

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

Maggie Shi [†]

April 5, 2022

Abstract

The 2010 Affordable Care Act expanded health insurance coverage to dependents up to age 26, allowing some parents to add adult children to their employer-sponsored plans. I consider a potential spillover of the dependent mandate policy on the labor supply of parents: did parents delay retirement to take advantage of the dependent mandate? I find that among parents aged 55-66, affected parents' retirement rate fell by 3.9 percentage points after policy enactment, causing them to delay retirement by 0.74 years on average. An estimated 296,000 parents delayed retirement in order to obtain coverage for their children.

^{*}I am grateful to Wojciech Kopczuk, Michael Best, and Adam Sacarny for their guidance and support. I also thank Bentley Macleod, Sandra Black, Claudia Halbac, Francois Gerard, Day Manoli, Pietro Tebaldi, Gautam Gowrisankaran and participants at the Columbia University Applied Microeconomics colloquium, 2019 Society of Labor Economists conference, and 2019 ASHEcon conference for insightful comments and suggestions.

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

1 Introduction

One of the defining features of the 2010 Affordable Care Act (ACA) was that it enabled a large number of Americans to gain health insurance coverage. Between 2010 and 2016, the number of uninsured adults under the age of 65 fell by 41 percent, as 20 million people gained access to insurance (Garfield et al. 2019). Many ACA provisions that increased insurance access, like the expansion of Medicaid and the introduction of the marketplace exchanges, gave individuals access to health insurance that was not tied to employment. Therefore, many speculated that the ACA would spill over to the broader economy by weakening “job lock,” which is the notion that the prevalence of offering health insurance as an employment benefit in the U.S. can distort labor supply by keeping individuals who value these benefits from leaving jobs that they otherwise would (Madrian et al. 1994; Swartz 2014; Baker 2016). However, there may be another provision in the ACA that could have also unintentionally *increased* job lock for some: the dependent coverage mandate, which required insurance plans covering dependents to cover them until age 26.

This paper explores whether the dependent mandate induced job lock for *parents* with children who were newly eligible for coverage under the mandate. In particular, I focus on whether it caused parents with children under the age of 26 to delay retirement in order to add them to their employer-sponsored health insurance. The dependent mandate may have affected parental labor supply because it was a well-publicized policy change that, for many eligible parents, was relatively straightforward to take advantage of. Takeup of the policy was high – about two million young adults were added to their parent’s employer-sponsored health insurance (Antwi et al. 2013), and for each young adult who gained insurance, there was a parent who added them. I focus in particular on the effects of the mandate on retirement, as retirement is an important and relevant dimension of labor supply for the cohort of parents affected by the mandate, most of whom were in their 50s and 60s.

To understand whether the dependent mandate shifted parental retirement patterns, I use an empirical strategy that compares parents with children who are young enough to be eligible under the dependent mandate (i.e., under 26 in 2010) to parents with children who are too old to be eligible (i.e., over 27). Using the Survey of Income and Program Participation (SIPP), I look at differences in retirement rates across the two groups of parents, before and after the implementation of the mandate in September

2010. I find that parents with eligible children were *less likely to be retired* by 2013, even after including a rich set of controls for parental factors that may play a role in the retirement decision. Importantly, the analyses control for parental age fixed effects to ensure comparisons of eligible and ineligible parents who themselves were the same age. I find that the gap in retirement rates between eligible and ineligible parents is driven by parents in their early 60s, and appears only after the mandate is enacted. I interpret this as evidence of parental job lock induced by the expansion of dependent coverage for young adults.

In this policy setting, the mandate’s eligibility cutoff age of 26 for dependents allows me to cleanly identify eligible and ineligible parents. The simplicity of the mandate, in which eligibility depends solely on dependents’ age, means that I only need to know the child’s birth year in order determine eligibility. This information is available in the SIPP. Assuming that among parents of the same age, having a child slightly older or younger than 26 in 2010 is as good as random, then my estimates recover the causal effect of dependent health insurance on parental retirement.

The lower retirement rate for eligible parents translates into an average retirement delay of 0.74 years. Extrapolating the effect to the entire U.S. population implies that about 296,000 parents delayed retirement in response to the dependent mandate. I also deploy a series of placebo tests and robustness checks. I consider three placebo tests: I compare the retirement rates of the eligible and ineligible parents *before* policy implementation, I compare the retirement rates of two ineligible groups, and I explore the effects of a placebo “policy” using an earlier SIPP panel. All of these placebo tests are insignificant. The difference in retirement rates appears only when comparing parents with differing eligibility for the dependent insurance mandate, and arises only after the mandate is implemented. The findings are also robust to alternative specifications such as changing the definition of the treated and control groups.

This paper contributes to the literature on job lock by providing evidence that it can arise even from *dependent* insurance. Previous studies have demonstrated that job lock can arise from *own* insurance (Madrian et al. 1994; Gruber and Madrian 1995; Gruber and Madrian 1997; Dave et al. 2015; Garthwaite et al. 2014). But there is less evidence on whether dependent insurance availability contributes to job lock, even though the two are often packaged together. In a related paper looking at a different setting – mothers of

young children who become eligible for children’s Medicaid – Grossman et al (2022) find evidence that increases in children’s Medicaid eligibility decrease maternal labor market participation, implying that some mothers faced job lock prior to their child becoming eligible for Medicaid. My findings complement these findings and suggest that even when the dependents in question are much older, dependent insurance availability can generate parental job lock.

I also contribute to a growing literature documenting the effect of the ACA dependent mandate, which has mostly focused on how it affected young adults.¹ The mandate increased young adults’ insurance coverage (Antwi et al. 2013), improved health in a variety of dimensions (Barbaresco et al. 2015; Robbins et al. 2015; Akosa Antwi, Moriya, et al. 2015; Akosa Antwi, Ma, et al. 2016), reduced healthcare debt (Blascak and Mikhed 2021), and increased wages (Dillender 2014). The evidence on job lock for young adults is mixed, with some papers finding a decreased probability of working full time and hours worked (Antwi et al. 2013), but others finding no effect on their job mobility and employment (Bailey and Chorniy 2015; Heim et al. 2018). My paper focuses on another party who is mechanically also affected by the dependent mandate – the parents who stay at their job in order to provide the insurance.

The rest of the paper is organized as follows: in Section 2, I describe the reform and lay out a conceptual framework. In Section 3, I discuss the data and identification strategy. I present my results in Section 4, and discuss their implications and conclude in Section 5.

2 Policy Change and Conceptual Framework

The dependent coverage mandate was one of the most popular and well-publicized components of the ACA (Goldman 2013). This mandate required insurance plans with dependent coverage to expand coverage to children up to age 26. For parents whose employer-sponsored insurance covered dependents, which is the case for 96% of firms (Kaiser Family Foundation 2020), adding their children to their insurance plan was relatively straightforward and low-cost. Before the dependent mandate, in many states insurance plans only

¹There is also a small literature documenting how the mandate affected firms and workers. Premiums for plans covering dependents increased (Depew and Bailey 2015), and workers in firms with employer-based coverage saw a reduction in wages (Goda, Farid, et al. 2016).

covered dependents up to age 18 or until they were out of college. Some individual states had their own dependent coverage mandates, but eligibility requirements varied widely by the dependent's marital status, student status, or own employer-sponsored insurance, and almost none required coverage up to age 26.²

In contrast to individual state policies, the ACA dependent coverage mandate was a national policy with no eligibility requirements besides age. Any young adult under 26 whose parent's health insurance covered dependents could now obtain coverage under her parent's plan. The policy was announced in March 2010 and insurers were required to comply by September 2010. In March 2010, the Internal Revenue Service also amended its rules to allow health benefits for dependents to be tax-exempt up until the dependents reached age 27 (Internal Revenue Service 2010). The dependent mandate had a significant impact on insurance rates for young adults. About 2 million young adults (7 percent of adults aged 19-26) added parental employer-sponsored insurance as a result of the mandate (Antwi et al. 2013). By 2016, 26 year olds were 20 percent more likely to be uninsured than 25 year olds (Barnett and Berchick 2017).

Of course, for every young adult who gained insurance through the mandate, there was a parent who added them to their insurance plan. Therefore in families where parents were on employer-sponsored insurance plans, the mandate established a link between the adult children's health insurance coverage and their parent's employment. The mandate's age cutoff at 26 presents us with a quasi-experimental setting to study the extent to which parents are willing to adjust their labor supply to gain insurance coverage for their adult children. Given that most of the affected parents were in their 50s or 60s,³ retirement is a highly relevant labor supply outcome of interest.

Next, I lay out a conceptual framework based on Gruber and Madrian (2004) to illustrate how the dependent mandate could induce job lock for parents and cause them to delay 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, Medicare, Medicaid, or non-employer private insurance). Let utility in period t be $U_t(w, H, L)$,

²See Depew (2015), Table 1 for an overview of state policies. Note that my results rely on a comparison of parents with children above and below age 26 in 2010. Since most of the pre-ACA state policies did not require coverage up to age 26, the existence of state provisions prior to 2010 does not directly affect the analysis.

³The average age of parents in my sample (parents of 23-29 year olds in 2010) was 58 at the end of the panel.

where w is the compensation 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:

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

Individuals simply compare working with health insurance, $U_t(w, 1, 0)$, to not working without health insurance, $U_t(0, 0, 1)$. If $U_t(w, 1, 0) > U_t(0, 0, 1)$, then they work this period. The periods can be at regular intervals (i.e., individual decides each year whether to continue working next year), or they can center around reference points like statutory retirement ages (i.e., individual decides whether work until the early retirement age or the full retirement age).

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

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

where $B > 0$ represents the value of dependent health insurance to the parent and $C > 0$ is the compensating differential by which an employer reduces the parent's compensation.⁴ If 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.⁵ If $w + B - C > 0$ (i.e., $B - C > 0$), then the reform increases the value of continuing to work for affected parents.

Parental labor supply is unaffected by the mandate if $B - C = 0$. This could happen in two ways. First, if parents do not value the dependent coverage benefit, $B = 0$ and employers would set $C = 0$. Alternatively, even if parents do value dependent coverage and $B > 0$, firms could still set $C = B$ and exactly compensate for it. In both cases, the workers' labor supply will be unaffected and the reform will not induce job lock.

⁴Goda et al. (2016) finds evidence that workers in firms offering dependent coverage saw annual wages decrease by \$1200.

⁵Of course, if premiums paid by the employer went up as well (which indeed seems to be true according to Depew and Bailey (2015)), this would not be a rent per se as the cost to employers of providing insurance to workers increases. For the sake of simplicity in this model, I do not include the employer's decision of whether or not to offer health insurance. Additionally, I assume that $C < B$, meaning employers do not impose compensating differentials higher than how much parents value dependent insurance.

Job lock occurs when $B - C > 0$, meaning workers value the dependent benefit and employers do not set compensating differentials to exactly offset how much the worker values it. Then, there could be workers for whom $U_t(w + B - C, 1, 0) > U_t(0, 0, 1) > U_t(w, 1, 0)$. Put another way, absent the policy these individuals would have retired in period t ($U_t(0, 0, 1) > U_t(w, 1, 0)$), but now they continue to work in order to take advantage of the policy ($U_t(w + B - C, 1, 0) > U_t(0, 0, 1)$). If we find causal evidence that affected parents delayed their retirement, this would imply that the additional value of the dependent insurance (net of the compensating differential) is high enough that they are willing to continue working when they otherwise wouldn't have. Thus if we observe job lock, we can conclude that $B - C > 0$. Note that separately identifying B and C is not required to establish job lock.

3 Data and Identification Strategy

The data 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 a variety of information from the past four months. The 2008 panel covers May 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.

Using the SIPP, I compare retirement rates for parents whose children are affected by the mandate to those whose children are too old to be affected by it. To define my treatment and control groups, I will use the age of a respondent's youngest child at the time of implementation. 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. 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, following Antwi et al. (2013), I do not assign parents of 26 year olds to either treatment or control as it is unclear what their eligibility is in 2010.

In order to have relatively comparable groups, I restrict my treatment group to parents

of 23-25 year olds and control group to parents of 27-29 year olds. In Table 1, I report summary statistics of individual characteristics by treatment status. In the main analysis, I restrict the sample of parents to those aged 55-66 at a given point in time because of small sample sizes at older ages for the treatment group and at younger ages for the control group.

While the treatment and control groups are similar in many dimensions, it is important to note that parents in the control group are older than those of the treatment group. This makes sense, since the control group comprises 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 that of the treatment group. Accordingly, we also see that control parents are less likely to be working in 2009 (Table 1). In order to address this, my empirical strategy holds age fixed when comparing the treated and control groups. All regression specifications will include age fixed effects, and all raw differences will be calculated by parental age cohort. Doing this ensures that the comparison of treatment and control groups always occurs *within* parental age cohort. Thus, differences in the age distribution of treatment and control parents are not driving the results.

It is also of note that the sample contains more women than men. Women comprise 66 percent and 62 percent of the treatment and control groups, respectively. This is because I do not directly observe every child's age for children no longer living in the household for both mothers and fathers. Instead, I use mothers' fertility history⁶ and consider the ages of the children a woman gave birth to. Fathers (or step-fathers) are then assigned to treatment or control depending on the fertility history of their wife.⁷ Thus, I can only add married men whose wives were also in the sample, which explains the gender imbalance. Note that both the treatment and control groups are unbalanced and have similar female shares, and thus are comparable to each other.

I use responses 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

⁶SIPP only asks for the oldest and youngest children that a woman gave birth to. Thus I also have no sense of exactly how many eligible or ineligible children the respondent has, although I do know the total number of children. I also do not observe ages of adopted children or stepchildren, who would both count as dependents for insurance purposes. Not accounting for them will bias my result toward zero.

⁷See Data Appendix D for a detailed description of the procedure to link husbands and wives.

looking for work” and the main reason for not having a job in the reference period (last four months) is retirement. In the main analysis I use the last period of the 2008 panel (which for most individuals is in 2013) rather than the date their child becomes eligible.⁸ I do this because it is unclear whether a parent who has decided to delay retirement because of the mandate would retire immediately once their child becomes ineligible – parents could be, for example, deciding between whether to retire immediately or to delay until a statutory retirement age, like the full retirement age. Additionally, individuals’ decisions to delay retirement may not be immediately observable in the data, it may take time for changes in retirement patterns to emerge. Finally, since I only know each child’s birth year rather than their birth date, I have an imprecise measure of when each child becomes ineligible.

Intuitively, the empirical strategy compares cross-sections of each group’s retirement rate by parental age. I first plot the retirement rates of each group, which represent the percent of each group that is retired at a given age in 2013. I then calculate the raw difference in retirement rates for each age separately. This strategy allows me to hold age fixed while comparing the treated and control groups, which is important since the retirement decision is tightly linked to age and the control group is on average older than the treated group.

Then I pool these within-age comparisons together in a regression framework (Equation 1). The regression specification includes a parental age fixed effect (Age_i), which allows for a within-age comparison and accounts for the different age distributions among the treatment and control groups. The specification also includes a rich set of controls, including a fixed effect for the year the individual is observed ($Year_i$)⁹ and individual controls \mathbf{X}_i for pre-reform monthly earnings, pre-reform working status, gender, race, Hispanic, marital status, education, pre-reform private health insurance coverage, number of children, pre-reform union status, pre-reform paid hourly status, and pre-reform worker class.¹⁰ Controlling for demographic and pre-reform employment characteristics

⁸In the 2008 SIPP panel, individuals exit the sample at different points in time. I restrict the analysis to individuals whose last observation is at least 1 year after policy implementation (i.e., whose last observation is on or after September 2011). This is true for 75 percent of individuals. Of this group, 83 percent of individuals’ last observation is on or after June 2013.

⁹The year an individual exits the 2008 panel can differ; 92% of individuals in the main sample exit in 2013, 6% exit in 2012, and 2% exit in 2011.

¹⁰Workers are separated into the following classes (*eclwrk1*): private for-profit employee, private not for profit employee, local government worker, state government worker, federal government worker, family worker without pay, and not in universe.

accounts for differential trends across different groups of people over time. Given that this time period also coincides with the recovery from the Great Recession which could have also affected individuals' retirement decisions, I also include as a control the state unemployment rate ($StateUR_i$).

$$Y_i = \beta_0 + \beta_1 Treat_i + \mathbf{X}_i \times \boldsymbol{\gamma} + StateUR_i + Year_i + Age_i + \varepsilon_i \quad (1)$$

I argue that the coefficient β_1 on the $Treat_i$ dummy (which is 1 if the child is 23-25 in 2010, and 0 if the child is 27-29 in 2010) captures the causal effect of the dependent mandate on the outcome variable. The identification assumption underlying this comparison is that absent the mandate, there would be no difference in outcomes between treated and control parents, after holding fixed parent's age and other characteristics. This would be true if having a child slightly older or younger than 26 in 2010 is as good as random. With this assumption, any observed difference reflects the causal effect of the mandate. This assumption is violated if, conditional on age and other control variables, parents of children above and below 26 retire at different rates for reasons unrelated to the dependent mandate. I then conduct a series of placebo tests and robustness checks to support this identification assumption and argue for a causal interpretation of the results.

An alternative research design one might consider using would be a difference-in-differences framework. However, using a difference-in-differences framework would require making a parallel trends assumption that the outcomes of the two groups would have evolved similarly in the post-policy period absent the mandate. In the context of retirement, where a large proportion of workers retire once they hit statutory ages like 63 and 66, it may not be reasonable to assume that the retirement rate of an older group would trend similarly to that of a younger group. This poses a problem when comparing parents of older children to parents of younger children, because on average the former is older than the latter. Indeed, if we consider the retirement rate for the two groups in 2008 in Figure C.1 (left panel), we see that in the post-policy period the control group's retirement rate increases faster than that of the treatment group. In a difference-in-differences framework, this would suggest that the mandate caused the treated group to delay retirement. But we also see a similar pattern if we compare two similarly-defined groups in 2004 data in the absence of the mandate (Figure C.1, right panel), violating the parallel trends assumption. Since a difference-in-differences strategy cannot simultaneously

control for age and time, it is difficult to disentangle whether the difference in retirement rates after the mandate is due to the different age compositions of the two groups or due to the mandate. But once I restrict to cross-sectional comparisons holding age fixed, then any difference in outcomes between the two groups cannot be due to differences in age composition between the two groups. Thus, the method of comparing cross-sections of retirement profiles is preferable to a difference-in-differences framework.

Finally, I quantify the delay in retirement induced by the mandate in years, rather than probability of being retired. One challenge to estimating this directly in a regression framework is that the 2008 SIPP panel ends in 2013, before many parents in the sample have retired. Retirement age is censored for most parents in my sample and thus I cannot use observed retirement age as an outcome variable directly in the regression without modelling censorship or duration. Instead, I opt for an alternate way of measuring retirement delay – I consider how much the dependent mandate “shifted” the treated group’s retirement profile to the right, relative to the control group’s profile.

In practice, I measure the horizontal distance between each group’s retirement cross-sections at different retirement rates. While the previous set of results looked at retirement rate as a function of age, with this measure I consider age as a function of retirement rate. This non-parametric method allows me to estimate the difference in retirement age between the two groups without having to introduce a more complicated model of duration or censorship. If we assume rank preservation, meaning that an individual would be in the same retirement age quantile regardless of whether she is in the treated or control group, then the horizontal distance between the two retirement profiles represents the retirement delay induced by the dependent mandate measured in years.

4 Results

4.1 Main Results

Figure 1 plots the average retirement rate in the post-policy period by age and treatment group. I subset to individuals who exit the panel at least one year after policy implementation (i.e., after September 2011). Age is defined as the person’s age in the last

observation in the SIPP panel, rounded to the closest integer.¹¹ The gap in retirement rates between the two groups represents the difference in the percent of individuals who are retired by a given age.

I next quantify the raw difference in retirement rates between the two groups, taking repeated bootstrap samples to estimate confidence intervals; Figure 2 plots the results by age. There are no differences in retirement rates for parents in their late 50s, but a gap emerges from age 60 on. The point estimates of the raw differences at each age are reported in Table B.1. After age 60, the individual age cohort estimates of the difference in retirement rates are all negative. For ease of interpretation and to improve power, I aggregate the estimates from separate age cohorts into groups (i.e., 55-60, 61-66, and 55-66), which are also reported in Table B.1. The group average difference is calculated by estimating separate differences for each age cohort and taking the mean of the estimates, rather than pooling individuals in each group together and calculating the difference. Taking the average of the estimates rather than pooling ensures that the estimates are not a result of differing age compositions between the treatment and control groups. On average for all individuals between 55 and 66, the retirement rate is 5.4 percentage points, or 25.9 percent, lower in the treated group than to the control group. This is driven by individuals in the 61-66 age group. The average difference is statistically significant for the overall group and for the 61-66 age group, as the 95% bootstrapped confidence intervals for the difference fall below 0.

Next, I run a regression of the treatment dummy on retirement, controlling for individual characteristics, state unemployment rates, and age and year fixed effects. Including these controls accounts for factors which contribute to an individual's retirement decision that are unrelated to the dependent mandate. The results are reported in Table 2. Even after taking into account these controls, the gap in retirement between treated and control remains negative and statistically significant. Being in the treated group makes an individual 3.9 percentage points (18.7 percent) less likely to be retired. Again, this gap is driven by the large effects for treated individuals aged 61-66.

Given the timing of the policy, one concern may be that the recession recovery may have contributed to the different retirement rates between treatment and control parents.

¹¹The specification is robust to finer definitions of age, such as rounding to the age by 0.5 years (Table B.2). I choose to round parent's age to the closest integer to match the precision of the measure of children's age, which is by the year.

For this to be true, treatment parents must have experienced a different labor force market recovery than control parents, conditional on parental age and other demographic controls. If the effect is driven mostly by the recession recovery and not the dependent mandate, then we should see the results change when adding in controls that should be correlated with the recession recovery. Table B.3 demonstrates that the results are not sensitive to the inclusion of recession recovery-related controls like state unemployment rate and pre-reform job characteristics. This suggests that the effect is not driven by the recession recovery, or other differential trends over time that would be correlated with employment characteristics such as changes in pensions.

I can extrapolate the regression results to the overall US population and calculate how many workers delayed retirement in response to the dependent mandate. Using the age profiles of all affected parents in my sample,¹² I calculate that there were about 7.6 million parents of 19-25 year olds who were between 55 and 66 in 2013. Combined with the estimates from the main regression results, this implies that about 296,000 parents delayed retirement in order to take advantage of the mandate.

Overall, the results on retirement demonstrate that among parents over 60, those eligible for the mandate were less likely to be retired than ineligible parents post-policy. Because the analysis holds age fixed, the difference in retirement rates is not due to the differing age composition between the treated and control groups. Since other covariates are relatively balanced across the two groups (Table 1), adding them as controls does not change the findings much from the comparison of raw differences. Thus, the difference in retirement rates between the two groups is not driven by parental age composition or other parental covariates.

The two groups do, however, differ mechanically in one important way – the children of treated parents are younger than the children of control parents. If parents of younger children are less likely to retire in their early 60s than parents of older children, then we would see this pattern in retirement rates even without the mandate. Since treatment and control are defined by children’s ages, there is no way to disentangle this by simply controlling for it.

In order to address the issue of differing children’s ages, I conduct a series of placebo

¹²Within SIPP, I calculate the percentage of each age cohort in 2010 that has a child aged 19-25. I then combine this with population count by age in the 2010 Census, and extrapolate to calculate the total number of individuals who are parents of a 19-25 year old and were between 55 and 66 in 2013.

tests to check whether differences in children’s ages could be driving the difference in retirement rates between the two groups. The first placebo test considers treated and control parents in the pre-period – that is, it takes a cross-section of retirement rates by age in *September 2010*, right before the dependent mandate was implemented. If control parents’ higher retirement rate is driven some unobserved, non-time-varying characteristic of the control parents, then their retirement rate should be higher before implementation of the mandate as well. Table 3 reports the regression results. For all groups, the difference in retirement rates between treated and control individuals is statistically insignificant. Thus, the observed difference in retirement rates between the treatment and control group emerges only after the dependent mandate is implemented.

The second placebo test considers a group of parents that should not be affected – parents whose children are 31-33 in 2010. This placebo group of parents has children with a similar age difference to the control group as the treated group does, albeit in the opposite direction. We would expect to see an effect with the placebo group if there are time-varying trends which cause differential retirement behavior in parents with children of different ages, such as the recession recovery affecting parents of older children differently than parents of younger children. If, for a reason unrelated to the mandate, parents of younger children tend to retire at lower rates than parents of older children conditional on parental age, then we would expect to see that parents in the control group should have *lower* retirement rates than parents in the placebo group (who are older than control children). Table 3 reports the regression results. For all groups, the coefficient is insignificant, albeit negative. However, note that the placebo group has children *older* than those of the control group. If the main results are driven by parents with older children being more likely to be retired, then the placebo parents would be more likely to be retired and thus we would expect the coefficient to be positive.

The third placebo test considers the effect of a placebo policy in 2004 using a previous SIPP panel. I define treatment and control groups using parents in the 2004 SIPP panel and compare their retirement rates at the end of the panel in 2007. Specifically, treatment is defined as having a child aged 23-25 in 2004 and control is defined as having a child aged 27-29 in 2004. This definition mirrors that of the 2008 panel because the “policy” occurs about 3 years before the last observation; in the 2008 panel, the policy occurs in 2010 and the last observation usually occurs in 2013. Similar to the first placebo test in

the pre-period, since there was no mandate in effect in 2007, any difference between the two groups would imply that the main results might be driven by something other than the dependent mandate. Again, if parents of young children have different retirement patterns than parents of older children, then we would expect to see a similar effect in the 2004 panel. Table 3 reports the regression results. I find that the coefficients are all statistically insignificant, meaning the difference in retirement rates is unique to the post-ACA period.

All three placebo tests are insignificant, indicating that the difference in retirement rates between treatment and control in Table 2 is not driven by the differing age compositions of each group’s children and any time-varying trends that may be correlated with children’s age. The difference only emerges in the post-policy period, and only for groups that were differentially affected by the reform (i.e., 23-25 vs. 27-29 opposed to 31-33 vs. 27-29). Taken together, the main results and placebo tests suggest that the gap in retirement between treated and control parents is caused by the dependent mandate, implying that treated parents eligible for dependent health insurance delayed retirement to take advantage of it.

Finally in Section A, I show that the results are robust to alternate specifications – specifically, I consider specifications that define cohorts by birth year rather than age, and specifications that vary the definition of treatment and control groups.

4.2 Quantifying Retirement Delay

The main results demonstrate that the ACA dependent mandate decreased the likelihood of retirement for eligible parents, especially for those in their 60s. Since almost all individuals will eventually retire, this decreased probability eventually translates into a delay in retirement age. I next calculate the length of retirement delay in years implied by my results. The challenge is that retirement age is censored for many individuals because the SIPP panel ends in 2013. So, I cannot use retirement age directly as an outcome variable in the regression framework used in the main results.

Instead, in Figure 3 I calculate how much “longer” it took the treated group to reach a given retirement rate compared to the control group. Rather than looking at how the mandate affected retirement at a given age, I consider how it affected the average age at a given retirement rate. Doing this gives the retirement delay induced by the mandate

as measured in years, without having to introduce a duration or censorship model of retirement age.

I construct a cross-section of retirement rates by age in the post-policy period, utilizing additional observations of an individual's monthly work history between policy implementation and the individual's last observation in the panel. Including this additional information allows for a more granular measure of retirement rate – by using the full set of post-policy observations, I measure retirement rate at every age in months, rather than rounding to age in years as in Figure 1. This additional information about each group's retirement profiles is otherwise difficult to use in a regression approach without a more complicated duration or censorship model. The bootstrap method for estimating confidence intervals is clustered at the individual level to ensure that inference is still at the individual level. Each point on this retirement profile at a given age represents the percent of the treated or control group that was retired by that age in the post-policy period. Thus, the horizontal distance between the two profiles at a given quantile represents how much longer it took the treated group to reach that retirement rate relative to the control group. For example, 10 percent of the control group had retired by age 58.73, while it took until age 60.25 for the treated group to reach 10 percent retired. Thus, the treated group required 1.52 more years to achieve a 10 percent retirement rate relative to the control group.

In order to estimate the horizontal distance between the two retirement profiles, I use a locally estimated scatterplot smoothing (LOESS) regression with smoothing parameter of 0.15 to fit a smooth curve to each group's retirement profile. I then calculate the horizontal distance between the smoothed retirement profiles at each quantile on the y-axis. In order to estimate bootstrapped confidence intervals, I resample the data at the individual level and recalculate the LOESS regression fitted line and horizontal distances. The left panel of Figure 3 plots the original cross-section and the average of fitted lines from 1000 resamples. The right panel plots the horizontal distance between the 1st and 50th quantile and corresponding 95% bootstrapped confidence intervals.

The difference is positive for almost every quantile of retirement rate and statistically significant at the 95th percent confidence level between the 9th and 18th quantiles. Between the 1st and 50th quantiles, the average difference between treated and control retirement ages was 0.74 years. The largest statistically significant difference was 1.7,

at the 9th quantile. If we assume rank preservation, meaning that an individual would be in the same retirement age *quantile* regardless of whether she is in the treated or the control group, then we can interpret this as a treatment effect at each quantile. This would imply that the mandate caused treated parents to delay their retirement by 0.74 years on average. If we do not assume rank preservation, this means that on average it took 0.74 years longer for the treated group to reach a given retirement rate between 1 percent and 50 percent relative to the control group.

5 Discussion

The main results demonstrate that the dependent insurance mandate increased job lock for parents – it led parents of eligible children to delay retirement in order to take advantage of the mandate. The estimates suggest a relatively large effect of the dependent mandate on parental retirement. Recent work on job lock from dependent insurance for young children also finds relatively large magnitudes – the estimates Grossman et al. (2022) would imply that having a child go from ineligible for Medicaid to eligible would increase the chance a mother is out of the labor force by 40 percent. The magnitude of my findings also echos earlier work on the effect of retiree health insurance (RHI) in inducing *earlier* retirement. Madrian et al. (1994) found that RHI availability led workers to retire 0.4-1.2 years earlier. Karoly and Rogowski (1994) found that the availability of RHI increased retirement rates by about 8 percentage points, or 50 percent of base-line. Other empirical papers in this literature also find relatively large effects (Blau and Gilleskie 2001; Shoven and Slavov 2014; Johnson et al. 2016).

There are also some additional context-specific factors surrounding the ACA and the dependent mandate that may contribute to the magnitude of the estimate. First, individuals did not directly pay the full cost of adding a dependent to their employer-sponsored plan (Goda, Farid, et al. 2016). The RHI literature has found that the effect of RHI on employment exit is three times larger when the firm pays the entire cost relative to when the individual shares the cost with the firm (Blau and Gilleskie 2001).

Second, the dependent mandate was well-publicized and popular, and had a straightforward eligibility cutoff of 26. 74 percent of the public felt favorably towards the mandate within a month of enactment, making it one of the most popular components of the ACA

(Kaiser Family Foundation 2010). There is evidence of large “statutory age effects” on retirement which are independent of and can be even larger than financial incentives, meaning that simply defining a salient retirement age can affect workers’ retirement behavior (Seibold 2021). Thus, the cutoff child’s age of 26 could have served as a salient “statutory age” around which parents planned their retirement.

Finally, the policy change occurred in the wake of the Great Recession, and more parents may have been on the margin of delaying retirement in this time period. Goda et al. (2011) find that retirement decisions were much more responsive to shifts in the S&P 500 index in this period relative to other periods, suggesting that individuals may have been more willing to delay retirement in this period. Young adults were also worse off in terms of insurance coverage due to the recession. Young adults disproportionately lost insurance in the recession – 2.5 million adults aged 19-34 (3.1 percent) became uninsured from 2007 to 2009, the largest increase in uninsured rate across all nonelderly adults (Holahan 2011). This means that more adult children may have relied on their parents for insurance coverage in this period. The coincidence of more parents on the margin of delaying retirement and more young adults losing health insurance in the Great Recession could explain the relatively large response to the ACA dependent mandate.

In summary, I use the ACA dependent mandate to estimate the effect of dependent insurance on parental retirement decisions in order to understand whether the mandate induced job lock for parents. I find that the availability of dependent coverage significantly decreased the likelihood of retirement for eligible parents. Dependent health insurance can generate job lock for parents, even when the dependents themselves are adults. On average, dependent insurance coverage reduced the retirement rate among 55-66 year old parents eligible for the mandate by 3.9 percentage points after policy enactment. This implies that in response to the mandate, eligible parents delayed retirement by 0.74 years on average to take advantage of the dependent mandate. Overall, about 296,000 parents delayed retirement to take advantage of the policy. Altogether, the results indicate that one unintentional effect of the ACA was that it induced job lock for some individuals – in particular, parents of young adults eligible for the dependent mandate. These findings demonstrate the potential spillovers of policies aimed at expanding insurance access onto the broader economy, and highlight the importance of taking these spillovers into account when making policy decisions.

References

- Akosa Antwi, Yaa, Jie Ma, et al. (Jan. 2016). “Dependent Coverage under the ACA and Medicaid Coverage for Childbirth”. eng. In: *The New England Journal of Medicine* 374.2, pp. 194–196. ISSN: 1533-4406. DOI: [10.1056/NEJMc1507847](https://doi.org/10.1056/NEJMc1507847).
- Akosa Antwi, Yaa, Asako S. Moriya, and Kosali I. Simon (Jan. 2015). “Access to health insurance and the use of inpatient medical care: evidence from the Affordable Care Act young adult mandate”. eng. In: *Journal of Health Economics* 39, pp. 171–187. ISSN: 1879-1646. DOI: [10.1016/j.jhealeco.2014.11.007](https://doi.org/10.1016/j.jhealeco.2014.11.007).
- 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.
- Bailey, James and Anna Chorniy (July 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](https://doi.org/10.1111/coep.12119).
- Baker, Dean (Apr. 2016). “Op-Ed: Will Obamacare end ‘job lock’?” In: *Los Angeles Times*.
- Barbaresco, Silvia, Charles J. Courtemanche, and Yanling Qi (Mar. 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](https://doi.org/10.1016/j.jhealeco.2014.12.004).
- Barnett, Jessica C and Edward R Berchick (Sept. 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 (Aug. 2021). “Health Insurance and Young Adult Financial Distress”. In: *Federal Reserve Bank of Philadelphia*. Vol. 19-54. Working papers.
- Blau, David M. and Donna B. Gilleskie (2001). “Retiree Health Insurance and the Labor Force Behavior of Older Men in the 1990s”. en. In: *The Review of Economics and Statistics* 83.1. Publisher: MIT Press, pp. 64–80.
- Dave, Dhaval et al. (Feb. 2015). “The Effect of Medicaid Expansions in the Late 1980s and Early 1990s on the Labor Supply of Pregnant Women”. In: *American Journal of*

- Health Economics* 1.2. Publisher: The University of Chicago Press, pp. 165–193. ISSN: 2332-3493. DOI: [10.1162/AJHE_a_00011](https://doi.org/10.1162/AJHE_a_00011).
- Depew, Briggs (Jan. 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](https://doi.org/10.1016/j.jhealeco.2014.11.008).
- Depew, Briggs and James Bailey (May 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](https://doi.org/10.1016/j.jhealeco.2015.01.004).
- Dillender, Marcus (July 2014). “Do more health insurance options lead to higher wages? Evidence from states extending dependent coverage”. In: *Journal of Health Economics* 36, pp. 84–97. ISSN: 0167-6296. DOI: [10.1016/j.jhealeco.2014.03.012](https://doi.org/10.1016/j.jhealeco.2014.03.012).
- Garfield, Rachel, Anthony Damico Published: Jan 25, and 2019 (Jan. 2019). *The Uninsured and the ACA: A Primer – Key Facts about Health Insurance and the Uninsured amidst Changes to the Affordable Care Act - How many people are uninsured?* en-US.
- Garthwaite, Craig, Tal Gross, and Matthew J. Notowidigdo (May 2014). “Public Health Insurance, Labor Supply, and Employment Lock”. en. In: *The Quarterly Journal of Economics* 129.2. Publisher: Oxford Academic, pp. 653–696. ISSN: 0033-5533. DOI: [10.1093/qje/qjw005](https://doi.org/10.1093/qje/qjw005).
- Goda, Gopi Shah, Monica Farid, and Jay Bhattacharya (Jan. 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](https://doi.org/10.3386/w21846).
- Goda, Gopi Shah, John B. Shoven, and Sita Nataraj Slavov (May 2011). “What Explains Changes in Retirement Plans during the Great Recession?” en. In: *American Economic Review* 101.3, pp. 29–34. ISSN: 0002-8282. DOI: [10.1257/aer.101.3.29](https://doi.org/10.1257/aer.101.3.29).
- Goldman, TR (Dec. 2013). *Progress Report: The Affordable Care Act’s Extended Dependent Coverage Provision | Health Affairs*. en. Library Catalog: www.healthaffairs.org.
- Grossman, Daniel S., Sebastian Tello-Trillo, and Barton Willage (Jan. 2022). *Health Insurance for Whom? The ‘Spill-up’ Effects of Children’s Health Insurance on Mothers*. Working Paper 29661. Series: Working Paper Series. National Bureau of Economic Research. DOI: [10.3386/w29661](https://doi.org/10.3386/w29661).

- Gruber, Jonathan and Brigitte Madrian (Dec. 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](https://doi.org/10.1016/S0047-2727(96)01621-0).
- (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: FrXlqoFwjjeQC. The Urban InSTITUTE. 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.
- Heim, Bradley, Ithai Lurie, and Kosali Simon (Oct. 2018). “Did the Affordable Care Act Young Adult Provision Affect Labor Market Outcomes? Analysis Using Tax Data”. en. In: *ILR Review* 71.5. Publisher: SAGE Publications Inc, pp. 1154–1178. ISSN: 0019-7939. DOI: [10.1177/0019793917744176](https://doi.org/10.1177/0019793917744176).
- Holahan, John (Jan. 2011). “The 2007–09 Recession And Health Insurance Coverage”. In: *Health Affairs* 30.1. Publisher: Health Affairs, pp. 145–152. ISSN: 0278-2715. DOI: [10.1377/hlthaff.2010.1003](https://doi.org/10.1377/hlthaff.2010.1003).
- Internal Revenue Service (May 2010). *Internal Revenue Bulletin: 2010-2*.
- Johnson, Richard W., Amy J. Davidoff, and Kevin Perese (June 2016). “Health Insurance Costs and Early Retirement Decisions:” en. In: *ILR Review*. Publisher: SAGE PublicationsSage CA: Los Angeles, CA. DOI: [10.1177/001979390305600410](https://doi.org/10.1177/001979390305600410).
- Kaiser Family Foundation (Apr. 2010). *Kaiser Health Tracking Poll – April 2010*.
- (Oct. 2020). *2020 Employer Health Benefits Survey - Section 2: Health Benefits Offer Rates*. en-US.
- Karoly, Lynn A. and Jeannette A. Rogowski (Oct. 1994). “The Effect of Access to Post-Retirement Health Insurance on the Decision to Retire Early”. en. In: *ILR Review* 48.1. Publisher: SAGE Publications Inc, pp. 103–123. ISSN: 0019-7939. DOI: [10.1177/001979399404800108](https://doi.org/10.1177/001979399404800108).
- 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](https://doi.org/10.2307/2534632).
- Robbins, Anthony S. et al. (Nov. 2015). “Association Between the Affordable Care Act Dependent Coverage Expansion and Cervical Cancer Stage and Treatment in Young

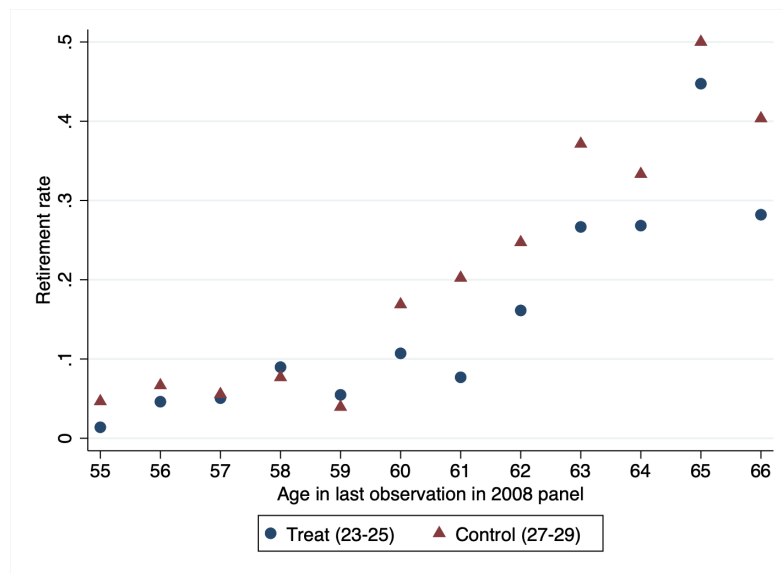
Women”. en. In: *JAMA* 314.20, pp. 2189–2191. ISSN: 0098-7484. DOI: [10.1001/jama.2015.10546](https://doi.org/10.1001/jama.2015.10546).

Seibold, Arthur (Apr. 2021). “Reference Points for Retirement Behavior: Evidence from German Pension Discontinuities”. en. In: *American Economic Review* 111.4, pp. 1126–1165. ISSN: 0002-8282. DOI: [10.1257/aer.20191136](https://doi.org/10.1257/aer.20191136).

Shoven, John B. and Sita Nataraj Slavov (2014). “The role of retiree health insurance in the early retirement of public sector employees”. en. In: *Journal of Health Economics* 38.C. Publisher: Elsevier, pp. 99–108.

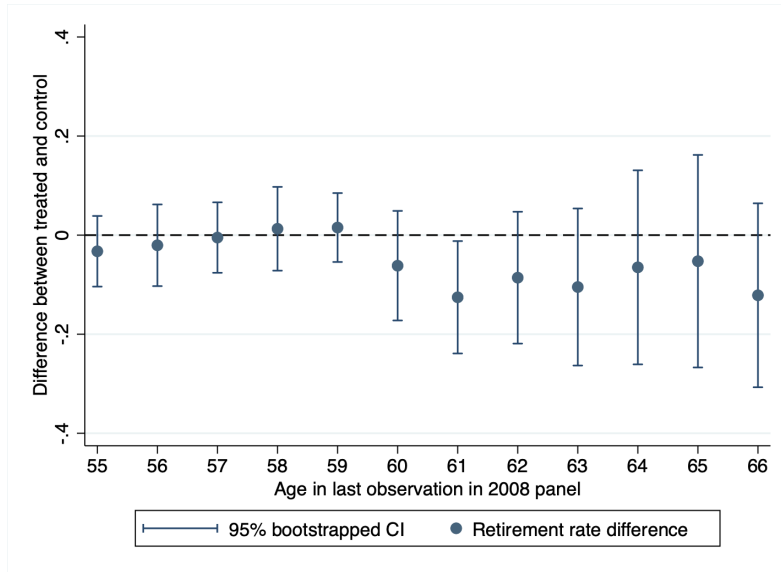
Swartz, Theda Skocpol and Katherine (Feb. 2014). *Obamacare cures 'job lock': Column*. en-US.

6 Figures



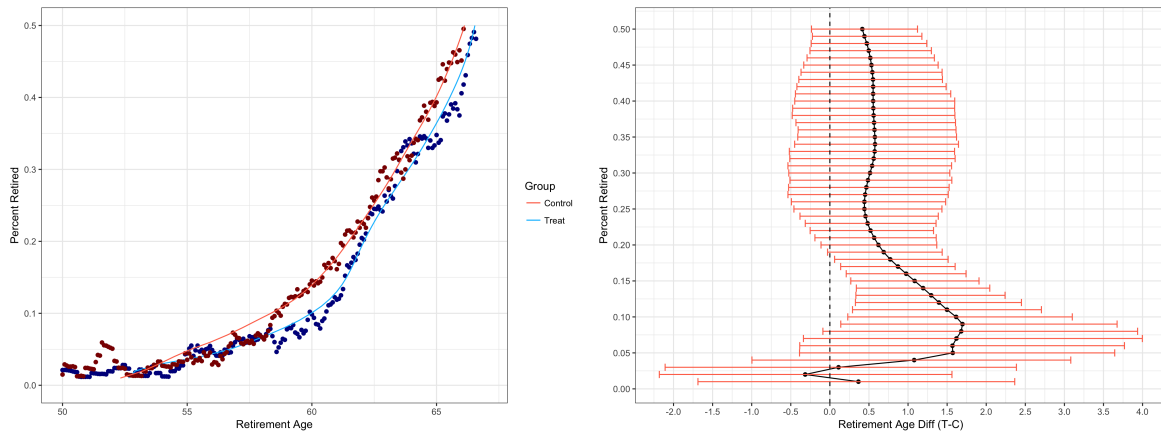
Treated is defined as having a child aged 23-25 in 2010, and control is defined as having a child aged 27-29 in 2010. Sample restricted to individuals who exit the panel on or after September 2011, and are 55-66 in their last observation. Data: 2008 SIPP.

Figure 1: Percent of individuals retired by age and treatment group in post-policy period



Difference calculated as treated average minus control average. Treated is defined as having a child aged 23-25 in 2010, and control is defined as having a child aged 27-29 in 2010. Sample restricted to individuals who exit the panel on or after September 2011, and are 55-66 in their last observation. 95% confidence intervals are estimated from non-parametric bootstrap sampling ($N = 1000$) with replacement from the initial sample, clustered by household. Data: 2008 SIPP.

Figure 2: Difference in average retirement rate between treated and control groups by age in post-policy period



Left: Cross-section of retirement rate is constructed using retirement rates for treated and control group individuals in the post-policy implementation period after September 2010. A LOESS regression (span = 0.15) is used to fit a bootstrapped smooth line to the retirement profile for the treated and control groups. Right: The horizontal distance between the fitted lines in cross-section of retirement rate (treatment minus control) is calculated between retirement rates of 1 percent and 50 percent. Confidence intervals are estimated from non-parametric bootstrap sampling ($N = 1000$). Data: 2008 SIPP.

Figure 3: Percent retired by age in post-policy period (left) and retirement age difference by quantile (right)

7 Tables

Table 1: Summary statistics by treatment and control

Variable	(1) Treat (23-25)	(2) Control (27-29)	(3) Diff (T-C)
Parent age in 2010	60.590 (3.204)	59.801 (3.296)	-0.789*** (0.165)
Monthly earnings (2009)	2,716.5 (3,503.4)	3,158.6 (3,946.3)	442.0** (189.3)
Paid job all of 2009	0.545 (0.498)	0.569 (0.496)	0.024 (0.025)
Female	0.660 (0.474)	0.617 (0.487)	-0.043* (0.024)
White (race)	0.835 (0.372)	0.829 (0.376)	-0.005 (0.019)
Hispanic (ethn.)	0.059 (0.235)	0.081 (0.273)	0.022* (0.013)
Married	0.788 (0.409)	0.797 (0.402)	0.009 (0.021)
HS	0.909 (0.288)	0.892 (0.310)	-0.017 (0.015)
College or higher	0.410 (0.492)	0.441 (0.497)	0.031 (0.025)
Priv. HI all of 2009	0.787 (0.410)	0.800 (0.400)	0.013 (0.021)
Num. children	2.304 (1.015)	2.478 (1.156)	0.174*** (0.055)
In union (2009, if in universe)	0.165 (0.372)	0.172 (0.378)	0.007 (0.024)
Paid hourly (2009, if in universe)	0.468 (0.499)	0.472 (0.500)	0.003 (0.032)
Observations	835	715	1,550

Treated is defined as having a child aged 23-25 in 2010, and control is defined as having a child aged 27-29 in 2010. Sample restricted to individuals who exit the panel on or after September 2011 and are aged 55-66 in their last observation.

Table 2: Effect of treatment on retirement rate in post-policy period, with age fixed effect and year fixed effect

Dep. variable: retired in post-policy period			
	(1)	(2)	(3)
	<i>Parent's age</i>		
	<i>55-66</i>	<i>55-60</i>	<i>61-66</i>
Treat	-0.0394** (0.017)	-0.00499 (0.017)	-0.0849*** (0.032)
Monthly earnings (thousands, 2009)	-0.00745*** (0.002)	0.000694 (0.002)	-0.0110*** (0.003)
Paid job all of 2009	-0.148*** (0.028)	-0.0813*** (0.029)	-0.231*** (0.047)
Female	0.00823 (0.022)	0.0158 (0.023)	0.0181 (0.037)
White (race)	0.0273 (0.022)	0.00468 (0.020)	0.0706* (0.042)
Hispanic (ethn.)	-0.0517 (0.036)	-0.0496* (0.029)	-0.0400 (0.075)
Married	0.128*** (0.030)	0.0920*** (0.030)	0.171*** (0.057)
HS only	0.000327 (0.030)	0.0150 (0.027)	-0.0399 (0.067)
College or higher	0.00219 (0.034)	0.00578 (0.032)	-0.0249 (0.068)
Priv. HI all of 2009	0.0685*** (0.025)	0.0378 (0.025)	0.0944* (0.050)
Num. children	-0.0154** (0.008)	-0.0134* (0.007)	-0.0175 (0.014)
State UR	0.00714 (0.007)	-0.00127 (0.007)	0.0196 (0.013)
Age FE	X	X	X
Year FE	X	X	X
Emp. Ctrls	X	X	X
Ctrl Avg	.21	.08	.33
Obs	1550	829	721

: $p < 0.1$, : $p < 0.05$, : $p < 0.01$, based on bootstrap standard errors (in parentheses) and asymptotic normal approximation. Regression includes age fixed effect, year of last observation fixed effect, state unemployment rate, demographic, educational, and pre-reform employment controls. Treated is defined as having a child aged 23-25 in 2010, and control is defined as having a child aged 27-29 in 2010. Paid job all of 2009 and private health insurance all of 2009 are dummy variables equal to 1 if the individual held a paid job every month in 2009 or had private health insurance every month in 2009. Pre-reform employment controls comprise of union status, paid hourly status, and worker class. Sample restricted to individuals who exit the panel on or after September 2011 and are aged 55-66 in the last observation in the panel. Standard errors are estimated from non-parametric bootstrap sampling ($N = 1000$) with replacement from the initial sample, clustered by household. Data: 2008 SIPP.

Table 3: Placebo Tests

A. Pre-implementation placebo

Dep. variable: retired in pre-policy period (September 2010)			
	(1)	(2)	(3)
	<i>Parent's age in last observation in 2008 panel</i>		
	<i>55-66</i>	<i>55-60</i>	<i>61-66</i>
Treat	-0.00121	-0.00794	0.0117
	(0.015)	(0.016)	(0.032)
Ctrl Avg	.15	.06	.29
Obs	1313	856	457

B. Ineligible Parents

Dep. variable: retired post-ACA			
	(1)	(2)	(3)
	<i>Parent's age in last observation in 2008 panel</i>		
	<i>55-66</i>	<i>55-60</i>	<i>61-66</i>
Placebo (31-33 in 2010)	-0.0290	-0.0218	-0.0376
	(0.020)	(0.023)	(0.030)
Ctrl Avg	.21	.08	.33
Obs	1507	639	868

C. 2004 Placebo Reform

Dep. variable: retired after placebo reform in 2004			
	(1)	(2)	(3)
	<i>Parent's age in last observation in 2004 panel</i>		
	<i>55-66</i>	<i>55-60</i>	<i>61-66</i>
Treat (23-25 in 2004)	0.0212	0.00303	0.0566
	(0.019)	(0.020)	(0.038)
Ctrl Avg	.17	.09	.28
Obs	1263	781	482

: $p < 0.1$, : $p < 0.05$, : $p < 0.01$, based on bootstrap standard errors (in parentheses) and asymptotic normal approximation. Regressions includes age fixed effect, year of last observation fixed effect, state unemployment rate, demographic, educational, and pre-reform employment controls. In Panel A, treated is defined as having a child aged 23-25 in 2010, and control is defined as having a child aged 27-29 in 2010. In Panel B, placebo is defined as having a child aged 31-33 in 2010 and control is defined as having a child aged 27-29 in 2010. In Panel C, treated is defined as having a child aged 23-25 in 2004 and control is defined as having a child aged 27-29 in 2004. Sample restricted to individuals aged 55-66 in September 2010 or 2004. Standard errors are estimated from non-parametric bootstrap sampling ($N = 1000$) with replacement from the initial sample, clustered by household.

A Robustness Checks

As a robustness check, I use an alternate definition of cohort – birth year. The main results define cohorts by the age of an individual in their last observation in the panel, but an alternative way to define cohorts would be by birth year. These results will not exactly mirror those of the age cohort definition because each individual’s last observation in the panel differs, so two individuals born in the same year may exit the panel at different ages.¹³ If we define cohorts by age, then the gap in retirement rates represents the difference in the percent of individuals who are retired by a given age. Defining cohorts by birth year changes the interpretation of the difference to be the difference in percent of individuals in the same birth cohort who are retired by the time they exit the panel. The differences by birth year are plotted in Figure C.2 and regression results are reported in Tables B.4. The results by birth year are similar to the results by age, with older individuals driving a gap in retirement between the two groups. For parents born in 1947-1958 (who would be 55-66 in 2013), the treated group is on average 3.3 percentage points (15.7 percent) less likely to be retired. The difference is larger for older parents born in 1947-1952, where the treated group is 7.4 percentage points (21.8 percent) less likely to be retired.

I also consider alternate specifications that widen the bandwidth for defining treatment and control. The main analysis uses a bandwidth of 3 years, but alternatively I could have chosen to compare different-sized bandwidths. To see whether this affects the results, I define treatment and control with varying bandwidths from 1 to 5 years. A one-year bandwidth would compare parents of 24 year olds in 2010 (treated) to parents of 27 year olds (control), and a five-year bandwidth would compare parents of 20-25 year olds to parents of 27-32 year olds. The regression coefficients of treatment for varying bandwidths is plotted in Figure C.3; the results are relatively stable across bandwidths.

¹³About 92% of individuals in the sample leave the sample in 2013, meaning that most of the time when two individuals in the same birth cohort exit at different ages, the ages at which they exit will be at most one year apart from each other.

B Appendix Tables

Table B.1: Difference in average retirement rates in post-policy period between treatment and control group by parent's age

Age	N	Diff. in retirement rate	95% bootstrap CI	Ctrl retire rate	$\frac{\text{Difference}}{\text{Control avg}}$
55	115	-0.033	(-0.104, 0.039)	0.047	-0.701
56	125	-0.021	(-0.103, 0.062)	0.067	-0.308
57	151	-0.005	(-0.076, 0.066)	0.056	-0.089
58	156	0.013	(-0.072, 0.097)	0.077	0.167
59	149	0.015	(-0.054, 0.085)	0.039	0.388
60	133	-0.062	(-0.172, 0.049)	0.169	-0.365
61	136	-0.125	(-0.239, -0.012)	0.202	-0.620
62	151	-0.086	(-0.219, 0.047)	0.247	-0.348
63	130	-0.105	(-0.263, 0.054)	0.371	-0.282
64	104	-0.065	(-0.261, 0.131)	0.333	-0.195
65	104	-0.053	(-0.267, 0.162)	0.500	-0.105
66	96	-0.121	(-0.307, 0.064)	0.404	-0.301
55-60	829	-0.015	(-0.049, 0.019)	0.079	-0.194
61-66	721	-0.093	(-0.164, -0.021)	0.331	-0.280
55-66	1550	-0.054	(-0.094, -0.014)	0.208	-0.259

Difference calculated as treated average minus control average. Treated is defined as having a child aged 23-25 in 2010, and control is defined as having a child aged 27-29 in 2010. Sample restricted to individuals who exit the panel on or after September 2011 and are aged 55-66 in their last observation. The average difference for age ranges (i.e., 55-60) is calculated by estimating separate differences for each individual age cohort and taking the mean of the estimates. Confidence intervals are estimated from non-parametric bootstrap sampling ($N = 1000$) with replacement from the initial sample, clustered by household.

Table B.2: Regression of retirement on treatment, with parent age fixed effects rounded to 0.5

Dep. variable: retired in post-policy period			
	(1)	(2)	(3)
	<i>Parent's age post-ACA</i>		
	<i>55-66</i>	<i>55-60</i>	<i>61-66</i>
Treat	-0.0412**	-0.00377	-0.0915***
	(0.017)	(0.019)	(0.032)
Age (round to 0.5) FE	X	X	X
Year FE	X	X	X
Emp. Ctrl	X	X	X
Ctrl Avg	.21	.08	.33
N	1550	714	721

: $p < 0.1$, : $p < 0.05$, : $p < 0.01$, based on bootstrap standard errors (in parentheses) and asymptotic normal approximation. Regression includes age (rounded to 0.5) fixed effect, year of last observation fixed effect, state unemployment rate, demographic, educational, and pre-reform employment controls. Treated is defined as having a child aged 23-25 in 2010, and control is defined as having a child aged 27-29 in 2010. Paid job all of 2009 and private health insurance all of 2009 are dummy variables equal to 1 if the individual held a paid job every month in 2009 or had private health insurance every month in 2009. Pre-reform employment controls comprise of union status, paid hourly status, and worker class. Standard errors are estimated from non-parametric bootstrap sampling (N = 1000) with replacement from the initial sample, clustered by household. Data: 2008 SIPP.

Table B.3: Regression of treatment on retirement post-ACA, including employment-related controls

Dep. variable: retired in post-policy period					
	(1)	(2)	(3)	(4)	(5)
Treat	-0.0382** (0.017)	-0.0370** (0.018)	-0.0378** (0.017)	-0.0378** (0.017)	-0.0385** (0.017)
State UR		X	X	X	X
Worker class			X	X	X
Hourly				X	X
Union					X
Obs	1550	1550	1550	1550	1550


: $p < 0.1$, : $p < 0.05$, : $p < 0.01$, based on bootstrap standard errors (in parentheses) and asymptotic normal approximation. Regressions all include age fixed effect, year of last observation fixed effect, demographic, educational, and pre-reform working and private health insurance controls. Treated is defined as having a child aged 23-25 in 2010, and control is defined as having a child aged 27-29 in 2010. Pre-reform employment controls comprise of union status, paid hourly status, and worker class. Sample restricted to individuals who exit the panel on or after September 2011 and are aged 55-66 in the last observation in the panel. Standard errors are estimated from non-parametric bootstrap sampling ($N = 1000$) with replacement from the initial sample, clustered by household. Data: 2008 SIPP.

Table B.4: Effect of treatment on retirement rate in post-policy period with birth year fixed effect and year fixed effect

Dep. variable: retired in post-policy period			
	(1)	(2)	(3)
	<i>Parent's birth year</i>		
	1947-1958	1947-1952	1953-1958
Treat	-0.0351*	-0.0801**	0.000757
	(0.018)	(0.034)	(0.017)
Birth Year FE	X	X	X
Year FE	X	X	X
Emp. Ctrl	X	X	X
Ctrl Avg	.21	.34	.08
Obs	1552	715	837

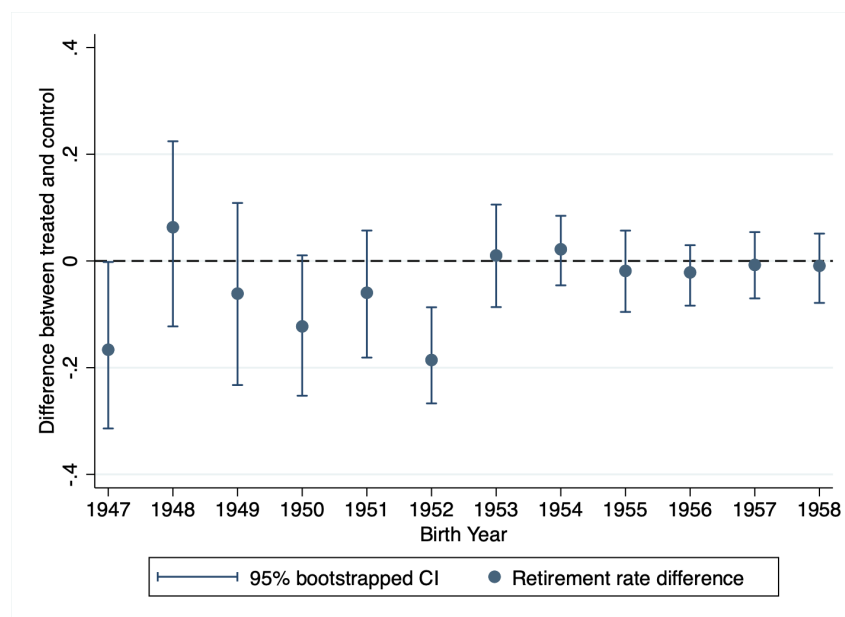
: $p < 0.1$, : $p < 0.05$, : $p < 0.01$, based on bootstrap standard errors (in parentheses) and asymptotic normal approximation. Regression includes birth year fixed effect, year of last observation fixed effect, state unemployment rate, demographic, educational, and pre-reform employment controls. Treated is defined as having a child aged 23-25 in 2010, and control is defined as having a child aged 27-29 in 2010. Paid job all of 2009 and private health insurance all of 2009 are dummy variables equal to 1 if the individual held a paid job every month in 2009 or had private health insurance every month in 2009. Pre-reform employment controls comprise of union status, paid hourly status, and worker class. Sample restricted to individuals who exit the panel on or after September 2011 and were born between 1947 and 1958. Standard errors are estimated from non-parametric bootstrap sampling ($N = 1000$) with replacement from the initial sample, clustered by household. Data: 2008 SIPP.

C Appendix Figures



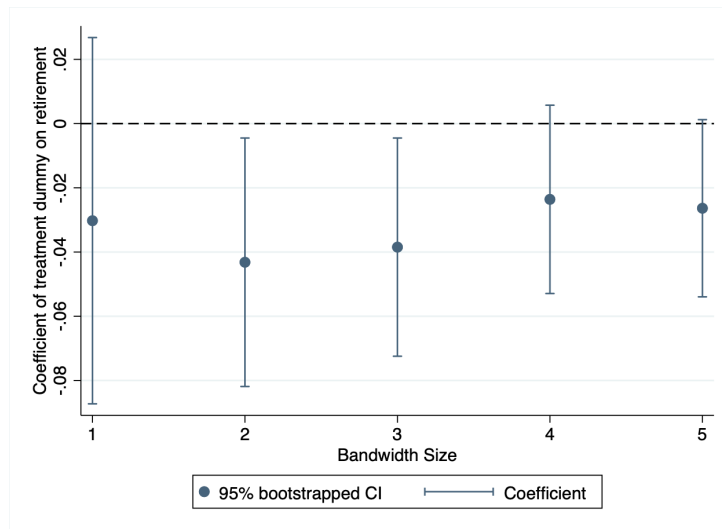
Treatment in 2008 SIPP (left) is defined as having a child aged 23-25 in 2010, and control is having a child aged 27-29. Treatment in the 2004 SIPP (right) is defined as having a child aged 23-25 in 2004, and control is having a child aged 27-29.

Figure C.1: Percent retired by treatment and control, 2008 SIPP (left) and 2004 SIPP (right)



Difference calculated as treated average minus control average. Treated is defined as having a child aged 23-25 in 2010, and control is defined as having a child aged 27-29 in 2010. Sample restricted to individuals who exit the panel on or after September 2011, and are 55-66 in their last observation. 95% confidence intervals are estimated from non-parametric bootstrap sampling ($N = 1000$) with replacement from the initial sample, clustered by household. Data: 2008 SIPP.

Figure C.2: Difference in retirement rate in post-policy period between treated and control groups by birth year



This figure plots the coefficients of a regression of a treatment dummy, where treatment is defined with bandwidths for child's age of varying size, on retirement. The treated group comprises parents whose children under 26 in 2010, and the control group comprises parents whose children are over 26 in 2010. Regression includes age fixed effect, year fixed effect, state unemployment rate, demographic, educational, and pre-reform employment controls. Sample restricted to individuals who exit the panel on or after September 2011 and are aged 55-66 in their last observation in the panel. 95% confidence intervals are estimated from non-parametric bootstrap sampling ($N = 1000$) with replacement from the initial sample, clustered by household. Data: 2008 SIPP.

Figure C.3: Coefficient of treatment definitions with varying bandwidths on retirement in post-policy period

D Data Appendix

The data used for the main analysis are waves 1-16 of the 2008 SIPP panel and 1-12 of the 2004 SIPP Panel, which can be downloaded from [NBER](#) (last accessed April 4, 2020). Specifically, I use the core files for each wave, as well as Topical Module 2 (Fertility History).

I exclude observations with interview status equal to "noninterview," who are missing a sample unit identifier, and who are listed as men but have a non-missing value for the year they gave birth to their last child (I take this to mean that either the sex or the birth year of their child is coded incorrectly).

I define an individual to be retired if their employment status (*rmesr*) is listed as "No job all month, no time on layoff and no time looking for work" and the main reason for not working in the reference period (*ersnowrk*) is that they are retired.

Ages For each mother, I collect the year of her last birth (*tlbirtyr*) if she had more than one birth and the year of her first (and only) birth if she had only one birth (*tfbrthyr*). For mothers who are married, I assign to their current husband the same value for child's birth year. For parent's ages, I define age at last observation in panel using birth year and birth month to calculate age in months, and then round to the closest integer.

Matching husbands to wives I use the person number of an individual's spouse (*epnsps*) within the same household (identified using *ssuid* and *eentaid*) to identify an individual's spouse. After linking husbands to their wives, I then assign the year of last birth of the wife to her husband. Using this method, I identify husbands in the sample for 88% women in the control group and 86% in the treated group for women who say they are married with a spouse present in their first observation.

Treatment, control, placebo, and robustness The treated group is defined as parents whose children's births were between 1985 and 1987, meaning their children were 23-25 in 2010. The control group is defined as parents whose children's births were between 1981 and 1983, meaning their children were 27-29 in 2010. *Treat* is a dummy which is 1 for individuals in the treated group and 0 for individuals in the control group. The first placebo selects parents who were aged 55-66 in September 2010 and defines treatment and control the same way as the main results. The second placebo

compares parents of children who were 31-33 in 2010 (born between 1977-1979) and 0 for parents of 27-29 year olds in 2010. The third placebo test defines treatment as having a child aged 23-25 in 2004 (born 1979-1981) and control as having a child aged 27-29 in 2004 (born 1975-1977). Bandwidth of 1 compares parents of 25 year olds to 27 in 2010, bandwidth of 2 compares parents of children aged 24-25 to 27-28, bandwidth of 3 compares parents of children aged 23-25 to 27-29 (main results), bandwidth of 4 compares parents of children aged 22-25 to 27-30, and bandwidth of 5 compares parents of children aged 21-25 to 27-31.

Other variables Monthly state unemployment rates are downloaded from the [Federal Reserve Economic Data](#) (last accessed: September 21, 2020). White is a dummy for whether an individual is white, defined using the variable *erace*. Female is a dummy for whether an individual is female, defined using the variable *esex*. Hispanic is a dummy for whether an individual is Hispanic, Spanish, Latino, defined using the variable *eorigin*. Married is defined as being never married, widowed, divorced, or separated and is defined using the variable *ems*. High school is defined as having *eeducate* as 39 and 40, and college is a dummy for having *eeducate* greater than or equal to 41. Private HI in 2009 is a dummy variable for having private health insurance for at least half of 2009, defined by taking the average of *ehimth* in 2009. Employed in 2009 is a dummy for working a paid job for all of 2009, defined by taking the average of *epdjbthn* for all of 2009. Age in 2009 is the difference between 2009 and birth year. Number of children is defined using *tfrchl* for men and *tmomchl* for women, and is set as 0 if missing. Monthly income in 2009 is defined as the mean of *tpearn* in 2009 and is missing if the individual did not work in 2009. Pre-implementation union status is defined using *eunion1*, paid hourly is defined using *epayhr1*, and worker class is defined using *ec1wrk1*; all are defined using the first job an individual reports.