Job Lock, Retirement, and Dependent Health Insurance: Evidence from the Affordable Care Act *

Maggie Shi †

January 14, 2019

Abstract

The 2010 Affordable Care Act introduced a mandate expanding dependent health insurance coverage to young adults up to age 26. I explore whether this expansion induced job lock for parents of eligible young adults, since many parents' health insurance is tied to employment. Using a difference-in-difference strategy comparing parents of children above and below the age cutoff, I find that the mandate reduced parents' retirement rate by 2.9 percentage points, causing them to delay retirement up to 1.7 years. Early retirees are very responsive, as well as individuals with Social Security eligible spouses, suggesting the mandate interacted with Social Security.

*I am grateful to Wojciech Kopczuk and Michael Best for their guidance and support. I also thank Francois Gerard, Adam Sacarny, Day Manoli, and Columbia University Applied Microeconomics colloquium participants for insightful comments and suggestions.

†Columbia University. Email: m.shi@columbia.edu

1 Introduction

For many Americans, health insurance coverage is inextricably tied to work. According to the Current Population Survey, employer-based health insurance covered 55.7 percent of the population for some or all of the 2016 calendar year (Barnett and Berchick 2017). Tying benefits to jobs may introduce inefficiencies in the labor market. Workers value health insurance differentially, but it is hard for employers to tailor compensation at an individual level to account for this variation. As a result, workers' labor supply and job mobility may be distorted if a worker stays at a job because she values the job's employer-based health insurance highly. This is a phenomenon known as "job lock."

Job lock poses a problem in many dimensions of labor supply. For one, it may limit job mobility. Instead of picking the job which maximizes her marginal productivity, a worker may pick a job in which she is less productive but which offers health insurance. Thus, job lock prevents optimal matches between firms and workers. Not only can healthcare benefits lock workers into one job versus another, but they can also lock workers into the labor force itself. Health insurance is an important factor in the retirement decision, since older workers are more likely to experience negative health shocks. Even if a worker would rather retire earlier, she may be too risk-averse to expose herself to a gap in insurance coverage before she is eligible for Medicare, and thus will delay retirement until 65. To add another dimension of complexity to the issue, many Americans are not insured through their own job, but rather through a spouse or parent's job. So job lock may be exacerbated when many individuals' insurance coverage is tied to a single job.

Identifying the existence and extent of job lock is an empirical challenge, since individuals with and without employer-sponsored insurance likely differ in many dimensions. So, simply comparing job mobility and labor supply of these two groups is insufficient. There is an extensive literature, discussed in Section 2, studying job lock and specifically its effect on retirement decisions. This paper contributes to this literature by using a policy change affecting older workers which exogenously increased the value of employer-sponsored health insurance for some workers but not others.

One of the most popular and well-publicized components of the 2010 Affordable Care Act (ACA) was the mandate to expand health insurance coverage for young adult dependents. The dependent coverage mandate required insurers offering dependent coverage to cover adult children up to age 26. As a result, many parents' employer-sponsored insurance plans suddenly expanded coverage to their young adult children. This policy change presents a quasi-experimental setting to study how employer-provided health benefits interact with these parents' labor supply. Specifically, by comparing parents of children older and younger than 26 in 2010, I leverage the eligibility requirements of the mandate to identify the causal effect of dependent health insurance on parental labor supply. Given that most of these parents were in their 50s or 60s, this provides an opportunity to study the effect of employer-provided health insurance on retirement behavior.

In this paper, I use a difference-in-difference strategy to find that expanding dependent insurance decreased the retirement rate for affected parents by 2.9 percentage points, or relative to the average post-policy rate had the mandate not been enacted. Comparing retirement profiles for the two groups, this translates into a delay in retirement of up to 1.7 years. Early retirees are particularly responsive – I find larger effects for parents too young to qualify for Social Security retirement benefits and parents who qualified for early Social Security (i.e., turn 62) in the post-policy period. Since only one parent has to delay retirement to take advantage of the mandate, I also look at how couples decide which parent delays. I find that age and Social Security are deciding factors; the younger parent is more likely to delay retirement, especially if the older parent qualifies for Social Security. I do not find any effect of the mandate on other labor force outcomes such as mobility, switching to part-time work, or involuntary job losses.

Before the ACA mandate, insurance plans often only covered dependents up to age 19 or until they were out of school. Some individual states had their own dependent coverage mandates, but eligibility requirements varied widely (e.g., by dependent's marital status, student status, or employer-sponsored insurance) and almost all did not require coverage up to age 26 (see Depew 2015, Table 1 for an overview of state policies). In contrast, the ACA dependent coverage mandate was a national policy with no eligibility requirements besides age. As a result, any young adult under 26 whose parent's health insurance covered dependents could obtain coverage under her parent's plan. So for dependents who were added to their parent's insurance after the ACA, the mandate implicitly tied their health insurance to their parents' job.

The policy was announced in March 2010 and insurers were required to comply by September 2010. The implementation was staggered as some insurers chose to enroll before the September deadline. But by September 2010, almost all insurers had complied. In March 2010, the Internal Revenue Service also amended its rules to allow health benefits for dependents to be tax-exempt up until age 27 (Internal Revenue Service 2010). The dependent mandate had a significant impact on insurance rates among young adults. Antwi, et al. (2013) estimate that about 2 million young adults added parental employer-sponsored insurance as a result of the mandate. In 2016, 26 year olds were 1.19 times more likely to be uninsured than 25 year olds (Barnett and Berchick 2017). In this paper, I study the impact of the mandate on another population: the parents of these young adults.

The paper is organized as follows: in Section 2 I review the related literature. In Section 3, I lay out a conceptual framework. In Section 4, I describe data and methodology. I present my results in Section 5, and conclude in Section 6.

2 Literature Review

Much of the empirical literature on employer-sponsored insurance and retirement looks at the availability of *retirement* insurance and thus suffers from endogeneity issues, as noted by Gruber and Madrian (2004). The most convincing studies use the Consolidated Omnibus Budget Reconciliation Act of 1985 (COBRA). COBRA was a national law which gave employees access to continuation coverage under employer-sponsored health insurance for up to 18 months after leaving a job. Under COBRA, former employees pay out of pocket for their own coverage after leaving, but they have access to the employer's lower group rate. Gruber and Madrian (1995) use variation in pre-COBRA continuation mandates across states to identify the effect of the national policy, which in most states expanded continuation coverage. They find that the retirement rate increased by 28 percent in states without mandates, whereas it stayed about constant in ones with mandates. In a following paper, Gruber and Madrian (1997) find that access to continuation coverage raises the odds of a job transition for working-age men. The identification strategy of COBRA papers relies on exogenous changes in the price of self-insuring, whereas my strategy relies on changes in the value of employer-sponsored insurance itself to individuals. Additionally, the identification strategy of COBRA studies relies on variation across states, whereas in my policy setting I can compare across *individual* workers and thus control for individual fixed effects.

There has also been substantial work on job lock for low-income mothers and secondary earners, who are not the focus on this paper; see Gruber and Madrian (2004) for a detailed review of the literature.

Another strand of the job lock and retirement literature uses structural models to estimate the relationship between retirement and health insurance. Using a dynamic programming model of retirement, Rust and Phelan (1997) showed that retiree health insurance decreases the probability of working full-time by up to 20 percent for older workers. French and Jones (2011) estimate a life cycle model using the rich Health and Retirement Study data and incorporate consumption smoothing through saving. They find that raising the Medicare age from 65 to 67 slightly increases employment probability for individuals aged 60-69. While not directly comparable to these results, my findings line up with these structural studies in that I also find that employer-sponsored insurance significantly affects retirement behavior.

To my knowledge, this is the first study looking at the effect of the ACA dependent coverage mandate on dependents' parents. There is already a substantial literature documenting how it affected young adults. Antwi et al. (2013) found that the mandate increased dependent coverage and decreased individual purchased/own-employer coverage. The evidence on job lock for young adults is mixed. Bailey and Chorniy (2015) find that the mandate did not increase job mobility among dependents. Antwi et al. (2013) find that the mandate decreased the probability of working full time and the number of hours worked, while Heim et al, (2017) report that employment and self-employment were largely unaffected. The ACA mandate improved health in certain dimensions like body mass index, self-assessed health, and cancer detection (Barbaresco, Courtemanche, and Qi 2015; Robbins et al. 2015). It also reduced debt from healthcare expenditures, delinquencies, and bankruptcy (Blascak and Mikhed 2018).

3 Conceptual Framework

Next, I lay out a simple conceptual framework based on Gruber and Madrian (2004) to illustrate how the dependent mandate could induce job lock for potential retirees. I model the decision to work or retire at a given age rather than the decision to switch jobs because, as shown later in my results, I find that the mandate had no effect on job switches. Additionally, the average age of parents in my sample was 58 at the end of the panel, so many were on the cusp of retirement.

Assume that workers with inelastic labor supply decide between two states: working with insurance, or not working without insurance. For simplicity, assume that health insurance is tied to working (i.e., there is no retirement insurance, Medicaid, or private insurance). Let utility at age t be $U_t(w, H, L)$, where w is the wage at the current job, H is a dummy variable for health insurance coverage, and L is inelastic leisure. Utility is increasing in all three arguments. Since health insurance is tied to working, define $H = \mathbb{1}\{L = 0\}$. w and L are related as follows:

$$\begin{cases} w > 0 & L = 0 \\ w = 0 & L = 1 \end{cases}$$

Since I eliminate the possibility of working at a job without health insurance, I sidestep modeling the employer's decision of whether to offer insurance. Each year, individuals simply compare working with health insurance, $U_t(w, 1, 0)$, to not working without health insurance, $U_t(0, 0, 1)$. If $U_t(w, 1, 0) \geq U_t(0, 0, 1)$, then they work this year.

Now, introduce a reform which changes the utility function as follows:

$$U'_t(w, H, L) = \begin{cases} U_t(w + B - CD, H, L) & child \leq 26 \\ U_t(w, H, L) & child > 26 \textit{ or no child} \end{cases}$$

where $B \geq 0$ represents the value of dependent health insurance to a parent and $CD \geq 0$ is the compensating differential by which an employer reduces the parent's compensation since she now values employer-sponsored insurance more. If the wage does not change ($CD = 0$), the reform increases the value of continuing to work.

In the case of no job lock, firms set CD such that $CD = B$ for each individual. Since workers now value employer-sponsored health insurance more, firms can lower their wage until their utility is back to the pre-reform level and capture this rent. As a result, affected workers' labor supply will not change and the reform will not induce job lock.

In reality, firms generally cannot tailor compensation packages on an individual basis, and thus cannot set $CD = B$. In the extreme case, assume $CD = 0$, meaning an individual's wage doesn't change at all after the reform. This is a plausible assumption since there are legal and logistical constraints on reducing compensation only for those parents affected by the mandate. When $CD = 0$, parents compare $U_t(w + B, 1, 0)$ to $U_t(0, 0, 1)$. If there is heterogeneity in the value of dependent insurance across parents, there will be individuals for whom previously $U_t(w, 1, 0) < U_t(0, 0, 1)$, but now $U_t(w + B, 1, 0) > U_t(0, 0, 1)$. Put another away, this means that absent the policy these individuals would have retired at age t, but now they continue to work. Thus, the mandate will induce job lock for these individuals.

This framework can incorporate Social Security retirement benefits by introducing another option: $U_t(S, 0, 1)$ if $t \geq A$, where S is the Social Security benefit and A is the eligibility age.[1] Since utility is increasing in wage, $U_t(S, 0, 1) > U_t(0, 0, 1)$. Assuming $CD = 0$, an SS-eligible individual

[1] If $A = 62$, then this simple framework captures how individuals first become eligible for early Social Security at age 62. But of course this is an extremely simplified version of Social Security, since S is actually a function of age t and also varies by an individual's work history and marital status.

now compares $U_t(S,0,1)$ to $U_t(w+B,1,0)$. If $U_t(S,0,1) >> U_t(0,0,1)$, we could observe individuals working while $t < A$ because $U_t(w+B,1,0) > U_t(w,1,0)$, but then retiring at age A since $U_t(w+B,1,0) < U_t(S,0,1)$ for $t \geq A$. In other words, individuals delay only up until they reach Social Security eligibility, at which point they retire and take their Social Security benefits[2]. However if $U_t(w+B,1,0) > U_t(S,0,1)$, then the mandate could cause some individuals to forgo Social Security.

4 Data and Identification Strategy

The data source I use is the 2008 panel of the Survey of Income and Program Participation (SIPP). SIPP surveys a nationally representative sample of 42,000 American households for about 4 years, asking them to recall the past four months. The 2008 panel covers August 2008 to December 2013, which spans the pre- and post-policy period of the ACA. It collects data on demographics, employment status, assets and earnings, health and disability, government program participation, and job benefits. SIPP provides a rich monthly snapshot of work history during this period for every individual in a household.

To define my treatment and control group, I will use the age of a respondent's youngest child. I define "treated" to mean that the respondent's youngest child is less than 26 in 2010, and "control" if the youngest child is older than 26. In order to have relatively comparable groups, I restrict my treatment group to parents of 23-25 year olds and the control group to parents of 27-29 year olds. It is unclear whether children who were 26 in 2010 would be eligible or not, since insurance companies varied in when they began to comply with the mandate. So, I do not assign parents of 26

[2]While individuals can technically claim SS while still working, they are subject to an earnings test. In 2018, Social Security benefits are reduced \$1 for every \$2 earned above \$17,040 (\$1420 a month).

year olds to either treatment or control. This leaves 1372 individuals in the treatment group and 1297 in control, with an average of 59.5 months of observations per individual (out of 64 total possible months). In Table 1, I report summary statistics of individual characteristics by treatment status.

Table 1: Summary statistics by treatment and control

Characteristic	**Type**	**Treat (23-25)**	**Control (27-29)**
White	percent	0.84	0.83
Female	percent	0.64	0.63
Hispanic	percent	0.09	0.07
Married	percent	0.78	0.78
High School	percent	0.89	0.90
College	percent	0.3	0.3
Private HI in 2009	percent	0.76	0.76
Employed in 2009	percent	0.74	0.70
Age in 2009	mean	53.75	56.31
Number of children	mean	2.30	2.36
Monthly income in 2009 (>0)	mean	$2970	$2556
Observations	#	72747	71580
Individuals	#	1372	1293
Months Observed	mean	59.28	59.79

While the treatment and control groups are relatively similar in most dimensions, it is important to note that the control group is older than the treatment group. This makes sense, since the control group consists of parents of older children. Age is an important factor in the retirement decision. Thus, the baseline retirement rate for the control group should be higher than the treatment group. However, in order for identification with a difference-in-difference strategy, we do not need that the two groups have the same baseline retirement rate – we just need that their retirement rates follow parallel trends (at this point in their retirement profiles), which I check for in the next section. Age may also play a role in other labor force decisions, but the same logic of parallel trends applies.

Something else to note is that the sample contains more women than men. Women comprise 64 percent and 63 percent of the treatment and control groups, respectively. This is because I do not directly observe every adult dependent's age, and must instead use mothers' fertility history. SIPP has limited information on family members *not* currently living in the household, which is likely the case for most adult dependents. But using a woman's fertility history, I can consider the ages of the children a woman gave birth to.[3] I then link fathers (or step-fathers) in through the mother's current husband. Thus, I only add married men whose wives were also in the sample, which explains the gender imbalance. Although men and women in this sample are not directly comparable to each other, the treatment and control groups are equally unbalanced in terms of gender ratio and thus are comparable to each other.

I use responses only from individuals with valid interview status and nonmissing identifiers. I consider whether a respondent has a job in the past month, whether she is looking for a job, and her reason for not having a job. I consider a respondent to be retired if her employment status for a given month is "no job all month, no time on layoff and no time looking for work" and the main reason for not having a job in the reference period (last four months) is retirement. Besides retirement, I also look at unemployment, non-retirement labor force exits, and changes in employer without leaving the labor force.

I use a difference-in-difference identification strategy to estimate the causal effect of the mandate, which relies on the parallel trends assumption: absent the policy, treatment and control group outcomes would be parallel to each other over this time period.[4] In Figure 1, I plot the retire-

[3]SIPP only asks for the oldest and youngest children that a woman gave birth to. Thus I also have no sense of how many eligible or ineligible children the respondent has, although I do know the total number of children. However, these characteristics are fixed over time and will be absorbed by individual fixed effects. I also do not observe ages of adopted children or stepchildren, who would both count as dependents for insurance purposes.

[4]Since retirement is an absorbing state, it would impossible for the retirement rate

ment rate per month for the treated and control groups. Similar figures for other outcomes like probability of not working and probability of changing employers are included in the appendix (Figures A.1, A.2, and A.3). Visually, the two trends seem parallel until after policy enactment (September 2010), when the retirement rate for the treated group flattens out, but the rate for the control group continues to climb. The gap in retirement rate between the two groups widens over time, even as some children in the treated group "age out" of eligibility (i.e., turn 26).

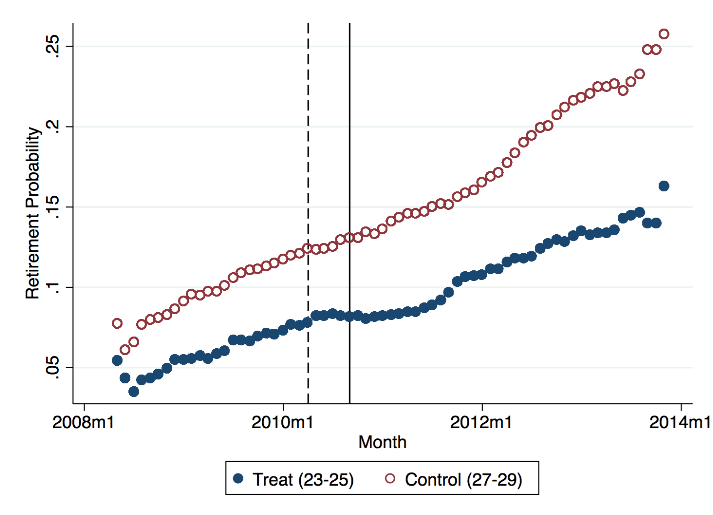

Figure 1: Monthly retirement rates for treated (parent of 23-25 in 2010) and control (parent of 27-29) groups. Dashed vertical line represents policy announcement in March 2010; solid vertical line represents policy enactment September 2010.

If the parallel trends assumption is met, I can implement my main difference-in-difference specification:

$$y_{it} = \beta_0 + \beta_1 Treat_i \times Enact_t + \beta_2 Seam_{it} + \alpha_i + \lambda_t + \varepsilon_{it} \quad (1)$$

y_{it} is the outcome variable for individual i in month t. In the main

of the two groups to be parallel indefinitely. However, given the age range of the two groups and the evidence in the graph, it may be safe to assume parallel trends in *this time period.*

specification, $Enact_t$ is a dummy variable which equals 1 after September 2010, when the mandate was enacted for all insurance plans, and 0 before. $Treat_i$ is a dummy variable which is 1 if the respondent's youngest child is 23-25 in 2010 and 0 if the respondent's youngest child is 27-29. α_i is an individual fixed effect and λ_t is a month-year fixed effect. $Seam_{it}$ is an individual-specific, time-varying dummy which equals 1 during a "seam" month between reference periods during which the respondent is actually surveyed. In a seam month, the individual is asked to recall information about the reference period (past 4 months), and it is a well-known phenomenon in longitudinal surveys that there is a higher likelihood of information changes during a seam month (U.S. Census Bureau 2001). Robust standard errors are clustered at the individual level. I use a linear probability model for ease of interpretation and because retirement rates are never "too close" to 0 or 1. None of the predicted retirement values from the main specification fall outside of (0,1)[5]. Later, I run this specification on a variety of subsamples to see which subsample is driving the response.

β_1 is the parameter of interest: it represents the average causal effect of dependent coverage expansion on the probability of exiting the labor force, retiring, or changing jobs. A negative β_1 means that the treated group is less likely to exit after the policy. With regard to retirement, it would mean that the treated group delays retirement after the policy is enacted.

One concern with this specification is that treatment is assigned based on *potential* eligibility through the child's age, rather than actual eligibility (i.e., whether the respondent had insurance which covered dependents). So the group of respondents actually receiving treatment is a strict subset of what I define as the treatment group. We do not have to worry about this in the control group. Since the mandate has an explicit eligibility cutoff at 26, no respondent in the control group is ever eligible for the dependent

[5]The predicted values from the main specification range from 0.054 to 0.205.

coverage expansion. As I am estimating an intent-to-treat effect, we should expect that, if anything, the estimate for β_1 is biased toward zero.

I also run more flexible specifications to see how β_1 varies over time and with the child's age. I replace $Enact_t$ with dummies for each month. $Month_t$ is a month-year dummy (equivalent to λ_t above):

$$y_{it} = \beta_0 + \sum_{t=0}^{T} \beta_{1t} Treat_i \times Month_t + \beta_2 Seam_{it} + \alpha_i + Month_t + \varepsilon_{it} \quad (2)$$

With this specification, β_{1t} represents the difference in labor force outcomes between treated and control groups in each month. If the parallel trends assumption holds, we should see that $\beta_{1t} = 0$ before policy enactment/announcement. While we expect to see a gap between treated and control emerge after September 2010, it is unclear whether it should increase or decrease in magnitude over time. On one hand, we expect it to increase as respondents' labor supply decisions made in 2010 are realized over time. On the other hand, it may decrease as children of the treated group "age out" of eligibility over the years.

We might also expect that β_1 varies with the age of the child. For a parent whose child is 20 in 2010, the opportunity cost of exiting the labor force in 2010 includes six potential years of dependent insurance; for a parent of a 25 year old, the opportunity cost includes only one potential year of dependent insurance. Thus, we should see that on average, parents of younger children are more attached to the labor force than parents of older (but still eligible) children. I explore this relationship with the following specification:

$$y_{it} = \beta_0 + \sum_{a=19}^{26} \beta_{1a} (ChildAge_a \times Enact_t) + \beta_2 Seam_{it} + \alpha_i + \lambda_t + \varepsilon_{it} \quad (3)$$

$ChildAge_a$ is a dummy variable which is 1 if the respondent's child is age a in 2010. β_{1a} is the causal effect of the mandate for parents of children of age a. I consider ages 19-26 because many insurance plans covered dependents only until age 19. The comparison group is parents with children aged 27-29 in 2010.

5 Results

5.1 Main Results

Table 2 presents the main difference-in-difference results from Equation 1 on various employment outcomes. Column 1 reports the result on a dummy variable for not having a job in a given month. The policy reduced the likelihood of *not* having a job by about 1.9 percentage points. There seems to be no effect on not having a job but looking for one (column 2). Instead, the response is driven by job exits after which the individual is no longer looking for a job (column 3). Put another way, we see an effect on labor force exits rather than unemployment. This is expected since job lock should not affect involuntary job losses which would contribute to unemployment. Breaking this down even further, the effect seems to be coming from reductions in retirement (column 5) rather than other labor force exits (column 4). From columns 6 and 7, we see that there is also no discernible effect on partial exits (i.e. having a job for some but not all weeks in a month) or employer changes. Therefore, the result in column 1 is driven entirely by a reduction in retirement. The policy reduced the likelihood of retirement by 2.9 percentage points. The observed average retirement rate after policy enactment for the treated group was 16.7 percent (although it increases steadily over time). So, the dependent coverage mandate reduced the retirement rate by about 14.8 percent (2.9/(2.9+16.7)) relative to the post-policy rate had the policy not been enacted. In comparison, Gruber and Madrian (1995) find that one year of continuation coverage *increased* the retirement rate by 1.1 percentage points (5.4 percent). As I will discuss in Section 6, the 2.9 percentage point reduction in retirement rates translates into an average retirement delay of 0.74 years (and up to 1.7 for some individuals).

Table 2: Difference-in-difference results from Equation 1

	(1) No job	(2) No job, looking	(3) No job, not looking	(4) Not looking, not retired	(5) Not looking, retired	(6) Partial LF exit	(7) Change employers
$Treat_i \times Enact_t$	-0.019*	0.003	-0.022**	0.007	-0.029***	-0.000	-0.000
	(0.011)	(0.006)	(0.010)	(0.009)	(0.009)	(0.001)	(0.002)
Seam	X	X	X	X	X	X	X
Indiv. FE	X	X	X	X	X	X	X
Month-year effect	X	X	X	X	X	X	X
Individuals	2678	2678	2678	2678	2678	2678	1900
Observations	144327	144327	144327	144327	144327	144327	92522

Standard errors in parentheses
(1) Outcome: No job (for at least some week of month)
(2) No job, looking for work
(3) No job, not looking for work
(4) No job, not looking, but not retired
(5) No job, not looking, and retired
(6) Partial exit (no job for part of but not entire month)
(7) Changed employers
* $p < .10$, ** $p < .05$, *** $p < .01$

In Table 3, I explore the retirement result further. In columns 2 and 3 I check that the results are robust to using $Enact_t$ (September 2010) and $Announce_t$ (March 2010) as the post-policy period. In column 3, the coefficient on $Treat_i \times Enact_t$ is significant and negative, while the coefficient on $Treat_i \times Announce_t$ is insignificant. Parents' retirement behavior diverged only when the policy was implemented by insurance companies, not when it was announced by the government. Thus, I use $Enact_t$ to separate the pre- and post-policy periods in the following analysis, although it is also robust to using $Announce_t$.

In column 4, I re-run the main specification on an alternative definition of retirement: not working because of retirement *or* because of a chronic health condition. We may worry that some older workers may not officially "retire" (i.e., state the reason they are not working as retirement), but stop working with no intention of returning to the workforce because of a chronic health condition. Chronic conditions are likely to lower productivity without completely preventing workers from working, unlike a short-term injury or illness. In this sense, they are similar to aging in that workers have some flexibility in deciding when to stop working. In both cases, workers weigh the benefits of working with reduced productivity against the cost of health deterioration over time. Using retired or chronic health condition as the outcome variable, I find a similar significant negative coefficient for β_1. For the rest of the analysis, I continue to use the main definition of retirement since adding chronic illness into the definition likely adds noise to the outcome variable, as seen by the slightly less significant coefficient in column 4 compared to the main result in column 1.

It is of note that the only labor force margin on which parents seem to respond is retirement. One potential reason could be that the average age of individuals considered is relatively high (about 54), and by this point in their careers, they are less likely to have a non-retirement labor force exit or a change in employers. This study differs from previous studies which found an effect of health coverage on such outcomes because previous studies focused on younger samples (Bansak and Raphael 2008; Gruber and B. Madrian 1997).

Figure 2 plots the β_{1t} coefficient from the flexible monthly specification in Equation 2. Prior to policy announcement and enactment, there is no significant difference between treatment and control groups, which supports the parallel trends assumption. After the mandate is introduced, a gap emerges between the two groups, which widens over time.

Table 3: Retirement difference-in-difference results from Equation 1

	(1) Retired	(2) Retired	(3) Retired	(4) Retired or chronic health
$Treat_i \times Enact_t$	-0.029***		-0.029***	-0.022**
	(0.009)		(0.008)	(0.010)
$Treat_i \times Announce_t$		-0.025***	-0.001	
		(0.008)	(0.006)	
Seam	X	X	X	X
Indiv. FE	X	X	X	X
Month-year effect	X	X	X	X
Individuals	2678	2678	2678	2678
Observations	144327	144327	144327	144327

Standard errors in parentheses
(1) Enactment, Sep. 2010
(2) Announcement, Mar. 2010
(3) Announcement and enactment
(4) Outcome variable: retired or disabled
* $p < .10$, ** $p < .05$, *** $p < .01$

Figure 2: Results from specification in Equation 2 using parents of 23-25 year olds as treated group. Coefficient on interaction of $Treat_i$ and $Month_t$. Dashed vertical line represents policy announcement in March 2010; solid vertical line represents policy enactment September 2010. Confidence intervals are 95%.

Figure 3 shows the results from a specification with interactions with child's age dummies, where the omitted group consists of parents of 27-29 year olds. As discussed in Section 4, we might expect heterogeneous treatment effects for parents of younger versus older (but still eligible) children. In general, as the age of the parent's youngest child gets closer to the cutoff of 26, β_{1a} decreases in magnitude. This is consistent with the idea that retiring while a child is younger has a higher opportunity cost for parents. Below age 25, the coefficients are significant and negative. The coefficients for 25-26 year olds are insignificant. However, the point estimates for parents of 19-24 year olds are not significantly different from each other.

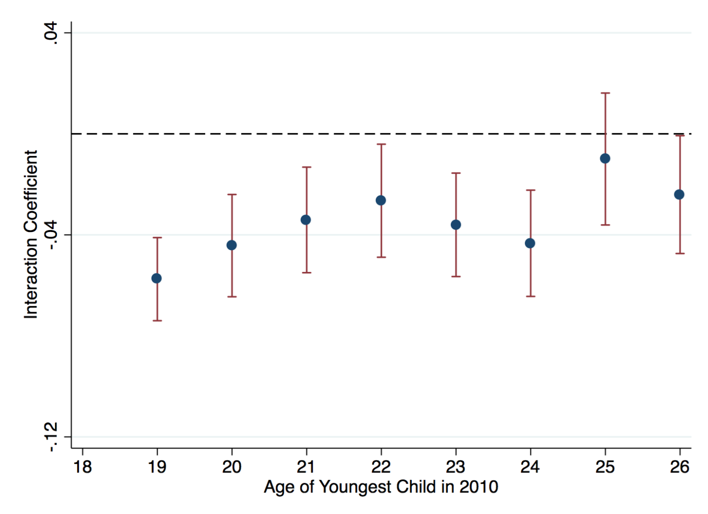

Figure 3: Coefficient on interaction of $Treat_i$ and $ChildAge_a$. Results from specification in Equation 3. Confidence intervals are 95%.

5.2 Placebo and Robustness Checks

In Table 4, I report results from running the main specification on four groups which, due to the specifics of the mandate, should be unaffected by the policy. One natural placebo group consists of parents of older children, like parents of 30-32 year olds (column 1). Comparing this placebo group to the original control group (parents of 27-29 year olds), we see that the coefficient on $Placebo_i \times Enact_t$ is insignificant. I run the same specification on subsamples which should be unaffected by the policy, parents covered by Medicaid (column 2) and parents who did not have insurance before the policy (column 3), and find no effect. Finally, I consider a counterfactual policy in September 2008, assigning "treatment" to parents of 23-25 year olds and "control" to parents of 27-29 year olds in 2008. Unfortunately, this is not a perfectly clean counterfactual since the 2008 "treatment" group slightly overlaps with the actual 2010 treated group (parents of 23 year olds in 2008 would be parents of 25 year olds in 2010), and the 2008 "control" treated group also overlaps with the actual control group (parents of 27 year olds in 2008 would be parents of 29 year olds in 2010). Using 2009 as a counterfactual policy would exacerbate this problem further. The cleanest counterfactual would be to use 2007 or earlier, but unfortunately the SIPP panel begins in 2008. Thus, we should expect that the coefficient in column 4 is biased *away* from zero relative to a "cleaner" 2007 counterfactual. In the Figures A.6 and A.7, I include results from the flexible monthly specification for each outcome in Table 4. Figures A.4 and A.5 check for parallel trends.

One concern may be that the treatment and control groups are not comparable to each other because control parents are older, and thus they may approach early or full Social Security age at different rates. In Figure 4, I vary the bandwidth used to define treatment and control groups to check that my result is robust to this. The main specification uses a bandwidth of 3, defining treatment as parents of 23, 24, or 25 year olds. To make the treatment and control groups more comparable, I could have chosen to use a bandwidth of 2 (24-25 vs. 27-28) or even 1 (25 vs. 27), at the expense of power. As I increase the bandwidth, the treated and control groups become less comparable. In Figure 4, we see that standard errors decrease as the

Table 4: Placebo tests and counterfactual

Outcome variable: monthly retirement rate

	(1) Placebo	(2) Medicaid subsample	(3) Uninsured subsample	(4) 2008 counterfactual
$Placebo_i \times Enact_t$	-0.003			
	(0.011)			
$Treat_i \times Enact_t$		-0.021	-0.031	
		(0.037)	(0.021)	
$Treat_{2008,i} \times Enact_{2008,t}$				-0.015
				(0.009)
Seam	X	X	X	X
Indiv. FE	X	X	X	X
Month-year effect	X	X	X	X
Individuals	2459	348	412	2621
Observations	134960	9531	22616	142118

Standard errors in parentheses

(1) Placebo group: parents of 30-32 year olds

(2) Subsample: parents covered by Medicaid

(3) Subsample: parents without insurance in 2009

(4) Counterfactual policy in Sep. 2008. Treated are 23-25 in 2008, control are 27-29 in 2008.

* $p < .10$, ** $p < .05$, *** $p < .01$

bandwidth increases, but the point estimate is consistently negative and relatively stable. If the negative main result was entirely driven by the age difference between treatment and control, then as the bandwidth increases and the age difference increases, the coefficient should become increasingly negative. Instead, we see that it is relatively stable between bandwidths 2 and 5.

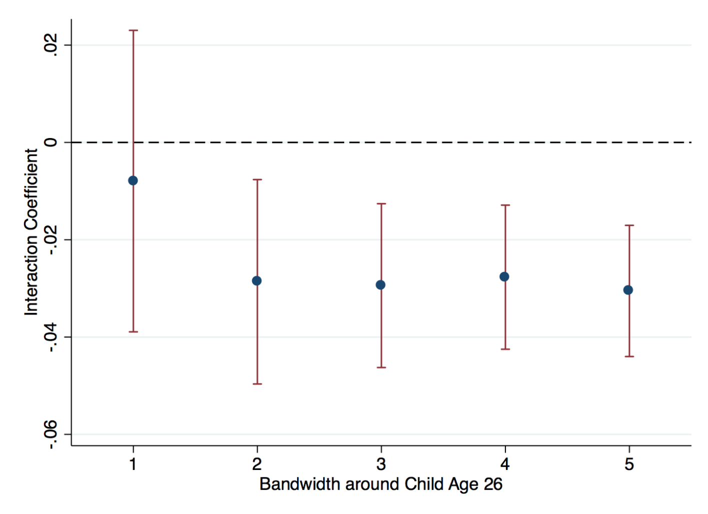

Figure 4: Coefficient on interaction of $Treat_i$ and $Enact_t$, by bandwidth for defining treatment and control groups. Confidence intervals are 95%.

5.3 Heterogeneity by Age and Social Security Eligibility

The marginal benefits and costs of retirement are shaped by a variety of incentives and preferences which vary over an individual's lifetime. In addition to the marginal earnings from continuing to work and the marginal costs of working as one ages, individuals must also consider incentives set by programs like Social Security. Social Security eligibility and benefit amounts are non-linear in age. Specifically, individuals gain early Social Security retirement benefits at age 62 and full benefits at age 66. It is a well-established empirical fact that individuals are more likely to retire at the early and full Social Security retirement eligibility ages (Rust and Phelan 1997). How does this change once the mandate is introduced, and is there heterogeneity depending on what Social Security benefits individuals qualify for?

I separate parents into groups depending on whether they cross either of these age thresholds (62 or 66) in the post-policy period, 2011-2013. Table 5 summarizes how I define Social Security cohorts: "Young" consists of parents who are younger than 62 in the panel, "Early" consists of parents who turn 62 in 2011-2013 and qualify for early SS, "Full" consists of parents who turn 66 in 2011-2013 and qualify for full SS, and "Old" consists of parents who turned 66 before 2011. Note that one group, parents who are 62 in 2010, does not fit neatly into any cohort as they turn 62 in 2010 and 66 in 2014. I also run a regression combining "Early" and "Full" together, and include this group as well.

Table 5: Social Security Cohort Definition

Age in 2010	Year turn 62	Year turn 66	Cohort
<59	**>2013**	>2017	Young
59	**2013**	2017	Early
60	**2012**	2016	Early
61	**2011**	2015	Early
62	2010	2014	
63	2009	**2013**	Full
64	2008	**2012**	Full
65	2007	**2011**	Full
>66	<2007	**<2011**	Old

I run the main specification on each Social Security cohort, and report results in Table 6. Column 4 combines the "Early" and "Full" cohorts, as well as parents who are 62 in 2010. The point estimates are negative for all columns, but only significant the "Young" and "Early" cohorts. The point estimate for "Early" individuals is substantially larger than that of "Young" individuals – almost 4 times larger.[6]

One question these results raise is, are the responsive "Early" individuals delaying until they reach early Social Security age at 62, or do they

[6]Alternatively, I define age cohorts by whether individuals turn 62 or 66 at some point in the panel (2008-2013), not just in the post-policy period (2011-2013). These results are reported in Table B.1 and are qualitatively similar. In A.8, I interact $Treat_i \times Policy_t$ with a dummy for parent's age in 2010. We see that in general, parents who were close to early retirement age were the most responsive, which lines up with the Social Security cohort results.

Table 6: Heterogeneity by SS age cohort

Outcome variable: monthly retirement rate

	(1) Young	(2) Early	(3) Full	(4) Early or Full	(5) Old
$Treat_i \times Enact_t$	-0.014*	-0.053**	-0.001	-0.020	-0.014
	(0.007)	(0.024)	(0.045)	(0.021)	(0.068)
Seam	X	X	X	X	X
Indiv. FE	X	X	X	X	X
Month-year effect	X	X	X	X	X
Individuals	1808	427	237	785	98
Observations	95216	23873	13001	43570	5541

Standard errors in parentheses

(1) Subsample: individuals who are $<$ 62 in 2013

(2) Subsample: individuals who turn 62 in 2011-2013

(3) Subsample: individuals who turn 66 in 2011-2013

(4) Subsample: individuals who turn 62 or 66 in 2011-2013

(5) Subsample: individuals who turned 66 before 2011

* $p < .10$, ** $p < .05$, *** $p < .01$

delay even further? If parents place high value on dependent health insurance, then they may be willing to forgo Social Security benefits to keep their insurance. In the context of the framework laid out in in Section 3, is $U_t(w + B, 1, 0) > U_t(S, 0, 1)$ when $t >= A$? In Figure 5, I plot the percent of treated and control parents who are retired at a given age in the post-policy period. If treated parents forgo Social Security for dependent health insurance, then we should see a gap in the retirement profiles of the two groups past age 62. However, if they delayed only up until they gain early Social Security, then we should see the gap close around 62. In Figure 5 we see that a gap emerges between the two groups between ages 58 and 62; treated parents in this age range are less likely to be retired than control parents. However, past 62 (the dashed line) the gap closes and retirement rates are similar across the two groups. This suggests that "Early" parents delayed up until they qualified for early Social Security, but not past that point.

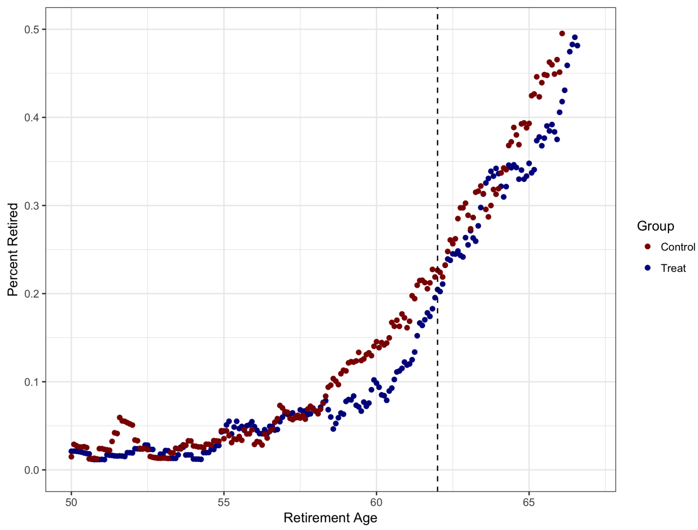

Figure 5: Percent retired by age in the post-policy period (2011-2013); dashed line at 62.

5.4 Heterogeneity Within Couples

From the perspective of individuals, the ACA dependent mandate effectively increased the benefit of continuing to work. From the perspective of couples, however, the mandate exogenously increased the benefit of *only one* of the two parents delaying retirement. Once a dependent is covered under one parent's insurance, there is no additional benefit to the other parent also delaying retirement (in terms of dependent coverage). Retirement is often a joint decision for married couples, as people take into account not only their individual circumstances, but also that of their spouse (e.g., Lee 2017). This leads to the question of how couples' joint retirement decision was affected by the mandate – if only one parent has to delay retirement for the dependent to benefit from the mandate, how did couples decide who would delay retirement?

I consider three factors which could affect this decision: gender, relative earnings, and relative age. Previous studies have found that women are more likely than men to spend transfers and savings on dependents (Duflo 2000; Ashraf, Karlan, and Yin 2010). We might expect to see the same phenomenon here: women may be more willing than men to adjust their labor supply to provide dependent health insurance. In this policy setting, delaying retirement to take advantage of dependent coverage is loosely analogous to spending (time, not money) on a dependent's health. In Table 7 column 1, I add an interaction with a dummy for female, $Female_i$, and restrict my analysis to married individuals. I find no statistically significant difference between married men and married women's responses.

The decision to retire depends crucially on earnings and age, so next I see if relative age and earnings *within* a couple affect which spouse responds to the mandate. I define the dummy variable $EarnMore_i$ to be 1 if individual i earns more than her spouse.[7] Relative earnings could matter because the higher-earner has greater bargaining power, but also simply because the higher earner derives greater returns to working. From the statistically insignificant coefficient on $Treat_i \times Enact_t \times EarnMore_i$ in column 2, we see that being the higher-earning spouse has no effect on retirement response.

I also define dummy variable $Older_i$, which is 1 if individual i is older than her spouse. The older spouse could have a higher marginal cost of working (say, if disutility from work increases with age or decreases with health), and she may also have a higher opportunity cost of continuing to work if she qualifies for Social Security benefits. In column 3, the coefficient on the interaction term $Treat_i \times Enact_t \times Older_t$ is positive and significant, meaning the older spouse is *less* likely to delay retirement (i.e., more likely to retire) than the younger spouse.

Older spouses may be less likely to delay retirement because they face higher disutility from continuing to work at an older age, or possibly because they qualify for Social Security while their younger spouse does not.

[7] I use average pre-policy earnings in 2009. Using post-policy earnings is unsuitable since the outcome variable is retirement, and retirement mechanically affects earnings.

Table 7: Regressions with within-household characteristic

Outcome variable: monthly retirement rate

	(1)	(2)	(3)
$Treat_i \times Enact_t$	-0.023*	-0.034***	-0.054***
	(0.012)	(0.013)	(0.011)
$Treat_i \times Enact_t \times Female_i$	-0.016		
	(0.013)		
$Treat_i \times Enact_t \times EarnMore_i$		-0.010	
		(0.012)	
$Treat_i \times Enact_t \times Older_i$			0.036***
			(0.013)
Seam	X	X	X
Indiv. FE	X	X	X
Month-year effect	X	X	X
Individuals	2148	1046	1046
Observations	112668	56414	56414

Standard errors in parentheses

(1)$Female_i$: 1 if female

(2) $EarnMore_i$: 1 if earned more than spouse in 2009 and both working.

(3) $Older_i$: 1 if older than spouse and both working in 2009.

* $p < .10$, ** $p < .05$, *** $p < .01$

As shown in Table 5, Social Security eligibility is an important factor in the retirement decision and size of the response. Next, I categorize individuals by their own (i) and their spouse's (s) age, relative to the early Social Security age of 62 (by 2013). Individuals can be in a marriage where both are too young to qualify for Social Security (column 1), where both are old enough to qualify by the end of the panel (column 2), where the individual is old enough but the spouse is not (column 3), and where the individual is not old enough but the spouse is (column 4). Note that the individuals in column 3 should be married to the individuals in column 4. I run the main specification on each group and report results in Table 8.

The most responsive subsample consists of individuals who themselves

Table 8: Couples with SS-eligible and SS-ineligible Spouse

Outcome variable: monthly retirement rate

	(1)	(2)	(3)	(4)
	$i < 62$	$i \geq 62$	$i \geq 62$	$i < 62$
	$s < 62$	$s \geq 62$	$s < 62$	$s \geq 62$
$Treat_i \times Enact_t$	-0.010	-0.017	-0.005	-0.048*
	(0.009)	(0.030)	(0.032)	(0.028)
Seam	X	X	X	X
Indiv. FE	X	X	X	X
Month-year effect	X	X	X	X
Individuals	1044	462	219	220
Observations	54820	26401	11634	11713

Standard errors in parentheses

(1) Subsample: individual and spouse ineligible for SS

(2) Subsample: individual and spouse eligible for SS by 2013

(3) Subsample: individual eligible for SS by 2013 but spouse ineligible

(4) Subsample: individual ineligible but spouse eligible by 2013

* $p < .10$, ** $p < .05$, *** $p < .01$

are too young to qualify, but are married to spouses who qualify for Social Security (column 4). Their spouses (column 3), in contrast, do not respond to the mandate. This coincides with the evidence from Table 7 that older spouses are *less* likely to delay retirement in response to the mandate, and the evidence from Figure 5 that parents did not forgo Social Security to take advantage of the mandate. Together, they suggest that some couples coordinated to take advantage of two public policies: the SS-eligible parent retires and claims Social Security, while the SS-ineligible parent delays retirement to hold on to dependent insurance.

6 Discussion

The main results show that the ACA dependent mandate significantly decreased the likelihood of retirement for eligible parents, especially those considering early retirement. Since almost all individuals will eventually

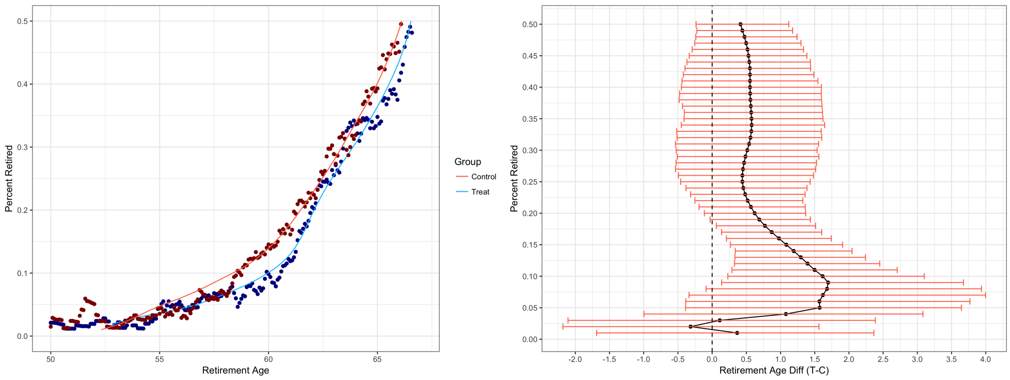

Figure 6: Percent retired by age in post-policy period and bootstrapped non-parametric regression fitted lines (n=1000) (left); Horizontal retirement age differences. Bootstrapped 90% confidence intervals (n=1000) (right)

retire, this decreased probability eventually translates into a *delay* in retirement age. In order to interpret the findings and benchmark them against other studies, I calculate the length of retirement delay implied by my results. One challenge is that retirement age is censored for many individuals because of the panel structure of the data. So, I cannot use retirement age as an outcome variable. Instead, I calculate the horizontal distance between the treated and control retirement profiles in Figure 5 at each quantile of retirement age, which corresponds to the percent retired at each age. In Figure 6, I non-parametrically estimate the retirement profiles for treated and control, and then predict the retirement age difference at each quantile of retirement age.

The difference in retirement age is positive for almost every quantile and statistically significant between the 9th and 18th quantiles. Between the 1st and 50th quantiles, the average difference between treated and control retirement ages was 0.74. The largest statistically significant difference was 1.7, at the 9th quantile. If we assume rank preservation[8], we can interpret this as a treatment effect at each quantile. This means that the mandate

[8]Meaning that an individual remains in the same retirement age quantile whether she is in treated or control.

caused individuals at the 9th quantile to delay retirement by 1.7 years. If we do not assume rank preservation, this still means that the mandate shifted the 9th quantile of retirement age up by 1.7 years.

If we assume rank preservation, then we can compare the treatment effect of this mandate to findings of other studies. Madrian et al. (1994) found that retirement health insurance availability led workers to retire 0.4-1.2 years earlier. French and Jones (2011) estimated that raising the Medicare age to 67 delayed retirement by only 0.07 years. Brown (2013) found that in response to a 10 percent increase in return to work, individuals delayed retirement by 0.17 years. So, my finding that some individuals delayed on average 0.74 years, and up to 1.7 years, in response to the dependent mandate is relatively large. One possible explanation could be that the dependent mandate was well-publicized and had a straightforward eligibility cutoff of 26. Over 70 percent of the public was aware of the provision within a month of enactment (Kaiser Family Foundation 2010). There is evidence of large "statutory age effects" on retirement which are independent of pure financial incentives, in line with a behavioral model of retirement where statutory ages serve as reference points (Seibold 2017). Thus, the cutoff of 26 could have been a salient "statutory age" around which parents planned their retirement.

These findings also raise additional questions, which are left for further work. While I show empirical evidence of how couples' behavior changes in response to the mandate, one could potentially combine this with a model of the joint retirement decision. This model would have to take into account the well-documented fact that couples tend to retire together. In the overall SIPP sample, we see that a large portion of couples retire within the same month, as shown in Figure A.9. While previous studies have modeled the joint retirement decision (Gustman and Steinmeier 2004; Gustman and Steinmeier 2000; Lee 2017), the quasi-experimental setting of the ACA dependent mandate presents an opportunity to get a better sense of how much a household values joint retirement.

Looking beyond the individuals affected by mandate, we may wonder about its impact on firms. Throughout this paper, I implicitly assumed that firms could not respond to the mandate by encouraging workers to retire or reducing their compensation. Both of these would attenuate my job lock finding. However, it has been shown that the mandate increased premiums for insurance plans covering dependents (Depew and Bailey 2015) and that workers in firms offering dependent coverage saw annual wages decrease by \$1200 (Goda, Farid, and Bhattacharya 2016), both of which would bias my findings toward zero. Given this paper's finding that firms not only had to pay for additional dependent coverage but also had to accommodate workers delaying retirement, it would be interesting to explore how each of these channels affected firms.

In summary, I use the ACA dependent mandate to estimate the effect of dependent insurance on parental retirement rates, and find that the ACA significantly decreased the likelihood of retirement for eligible parents. On average, dependent insurance coverage reduced the retirement rate for the treated group by 2.9 percentage points after policy enactment, which is 14.8 percent of the retirement rate absent the policy. Assuming rank preservation, this means that in response to the policy, individuals delayed retirement by 0.74 years on average, and up to 1.7 years. I did not find that the mandate affected other labor force outcomes like non-retirement exits and employer changes. The effect was particularly large for early retirees, although I find evidence that parents do not forgo Social Security benefits in response to the mandate. I find that Social Security eligibility plays a role in a couple's joint decision about which spouse delays retirement to take advantage of the mandate. Couples with one SS-eligible spouse and one SS-ineligible spouse tend to coordinate so that the former retires to claim Social Security and the latter delays retirement. Altogether, the findings raise the point that it is important to consider how various government programs interact with each other to shape individual and household incentives. Specifically, in this case, individuals and couples faced the tradeoff of whether to retire and claim Social Security, or delay retirement to obtain dependent health insurance. Overall, the mandate reshaped both retirement patterns for individuals and joint retirement patterns for married couples.

References

Antwi, Yaa Akosa, Asako S. Moriya, and Kosali Simon (2013). "Effects of Federal Policy to Insure Young Adults: Evidence from the 2010 Affordable Care Act's Dependent-Coverage Mandate". In: *American Economic Journal: Economic Policy* 5.4, pp. 1–28. ISSN: 1945-7731. URL: http://www.jstor.org/stable/43189352.

Ashraf, Nava, Dean Karlan, and Wesley Yin (2010). "Female Empowerment: Impact of a Commitment Savings Product in the Philippines". In: *World Development* 38.3, pp. 333–344. ISSN: 0305-750X. DOI: 10.1016/j.worlddev.2009.05.010. URL: http://www.sciencedirect.com/science/article/pii/S0305750X09001910.

Bailey, James and Anna Chorniy (2015). "Employer-Provided Health Insurance and Job Mobility: Did the Affordable Care Act Reduce Job Lock?" en. In: *Contemporary Economic Policy* 34.1, pp. 173–183. ISSN: 1465-7287. DOI: 10.1111/coep.12119. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/coep.12119.

Bansak, Cynthia and Steven Raphael (2008). "The State Children's Health Insurance Program and Job Mobility: Identifying Job Lock among Working Parents in Near-Poor Households". In: *Industrial and Labor Relations Review* 61.4, pp. 564–579. ISSN: 0019-7939. URL: http://www.jstor.org/stable/25249173.

Barbaresco, Silvia, Charles J. Courtemanche, and Yanling Qi (2015). "Impacts of the Affordable Care Act dependent coverage provision on health-related outcomes of young adults". In: *Journal of Health Economics* 40, pp. 54–68. ISSN: 0167-6296. DOI: 10.1016/j.jhealeco.2014.12.004. URL: http://www.sciencedirect.com/science/article/pii/S0167629614001519.

Barnett, Jessica C and Edward R Berchick (2017). "Health Insurance Coverage in the United States: 2016". en. In: *Current Population Reports*. Washington, DC: U.S. Government Printing Office, pp. 60–260.

Blascak, Nathan and Vyacheslav Mikhed (2018). "Did the ACA's Dependent Coverage Mandate Reduce Financial Distress for Young Adults?" In: *Federal Reserve Bank of Philadelpha*. Vol. 18-03. Working papers. URL: https://www.philadelphiafed.org/-/media/research-and-data/publications/working-papers/2018/wp18-03.pdf.

Brown, Kristine M. (2013). "The link between pensions and retirement timing: Lessons from California teachers". In: *Journal of Public Economics* 98, pp. 1–14. ISSN: 0047-2727. DOI: 10.1016/j.jpubeco.2012.10.007. URL: http://www.sciencedirect.com/science/article/pii/S0047272712001156.

Depew, Briggs (2015). "The effect of state dependent mandate laws on the labor supply decisions of young adults". In: *Journal of Health Economics* 39, pp. 123–134. ISSN: 0167-6296. DOI: 10.1016/j.jhealeco.2014.11.008. URL: http://www.sciencedirect.com/science/article/pii/S0167629614001465.

Depew, Briggs and James Bailey (2015). "Did the Affordable Care Act's dependent coverage mandate increase premiums?" eng. In: *Journal of Health Economics* 41, pp. 1–14. ISSN: 1879-1646. DOI: 10.1016/j.jhealeco.2015.01.004.

Duflo, Esther (2000). "Child Health and Household Resources in South Africa: Evidence from the Old Age Pension Program". en. In: *American Economic Review* 90.2, pp. 393–398. ISSN: 0002-8282. DOI: 10.1257/aer.90.2.393. URL: https://www.aeaweb.org/articles?id=10.1257/aer.90.2.393.

French, Eric and John Bailey Jones (2011). "The Effects of Health Insurance and Self-Insurance on Retirement Behavior". en. In: *Economet-*

rica 79.3, pp. 693–732. ISSN: 1468-0262. DOI: 10.3982/ECTA7560. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA7560.

Goda, Gopi Shah, Monica Farid, and Jay Bhattacharya (2016). *The Incidence of Mandated Health Insurance: Evidence from the Affordable Care Act Dependent Care Mandate*. Working Paper 21846. National Bureau of Economic Research. DOI: 10.3386/w21846. URL: http://www.nber.org/papers/w21846.

Gruber, Jonathan and Brigitte Madrian (1997). "Employment separation and health insurance coverage". en. In: *Journal of Public Economics* 66.3, pp. 349–382. ISSN: 0047-2727. DOI: 10.1016/S0047-2727(96)01621-0. URL: https://www-sciencedirect-com.ezproxy.cul.columbia.edu/science/article/pii/S0047272796016210.

— (2004). "Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature." en. In: *Health Policy and the Uninsured*. Ed. by Catherine G. McLaughlin. Google-Books-ID: FrXlqoFwjeQC. The Urban Insitute. ISBN: 978-0-87766-719-3.

Gruber, Jonathan and Brigitte C. Madrian (1995). "Health-Insurance Availability and the Retirement Decision". In: *The American Economic Review* 85.4, pp. 938–948. ISSN: 0002-8282. URL: http://www.jstor.org/stable/2118241.

Gustman, Alan L. and Thomas L. Steinmeier (2000). "Retirement in Dual-Career Families: A Structural Model". In: *Journal of Labor Economics* 18.3, pp. 503–545. ISSN: 0734-306X. DOI: 10.1086/209968. URL: https://www.journals.uchicago.edu/doi/abs/10.1086/209968.

— (2004). "Social security, pensions and retirement behaviour within the family". en. In: *Journal of Applied Econometrics* 19.6, pp. 723–737. ISSN: 1099-1255. DOI: 10.1002/jae.753. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.753.

Heim, Bradley, Ithai Lurie, and Kosali Simon (2017). *Did the Affordable Care Act Young Adult Provision Affect Labor Market Outcomes? Analysis Using Tax Data*. Working Paper 23471. National Bureau of Economic Research. DOI: 10.3386/w23471. URL: http://www.nber.org/papers/w23471.

Internal Revenue Service (2010). *Internal Revenue Bulletin: 2010-2*. URL: https://www.irs.gov/irb/2010-20_IRB.

Kaiser Family Foundation (2010). *Kaiser Health Tracking Poll – April 2010*. URL: https://www.kff.org/health-reform/poll-finding/kaiser-health-tracking-poll-april-2010/.

Lee, Ajin (2017). "Late Career Job Loss and Retirement Behavior of Couples". en. In: *Research on Aging* 39.1, pp. 7–28. ISSN: 0164-0275. DOI: 10.1177/0164027516656136. URL: https://doi.org/10.1177/0164027516656136.

Madrian, Brigitte C., Gary Burtless, and Jonathan Gruber (1994). "The Effect of Health Insurance on Retirement". In: *Brookings Papers on Economic Activity* 1994.1, pp. 181–252. ISSN: 0007-2303. DOI: 10.2307/2534632. URL: https://www.jstor.org/stable/2534632.

Robbins, Anthony S. et al. (2015). "Association Between the Affordable Care Act Dependent Coverage Expansion and Cervical Cancer Stage and Treatment in Young Women". en. In: *JAMA* 314.20, pp. 2189–2191. ISSN: 0098-7484. DOI: 10.1001/jama.2015.10546. URL: https://jamanetwork.com/journals/jama/fullarticle/2471561.

Rust, John and Christopher Phelan (1997). "How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets". In: *Econometrica* 65.4, pp. 781–831. ISSN: 0012-9682. DOI: 10.2307/2171940. URL: https://www.jstor.org/stable/2171940.

Seibold, Arthur (2017). "Reference Dependence in Retirement Behavior: Evidence from German Pension Discontinuities". In:

U.S. Census Bureau (2001). *Survey of Income and Program Participation Users' Guide.*

A Appendix - Figures

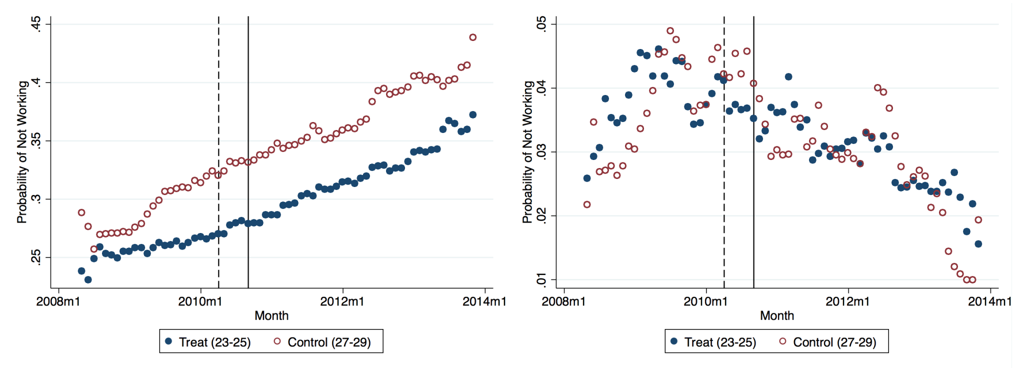

Figure A.1: Monthly probability of not working (left); not working and looking for work (right). Dashed vertical line represents policy announcement in March 2010; solid vertical line represents policy enactment September 2010.

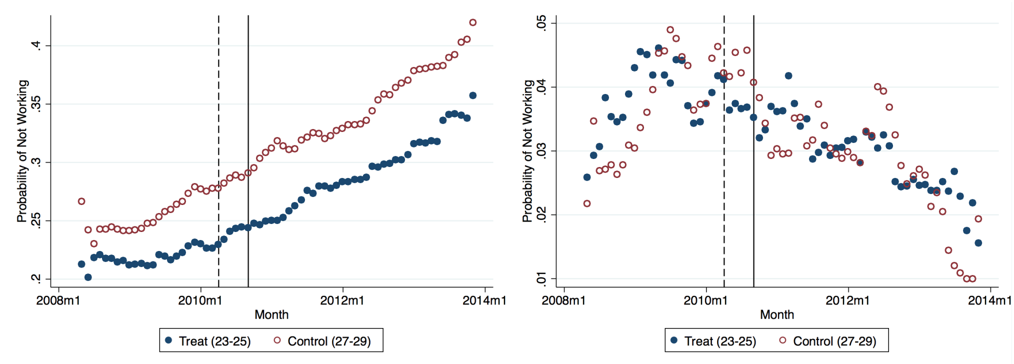

Figure A.2: Monthly probability of not working and not looking for work (left); partial LF exit (right). Dashed vertical line represents policy announcement in March 2010; solid vertical line represents policy enactment September 2010.

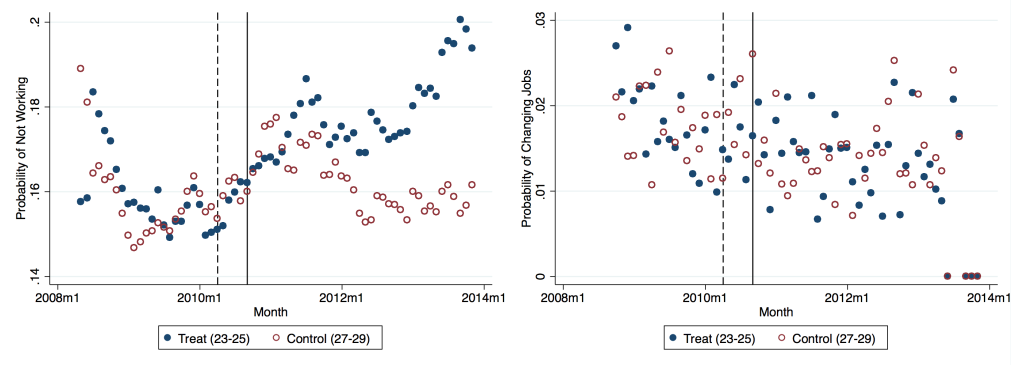

Figure A.3: Monthly probability of not working, not looking for work, and not retired (left); job change (right). Dashed vertical line represents policy announcement in March 2010; solid vertical line represents policy enactment September 2010.

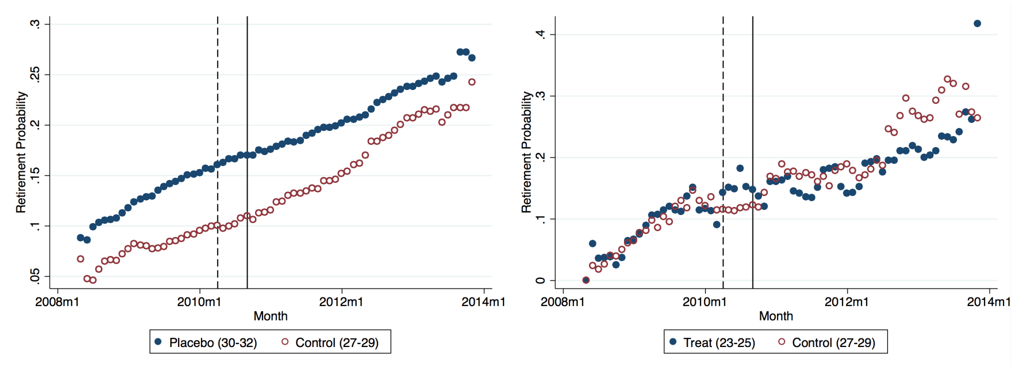

Figure A.4: Monthly probability of retirement for placebo group (left); Medicaid recipients (right). Dashed vertical line represents policy announcement in March 2010; solid vertical line represents policy enactment September 2010.

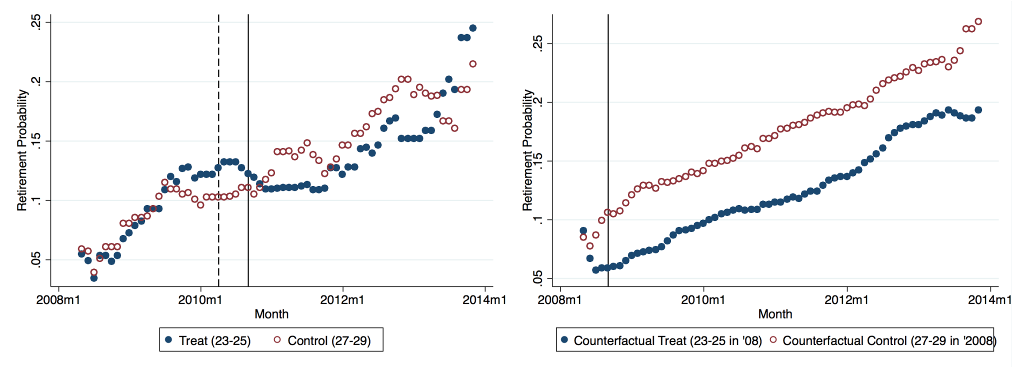

Figure A.5: Monthly probability of retirement for individuals without insurance in 2009 (left); Counterfactual (September 2008) (right). Dashed vertical line represents policy announcement in March 2010; solid vertical line represents policy enactment September 2010 (or September 2008 for counterfactual).

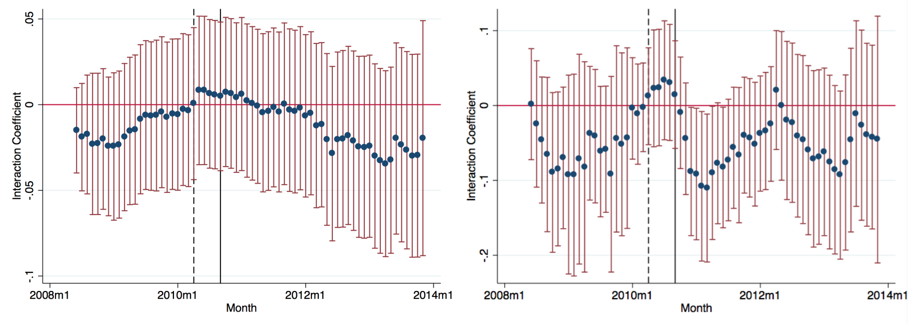

Figure A.6: Monthly coefficients for placebo (left); Medicaid recipients (right)

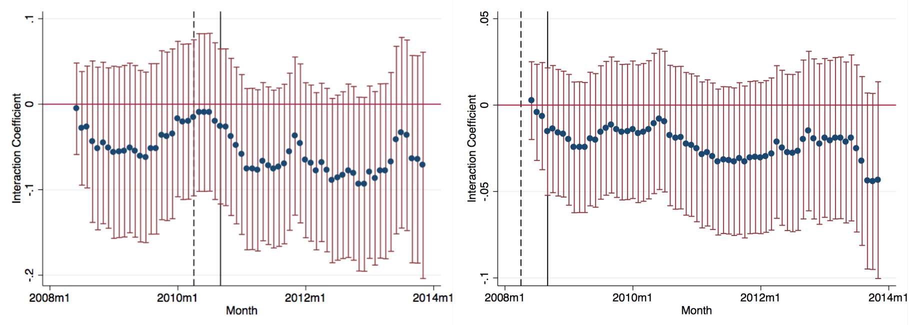

Figure A.7: Monthly coefficients for parents with no insurance in 2009 (left); 2008 counterfactual policy (right)

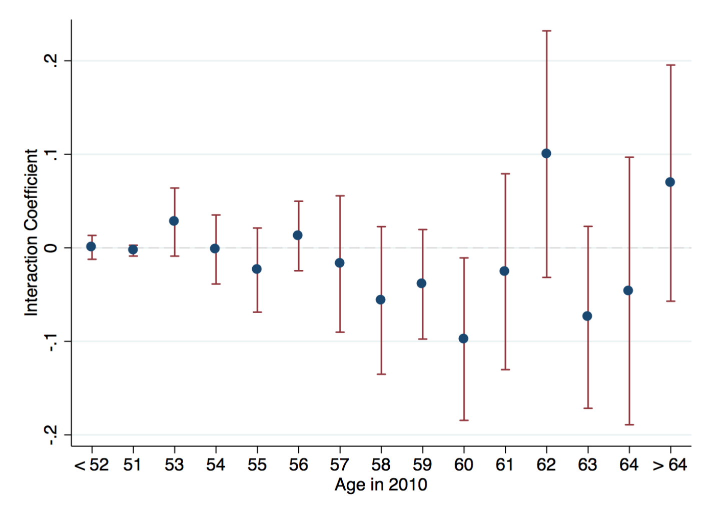

Figure A.8: Interaction coefficient by parent's age in 2010. Confidence intervals are 95%.

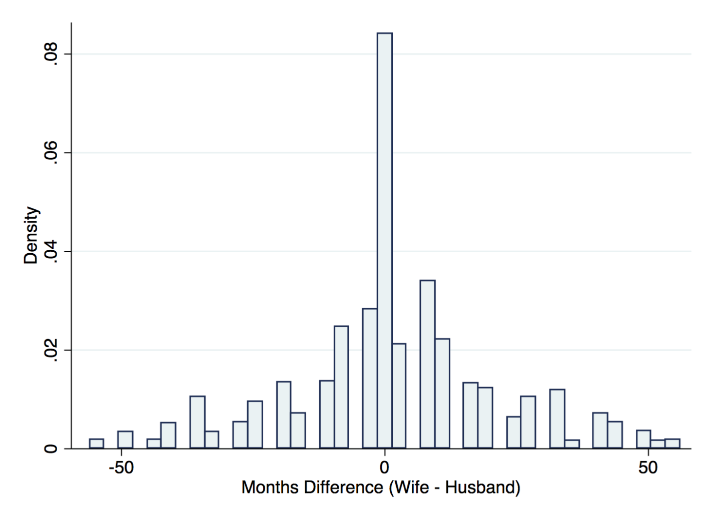

Figure A.9: Months between wife and husband's retirement dates if both retirements observed in SIPP

B Appendix - Tables

Table B.1: Regression coefficients by 2008-2013 age cohort. Outcome variable: monthly retirement rate

	(1) Young	(2) Early	(3) Full	(4) Early or Full	(5) Old
$Treat_i \times Enact_t$	-0.014*	-0.036*	0.076	-0.027	0.103
	(0.007)	(0.021)	(0.085)	(0.021)	(0.093)
Seam	X	X	X	X	X
Indiv. FE	X	X	X	X	X
Month-year effect	X	X	X	X	X
Individuals	1808	668	107	838	45
Observations	95216	37141	6101	46682	2429

Standard errors in parentheses

(1) Subsample: individuals who are < 62 in 2013

(2) Subsample: individuals who turn 62 between 2008 and 2013

(3) Subsample: individuals who turn 66 between 2008 and 2013.

(4) Subsample: individuals who turn 62 or 66 between 2008 and 2013

(5) Subsample: individuals who turned 66 before 2008.

* $p < .10$, ** $p < .05$, *** $p < .01$