-provided health insurance plan (either tied or retiree coverage), and who do not retire before age 55. This gives me a sample of 1,405 respondents. The youngest respondents in the first wave are age 50, allowing me to track all workers' retirement decisions through at least age 70. Critically, this allows me to track workers through almost the entirety of the retirement window, including the age of Medicare eligibility.

3.1 Defining Health Insurance Type

The first step is to classify workers by health insurance type. The HRS asks respondents whether their employer-provided insurance continues after retirement, though it is not consistently worded over time. I argue that the questions in wave 3 and beyond are ill-suited for the purpose of identifying health plan type (see Appendix B), and so use only the responses from the wave one questions⁷, a choice that makes my results directly comparable to most prior literature (for

⁷If a respondent reports being employed and covered by employer-provided insurance, but does not report having either tied or retiree insurance, I extract the information, if reported, from wave 2 (the phrasing of the question is

example, Blau and Gilleskie 2008 and French and Jones 2011).

A respondent has retiree coverage if he affirms that his employer health insurance continues after retirement. The respondent is considered to have “tied” coverage if he affirms that his employer health insurance does not continue after retirement, *and* he is not covered under a spouse’s retiree coverage. In effect, I am comparing retiree workers to tied-only workers.⁸

Rows 1 and 2 in Table 1 summarize the types of coverage in my sample. For context, I also show all other types of coverages for workers aged 50-55 in the baseline HRS sample, though they are not included in the study.

Table 1: Summary of Coverage Types, Age 50-55 at Baseline

Type	Frequency	Percent	Cumulative
Retiree	961	40.98	40.98
Tied	444	18.92	59.90
Coverage from spouse	426	18.17	78.07
Public insurance	104	4.43	82.50
None	181	7.72	90.22
Unknown	229	9.77	100
Total	2345	100	

Note: This table illustrates the distribution of insurance types at baseline. Data are from wave one of the HRS and include all respondents between age 50 and 55 who work.

3.2 Defining Retirement

I determine the respondent’s age of retirement from the response to the question: “In what month and year did you retire?” using the first available wave in which a response is given. To be classified as “retired,” a respondent must report that he considers himself to be either partially or fully retired.⁹

3.3 Defining sickness

“Sickness” can be defined in a variety of ways, but for my main analysis I use the number of reported chronic conditions at baseline. I show below that the results are qualitatively similar if identical in both waves).

⁸To the extent that tied firms attract workers with spouse-retiree coverage then the effects of the tie will be dampened by availability of the spouse’s plan. The effect of hidden (to the firm) coverage would depend on the spouse’s age and plan rules. To obtain the pure effects of the tie, I look at tied-only workers.

⁹If a respondent reports that he is partially retired, then it is unlikely that he still is working at his main job, that is, the job in which he was working fulltime at baseline. In other words, I care only about retirement from the firm, and not about retirement from the workforce. Further, if the respondent confirms retirement but does not report a year, I assume that he retired one year prior to the first wave in which he confirmed retirement.

use either self-reported health, or a more complex variable that represents projected health costs from ages 55-65.¹⁰

The seven chronic conditions are hypertension, diabetes, cancer (except skin cancer), chronic lung disease (except asthma), heart problems (including heart attack), stroke, psychiatric problems, or arthritis. The questions follow the pattern ‘has a doctor ever told you that you have arthritis, etc.?’ I simply sum the number of these pre-determined conditions for which the respondent answers in the affirmative. Table 2 shows the distribution of chronic conditions at baseline for workers in firms that offer tied or retiree coverage.

Table 2: Number of Chronic Conditions at Baseline

Number of Conditions	Frequency	Percent	Cumulative Percent
0	684	48.75	48.75
1	482	34.35	83.11
2	191	13.61	96.72
3	34	2.42	99.14
4 or more	12	.86	100
Total	1,403		

Note: This table illustrates the distribution of the number of health conditions at baseline. Data are from wave one of the HRS and include all respondents between age 50 and 55 who work at a firm which offers tied or retiree insurance.

About half of the respondents report having no chronic conditions, so I consider these as “healthy” workers and those with one or more conditions as “sickly.”

3.4 Worker Characteristics: Tied versus Retiree Firms

Table 3 shows means for selected characteristics of workers with tied and retiree insurance. The final column reports the difference between groups. Personal characteristics such as education, race, marriage status, and most measures of household wealth are also similar across tied and retiree firms, though workers in retiree firms are slightly more risk averse. Similarly, workers’ health is similar across firm types at baseline using either number of reported conditions or self-reported health. In short, workers do not seem to be sorting into retiree or tied firms based on personal characteristics.¹¹

¹⁰Wallace and Herzog (1995) provide a detailed explanation of health measures in the HRS. In addition to directly capturing major ailments, the chronic condition indicators are associated with measures of health care utilization, difficulties performing routine tasks, and other health measures. Similarly, self-reported health is “related to all physical as well as mental health measures,” particularly those related to pain, mobility, and depression.

¹¹Risk aversion is based on questions which ask the respondent is asked to choose between pairs of jobs where one guarantees current family income and the other offers a chance to increase income but also carries the risk of loss of income. If a respondent is not willing to take any option that involves the possibility of losing income they are

In contrast, there are notable differences in the characteristics of firms offering retiree or tied insurance. Retiree workers are more likely to be in a union, have a defined benefit pension, and have more vacation and sick time than tied firms. Retiree firms are much more likely to be large and to have workforces with more tenure (at baseline).

Table 3: Characteristics of workers employed at firms offering tied and retiree health insurance

Characteristic	Difference		
	Tied	Retiree	(Retiree minus Tied)
<i>Health Variables</i>			
Subjective health: excellent	0.30	0.29	-0.009
Number of chronic conditions	0.73	0.75	0.028
<i>Demographics</i>			
College graduate	0.10	0.10	.006
Black	0.17	0.16	-0.007
Married	0.70	0.73	0.034
Risk aversion	0.56	0.62	0.056**
Female	0.47	0.41	0.071**
Number of children	2.99	2.93	-0.061
<i>Household Wealth</i>			
Household income	56,104	59,540	3,436
Non-housing wealth	47,942	39,399	-8,543
Housing wealth	70,361	65,370	-4,991
IRA wealth	20,319	18,917	-1,403
<i>Employment Characteristics</i>			
At least 500 employees	0.41	0.58	0.165***
Covered by defined benefit	0.33	0.59	0.257***
Union membership	0.22	0.38	0.152***
Number sick days allowed per year	6.57	7.99	1.43***
Number sick days taken per year	3.22	5.50	2.275**
Weeks of vacation time per year	2.90	3.52	0.6249***
Tenure	12.33	16.23	3.89***

Note: Significance levels based on t-tests for the equality of means of the tied and retiree groups. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Data are from the first wave of the Health and Retirement Study. Sample includes all workers ages 50-55 at baseline who work in retiree or tied firms.

4 Retention Rates: Healthy versus Sickly in Tied and Retiree Firms

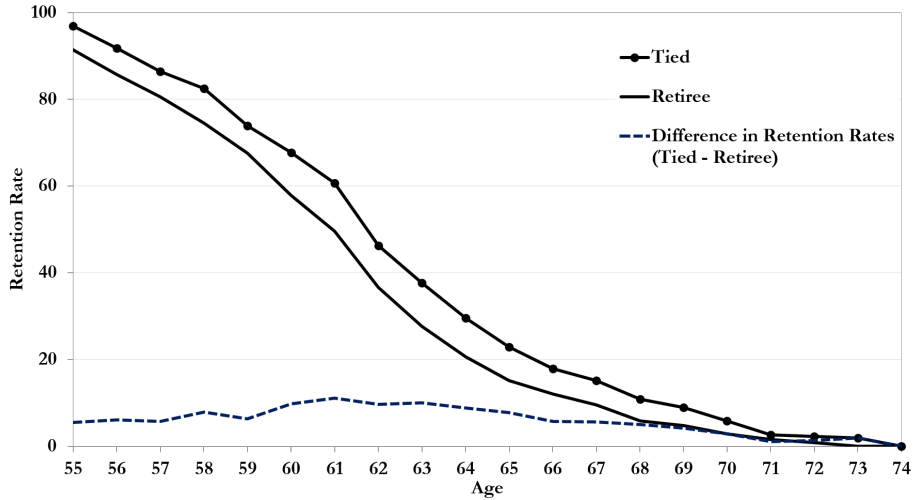
I now turn my attention to my main focus: retention rates among sickly and healthy older workers in tied versus retiree firms. I first show the results using a series of graphs, then report regression estimates of retirement, holding constant a number of worker and firm characteristics, including classified as “risk averse.”

those in Table 3. These estimates, along with various sensitivity tests, confirm the results shown in the graphs.

Figure 3 shows retirement patterns of relatively sick workers in tied and retiree firms (those with one or more chronic conditions at baseline). The line with markers denotes retention rates in tied firms, while the solid line with no markers shows the rates in retiree firms. The dashed line shows the excess retention rate generated by the tied contract (i.e., it is the vertical difference between the tied and retiree schedules). Sickly workers have higher retention rates in tied versus retiree firms.

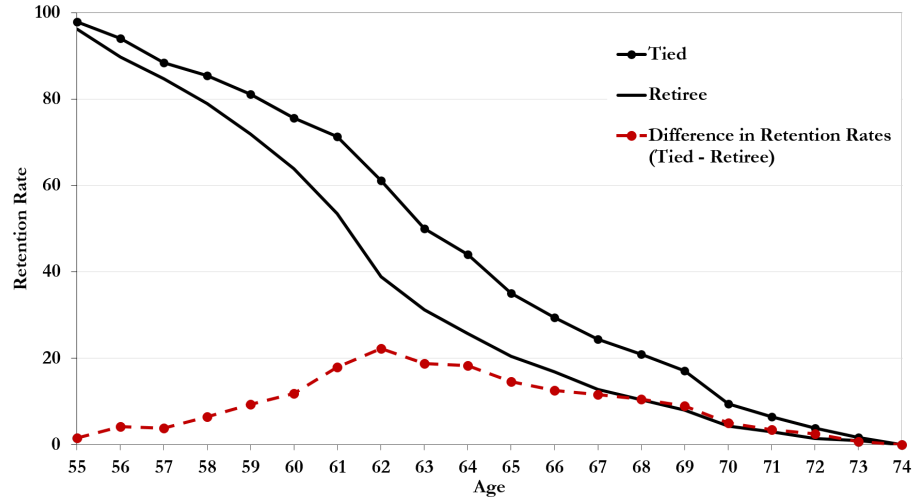
Figure 4 shows the same patterns, except now using the population of healthy workers—those with no chronic conditions at baseline. The results are qualitatively similar to those above. That is, healthy workers are less likely to retire from tied firms, particularly in their early to mid-sixties.

Figure 3: Comparing retention rates in tied and retiree firms for relatively sick workers



Note: This figure illustrates retention rates in tied and retiree firms for workers who have at least one chronic condition at baseline. Data are from the HRS cohort of the Health and Retirement Study and include all workers who are employed in a tied or retiree firm until at least age 55.

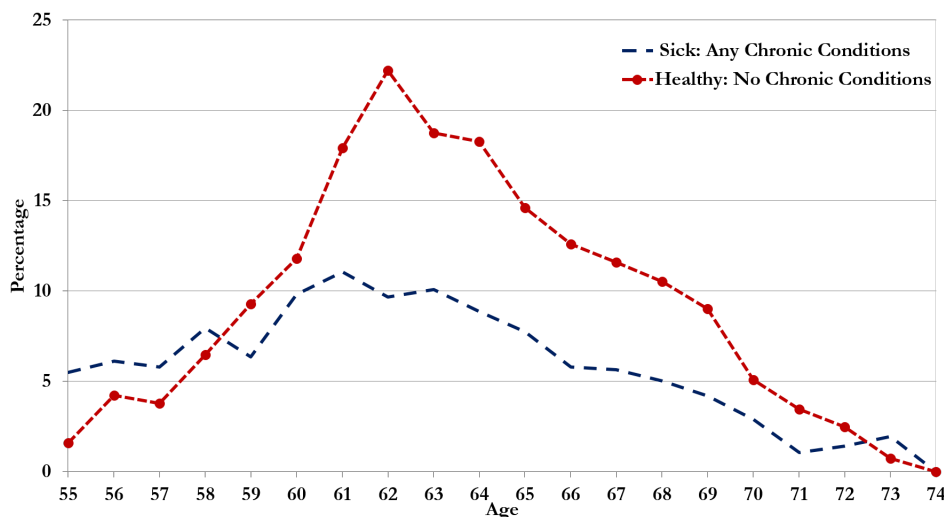
Figure 4: Comparing retention rates in tied and retiree firms for relatively healthy workers



Note: This figure illustrates retention rates in tied and retiree firms for workers who have no chronic conditions at baseline. Data are from the HRS cohort of the Health and Retirement Study and include all workers who are employed in a tied or retiree firm until at least age 55.

While these figures each reproduce the well-known effect of the tied contract, Figure 5 shows the selection effect. In particular, Figure 5 reproduces the excess retention rates from the tied contract for healthy (with markers) versus sickly workers (no markers) from the prior two figures. The effect of the tie is decidedly higher for the healthier workers, especially after age 60. That is, the tied contract generates *favorable* selection on health. In effect, the data are consistent with the implications of the simple model represented in Figure 2.

Figure 5: Excess retention rates of tied contracts: Healthy and sick workers



Note: This figure illustrates excess retention rates generated by tied contracts (relative to retiree contracts) for relatively healthy and sickly workers. Data are from the HRS cohort of the Health and Retirement Study and include all workers who are employed in a tied or retiree firm until at least age 55.

4.1 Illustrating the Reduced Form

I now show that the findings illustrated in Figure 5 persist when controlling for a rich set of employee and firm characteristics. I begin by estimating the following probit model for early retirement for a worker, i :

$$Pr(\text{early retirement}) = \alpha + \beta_1 X_i + \beta_2 Tied_i + \beta_3 Sick_i + \beta_4 Tied_i * Sick_i. \quad (7)$$

I define early retirement as retiring prior to age 63, though the results persist with alternate reasonable choices.¹² Controls in vector X_i include a host of demographic and job characteristics of respondents at baseline such as pension information, wealth, education, job tenure, marriage status, and spouse's health. $Tied_i$ is a binary variable equal to one if a worker has tied insurance (and is not covered by as spouse's retiree plan). $Sick_i$ is a dummy variable indicating whether a worker has any chronic conditions. β_4 , which captures the interaction between tied health insurance and sickness, is the central focus of this analysis.¹³ The omitted group is workers with retiree insurance. Standard errors are clustered at the household level to allow for serial correlation within households.¹⁴

¹²I prefer using age 63 since this means workers are retiring before they would be eligible for COBRA coverage that would extend through Medicare eligibility.

¹³I calculate the coefficient on the interaction term using the method suggested in Ai and Norton (2003).

¹⁴The validity of this empirical strategy is based on the assumption that workers do not sort into tied and retiree firms in a way that is correlated with retirement decisions. This assumption is supported by the limited differences

The results, reported in column (1) of Table 4, confirm the earlier findings. The coefficient on *Tied insurance* in the first row confirms the usual result that the tied contract reduces the likelihood of early retirement. The coefficient on the interaction between sick and tied is positive, large and statistically significant. Tied contracts increase the probability of early retirement for sickly workers by about 14 percent relative to healthy workers.

The results in column (1) assume that all worker characteristics impact retirement decisions identically for the healthy and sickly populations. I relax this assumption by estimating the probit model for each group separately (omitting the “sick” dummy and interaction term). The results in columns (2) and (3) are qualitatively similar to those presented in the first column. Tied contracts significantly reduce the probability of early retirement for the relatively healthy workers, while the coefficient on *Tied insurance* is small and insignificant for the relatively unhealthy group.¹⁵

in worker characteristics observed at baseline, as well as prior literature. For instance, Madrian (1994) reports that many workers, especially younger ones, do not know whether their plan is retiree or tied.

¹⁵The Affordable Care Act will likely impact these retirement patterns, though it’s effect is not clear. Healthy workers respond most to the potential loss of insurance, suggesting they would retire earlier in response to increased access to health insurance on the individual market (via health exchanges). However, the limitations on risk-adjustment in the exchanges reduce prices disproportionately for the sick. These possible effects are empirically testable when data become available, but are left for future work.

Table 4: The Relationship Between Tied Insurance and the Probability of Retiring Before Age 63 by Health Status

	Outcome: Probability of Retiring Before 63					
	(1)		(2)		(3)	
	Full Sample		Healthy (0 Conditions)		Sick (1+ Conditions)	
	Coeff.	Std Err	Coeff.	Std Err	Coeff.	Std Err
<i>Insurance</i>						
Has tied insurance	-0.187***	(0.044)	-0.194***	(0.046)	-0.026	(0.044)
Sick	0.013	(0.034)				
Tied insurance x Sick	0.144***	(0.056)				
<i>Job Characteristics</i>						
Managerial position	-0.131***	(0.037)	-0.067	(0.053)	-0.194***	(0.053)
Self employed	-0.116	(0.084)	-0.206*	(0.116)	-0.006	(0.135)
Owns a business	-0.059	(0.050)	-0.057	(0.075)	-0.075	(0.067)
Salaried	-0.038	(0.034)	-0.069	(0.050)	-0.013	(0.046)
Wage	-0.120	(0.131)	-0.047	(0.043)	-0.381	(0.210)
Works at large firm (500+ emply.)	0.044	(0.030)	0.055	(0.044)	0.021	(0.040)
Tenure	0.002	(0.002)	0.001	(0.002)	0.003	(0.002)
<i>Pension Information</i>						
Has defined benefit pension	-0.104	(0.066)	-0.080	(0.094)	-0.135	(0.094)
DB pension x Tenure	0.011***	(0.003)	0.011**	(0.004)	0.011**	(0.004)
Has defined contribution pension	0.027	(0.044)	-0.012	(0.062)	0.071	(0.061)
Earliest age for receiving pension						
By age 55	0.068	(0.060)	0.050	(0.085)	0.111	(0.089)
Between ages 55-62	0.043	(0.059)	-0.050	(0.083)	0.127	(0.086)
<i>Wealth Characteristics</i>						
Household income†	0.006	(0.004)	-0.001	(0.005)	0.016**	(0.006)
Household IRA wealth‡	0.029	(0.028)	0.044	(0.041)	0.008	(0.048)
Negative non-housing wealth	-0.098**	(0.036)	-0.082	(0.054)	-0.117*	(0.048)
Household DC pension wealth‡	0.031	(0.031)	0.034	(0.049)	0.018	(0.043)
<i>Demographics</i>						
Female	-0.040	(0.037)	-0.085	(0.053)	-0.014	(0.050)
Non-white	-0.023	(0.037)	-0.047	(0.053)	0.002	(0.052)
Has a spouse	0.029	(0.044)	0.167**	(0.064)	-0.111	(0.062)
Spouse has 1+ conditions	-0.012	(0.033)	-0.084	(0.047)	0.069	(0.047)
Spouse not working at baseline	-0.035	(0.037)	-0.078	(0.054)	0.005	(0.056)
Spouse is younger	-0.051	(0.038)	-0.132*	(0.053)	0.036	(0.054)
<i>Education Dummies</i>						
	Y	Y	Y	Y	Y	Y
Observations	1,426		684		717	

Note: This table reports the results of a Probit regression for early retirement against insurance type by health status. Coefficients represent marginal effects evaluated at the mean. Data are from the HRS cohort of the Health and Retirement Study and include all workers who are employed in a tied or retiree firm until at least age 55. Standard errors are in parentheses and clustered at household level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. †: measured in tens of thousands, ‡: hundreds of thousands.

4.2 Sensitivity

In this section I re-estimate the results using alternative specifications for the key variables. First, I consider alternative definitions of sickness. Next, I restrict the sample to those for whom I can affirm they did not change jobs after baseline and prior to their age of retirement.

4.2.1 Alternative Measures of Health

Above, I defined healthy and sickly based on chronic health conditions reported at baseline. In this section I report the results using alternate measures of health.

I first classify workers based on self-reported health at baseline. Respondents are asked to describe their health as excellent, very good, good, fair, or poor. I define as healthy, the 64 percent of respondents who classify their health as either excellent or very good.

I re-estimate the models reported in columns (2) and (3) of Table 4, using the self-reported definition. Columns (1) and (2) of Table 5 report the results. For brevity, I include only the coefficient on the tied variable. The results are similar to those presented above. The tied contract reduces the likelihood of early retirement by 16.6 percent in the healthy population, compared to a reduction of only 5.6 percent for the sickly group (which, again, is not significantly different from zero).

For a second alternative measure of health, I create a “sickness index” that incorporates additional baseline health information, and perhaps better captures expected medical costs. I regress (a proxy of) medical costs incurred by respondents aged 56-65 against a host of baseline health characteristics, then use these estimates to create a measure of predicted health expenses for my baseline sample, aged 50-55. For a more detailed explanation of this measure see Appendix C.

As above, I split the sample into two groups: those above the median \hat{S} are considered relatively sick, and those below are relatively healthy. I then estimate the effect of the tied contract on the probability of retiring before age 63 for each group. The results, reported in columns (3) and (4) of Table 5, show that tied contracts disproportionately reduce the probability of early retirement for the healthy group. Specifically, for the healthy workers, tied contracts are associated with a nearly 14 percent reduction in the probability of early retirement. For the sicker workers, the coefficient on tied insurance is about half as large, but for the first time, marginally significant (at the 10 percent level).

Table 5: The Relationship Between Tied Insurance and the Probability of Retiring Before Age 63: Alternative Health Measures

	Outcome: Probability of Retiring Before 63			
<i>Health Measure:</i>	(1)	(2)	(3)	(4)
	Self-Reported Health		Projected Sickness Index	
	Healthy	Sickly	Healthy	Sickly
Has tied insurance	-0.166*** (0.044)	-0.056 (0.045)	-0.136** (.045)	-0.073* (.044)
Controls	Y	Y	Y	Y
Observations	726	675	723	682

Note: This table reports the results of a Probit regression for early retirement against insurance type. Data are from the HRS cohort of the Health and Retirement Study and include all workers who are employed in a tied or retiree firm until at least age 55. Coefficients represent marginal effects evaluated at the mean. This analysis includes the full set of controls listed in Table 4. Standard errors are in parentheses and clustered at household level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.2.2 Ensuring No Switching of Insurance

In this section I address the possibility that some workers might switch firms after baseline and before retirement. If so, the insurance type may be mislabeled for some portion of the sample. I therefore restrict my sample to workers who report that their tenure immediately prior to retirement is sufficiently long to ensure that they did not change employer after baseline. The results in Table 6 are similar to those reported above. Healthy workers are 11 percent less likely to retire early in a tied contract, while the coefficient for sick workers is very small and insignificant.

Table 6: The Relationship Between Tied Insurance and the Probability of Retiring Before Age 63: No Changing of Firms

	Outcome: Probability of Retiring Before 63	
	(1)	(2)
	Healthy (No Conditions)	Sick (1+ Conditions)
Has tied insurance	-0.11** (0.050)	-0.018 (0.043)
Controls	Y	Y
Observations	548	607

Note: This table reports the results of a Probit regression for early retirement against insurance type. Coefficients represent marginal effects evaluated at the mean. This analysis includes the full set of controls listed in Table 4. Standard errors are in parentheses and clustered at household level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.3 The Post-Age-65 Effect

One interesting finding reported in Figures 3, 4, and 5 is that tied contracts appear to retain workers even past age 65 - the age of Medicare eligibility. This effect has not gone unnoticed in the literature. French and Jones (2011) suggest that the effect could be explained by self selection; that is, workers with higher preference for leisure might self-select into retiree firms, leaving a disproportionate number of ‘worker-types’ in tied firms. Rust and Phelan (1997) propose that the post-65 effect is explained by the relative generosity of private plans compared to Medicare.

Madrian and Beaulieu (1998) suggest an alternate theory. Unlike most private insurance, Medicare does not cover dependents. Even though workers themselves qualify for Medicare at age 65, their spouses do not until they themselves reach age 65. In this way, tied contracts can potentially defer the retirement even past age 65 for workers with younger spouses who do not have their own insurance. The same is true for those respondents with children who are covered under their health plan. Though respondents with minor children are relatively rare in my sample, about 400 have an insurance-dependent younger spouse. About half of these spouses are more than four years younger, suggesting tied contracts may retain some workers considerably past age 65.

I split the sample into two groups: those with an insurance-dependent younger spouse, and those without. Figure 6 illustrates the excess retention generated by tied contracts relative to retiree contracts for workers with and without a younger, dependent spouse. As expected, tied firms experience later retirements in both cases, but tied contracts appear to generate much higher post-65 retention among those with an insurance-dependent younger spouse. For workers without dependent younger spouses, the extra retention generated under the tied contract decreases rapidly as workers reach age 65. For workers with a dependent, tied contracts clearly generate higher retention rates than retiree contracts, even well past Medicare eligibility.


I now test this theory controlling for a full set of worker and spouse characteristics. I create a dummy variable ‘Dependency’ that is equal to one if the respondent has a spouse who is both younger, and entirely dependent on the respondent for health insurance coverage. I then estimate the following probit for retiring after age 65 (separately for those with and without a dependency):

$$\Pr(\text{Retire post } 65_i) = \alpha + \beta_1 X_i + \beta_2 Z_i + \beta_3 Tied_i$$

where X_i is a vector of worker characteristics, Z_i is a vector of spouse characteristics, and $Tied_i$ is an indicator equal to one if a worker has tied health insurance.

Table 7 reports the results. Column (1) reports the results for respondents with an insurance-dependent younger spouse, and column (2) reports the results for those without. The results are consistent with the Medicare-dependency hypothesis. Among respondents with an insurance-dependent younger spouse, the tied insurance increases the probability of retiring after age 65 by 13.3 percent. In the population with no such spousal dependency, the marginal impact of the tied

Figure 6: Comparing the difference between retention rates in tied firms and retiree firms for workers with and without dependents.



Note: This figure illustrates the excess retention rates generated by tied contracts relative to retiree contracts for two groups: those with a dependent (i.e., has a younger spouse who relies on the respondent for health insurance) and those without a dependent. Data are from the HRS cohort of the Health and Retirement Study and include all workers who are employed in either type of firm until at least age 55.

contract is small and insignificantly different from zero. Indeed, the latter result implies that after excluding dependencies from the data, there is no remaining post-65 effect to explain.

Table 7: The Relationship Between Tied Insurance and the Probability of Retiring After Age 65

	(1)	(2)
	Has a Dependent	Has No Dependent
Has tied insurance	0.133*** (0.044)	0.028 (0.029)
Controls	Y	Y
Observations	460	944

Note: This table reports the results of a Probit regression for retirement post age 65 against insurance type. Coefficients represent marginal effects evaluated at the mean. This analysis includes the full set of controls listed in Table 4. Standard errors are in parentheses and clustered at household level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.4 Completed Tenure in Tied and Retiree Firms

The late-career retention effects of tied contracts, both before and after 65, have important implications for overall tenure profiles of workers. At baseline, workers in retiree firms have higher

tenure than those in tied firms. As shown in Table 8, this is true of both healthy and sick workers.¹⁶ By retirement, however, the healthy workers almost “catch up” to their healthy counterparts in retiree firms. The catch up is more muted among the sick workers who, as previously showed, respond less to the tied incentives.

Table 8: Tenure at baseline and retirement

	Tied	Retiree	Difference (Retiree minus Tied)
<i>At Baseline</i>			
Healthy (0 Conditions)	13.6	16.2	2.6**
Sick (1+ Conditions)	13.0	17.4	4.4***
<i>At Retirement</i>			
Healthy (0 Conditions)	22.8	23.7	0.9
Sick (1+ Conditions)	20.8	24.2	3.4***

Note: Significance levels based on t-tests for the equality of means of the tied and retiree groups. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Data are from the first wave of the Health and Retirement Study. Sample includes all workers ages 50-55 at baseline who work in retiree or tied firms and do not switch firms prior to retirement.

In the end, tied and retiree firms generate similar levels of completed tenure in their workforces, especially amongst their healthier workers. The major difference is that, compared to tied firms, workers in retiree firms attain their completed tenure at younger ages. In other words, retiree firms hire workers at relatively young ages, encourage long tenure, but also encourage retirement at relatively early ages. Tied firms accomplish similar tenure profiles, but do so by hiring older workers and encouraging their retirement at older ages.

5 Conclusion

The results in this paper are consistent with rational choices of firms and workers. Tied firms clearly encourage their workers to retire at older ages. It is tempting to think that this contract would encourage the sickest to stay because of the relatively high value they attach to the health insurance, but this argument disregards other important influences of health, like the disutility of work when sick and old. In fact, the data show that tied firms are successful in retaining their healthiest workers; that is, the contract actually produces favorable selection.

Both retiree and tied firms achieve remarkably similar levels of completed tenure, particularly amongst healthy workers. Retiree firms produce their tenure profiles by hiring young workers and encouraging their early retirement. Tied firms hire older workers and produce longer tenure by

¹⁶In this section I focus on the roughly 80 percent of workers in my sample who do not change firm before retirement.

encouraging their workers to retire at older ages. In short, tied firms have older workforces than retiree firms.

It could be that retiree firms use production functions that are suitable to workers who attain experience before they get ‘too’ old. And that tied firms presumably use production functions that are more suited to workers with more labor market experience (age) and firm-specific experience (tenure). The underlying differences in workforce characteristics, and presumably the differences in underlying production functions across firms are issues not often pursued in the literature. But addressing the question, ‘why do firms offer the insurance contract they do?’ could represent a fruitful area of future research.

References

- Chunrong Ai and Edward C. Norton. Interaction terms in logit and probit models. *Economic Letters*, 80(1):123–129, 2003.
- David M. Blau and Donna B. Gilleskie. The role of retiree health insurance in the employment behavior of older men. *International Economic Review*, 49:475–514, 2008.
- Maria D. Fitzpatrick. Retiree health insurance for public school employees: Does it affect retirement? Working Paper 19524, National Bureau of Economic Research, October 2013.
- Eric French and John Bailey Jones. The effects of health insurance and self-insurance on retirement behavior. *Econometrica*, 79:693–732, 2011.
- Johnathan Gruber and Brigitte Madrian. Health insurance, labor supply, and job mobility: A critical review of the literature. Working Paper 8817, National Bureau of Economic Research, February 2002.
- Alan Gustman and Thomas Steinmeier. Employer-provided health insurance and retirement behavior. *Industrial and Labor Relations Review*, 48:124–140, 1994.
- Kanika Kapur and Jeannette Rogowski. How does health insurance affect the retirement behavior of women? *Inquiry*, 48:51–67, 2011.
- Lynn Karoly and Jeannette Rogowski. Health insurance and retirement behavior: Evidence from the health and retirement survey. *Journal of Health Economics*, 19:529–39, 2000.
- Edward P. Lazear. Why is there mandatory retirement? *Journal of Political Economy*, 87(6):1261–84, 1979.
- Greg Leiserson. Retiree health insurance and job separations: Evidence from pennsylvania state employees. *Manuscript. Massachusetts Institute of Technology*, 2013.
- Brigitte Madrian. The effect of health insurance on retirement. *Brookings Papers on Economic Activity*, 1:181–232, 1994.
- Brigitte Madrian and Nancy Beaulieu. Does medicare eligibility affect retirement? In *Inquiries in the Economics of Aging*, pages 109–131. University of Chicago Press, 1998.
- James Marton and Stephen A Woodbury. Retiree health benefits as deferred compensation: evidence from the health and retirement study. *Public Finance Review*, 41(1):64–91, 2013.
- Steven Nyce, Sylvester J. Schieber, John B. Shoven, Sita Nataraj Slavov, and David A. Wise. Does retiree health insurance encourage early retirement? *NBER*, Working paper, 2011.

John Rust and Christopher Phelan. How social security and medicare affect retirement behavior in a world of incomplete markets. *Econometrica*, 65:781–831, 1997.

Erin Strumpf. Employer-sponsored health insurance for early retirees: Impacts on retirement, health, and health care. *International Journal of Health Care Finance and Economics*, 10: 105–147, 2010.

Robert B. Wallace and A. Regula Herzog. Overview of the health measures in the health and retirement study. *The Journal of Human Resources*, 30:S84–S107, 1995.

A Adding Noise to the Model for the Empirical Work

In the model presented earlier, retirement decisions are entirely dependent on sickness draw. Disutility from work in period two is linear in S : $\beta + \gamma S$, where β measures disutility for a worker in perfect health. I can create differences in retirement decisions for workers with identical health by allowing period-two workers also to make a random draw for their value of β . Assume that this draw comes from a uniform distribution centered on zero (so half of workers enjoy their work when healthy; the other half do not). Figure 7 shows the model for a retiree firm with this modification.

The solid horizontal line represents the period-two wage. I portray sickness-related disutility by the upward-sloping schedule, labeled γS . A vertical line represent the distribution of β at any given sickness level along this schedule. Thus, the period-two disutility of work for individual j has two components: His sickly draw, S_j , determines his location along the upsloping schedule, and his β -draw determines his location on the vertical segment at S_j .

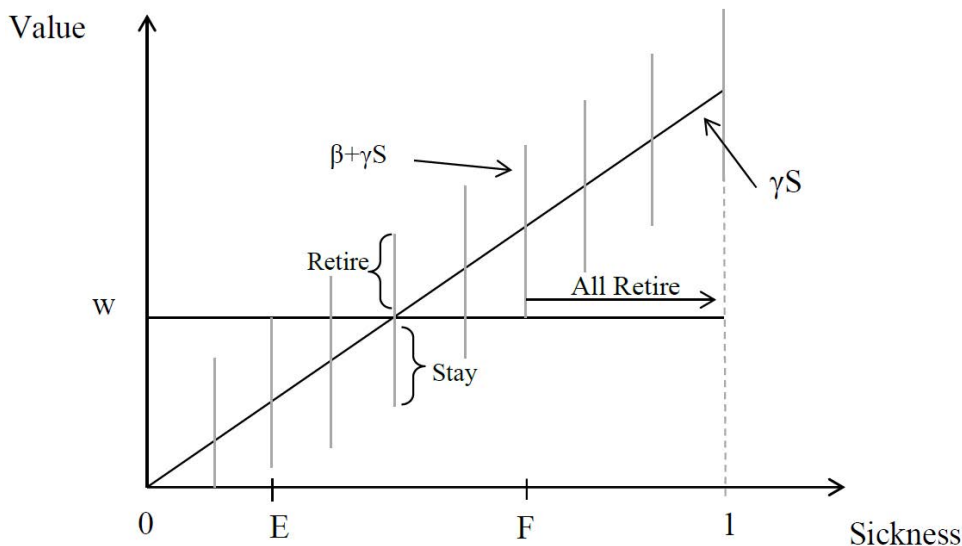


Figure 7: Retirement in a retiree contract with random draws of sickness and disutility.

In a retiree firm, workers retire if the disutility of work is more than the wage. All retiree workers with a sickness draw to the left of point E stay regardless of their β -draw; all workers with a sickness draw to the right of point F retire regardless of their β -draw; those with sickness draws between points E and F make their retirement decision depending on both their S -draw and their β -draw (as depicted).

Figure 8 shows the effect of the tied contract. The perceived wage now is $w + hS$, which is depicted by the upward-sloping, dashed-line schedule. In comparison to retiree workers, it is apparent that more workers with sickly draws between points G and H now decide to stay. In particular, workers in this range with a β -draw that puts them on a vertical segment between the

flat retiree wage schedule, w , and the upward-sloping dashed-line compensation schedule postpone retirement because of the tie. Notably, while the tie encourages more relatively healthy workers to stay, it exerts a relatively weak effect on the the sickliest workers.¹⁷




Figure 8: Retirement in a tied contract with random draws of sickness and disutility.

¹⁷The empirical impact of the tie depends on a variety of parameters and functional forms. For example, if the cost function $C(S)$ was concave in S instead of linear, then many more healthy workers would stay owing to the tie compared to those who are sickly.

B Determining Retiree Insurance in HRS

In each wave of the HRS respondents are asked a question related to the availability of health insurance in retirement. The wording changes in wave 3. In waves 1 and 2 workers are asked “Is this health insurance available to people who retire?” In wave 3 and beyond, workers are asked “If you left your current employer now, could you continue this health insurance coverage up to the age of 65?” The alteration makes it much less useful for the purpose of identifying retiree and tied coverage.

Trends in Retiree Reporting

To illustrate the problem, I select all workers who affirm that they have retiree health insurance in wave 1, then track their answer to the retiree question in subsequent waves. Table 9 shows the results.

In waves 1 and 2, when the question is phrased identically, the answers are remarkably consistent. When the question changes in wave 3, retiree health coverage rates fall precipitously, then partially rebound in the subsequent waves (the question is only asked of workers under age 65, so sample sizes are quite small in the last few waves reported).

Table 9: Consistency of Retiree Insurance Responses

	Fraction responding ‘yes’ to retiree insurance question
Wave 1	1.00
Wave 2	0.98
Wave 3	0.63
Wave 4	0.71
Wave 5	0.73
Wave 6	0.84
Wave 7	0.84
Wave 8	0.92

It is clear that the change in wording makes a substantive difference in retiree rates, but why? Recall that the wave 3 question asks “If you *left your current employer now*, could you continue this health insurance coverage up to the age of 65?” The fact that this question focuses so clearly on current access to insurance is critical. In most firms, workers are not eligible for full retirement benefits (pension or health insurance) unless they have vested and reached a required minimum age. Workers typically are entitled to retiree health coverage only if they ‘retire’ and not if they ‘quit.’

Consider a worker at baseline, aged 50-55, who is eligible for full or reduced pension benefits

at age 60. Assume that he is covered by retiree health insurance at the age of eligibility for full or reduced pension benefits. If asked at baseline whether “this insurance is available to people who retire” (waves 1 and 2 version), the respondent would answer ‘yes.’ If asked whether he would be covered by his health insurance if he “left [his] employer now,” however, the answer is ‘no’ because he has not reached retirement eligibility (age 60 in this example). Fortunately, the data allow me to test this theory.

While the HRS does not ask about eligibility related to health insurance, it asks if they are eligible for pension benefits, which should serve as a good proxy for retiree health eligibility. So, if I restrict the sample to pension-covered workers, I can test whether eligibility is driving the change in coverage answers starting in wave 3.

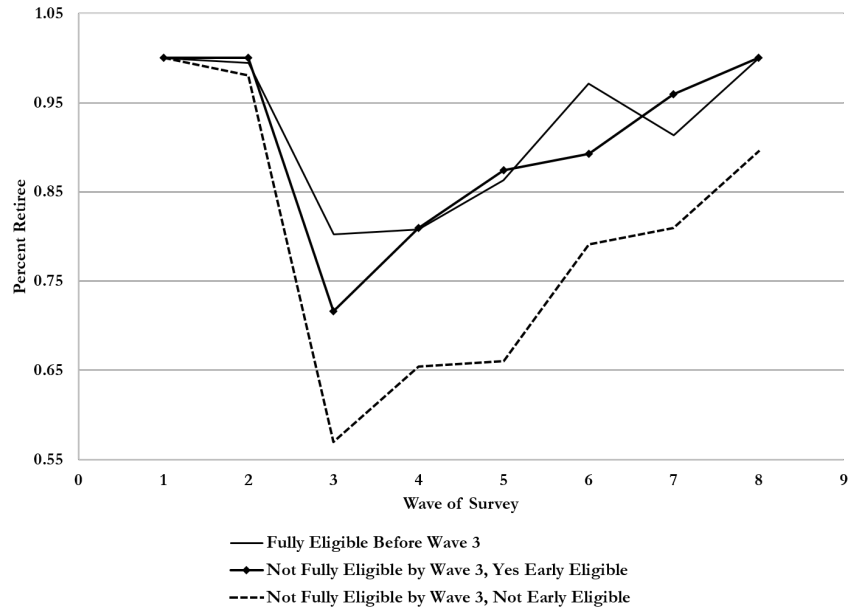
Workers are asked at what age they could first start to receive pension benefits, and when they would be eligible for full, unreduced, pension benefits. Retiring in the period between first eligibility and full eligibility means that the worker accepts some reduction in pension benefits.

I split retiree workers into three groups: (1) Those fully eligible for their pension by wave 3, (2) workers not fully eligible, but meet early eligibility at wave 3, and (3) workers neither fully nor partially eligible at wave 3.¹⁸ If workers are interpreting this question as I suspect, then the drop in reporting retiree insurance should be considerably more pronounced in the group who is not yet eligible for a reduced or full pension.

Figure 9 illustrates the patterns for the three groups. While the affirmative answer to the retiree health coverage question falls in wave 3 for all groups, it falls disproportionately for those not yet eligible for (early or full) pension (and presumably for retiree health coverage). Over time, the portion of wave-1 ineligible workers that become eligible increases because all are getting older and gaining more service, which explains why rates for this group gradually rebound over subsequent waves.

¹⁸I assume workers are early retiree eligible if “age at first eligible for pension benefit” < Age in wave 3 < “age fully eligible for pension benefits.”

Figure 9: Consistency of Retiree Self-Reporting Across Waves by Full Pension Eligibility Age



In short, while the retiree questions in wave 3 and beyond (and for cohorts entering the HRS after wave 3) may be useful for other purposes, it is ill suited to determine whether or not a worker is covered by retiree or tied coverage. This is why I focus my analysis on the original HRS cohort only, and why I rely on the wave 1 retiree insurance question.

C Regression for Creating Sickness Index

In this section I construct a sickness index that reflects projected health costs, conditional on respondents' health and other characteristics at baseline. Specifically, for every 55 year old in the the wave-one cohort, I track hospital stays for these workers (and their spouses) over subsequent waves through age 64. Hospital stays are easily reported and reflect high cost health care usage, so I use them as a proxy of overall medical expenses. This gives me ten years over which to count hospital stays. I then can sum the total number of nights in the hospital over those ten years, creating a variables S_{55-64} for each respondent.

I regress this measure against an entire set of baseline characteristics. The results of this regression are reported in Table 9, below. Using these values, I calculate a predicted value of \hat{S}_{55-64}^i for each worker in a tied or retiree firm age 55. This gives me a measure of expected health expenses in the years prior to Medicare eligibility, given respondents' baseline current characteristics.

This approach assumes that respondents can estimate future medical costs based on observations that they have made over time of a whole sequence of older people that they have observed before they, themselves, reach these ages; that is, they have some estimate of the effects of smoking, obesity, etc. on health costs over the ages 55-65. The measure has the advantage of incorporating more variables in respondents projected costs, and allowing the weights on medical conditions to vary, though it incorporates much more estimation error than relying simply on subjective health assessment at age 55. It turns out that either approach gives the same qualitative result.

Table 10: Regressing future hospital stays between age 55-65 against baseline characteristics

	Estimate	Standard Error
Self-reported Health: Poor	0.095	(0.087)
Self-reported Health: Excellent	-0.117*	(0.054)
Health Worse Than Last Year	0.058	(0.092)
Has Limitation that Interferes With Work	0.053	(0.100)
Can Climb Stairs Easily	-0.054	(0.051)
Can Jog a Mile	-0.140**	(0.049)
Obese	0.130*	(0.065)
Overweight	0.061	(0.053)
Smoker	0.185***	(0.054)
Drinks per Day: 1	0.018	(0.051)
Drinks per Day: 2	-0.079	(0.081)
Drinks per Day: 3	-0.147	(0.124)
Drinks per Day: 4	0.045	(0.202)
High Blood Pressure	0.296***	(0.060)
2+ Nights in Hospital in Past Year	0.461***	(0.088)
More than 7 Doctors Apts in Past Year	0.288***	(0.061)
Ever Applied for SSDI	0.333	(0.251)
Diabetes	0.332***	(0.098)
Cancer	0.428***	(0.120)
Lung Disease	0.005	(0.111)
Heart Disease	0.192*	(0.090)
Congestive Heart Failure	1.225**	(0.394)
Cholesterol	-0.100	(0.057)
Stroke	-0.022	(0.241)
Arthritis	0.061	(0.052)
Pain	0.021	(0.093)
Psychiatric Problem	0.123	(0.091)
Ulcer	0.139	(0.086)
Female	-0.222***	(0.050)
Parent Died Before age 65	0.0452	(0.046)
Household Income	-0.0005	(0.0004)
Has Spouse	-0.044	(0.059)
<i>N</i>	2345	

Note: This table reports the results of regression future hospital stays against baseline characteristics of individuals. Standard errors in parentheses, clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$