Retirement Lock and Prescription Drug Insurance: Evidence from Medicare Part D

Gal Wettstein*

11/11/15

Abstract

This paper examines whether the lack of an individual market for prescription drug insurance causes individuals to delay retirement. I exploit the quasi-experiment of the introduction of Medicare Part D, which provided subsidized prescription drug insurance to all Americans over age 65 beginning in 2006. Using a differences-in-differences design, I compare the labor outcomes of individuals turning 65 just after 2006 to those turning 65 just before 2006 in order to estimate the causal effect of eligibility for Part D on labor supply. I find that individuals at age 65 who would have otherwise lost their employer-sponsored drug insurance upon retirement decreased their rate of full-time work by 8.4 percentage points due to Part D, in contrast to individuals with retiree drug insurance even after age 65 for whom no significant change was observed. This reduction was composed of an increase of 5.9 percentage points in part-time work and 2.5 percentage points in complete retirement. I use these estimates to quantify the extent of the distortion due to drug insurance being tied to employment, and the welfare gains from the subsidy correcting that distortion. The results suggest that individuals value \$1 of drug insurance subsidy as much as \$3 of Social Security wealth.

*I am very grateful to my advisors, David Cutler, Edward Glaeser, Nathaniel Hendren, and Lawrence Katz, for their invaluable feedback and advice. I would also like to thank Mitra Akhtari, Alberto Alesina, Raj Chetty, Itzik Fadlon, Peter Ganong, Simon Jäger, Marc Kaufman, Divya Kirti, Danial Lashkari, Jing Li, Filippo Mezzanotti, Arash Nekoei, Mikkel Plagborg-Møller, Daniel Sullivan, Patrick Turley, and participants in the labor and public finance seminars at Harvard for helpful discussions and comments. I thank the Boston College Center for Retirement Research for their generous financial support.

1 Introduction

Do Americans work in order to maintain health benefits? In this paper I address this question by focusing on retiree prescription drug insurance and utilizing the 2006 introduction of Medicare Part D as a quasi-experiment. Stand-alone prescription drug insurance is almost non-existent on the individual market for those below age 65, and before Part D's introduction Medigap policies covering drugs for those over 65 offered limited coverage and were rarely taken up (Pauly and Zeng, 2004).[1] Thus the majority of Americans acquire their health insurance through an employer, and virtually all employer plans cover prescription drugs.[2] Therefore, individuals dependent on their employers for insurance may be "retirement locked": prevented from optimally retiring due to this extraneous consideration. The extent of retirement lock is important for many reasons, not least its role in the design of policies, such as the Affordable Care Act (ACA), which weaken the link between employment and insurance, impacting both the benefits of such policies and their costs.

This paper addresses the question of retirement lock by exploiting the quasi-experiment induced by the 2006 introduction of Medicare Part D. Part D expanded traditional Medicare in 2006 to give everyone over age 65 access to subsidized prescription drug insurance, indirectly inducing a sharp change in the incentives of individuals regarding whether to retire. Whereas before 2006 prescription drug insurance was available almost exclusively through employer-sponsored insurance (ESI) irrespective of age, after 2006 it became available to everyone over age 65 regardless of availability of ESI. I examine the effect of Part D using a differences-in-differences design: I estimate the causal effect of Part D by comparing labor outcomes of individuals reaching age 65 before 2006 to those reaching age 65 after 2006. I find that eligibility for Part D substantially decreased the labor supply of those who would have previously been dependent on their employers for drug insurance.

In order to focus on individuals who were potentially retirement locked to begin with, I consider those who had retiree health insurance until age 65 – such individuals continue to benefit from their employer coverage even if they retire. I divide this population into two groups: those who would be covered by their employer plan only until age 65 if they retired, and those who had retiree coverage after age 65 as well. The former constitute the "treatment group" – for them retiring implies a loss of drug coverage at age 65 before 2006, but no such constraint exists after 2006. Those with retiree coverage after age 65 were not retirement locked before or after 2006 – Medicare Part D should not change their retirement decisions through retirement lock. They are therefore a "control group" in a triple-differences design. If relaxation of retirement lock is the sole mechanism by which the labor supply of the treatment group is affected by Part D, it should exhibit a reduction in labor supply at

[1] In 2005 only 3.2% of Medigap policyholders in federally standardized plans chose plans offering any drug coverage at all (America's Health Insurance Plans, 2006).

[2] In 2014 about 70% of Americans were eligible for health insurance from their employer, and 99% of employer plans also covered prescription drugs (Kaiser Family Foundation, 2014).

age 65 in 2006, while there should be no change for the control group.

I find results consistent with these predictions. Those in the group with retiree coverage only to age 65 reduced their rate of full-time work by 8.4 percentage points more at age 65 after 2006 than they did at age 65 before 2006; for the group with retiree coverage over age 65 I observe a statistically insignificant 2 percentage point *increase* in full-time work. On a baseline of 35 percentage points of full-time work, this amounts to a 24% reduction in the rate of full-time work upon eligibility for Medicare Part D among the treated.[3] This drop in full-time work was largely composed of an intensive margin response of a shift to part-time work – an increase of 5.9 percentage points – with a smaller but substantial share accounted for by the extensive response of full retirement – an increase of 2.5 percentage points.

To interpret this effect I compare the reduction in labor supply due to Part D to that predicted to result from an increase in Social Security benefits. I find that a $1 subsidy to drug insurance leads to a labor response equivalent to $3 of Social Security. These substantial estimated behavioral responses to the relaxation of retirement lock suggest potential inefficiency in the existing individual drug insurance market in the absence of Part D. Using a simple model of labor responses to Medicare Part D's introduction I map the observed changes in labor supply due to the subsidy to individuals' willingness to pay for the subsidy out of retirement income. This implies a willingness to pay of $3 for every dollar of the subsidy among retirees.

The large estimated willingness to pay suggests the potential for large welfare gains from a subsidy to drug insurance for retirees. However, because the provision of insurance allows individuals to retire, this increases the government costs of Part D because of foregone tax revenue from those who would otherwise be working (i.e. a fiscal externality). To assess this cost I estimate the fiscal externality due to Part D using the labor supply responses of the treated. I find a large fiscal externality of 68 cents on the dollar for every dollar of subsidy. However, the valuation of the subsidy is larger than the cost, leading to a marginal value of public funds of Part D of $1.80 per dollar of subsidy, or a net social gain of 80 cents on the dollar.

The differences-in-differences approach I take allows me to non-parametrically account for the myriad changes which might otherwise affect the labor supply of 65 year-olds, such as health status and age dependent factors (e.g., pensions and full social security eligibility). It requires me to assume only that there was no sharp change in these factors in 2006. The fact that the magnitude of retirement lock can thus be cleanly estimated in a reduced form way, independent of strong modeling assumptions, is an advantage of this approach. It therefore complements past efforts to structurally estimate the effect of health insurance availability on retirement behavior.[4]

[3] This 35 percentage point baseline is the rate of full-time work for individuals aged 65-68 in the years 2006-2010, net of the estimated effect of Part D.

[4] There exists a rich literature attempting to structurally estimate the effect of health insurance availability on retirement. The conclusions of these papers are diverse, with some finding little effect of employer-insurance on retirement (e.g., Gustman and Steinmeier, 1994, Lumsdaine et al., 1994), while others find significant effects (for example, Rust and Phelan, 1997, Blau and Gilleskie, 2006, French and Jones, 2011).

My reduced form approach to estimation of retirement lock is most closely related to a number of previous papers which look at quasi-experiments estimating conceptually similar effects.[5] The predominant source of exogenous variation in this literature has been based on continuation of coverage laws (COBRA). This literature tends to find significant effects of relatively small magnitude (Madrian et al., 1994, Gruber and Madrian, 1995). However, the variation induced by COBRA can by necessity only identify the effect of a year or two of continued coverage; and the law still requires individuals to pay for coverage with after-tax dollars, making it less generous than employer sponsored insurance. Thus, both within the structural and reduced form attempts to estimate the extent of retirement lock there have been inconclusive results, along with a limited set of policy variations allowing clean identification, as outlined in Gruber and Madrian [2004].[6]

My approach to welfare is similar to that of Gruber and Madrian [1995]. There the authors provide a sense of the scale of retirement lock by comparing its impact to retirement wealth. They find that one year of continuation of coverage has the same effect on retirement as $13,600 of pension wealth, substantially higher than the $3,600 they estimate to be the cost of such coverage.[7] I formalize this comparison in a way which allows identification of both the distortion in labor supply induced by the inefficiency of the individual insurance market, and the willingness to pay of individuals for correcting this inefficiency. Such inference of welfare from labor market responses is related to Shimer and Werning [2007], Chetty [2008], Hendren [2013a], and Fadlon and Nielsen [2015].

This paper also contributes to the literature on Medicare Part D itself, particularly regarding welfare analysis of the program. An overview of early results on the structure and the effects of Part D is available in Duggan et al. [2008]. A great deal of research quantifies the effect of Medicare Part D on health expenditures and outcomes: for example, Engelhardt and Gruber [2011] find that Medicare Part D increased prescription drug coverage and utilization among the elderly, while reducing their out-of-pocket spending substantially.[8] They estimate

[5]A number of reduced form analyses not relying on quasi-experiments are also relevant here. These include Karoly and Rogowski [1994], Rogowski and Karoly [2000], Blau and Gilleskie [2001], Nyce et al. [2013], and Shoven and Slavov [2014]. These studies tend to find large effects of availability of retiree health insurance on retirement. The current paper's identification strategy circumvents some of the concerns raised by the lack of exogenous variation in insurance coverage in these studies, such as potential unobserved correlation of employer coverage with employer pension plans, or selection of individuals with particular preferences into matches with employers who provide health insurance (Gruber and Madrian, 2004).

[6]Two recent papers estimate the effect of health insurance on employment using variation other than the introduction of COBRA: Baicker et al. [2014] and Garthwaite et al. [2014] use exogenous enrollment changes in Medicaid. However, these papers do not focus on typical individuals near retirement, but rather on prime working age individuals who are in addition quite poor (on the margin of Medicaid eligibility). Furthermore, the two papers come to different estimates, with the latter finding substantial employment lock and the former finding only small and insignificant effects. I find that this divergence may be partially reconciled by the fact that most individuals who reduce their labor supply due to availability of subsidized individual insurance do so on an intensive margin. Baicker et al. [2014] observe only employment, without the ability to differentiate full- and part-time work.

[7]The authors speculate that this may be because policies available on the individual market generally exclude preexisting conditions from coverage, or because a number of early retirees are refused coverage at any price.

[8]Other papers in this literature include Lichtenberg and Sun [2007], Kaestner et al. [2014], Abaluck et al.

the welfare benefits of Medicare Part D by focusing on the gains due to increased insurance. These same authors also estimate large crowd-out of private insurance by the new program, cases in which there was no net gain in insurance per se. This paper complements such calculations by considering gains in welfare precisely among those whose private (employer) prescription drug insurance is potentially crowded out by Part D.[9] Rather than the null effect on welfare implied by the idea of crowd out I show substantial welfare gains from Medicare Part D; however, these gains accrue mostly along the margin of avoided labor disutility, rather than of a less risky distribution of health expenditures.

The structure of the paper is as follows: section 2 presents a simple conceptual model of retirement lock. Section 3 provides the institutional details of Medicare Part D. Section 4 describes the data and the identification strategy. Section 5 contains the main empirical results. Section 6 contains some robustness checks for these results. Section 7 discusses the implications of these results for welfare. Section 8 concludes.

2 Conceptual Framework of Retirement Lock

In this section I develop an extensive margin model of labor supply, which serves two purposes. First, it formally states what is meant by "retirement lock". I define this concept as the distortion arising in labor supply due to inefficiency in the individual insurance market. Second, I develop a framework for thinking about the subsidized prescription drug insurance offered through Medicare Part D and its effect on labor supply. I show that a negative labor supply effect of the policy does not in itself provide evidence of retirement lock; and provide a test that can provide such evidence by comparing the effect on labor of a subsidy to individual market insurance to that of increasing retirement income.

Individual preferences I assume individuals derive utility from consumption, c_i, and separable disutility from labor, v_i, such that:

$$U_i(c_i, l_i) = u_i(c_i) - v_i * l_i, \tag{1}$$

where $l_i = 1$ indicates full-time work and $l_i = 0$ otherwise. v_i is distributed according to a cumulative density function $G(v_i)$, with a probability density function $g(v_i)$. The realization of v_i is known to individuals at the time they make their labor and insurance choices. u_i is individual i's utility of consumption; $u'(c) > 0, u''(c) < 0$.

Individual budget Individuals' gross income is a function of their labor, $I(l_i)$ such that $I(0) < I(1)$. Individuals also face stochastic drug costs, Y_i. They can purchase insurance

[2015], and Ayyagari and Shane [2015].

[9] 97% of the treatment group had some form of prescription drug coverage before becoming eligible for Medicare Part D; see table 1.

against these costs at the quantity of x_i, leading to out-of-pocket costs of $y_i(Y_i, x_i)$ so that for all i and for each realization of Y_i: $\frac{dy_i}{dx_i} \equiv y_i' < 0, \frac{d^2 y_i}{dx_i^2} \equiv y_i'' > 0$.

$x_i(p)$ is i's demand for insurance as a function of the price of a unit of insurance. To capture the intuition of insurance being more expensive or of poorer quality on the individual market relative to the group market, the price of a given quantity of insurance will be permitted to differ based on whether the individual works full-time or not. In particular, the price of insurance will be $p(l_i)$ so that $p \equiv p(1) < p(0) \equiv P$.[10] In addition, I consider a stylized policy like Medicare Part D, of subsidizing the price of insurance only on the individual market by s – i.e., $s(1) = 0, s(0) = s$. Thus the consumer price on the individual market for a unit of insurance will be $P - s$, while the price for individuals getting their insurance on the group market is p.[11]

In sum, for each realization of Y_i and choice of (l_i, x_i), consumption for individual i is given by:

$$c_i = I(l_i) - y_i(Y_i, x_i) - (p(l_i) - s(l_i)) * x_i(p(l_i) - s(l_i)). \tag{2}$$

Optimal labor choice Individuals maximize their expected utility with respect to Y_i (noted by E_Y) over their choice of labor and the quantity of insurance they buy. An individual will work full-time if her expected utility of consumption from working minus her disutility of labor is greater than her expected utility of consumption when not working. Equivalently, there will be a cutoff level of labor disutility below which individuals choose to work full-time and above which they choose not to. That is, i works full-time if and only if:

$$E_Y[u_i(c_{1i}^*) - u_i(c_{0i}^*)] \equiv \overline{v}(s) > v_i, \tag{3}$$

where c_{1i}^*, c_{0i}^* are the values of consumption after having optimally chosen the level of insurance conditional on labor choice. $\overline{v}(s)$ is the cutoff value of labor disutility above which individuals choose to stop working full-time. An individual with labor disutility $v_i = \overline{v}(s)$ is precisely indifferent between the expected value of full-time work, with its higher income and lower price of insurance, and the expected value of retirement, with its lower income and higher price of insurance.[12]

[10] There are a number of reasons why the price of insurance on the individual market might be higher than on the group market. First, health insurance markets in general suffer from adverse selection (Hackmann et al., 2012, Hendren, 2013b). This is particularly true of prescription drug insurance, due to the persistence of drug expenditures over time (Pauly and Zeng, 2004). Second, there are fixed costs in contracting with an insurer. This is the result of administrative costs as well as the complexity of the choice problem which is particularly difficult for the elderly in the context of drug insurance (Abaluck and Gruber, 2011). Third, the exemption of employer sponsored insurance from the income tax leaves it cheaper in after-tax dollars than individual market alternatives. Fourth, the difficulty of forming long-term insurance contracts which do not result in premium increases following a negative health event makes risk pooling an integral part of insurance (Cutler, 1994).

[11] Medicare Part D also subsidized the group market at a lower rate. For simplicity I assume this subsidy was 0. What matters for this analysis is the change in the differential subsidy.

[12] In principle, all individuals could be made indifferent between working and retiring if employers could offer worker-specific $I(1)$. Two frictions preventing this are noted by Gruber and Madrian [2004]: the

Benchmark optimal insurance choice Individuals choose the amount of insurance to purchase conditional on their choice of labor. For a given l_i the first order condition for the optimal choice of x is:

$$\frac{dE_Y[u_i(c_i)]}{dx} = E_Y[u' * \frac{dc}{dx}] = -E_Y[u' * (y_i' + p(l_i))] = 0. \tag{4}$$

A market in which this condition holds can be thought of as "constrained efficient": given the possible insurance contracts in the market individuals will choose x so that in expectation the utility lost due to the dollars spent on an additional increment of x will equal the utility from the dollars saved on drug expenditures from that additional insurance.

There are numerous reasons to think that this first order condition does not hold in this form in practice. I will show below that inefficiency in the labor market will only occur if the insurance market is indeed inefficient; i.e., if this first order condition does not hold. Trivially, if the insurance market does not exist, or does not exist for some individuals (such as those with preexisting conditions, see Hendren, 2013b), then the first order condition for insurance will not hold for every i. This is a close approximation to the prevailing drug insurance market for those under age 65, for example, or to the market for those over 65 before Medicare Part D, due to adverse selection (Pauly and Zeng, 2004).

Analysis of changes in the level of s Define the marginal utility of consumption of the i'th individual as a retiree, given a subsidy of s: $u_{0i}'(s) \equiv u_i'(c_{0i}^*)$. The change in the cutoff disutility of labor when the subsidy is increased is given by differentiating equation (3): $\frac{d\overline{v}(s)}{ds} = -E_Y[u_{0i}'(s) * \frac{dc_i}{ds}]$. Therefore the change in the actual share of individuals working full-time will be:

$$\frac{dG(\overline{v}(s))}{ds} = -g(\overline{v}(s)) * E_Y[u_{i0}'(s) * \frac{dc_i}{ds}] < 0. \tag{5}$$

$E_Y[u_{0i}'(s) * \frac{dc_i}{ds}]$ must be weakly positive by revealed preference: individual welfare cannot decrease when an (unfunded) subsidy is increased. Therefore, labor supply would decline with increases in the subsidy regardless of whether or not there were any inefficiency in the insurance market. Both a substitution effect of giving another dollar conditional on retirement, and an income effect of making individuals richer work in the same direction in this case. To find evidence of inefficient labor supply due to retirement lock the bar is higher – there must be a decline in labor supply beyond what would result from a mere increase in retirement income due to the subsidy.

We can decompose retirement income, $I(0)$, into Social Security benefits, b, and other income. If instead of increasing s we increase b, the cutoff labor disutility change will be

first is the administrative cost of designing worker-specific contracts. The second is preference revelation constraints, where employers do not know the individual valuations of insurance and of leisure. In this model this latter point can be supported by assuming employers do not know each individual's v_i and, potentially, heterogeneity in the distribution of Y_i and preference parameters such as risk aversion. There is some evidence that while employers can offset the value of benefits by reducing compensation for groups of workers (e.g., Gruber, 1994), they cannot do so at an individual level (for example, Chetty et al., 2011).

(suppressing the arguments of $\overline{v}$): $\frac{d\overline{v}}{db} = -E_Y[u'_{0i}(s)]$. Such a change leads to a corresponding change in the share of full-time workers of:

$$\frac{dG(\overline{v})}{db} = -g(\overline{v})E_Y[u'_{0i}(s)]. \tag{6}$$

Note further that an increase of 1 in s corresponds to an increase of x_i dollars to individual i, or one dollar per unit of insurance. What I look for to provide evidence of retirement lock is a large ratio of the effect on labor supply of an increase of one dollar of subsidy to retiree insurance versus the effect of an increase of one dollar in retirement income:

$$\frac{\frac{dG(\overline{v}(s))}{ds}/x_i}{\frac{dG(\overline{v})}{db}} = \frac{E_Y[u'_{0i}(s) * \frac{dc_i}{ds}]/x_i}{E_Y[u'_{0i}(s)]}. \tag{7}$$

It is helpful here to illustrate the benchmark expected magnitude of this ratio if indeed individuals faced an efficient individual market for insurance.

Claim. In the presence of efficient insurance markets the effect of a dollar's worth of subsidy on labor supply is equal to the effect of a dollar of retirement income.

Proof. Plugging in the first order condition from equation (4) into equation (7) gives:

$$\frac{\frac{dG(\overline{v}(s))}{ds}/x_i}{\frac{dG(\overline{v})}{db}} = 1 - \frac{dx_i(P-s)}{ds} * \frac{E_Y[u'_{0i}(s) * (((P-s)+y')]}{x_i E_Y[u'_{0i}(s)]} = 1 \tag{8}$$

which gives the result.

This result is intuitive: if individuals can optimize their choice of insurance in an efficient market then they value a dollar of subsidy to insurance as exactly one dollar. If markets are efficient then compensation provided in the form of some good, in this case insurance, is equivalent to compensation in dollars, because the good can be exchanged for other consumption on a dollar-to-dollar basis.

Furthermore, note that $\frac{\frac{dG(\overline{v}(s))}{ds}/x_i}{\frac{dG(\overline{v})}{db}} = \frac{Cov(u'_{0i}(s), \frac{dc_i}{ds}) + E_Y[u'_{0i}(s)] * E_Y[\frac{dc_i}{ds}]}{x_i E_Y[u'_{0i}(s)]}$. All else equal, the larger the covariance of marginal utility of consumption and the gain in consumption from increasing the subsidy, the greater the effect of the subsidy. This is precisely the insurance value of the subsidy: individuals value it more the more it tends to increase consumption when marginal utility is otherwise high. When the insurance market is efficient this gain in consumption from one dollar of subsidy to insurance is precisely one dollar of consumption, leaving the covariance 0 and changing labor supply in exactly the same way as a change in income would.

Retirement Lock I define the distortion due to retirement lock, R, to be the extent to which labor responds to the insurance subsidy above and beyond its response to equivalent retirement income, or the excess of the ratio $\frac{\frac{dG(\overline{v}(s))}{ds}/x_i}{\frac{dG(\overline{v})}{db}}$ above 1:

$$R \equiv \frac{\frac{dG(\overline{v}(s))}{ds}/x_i}{\frac{dG(\overline{v})}{db}} - 1 \tag{9}$$

The numerator is the change in labor due to a \$1 increase in subsidy to insurance; the denominator is the change in labor from a \$1 increase in Social Security. R measures the extent to which individuals work in order to avoid having to acquire their insurance on a dysfunctional individual market, above and beyond how much they are willing to work for income. A positive value indicates individuals work more for a dollar's worth of insurance than for a dollar of income, a situation which cannot arise if markets are efficient. In Section 7 I quantify this distortion in monetary terms using a calibration based on my empirical estimates of $\frac{dG(\overline{v}(s))}{ds}$.

To tie this model to the empirical estimates in Section 5 note that from equation (7) it follows that the following relation of labor market responses to a ratio of expected marginal utilities holds:

$$\frac{s\frac{dG(\overline{v}(s))}{ds}}{x_i b\frac{dG(\overline{v})}{db}} = \frac{s}{b}\frac{E_Y[u'_{0i}(s) * \frac{dc_i}{ds}]}{x_i E_Y[u'_{0i}(s)]} \tag{10}$$

For a small s the key quantity $s\frac{dG(\overline{v}(s))}{ds} \approx \frac{\Delta G(\overline{v}(s))}{\Delta s}$, which is precisely what I will estimate in the empirical section.[13] To do so, I now turn to the institutional details of Medicare Part D which will be relevant to the empirical design.

3 The Medicare Part D Program

This section provides some institutional details regarding the Medicare Part D program: a change to traditional Medicare which took place in 2006 which provided a subsidy for prescription drug insurance plans for individuals over age 65. These details inform the identification strategy detailed in the next section.

Medicare provides universal health insurance coverage to Americans over age 65. When the program was started in 1966 it did not cover prescription drugs. However, the past 30 years have seen the share of health expenditures going towards prescription drugs increase substantially. In 1982 prescription drugs accounted for about 4.5% of health expenditures, while by 2005 that share had more than doubled, to about 10.1% (Duggan et al., 2008).

To address the lack of insurance for such large health expenditures among the elderly the administration and Congress passed a bill which, beginning January 1st, 2006, provided subsidized prescription drug insurance to everyone eligible for Medicare. This essentially meant that every American over age 65 would have access to prescription drug insurance.

[13] The key intuition that a labor response to a subsidy which is larger than a response to equivalent income implies a high valuation of the policy change can be derived from a simpler model with even less structure. Without specifying either that the policy change is small or imposing any structure on how insurance works an analysis of labor responses based on the equivalent variation of Medicare Part D can quantify the welfare value of the program. For such an analysis see Appendix D.

By 2014 the annual cost of this program had reached \$79 billion (Medicare Board of Trustees, 2014). This made Medicare Part D the largest expansion of a public health insurance program since the start of Medicare itself, a position it retained until the ACA's passage in 2010.

Medicare Part D works by allowing anyone eligible for Medicare to choose between three subsidized insurance options: a stand-alone prescription drug plan, offering only prescription drug benefits; a Medicare Advantage plan, offering the full range of Medicare benefits including prescription drugs; and the option of remaining on an employer/union health insurance plan provided that plan's prescription drug coverage was at least as generous as the standard Part D plan. All basic Part D plans are actuarially equivalent.

Those choosing the option of staying on an employer plan would still receive a subsidy from the government, which covers 28% of employer costs between the deductible of \$310 and an upper limit of \$6,350 in 2014, for a maximum subsidy of \$1,691. This subsidy is intended to discourage employers from dropping their coverage for elderly employees, knowing the government would replace it. It is noteworthy in order to interpret the results estimated below. It implies that virtually all the change in the insurance environment for individuals with employer sponsored insurance stems from introducing and subsidizing an individual market alternative to employer insurance, not from the loss of employer insurance due to a change in the worker's compensation package as a result of the change in policy.[14]

In sum, whereas before 2006 access to prescription drug insurance had been almost exclusively restricted to those with employer sponsored insurance, from 2006 onward everyone over age 65 had the option of purchasing subsidized prescription drug insurance. This sharp change forms the basis of my identification strategy, to which I turn in the next section.

4 The Health and Retirement Study Data and Empirical Strategy

This section describes the data used to estimate the effect of Medicare Part D eligibility on labor supply and how I go about estimating that effect. The rich data available in the Health and Retirement Study (HRS) provide detailed information on employment status, permitting differentiation of full-time and part-time work. This is crucial for my analysis. They also allow identification of the insurance status of individuals, enabling me to construct treatment and control groups to be used in a differences-in-differences and a triple differences design. This design recovers the causal effect of Part D on labor supply and reveals the extent to which individuals work solely in order to retain their group drug insurance. The triple differences with a control group demonstrates that it is Part D's relaxation of retirement lock that drives the effect on the treated.

[14] There has been a long-term trend of employers offering less retiree coverage since at least the 1980's; the share of employers who offer retiree coverage out of employers who offer health benefits to active workers has fallen from 66% in 1988 to 25% in 2014 (Kaiser Family Foundation, 2014). However, there was no sharp change in this trend around 2006, nor has there been any change in the share of employer plans which cover prescription drugs.

The data I use are primarily from the RAND version of the HRS (RAND HRS Data, 2014).[15] The HRS is a longitudinal survey of roughly 20,000 Americans over the age of 50 and their spouses conducted every two years since 1992. As Medicare Part D began January 1st, 2006, I restrict the sample to years 2000-2010. Because eligibility for Part D, as for Medicare in general, begins at age 65, I further restrict the sample to individuals aged 55-68.[16]

Retirement lock is not expected to operate on all individuals. In particular, for those individuals provided with retiree health insurance from their employer without an age limit the retirement decision is completely divorced from considerations of health or prescription drug insurance. These individuals will have such insurance irrespective of whether they work or not. Similarly, individuals who have no employer sponsored insurance whatsoever should not be expected to have any labor supply response, as they will not have prescription drug insurance regardless of whether or not they work.

To estimate the effect of Part D on those affected by the new policy with respect to their labor supply decisions I define a "treatment" group of individuals who would have retiree health insurance from their employer should they retire, but only until age 65. For the precise method of defining this group based on HRS data see the Data Appendix. Before 2006 such individuals could generally retire at any age before 65 and keep their health and prescription drug insurance. However, upon reaching age 65 they would have lost the latter. Non-prescription drug health insurance was guaranteed to them at that age by Medicare, but Medicare did not cover prescription drugs. Therefore, if maintaining prescription drug coverage were sufficiently important for them, members of the treatment group would have had to keep working, most likely at full time, or else lose drug coverage at age 65.

In contrast, from 2006 onward Medicare began to cover prescription drugs as well. As a result, members of the treatment group were now released from the potential retirement lock imposed by their employer sponsored prescription drug coverage in the past, and could choose when to retire without having to take into account possible loss of drug insurance. They could now retire at any age and maintain continuous coverage of both health and prescription drug insurance until age 65 (from their retiree health insurance) and from age 65 on (when Medicare would cover both health and prescription drug insurance).

This sharp change in the chaining of the labor supply decision to availability of prescription drug insurance at age 65, in year 2006, motivates a differences-in-differences design for the treatment group. The average change in outcomes for individuals just over age 65 (ages 65-68) relative to individuals just under age 65 (aged 55-64) reveals the life-cycle-driven changes in the outcome at age 65. Comparing this mean change at age 65 just after 2006 (years 2006-2010) to the mean change that prevailed just before Medicare Part D (years

[15] For information on prescription drug coverage and out-of-pocket spending I refer to the raw HRS data (Health and Retirement Study, 2013).

[16] The HRS asks questions about potential retiree insurance over age 65 only of respondents below age 65 at the time of the survey, and these questions were first asked in the 1996 wave of the survey. Those older than 68 in 2000 would have been too old to be asked these questions in any wave in which they were observed in the data. For details see the Data Appendix.

2000-2004) identifies the effect of Part D's introduction on individuals aging into eligibility for the program. Assuming no other sharp and systematic changes to the environment of individuals with respect to labor outcomes occurred in 2006, this effect can be attributed to Medicare Part D itself.

This latter assumption is equivalent to assuming that in the absence of Part D, the change in outcome at age 65 before 2006 would have been similar to the change in that outcome after 2006. A test of this assumption is that the outcome changes before age 65 are parallel before and after 2006. I show this to be the case below.

One other assumption in this identification strategy is that Medicare Part D had no effect on the incentives to retire of individuals under age 65. If a substantial share of people under age 65 continued working for the option value of having a job after age 65 which would provide prescription drug coverage then the differences-in-differences estimator would understate the true effect of Medicare Part D. In such a case both those over 65, and those under 65, would reduce their labor supply in 2006 due to Part D. This potential bias does not seem to be quantitatively important: in practice the full-time work rates of the treatment group before age 65 rise in 2006, rather than fall, in continuation of long-term trends in labor supply since the mid 1980's. For further details see section 5.2, and figure 8.

Confining the treatment group to those who had retiree health insurance if and only if they were younger than age 65 has another advantage in that it suggests a natural control group: individuals who have retiree health insurance up to any age. Including this latter group in the analysis leads to a triple-differences design (as in, e.g., Gruber, 1994), whereby the control group serves two purposes. The first is to absorb any residual labor market shocks post 2006 which might differentially affect individuals aged 65-68 differently than individuals aged 55-64. It will be apparent in the next section that this is not a major concern. The second and more useful role the control group will play is to demonstrate that Medicare Part D did not have any significant effect on the labor supply decisions of individuals who were not subject to retirement lock to begin with. This serves to establish the mechanism of the effect on the treated: any reduction in their labor supply can be more confidently attributed to Medicare Part D, and specifically to its relaxation of their retirement lock.

Estimation Equation The following equation will form the basic specification for the analysis in the next section:

$$\begin{gathered} y_{i,t,a} = \beta_1 * Post2006_{i,t} * Over65_{i,t,a} * Treat_{i,t} + \beta_2 * Post2006_{i,t} * Over65_{i,t,a} + \\ \beta_3 * Post2006_{i,t} + \beta_4 * Over65_{i,t,a} + \beta_5 * Treat_{i,t} + \\ \beta_6 * Treat_{i,t} * Post2006_{i,t} + \beta_7 * Treat_{i,t} * Over65_{i,t,a} + \\ \alpha_a + \gamma_t + \delta_a * Treat_{i,t} + \zeta_t * Treat_{i,t} + \mu_i + \sum_{j=1}^{k} \theta_j X_{j,i,t,a} + \varepsilon_{i,t,a}, \end{gathered} \tag{11}$$

where i indexes individuals, t indexes years and a indexes age. $y_{i,t,a}$ is an outcome variable such as an indicator of full-time work; $Post2006_{i,t}$ and $Over65_{i,t,a}$ are dummies equal to 1 if and only if the observation is observed at year 2006 or later, and at age 65 or over, respectively; and $Treat_{i,t}$ is a dummy equal to 1 if and only if the individual would be eligible for retiree health insurance should she retire, and this insurance is limited to those younger than age 65. All specifications further include a full set of age and year fixed effects, as well as their interactions with $Treat_{i,t}$.[17] μ_i is an individual fixed effect which is included in all specifications unless otherwise noted. Thus, β_1 gives the causal effect of meeting the eligibility criteria for Medicare Part D on y for those in the treatment group, while β_2 gives the causal effect for those in the control group.[18]

$X_{j,i,t,a}$ is a vector of additional controls. They generally include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Additional health controls are also included except where stated otherwise, including a set of dummies for self-reported health on a scale of 1-5 from poor to excellent; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. All monetary variables are inflated to 2010 prices by the consumer price index. All standard errors are clustered at the individual level.[19]

In specifications without individual fixed effects some other demographic controls are included instead: gender, a full set of dummies for years of education, veteran status, and dummies for race (white, African American, or other) and religion (Protestant, Catholic, Jewish, None, or other).

The main outcome variables of interest are a full-time work indicator and an indicator of part-time work. Individuals are considered full-time workers if they report working more than 35 hours a week for more than 36 weeks a year. If they work less than that they are considered part-time workers. Hours from both main and secondary jobs are counted. In addition, some specifications have as their outcome variable an indicator of job switching: it is 1 if tenure with the current employer declines from more than two years to less than two years between two consecutive survey waves, and 0 otherwise. This indicates a change of a relatively long-term employer at the finest resolution available in the bi-annual HRS survey.

Furthermore, self-reported annual labor earnings are also analyzed. To construct these I use the RAND variable on earnings which sums up individual responses to questions in the HRS regarding wages and salaries, bonuses, overtime pay, commissions, tips, second job and military reserve earnings and professional practice or trade income. As with all monetary variables, earnings are inflated to 2010 dollars using the consumer price index. Furthermore,

[17] A dummy for age 68, for year 2010 and for their interactions with being in the treatment group are omitted to avoid perfect multicollinearity and provide the baseline.

[18] The HRS does not survey a random sample of the US population, but rather oversamples minorities and some states. Because individuals are sampled at different years and weighted to match different populations (based on the CPS) the results presented below are not weighted. However, all results are virtually identical when weighted by the HRS sampling weights at the wave when they were first sampled.

[19] Where possible, results are also robust to clustering at the household level.

I top-code earnings at $100,000. This is the 95th percentile of earnings in the sample for individuals working full-time.

Descriptive statistics are presented in table 1 for the pre-treatment sample: individuals aged 55-64, in the years 2000-2004.[20] Column 1 provides statistics for demographic variables, prescription drug insurance and utilization, and the main outcome variables of full- and part-time work and labor earnings for the treatment group, as well as the number of individuals included in the group; column 2 does the same for the control group. There are about 4000 unique individuals in each group, and the two groups are very similar in their demographic characteristics: about 50% women, have a mean age of 62 and between 13 and 14 years of education on average. Likewise, the groups are similar in their coverage for prescription drugs, which is almost universal (both groups before age 64 have employer-sponsored health insurance which almost invariably also includes drug coverage), and in their part-time work rates. They differ in their full-time work rates; however as discussed above it is parallel trends, rather than identical levels, which are the identifying assumption of the triple-differences estimation strategy.

The distribution of the treatment and control groups' occupations and industries (among those still working) are also very similar, and there is no substantial change in these respective distributions from before Medicare Part D's introduction to after it. These distributions for each experimental group, in years 2004 and 2006, are presented in figure 1 (occupations) and figure 2 (industries). Both treatment and control groups are predominantly in managerial, clerical and professional occupations (together accounting for over half of each group), with sales accounting for an additional 10% of each group. The remaining 30-40% are roughly uniformly distributed across a variety of occupations. With respect to industry, both treatment and control groups are most likely to work in professional services (between 30% and 40%), with public administration (between 5% and 12%), manufacturing (around 15%) and retail (about 15%) making up the bulk of the remainder.

5 Estimation of Prescription Drug Insurance Retirement Lock

5.1 Take-up of Medicare Part D

Before estimating the effect of Medicare Part D on labor supply, it is helpful to see that the program was, in fact, taken up by the treated individuals. Figure 3 shows the rates of public insurance for prescription drugs by age, before and after 2006, in the sample of individuals who have retiree health insurance at least till age 65.[21] Before 2006 public prescription drug insurance was limited to those on Medicaid, on Disability Insurance or veterans receiving

[20] Except for statistics on age and number of unique individuals, which are not limited to observations of less than 64 years of age, before 2006 but rather encompass the entire sample.

[21] Very similar figures result from restricting the sample to only the treatment group, or only the control group, with take-up rates of Part D substantially higher for the former.

health insurance through the Civilian Health and Medical Program of the Uniformed Services or the Department of Veterans Affairs. As is clear in the figure, a very small share of the sample had such insurance, thus very few benefited from public prescription drug insurance before 2006. In stark contrast, with the beginning of Medicare Part D in 2006 individuals aged 65 or older became eligible for public prescription drug insurance through Medicare, explaining the large increase in the share of the sample having public insurance at age 65 post 2006. This figure therefore demonstrates the conceptual "first stage" of the Part D quasi-experiment, showing that individuals effectively assigned to the "treatment" of eligibility for Medicare Part D did in fact take up the treatment.

5.2 Differences-in-Differences Estimates of Retirement Lock

The left-side panel of figure 4 depicts the full-time work rate of individuals in the treatment group at different ages. In the blue squares are the full-time work rates of individuals at the age along the x-axis before 2006. In the red circles are the corresponding values after 2006. Note the drop in the full-time work rate both before and after 2006 at age 65, and, to a lesser extent, at age 62. These drops correspond to eligibility for Social Security full and early retirement ages, respectively.

Of particular interest, however, is the noticeably larger decline in the full-time work rate at age 65 after 2006, relative to before 2006. This is a visual representation of the differences-in-differences estimation of the effects of Medicare Part D eligibility on full-time work. Also of note is the parallel movement of the curves in blue squares and red circles before 2006. The identifying assumption of differences-in-differences is that absent the treatment, treatment and control groups will move in parallel. These parallel pre-trends are a test of this identifying assumption.[22] Both of these qualities are easier to observe in the right-side panel of figure 4, where the means of the post-2006 period are adjusted to match the means of the pre-2006 period for ages 55-64. This is a graphical representation of the first difference of the differences-in-differences. The trends for ages 55-64 line up very closely for the pre-2006 and post-2006 periods, and the differences-in-differences estimator is the difference in means between the post-2006 and pre-2006 periods for ages 65-68.

Table 2 estimates equation (11) solely for the treatment group, the regression equivalent of this differences-in-differences analysis with additional controls. This estimation shows a reduction of 7 percentage points in full-time work as a consequence of eligibility for Medicare Part D among individuals who have prescription drug insurance through retiree plans only until age 65. At a baseline mean rate of full-time work of 0.40, this represents a decline of

[22]It is apparent in figure 4 that post 2006 the level of full-time work is higher in ages 55-64 than it was in the years 2000-2004. While identification requires only parallel trends, not identical levels of the outcome, one might be concerned as to what drives that difference in levels. In this case, there has been a long-term trend of increasing labor supply among the elderly since the mid 1980's, long predating Medicare Part D. To see this please refer to figure 5, which shows the labor force participation rate of individuals aged 55-64 from the Current Population Survey. As a result of this secular trend the levels of full-time work are higher in the years 2006-2010 than they were in the years 2000-2004. This is therefore not directly related to Medicare Part D, nor is it an artifact of the HRS data.

18% in the share of full-time workers upon eligibility for Part D.

5.3 Triple-Differences Estimates of Retirement Lock: Full-Time Work

As an additional control I estimate a similar specification using a control group of individuals whose labor decisions are not tied to their prescription drug insurance: those with retiree health insurance till any age. Figure 6 shows there is no substantial differential change in the full-time work of this group at age 65, before and after implementation of Medicare Part D in 2006. This validates the differences-in-differences estimation above. It is also reassuring for the interpretation of Medicare Part D's labor supply effect as one driven by relaxation of retirement lock: where there is no retirement lock there is also no effect on full-time work rates.

A different way of looking at this placebo test is in the form of a triple differences estimation. Figure 7 is a graphic representation of the triple differences. The red circles now represent the treatment group, while the blue squares depict the control group's full-time work rates at every age. The left panel shows these for the years 2000-2004, while the right panel does the same for the years 2006-2010. In this figure one can see that while the control group has no sharp drop in full-time work rates at age 65 either before or after Medicare Part D, the treatment group has a substantially larger drop post-2006 relative to pre-2006. Furthermore, one can also see the parallel movements of full-time work rates between the treatment and control groups, complementing the parallel movement within each group in the pre- and post-2006 periods noted in figures 4 and 6. It is of particular interest to note that in the post-2006 period the treatment and control groups behave remarkably similarly after age 65, consistent with both groups at this point facing a similar detachment of the labor supply decision and their insurance environment.

Instead of pooling all three pre-Part D survey years and all three post-Part D survey years as figure 7 does, figure 8 shows the same information on full-time work rates by age and by treatment group at a yearly level. In the interest of clarity and reduction in sampling noise I have pooled every two consecutive ages in this figure. Figure 8 serves to illustrate two main points: the first is that the treatment and control groups have parallel pre-trends every year, not just averaged out over the pre- and post-Part D years. Second, it allows us to ascertain that the pivotal year in which the full-time work rates of the treatment group begin to decline much more sharply at age 65 is in fact 2006. Whereas the decline in the years 2000-2004 is around 23 percentage points (averaged over the three years), the fall at age 65 in 2006 is around 28 percentage points, a relative increase in magnitude of 22%. This gap only increases further in 2008 and 2010, consistent with some labor market frictions and delayed responses. A more complete discussion of this last point is deferred to the robustness checks in the next section. This difference in the decline of the full-time work rate at age 65 in the pre- and post-Part D periods for the treatment group will prove statistically significant and robust to various controls. To show this I turn now to regression results.

Results of the triple differences estimation can be seen in table 3. Column 1 shows the results without demographic and health controls, and column 2 shows the baseline specification of equation (11). The estimate of the effect on full-time work is quite robust, and in the baseline specification indicates a reduction of 8.36 percentage points in the rate of full time work for the treated group. This reduction is large relative to the baseline rate of full-time work, 0.349 (evaluated at the means of all controls); thus a reduction of 8.36 percentage points corresponds to a drop of 24% in treated individuals working full time.[23]

Reassuringly, the effect of eligibility for Part D on the control group is not statistically significant in any specification. For example, there is an insignificant point estimate of a 2 percentage point *increase* in full-time work for the control group in the baseline specification. This formalizes the visual impression from figure 6 that Part D eligibility has no effect on labor outcomes for individuals who were not retirement locked to begin with. Furthermore, it can isolate potential labor market shocks which might affect individuals at age 65 differentially post- and pre-2006, threatening the validity of the differences-in-differences design. The fact that no significant effect is seen for the control group helps allay concerns that the results in the treatment group are influenced by other unobserved changes rather than the relaxation of retirement lock due to Part D.

Table 4 contains some variations on this specification with the estimated effect on the treated remaining extremely robust and uniformly insignificant effects persisting on the control group. Column 1 excludes individual fixed effects, and instead includes richer demographic controls; column 2 includes interactions of the age and year fixed effects with demographic characteristics; column 3 excludes from estimation individuals younger than age 62, to verify that results are not driven by younger workers who may be less comparable to the treated group of over 65-year-olds; and column 4 excludes individuals who are on Medicaid or Veteran Affairs insurance, as these individuals would have had prescription drug insurance before Medicare Part D.

5.4 Triple Differences Estimates of Retirement Lock: Part-Time Work

Having established this effect on full-time work I now turn to consider what kind of work or retirement these individuals are replacing their full-time work with. Individuals may wish to slowly phase from full-time work to complete retirement; this is both optimal in various models of life-cycle behavior (e.g., Rust, 1990), and there is evidence that individuals also choose to act in this manner in practice (Ruhm, 1990, Peracchi and Welch, 1994). However, just as the prospect of losing employer health insurance may prevent individuals

[23]This reduction is also very large relative to the effect of wealth in the regression. Mean non-housing household wealth in the sample is about $380,000. At this mean, and using the fifth-order polynomial of wealth controlled for in the regression, an increase of $10,000 of wealth is predicted to reduce the rate of full-time work by 0.09 percentage points, almost two orders of magnitude smaller than the effect of Part D. The effect of wealth estimated here is likely biased due to measurement error, reverse causality, and omitted variables. For a more careful comparison of the effect of Part D to Social Security wealth see Section 7.

from completely retiring, it may also prevent them from reducing their labor supply gradually, as the vast majority of employers do not offer health insurance to part-time workers.[24] It is therefore of interest to explore how much of the reduction in full-time work estimated above is due to individuals shifting to part-time work, and how much of it is due to individuals shifting into complete retirement.

Figure 9 shows the differences-in-differences plot of part-time work for the treated group, with every two consecutive ages pooled in order to reduce noise. It is readily apparent that before age 65 the changes in part-time work rates over ages in the 2006-2010 period move in parallel to those in the 2000-2004 period. It is also clear that at age 65 there was a large increase in part-time work rates after 2006 (of roughly 6 percentage points), while there was no sharp change before 2006.

Column 3 of table 3 mirrors this graphical evidence, showing the results of the baseline specification with the dependent variable now being the rate of part-time work.[25] As expected, there is an increase in part-time work among the treated group, with an increase of 5.9 percentage points in part-time work with the relaxation of prescription drug insurance retirement lock. Over a baseline rate of part-time work of 16.2 percentage points, this represents an increase of 36%. As with full-time work, the control group shows no significant or systematic change in part-time work.

Column 4 of table 3 shows the effect of Part D eligibility on any work; this is the residual of the effect on full-time work after accounting for the increase in part-time work. It indicates that participation declined by 2.5 percentage points with Part D. According to these estimates 70% of those leaving full-time work do so in order to go into part-time work. Only 30% of people leaving full-time work as a result of the relaxation of retirement lock do so in order to fully retire.

5.5 Job Lock and the Transition from Full-Time to Part-Time Work

There are two ways in which one might go from full-time to part-time work. The first is to simply reduce hours while staying in essentially the same job. The second is to switch jobs, to one that involves fewer hours of work. Previous literature has found this latter to be a common choice (Ruhm, 1990). Table 5 shows to what extent these two mechanisms operate.

Column 1 again reproduces the basic specification of part-time work from column 3 of table 3. Column 2 of table 5 shows the increase in job-switching for the treated upon Part D eligibility. This is essentially an estimation of job lock in the more traditional sense of job mobility: eschewing movement between jobs due to concerns about employer-sponsored insurance coverage, as defined, for example, in Gruber and Madrian [2004]. This estimate indicates that individuals increase the rate at which they move between employers by 4.4

[24] In 2014 only 24% of employers who provided health insurance to some workers extended that offer to part-time workers (Kaiser Family Foundation, 2014).

[25] Results are robust to other specifications such as a differences-in-differences estimation (with no control group of individuals with retiree insurance past age 65), omitting individual fixed effects and omitting demographic and health controls.

percentage points when no longer faced with prescription drug-induced job lock. The baseline rate of job switching in any two-year wave of the HRS is 3.5 percentage points; thus this estimate represents a very large semi-elasticity of job switching with respect to Part D eligibility of 1.25.

This job lock estimate includes job switches between two full-time jobs and between two part-time jobs, as well as the movements between full and part-time jobs which are the focus here. To decompose the full- to part-time movements into those entailing job switches and those only involving a reduction of hours, column 3 of table 5 takes as its outcome variable the interaction of part-time work and job switching. Thus the dependent variable here equals 1 if the individual works part-time and has switched employers since the previous survey wave, and equals 0 otherwise.

The resulting estimate shows that Part D eligibility increases part-time work associated with job switching among the treated by 4.1 percentage points. In other words, almost all (93%) of the job switches are a result of scaling back work from full to part-time. More to the point, this estimate also indicates that about 69% of the increase in part-time work among the treated is due to a change in jobs, while only 31% is due to a reduction of hours on the same job.

The 125% estimated increase in job turnover upon introduction of Medicare Part D is much larger than common estimates in previous literature. For example, Madrian et al. [1994] find an increase of 25% in job turnover due to introduction of COBRA. This difference can be attributed to two main differences between my setting and that in previous work. First, the nature and scale of the policy reform are substantially different. Medicare Part D provides prescription drug insurance in perpetuity, whereas COBRA provides health insurance, and only lasts for 18 months.

Second, the quality of the job turnover in my setting is very different. A large bulk of the changes in jobs here is accounted for by a reduction in work intensity, moving from full-time to part-time work. For my treated group of over 65-year-olds this is evidently an attractive option, but it may be much less attractive for the prime working-age males which have been the focus of most previous work in this area.

5.6 Effect on Earnings and Wages

Earnings There is a statistic which captures both the decline in full-time work and the increase in part-time work, and that is individual annual labor earnings. The advantages that labor earnings has as a summary statistic of the two main and partially offsetting effects of Medicare Part D on labor supply are paired with the two notorious problems of survey measures of earnings. Reported earnings are often inaccurately reported, and they tend to be very right-skewed. To ameliorate this issue I top-code earnings at the 95th percentile among full-time workers, which is $100,000 in my sample. The results are reported in table 6. Column 1 provides a parsimonious specification excluding individual fixed effects; column 2 shows the baseline specification; and column 3 shows an enhanced baseline specification allowing for differential time-trends and age-trends by demographics (gender, years of education and a quadratic in household non-housing wealth), as well as excluding individuals on Medicaid and those covered by veteran's insurance.

All three specifications indicate substantial declines in annual labor earnings, although the estimates are very noisy and not always statistically significant. The baseline specification indicates a (statistically insignificant) reduction of $1,477, albeit with a large standard error of about 1,900. The other specifications yield larger estimates which are significant, but not statistically different from the baseline result.

Wages In equilibrium labor outcomes are determined not only by labor supply but also by labor demand.[26] It would be helpful to rule out that the shift from full-time work to part-time work and retirement is driven by a negative labor demand shock, rather than a change in labor supply. The primary evidence on this point comes from the control groups: a general shock to labor demand would be expected to impact the labor outcomes of individuals both above and below the age 65 cutoff for Medicare eligibility. This kind of shock should be absorbed by the differences in differences estimator. Furthermore, the existence of the control group of individuals who were not retirement locked to begin with allows me to test whether any age-specific shock to over 65 year-olds after 2006 remains. In none of the regressions above has there been any systematic or statistically significant effect of Part D eligibility on individuals with retiree health insurance unlimited by age. A negative labor demand shock would have been expected to lower the equilibrium labor of this group, as well as the treated group of individuals with retiree health insurance only till age 65.

Nevertheless, there is still a possibility that a negative labor demand shock for the particular kinds of workers who are over age 65 and have their retiree health insurance limited to pre-age 65 is confounding my estimates of retirement lock. One way to allay this concern is by looking at wages, as in Garthwaite et al. [2014]. Column 4 of table 6 shows the effect of Part D eligibility on wages. Conditional on positive wages there is no significant effect on the wages of the treated group (or on the control group), with a point estimate of a reduction of less than 1 cent per hour for the treatment group. Large standard errors preclude me from saying conclusively that there was no change in wages. However, the small point estimates do not suggest that the fall in full-time work for the treated group at age 65 in 2006 is driven by a fall in demand for their labor.[27]

5.7 Heterogeneity in the Treatment Effect

Heterogeneity by Health Status In this section I examine whether there is more retirement lock for workers who use significantly more prescription drugs (for previous work using

[26] In the model in Section 2 a decline in demand for labor would map into a decline in $I(1)$. It is straightforward to see that this would reduce $\overline{v}$, and thus also reduce labor supply.

[27] Similarly insignificant effects are found when the dependent variable is log-wages.

similar heterogeneity by health status to identify job lock due to health insurance see Kapur, 1998). Holding risk aversion constant, for individuals who have experienced negative health shocks such insurance is more valuable, both because they are more likely to use this insurance again (Pauly and Zeng, 2004) and because they would have found it more expensive than others to purchase insurance on the private market (if any insurer were willing to cover them). Their demand for insurance is therefore higher and the supply of such insurance on the individual market is slimmer- raising the relative value of employer sponsored insurance.

I first define two groups based on plausibly exogenous, physician-diagnosed health conditions.[28] The first group is the "sick" group, comprised of individuals who had at least one of the following conditions: cancer, heart disease, lung disease, stroke, arthritis or psychiatric conditions. Roughly two-thirds of the sample fall in this group. The second group is the "healthy" group, of individuals who do not have any of those conditions. The first group is more likely than the latter to require a greater quantity of expensive prescription drugs, and to face a larger risk of drug expenses: mean monthly out-of-pocket spending on drugs in the sick group is \$80 with a standard deviation of 466, while for the healthy it is \$34 with a standard deviation of 125.

The basic full-time and part-time work specifications can be estimated for each of these groups separately (excluding health status controls). Figure 10 shows the differences-in-differences plot for full-time work broken down by health status. While there seems to be no substantial difference in the evolution of full-time work over age before and after Medicare Part D for the healthy, there is a very large decline in full-time work for the sick after 2006.

Table 7 shows the regression results of this estimation. The first two columns give the estimates on full-time work for the sick and healthy groups, respectively. Columns 3 and 4 do the same with part-time work as the outcome. Reflecting the impression from figure 10, for both outcomes the entire retirement lock effect is concentrated in the sick group. This group experiences a 12.2 percentage point drop in the full-time work rate, while they experience a 9.9 percentage point increase in part-time work. For the healthy group there are no statistically or economically significant changes in any direction (likewise for the control group in all these regressions). This pattern is consistent with Medicare Part D being the driving force behind the observed effects, reassuring that we really are estimating the relaxation of retirement lock due to the publicly provided insurance.

Heterogeneity by Spousal Health Status Availability of spousal health insurance has also been used in the past to estimate job-lock (for example, Madrian and Beaulieu, 1998). With respect to spouses the most obvious difference between employer plans and Medicare Part D is that the latter does not provide coverage to spouses. In sharp contrast, the vast

[28] The HRS contains data on whether individuals use prescription drugs regularly, however this cannot be used in order to examine heterogeneity directly as it is endogenously determined based on insurance. Indeed, previous work has found that Part D eligibility increased prescription drug utilization (Lichtenberg and Sun, 2007, Engelhardt and Gruber, 2011, Ayyagari and Shane, 2015).

majority of employer plans do cover spouses.[29] Therefore, while Part D relaxed the retirement lock of unmarried individuals, or those whose spouses were unlikely to need expensive drugs, those who work predominantly in order to ensure their spouses are covered might remain locked, unable to retire without shouldering the cost of their spouses' drug coverage.

That is indeed what is observed in the data. Table 8 does the same as table 7, but instead of breaking the sample down by whether the observed individual is sick or not, now the sample is divided into those who have sick spouses or not. Single individuals are placed in the group without sick spouses. Columns 1 and 2 show the effect of Part D eligibility on full-time work for individuals who do not have a sick spouse, or do, respectively. Columns 3 and 4 do the same for part-time work. All these specifications control for the respondent's own health status. As expected, responses are larger in magnitude for those without a sick spouse. The full-time work rate of individuals without sick spouses declines by 17 percentage points, versus (a statistically insignificant) 1.4 percentage points for those with sick spouses. Regarding part-time work, individuals without sick spouses have an increase of 8.9 percentage points, 50% larger than the (statistically insignificant) 6 percentage point increase estimated for individuals with sick spouses.[30]

6 Robustness Checks

This section demonstrates that the results in Section 5 are robust to a number of perturbations of the sample and design.

6.1 Alternative Measurements of Labor Supply

Until now the measures of labor force status have been based on average hours of work per week and number of weeks worked per year (as described in Section 4; for further details on their construction see the Data Appendix). An interesting question in its own right and a natural robustness check for previous results is to consider the effect of Part D eligibility on the average of hours of work per week itself, as a measure of work intensity.

The results of using this variable as the outcome for the basic specification of equation (11) are in columns 1 and 2 of table 9. Column 1 shows the effect unconditional on working, with hours worked for individuals who do not work set to 0. Column 2 does the same, conditional on working. In both there is a large negative effect of Part D eligibility on average hours of work a week, of between 2.7 and 4.9 hours a week less for the treated individuals upon eligibility. Column 3 constructs a new full-time work variable purely from reported average hours a week, with the variable equal to 1 if average hours a week are more than 35, and 0

[29] In 2014 96% of employers who offered health plans also covered employees' spouses (Kaiser Family Foundation, 2014).

[30] Qualitatively similar results are obtained when the groups are defined as having a spouse needing drug insurance – being married to a spouse who is sick and also not eligible for Medicare Part D – or not having a spouse needing drug insurance – the complement of the former group.

otherwise. The estimated effect of Part D is remarkably similar to results in the previous section, with a fall of 7.7 percentage points in full-time work for the treated.

6.2 Alternative Control Group: No Employer Sponsored Insurance

Thus far all the triple-differences regressions have used a control group of individuals who have retiree health insurance until any age. They are similar to the treatment group of individuals who have retiree insurance only until age 65, but differ in their prescription drug-induced retirement lock. A group less comparable to the treatment group, but equally unaffected by the relaxation of retirement lock is composed of workers who do not have any employer-sponsored health insurance.

Those without any employer-sponsored insurance are less similar to the treatment group than those with retiree insurance to any age on virtually every observable, from gender distribution to income (see columns 1 and 3 of table 1). This second control group nevertheless allows me to test the robustness of the main results by comparing the treated group to a different, yet still untreated (with respect to retirement lock), control group.

Figure 11 shows the pre-trends of full-time work for the treatment group, who have retiree health insurance until age 65, in the red circles; and for this alternative control group of individuals with no ESI whatsoever, in the green squares. The gap between the two groups' mean full-time work rates before 2006 is larger than it was when using the original control group (as can be seen in figure 7). Nevertheless, the trends are roughly parallel, which is the relevant test of the identifying assumption of the triple differences estimation.

Table 10 confirms that the qualitative results hold using this alternative control group. While the precise numbers are naturally slightly different, they are of the same sign and order of magnitude. This estimation indicates a 6.7 percentage point decline in full-time work and a (statistically insignificant) 2.5 percentage point increase in part-time work for the treated in the baseline specification. As above, there are no statistically significant effects for the control group.

6.3 Excluding the Great Recession

The Great Recession which began in December 2007 and ended in June 2009 was a huge negative shock to the labor market (Elsby et al., 2010). One might be concerned that such a large macro shock to the labor market may confound estimates of Medicare Part D's effect on labor supply, as the period of treatment starts in 2006 and persists until 2010.

Recessions in general and the Great Recession in particular had differential effects on different demographic groups (Elsby et al., 2010); in particular men have usually been more strongly hit than women. To the extent that this is true the specifications including differential time and age trends for different demographic groups, including by gender, should have absorbed such specific shocks (see column 2 of table 4). To the extent that having retiree health insurance might have mediated such shocks, use of the control group of those with retiree health insurance to any age should have simultaneously absorbed such an idiosyncratic shock, as well as tested for its existence, insofar as having retiree health insurance till any age is similar to having retiree health insurance only till age 65. As stated above, such tests were never significant at standard significant levels, and so there is no substantial evidence of such residual shocks.

Nevertheless, to guarantee that the Great Recession does not drive the results I can utilize the fact that the treatment period includes observations from before and after the recession. Table 11 shows results excluding some of the later sample years entirely. Columns 1 and 2 show results for full-time and part-time work, respectively, when the only treatment period is 2006 itself (before the recession). While the standard errors are large due to the small sample size, leading to statistical insignificance, the effects are still economically large. In particular, they indicate a 5.6 percentage point reduction in full-time work for the treated.

Two things are worth noting here. The first is that the magnitude of the effects in 2006 seems smaller than for the entire post-2006 period, with a 5.6 percentage point effect that is smaller (albeit by less than one standard deviation) than the 8.4 percentage point effect estimated on the basis of the full sample. While this may only be a result of statistical noise and the large standard errors arising from the smaller sample (the standard errors are about 25% larger when excluding years 2008-2010), it is also consistent with a certain amount of labor market frictions. Medicare Part D went into effect at the beginning of 2006 but it may have taken time for individuals to change their labor supply. In particular, as the HRS is a survey, individuals surveyed during 2006 may have been contacted before they had time to adjust their working arrangements in response to the reduced retirement lock.

The small effect in 2006 relative to 2008-2010 is also consistent with the pattern which can be seen in figure 8. The drop in full-time work for the treatment group (in red circles) at age 65 goes from about 22.6 percentage points in 2004 to about 28.5 in 2006, indicating a difference-in-differences of about 6 percentage points. However, this drop increases further in 2008 to around 41 percentage points, consistent with an 18 percentage point fall in the full-time work rate relative to 2004, perhaps influenced by the recession (the figure does not control for year fixed effects, with every year graphed individually). This declines to around a 38.5 percentage point drop in full-time work in 2010, slightly smaller than the difference in 2008 but still substantially larger than in 2006.

The second observation regarding columns 1 and 2 is that 2006 seems to be the year at which a substantial change in retirement behavior takes place for the treated group, with the large (though insignificant) decline in full-time work in column 1. This provides further support for the visual impression that figure 8 imparts, of 2006 being the pivotal year. This is helpful in ruling out alternative explanations for the results: while the effect size increases in later years, the year Medicare Part D was implemented does seem to break with preexisting trends regarding the drop in full-time work at age 65. In addition to ruling out that the Great Recession is driving the results, this also rules out any other potential mechanism which does not take place in 2006.

Columns 3 and 4 of table 11 show results only excluding observations surveyed during the Great Recession (i.e., observations from 2008). Thus the pre-Part D period consists of years 2000, 2002 and 2004; and the post-Part D period here consists of observations from before the recession, in 2006, and from after the recession ended in 2010. Here once again there is a large and statistically significant drop in full-time work for the treated of about 9.3 percentage points, and a concurrent rise in the part-time work rate of (statistically insignificant) 4.3 percentage points. These numbers are economically meaningful even where they do not meet statistical significance.

7 Welfare Implications of Medicare Part D

7.1 A Test of Distortion due to Retirement Lock

The estimates in Section 5 show that Medicare Part D had a large effect on the full-time work rate of individuals without retiree health insurance after age 65. However, merely observing a reduction in labor supply in response to the subsidy is not sufficient evidence for concluding that any labor supply distortion existed before the policy change. This is because implicit in this policy are also substantial incentives to retire irrespective of any inefficiency in drug insurance markets. As described in Section 2, the subsidy has both an income and a substitution effect which both lead to lower labor supply. Evidence of retirement lock should therefore meet a higher bar: the effect of the Part D subsidy on labor should be larger than an equivalent increase in retirement income, such as Social Security, which involves the same income and substitution effects, but does not address any potential insurance market distortion. In terms of the model in Section 2, R as defined in equation (9) must be positive.

To measure the effect of a dollar of Part D subsidy on labor supply it is necessary to establish how many dollars of subsidy are actually given by the program to the average enrollee. In 2006 the benefits per capita from Medicare Part D were \$1,708, and these are projected to increase to \$3,188 a year by 2023[31] (Medicare Board of Trustees, 2014). These benefits include the premiums enrollees pay themselves. I therefore subtract from these benefits the per capita premium paid by the enrollees to get the net subsidy per capita in each year. The sum of these net benefits for those 17 years from 2006 to 2023, discounted at a rate of 3% annually, is \$25,000 in net present value in 2006.

In the model in Section 2 Part D was conceptualized as a subsidy per "unit of insurance". Such a unit of insurance is not observed, but the total net present value of the lifetime subsidy, $s * x$, is shown by the above calculation to be $s * x = 25,000$. From the estimates in Section 5 we know that $s\frac{dG(\overline{v}(s))}{ds} \approx \frac{\Delta G}{\Delta s} = 0.084$. It is now possible to calibrate R by comparing $s\frac{dG(\overline{v}(s))}{ds}$ to the effect on labor supply of increasing lifetime discounted Social Security wealth by b dollars, as in equation (10).

[31] Individuals reaching age 65 in 2006 had a mean life expectancy of roughly another 17 years.

To get the effect of Social Security wealth on labor supply I turn to the literature.[32] Much of that literature finds either relatively small or statistically insignificant effects (e.g., Burtless, 1986, Krueger and Pischke, 1992, Costa, 1998). For a comprehensive overview of this literature see Krueger and Meyer [2002] and Feldstein and Liebman [2002]. Recent analysis of exogenous changes in Social Security due to changes in the calculation of benefits (the Social Security "Notch") using administrative micro-data provides the most precise estimate available, to my knowledge (Gelber et al., 2015). These authors estimate that a $10,000 increase in lifetime Social Security wealth (discounted at 3% annually) would lead to a decline of labor participation of 1.1 percentage points. In terms of the model, this corresponds to $b\frac{dG(\overline{v})}{db} = 0.011$, where $b = 10000$.

Rewriting R and plugging in the estimates yields:

$$R = \frac{s\frac{dG(\overline{v}(s))}{ds}/sx_i}{b\frac{dG(\overline{v})}{db}/b} - 1 = 2.03 > 0.$$

R is estimated to be 2.03, substantially larger than 0.[33] In other words, the effect of a dollar of drug insurance subsidy on labor supply is 3 times as large as the effect of a dollar of Social Security.

This calibration, based on Gelber et al. [2015], is a conservative one in that most of the literature on Social Security finds even smaller effects on labor participation. As this parameter enters into the denominator of R, smaller estimates of the effect will increase the estimated retirement lock distortion. For example, based on estimates in Hurd and Boskin [1984] $R = 2.25$.[34] Similar exercises can be done using simulations of potential policy changes in the structural literature estimating the effects of Social Security on retirement. All that is required is a way of mapping the simulated policy change to a dollar-denoted change in Social Security generosity. For instance, based on simulations in Samwick [1998] I find $R = 5.7$[35]; and using estimates from Van der Klaauw and Wolpin [2008] I find $R = 2.07$.[36]

[32] The relation of Social Security to retirement has been extensively studied. A very partial list includes Hurd and Boskin [1984], Gustman and Steinmeier [1985], Burtless [1986], Krueger and Pischke [1992], Rust and Phelan [1997], Costa [1998], Samwick [1998], Coile and Gruber [2007], Van der Klaauw and Wolpin [2008], Gelber et al. [2015], and Gustman and Steinmeier [2015].

[33] Assuming there is only sampling error in my own estimate of $s\frac{dG(\overline{v}(s))}{ds}$, R is significantly larger than 0 at a 90% level of significance. Its 95% confidence interval is [-0.2,4.24].

[34] Hurd and Boskin [1984] find that $10,000 in 1969 would have led to a decrease of 7.8 percentage points in labor participation, using a 6% discount rate. When the Part D benefits are discounted at this rate and the 1969 dollars are inflated to 2010 dollars this implies $R = 2.25$.

[35] Samwick [1998] estimates that a 20% reduction in Social Security PIA would decrease retirement by 1 percentage point; in that sample this corresponds to a decrease in Social Security wealth of about $20,000 in 2010 dollars.

[36] Van der Klaauw and Wolpin [2008] consider a counterfactual policy reducing Social Security benefits by 25%. To get the dollar value of such a counterfactual I average the expected Social Security benefits of married men and women in their sample, inflate them to 2010 dollars, and calculate 25% of the total annual benefits. The result is a policy change which reduces annual benefits by $2,667. The authors estimate such a policy variation would lead to an increase in full-time work of 7.4 percentage points for men and 1.8 percentage points for women, at ages 62-69, which I average to get a 4.6 percentage point increase overall. The policy change considered is a change in an annual flow of benefits so I compare it to the annual net subsidy of Medicare Part D which in 2010 was $1,588 (Medicare Board of Trustees, 2014). Thus from Van der

This positive R is evidence of a lack of an efficient individual drug insurance market: if it were possible to buy a dollar's worth of insurance in exchange for a dollar, providing a dollar of insurance should have had precisely the same effect as providing a dollar of income, as the two could be exchanged on the market. The constraints on individuals' ability to freely purchase insurance cause a dollar of insurance to have an outsize effect on labor relative to a dollar of retirement income.

7.2 Welfare Implications of Medicare Part D

My approach to analyzing the welfare implications of Medicare Part D is in the spirit of Baily [1978], Chetty [2006], and Chetty [2009]: I show that welfare-relevant statements can be made by calibrating the model in Section 2 with a small number of sufficient statistics.

Some previous work on Medicare Part D has followed a more structural approach and found modest welfare gains from Medicare Part D, concentrated particularly at the high end of prescription drug consumers (Engelhardt and Gruber, 2011). These authors examine the distribution of out-of-pocket spending on prescription drugs with and without Part D coverage and calculate the utility gains from the reduction of risk from the added insurance under a CRRA utility function with various risk-aversion parameters.

This approach does not account for welfare gains among individuals who were insured both before and after Part D. Such individuals may replace their private insurance with public insurance, but there is no added insurance gained by this, merely crowd-out of the private insurance. In the limit, where the added public insurance completely crowds out preexisting private insurance (and is of similar quality), there would be no welfare gain from insurance whatsoever (and potentially a deadweight loss if the public insurance is funded through distortionary taxes).

However, the results in Section 5 suggest that there may be large welfare gains to individuals for whom public insurance crowds out private, employer-sponsored insurance. These gains do not come only from a better distribution of out-of-pocket spending, but rather from the flexibility of labor supply afforded by the public alternative to the employer-sponsored insurance. Thus such welfare gains from relaxation of retirement lock would be completely overlooked by an analysis which focuses on reductions in out-of-pocket spending.

Figure 12 demonstrates the relatively small decreases in out-of-pocket spending on prescription drugs for the treatment group at every percentile of the distribution of out-of-pocket spending, between the median and the 95th percentile.[37] Similar to the approach of Engelhardt and Gruber [2011], this is done by estimating quantile regressions for each percentile, based on a specification similar to the baseline specification in equation (11).[38]

Klaauw and Wolpin [2008] a change of $2,667 of Social Security leads to a change of 4.6 percentage points in full-time work; while I find that a $1,588 subsidy from Medicare Part D results in an 8.4 percentage point change in full-time work. The value of R is easy to compute from here to be 2.07.

[37] Below the median the effects are very small, while above the 95th percentile the standard errors become very large.

[38] The estimation equation here is simplified in order to reduce computational complexity by excluding

It is readily apparent that the estimated effect of Medicare Part D eligibility on out-of-pocket spending is quite small for the treatment group. At the median there is no estimated reduction in out-of-pocket spending from Part D eligibility, in sharp contrast to a \$180/year reduction in Engelhardt and Gruber [2011]. While by the 90th percentile I estimate a (statistically insignificant) \$440/year effect, this is still substantially smaller than the \$800/year estimate for the 90th percentile in Engelhardt and Gruber.[39] These relatively small effects in the current setting are consistent with the notion that the treatment group is in fact mostly crowding out their employer insurance with the public insurance from Part D. Large reductions in out-of-pocket spending should not be anticipated here because the individuals in question are not necessarily gaining much in terms of prescription drug insurance. Their gains in welfare arise from increased leisure, not from reduced risk. A similar effect has been noted before in Gruber [1996], Greenberg [1997], Greenberg and Robins [2008], Fadlon and Nielsen [2015].

The intuition for linking reductions in labor supply to utility stems from equation (5) in Section 2. This equation states that the reduction in labor supply resulting from a subsidy to the prescription drug insurance of individuals working less than full time is proportional to the marginal utility of consumption of individuals in that group. Relating the marginal utility of consumption to the change in labor supply is the key which permits me to look at the change in welfare due to the increase in leisure.

I proceed through the welfare analysis of Medicare Part D in three steps: first I calibrate the willingness to pay of retirees for a subsidy to their insurance. Then I estimate the total fiscal costs of such a subsidy, including the behavioral responses to it. Finally I combine those two quantities to estimate a marginal value of public funds for Part D.

7.2.1 Willingness to Pay for Medicare Part D

This section quantifies the value of Medicare part D to its beneficiaries by estimating an individual's willingness to pay for the subsidy out of her own income (as outlined in Hendren, 2013a). Consider the thought experiment of asking a retiree how much she would pay for a dollar of subsidy towards prescription drug insurance on the individual market. The value of such a dollar is precisely $E_Y[u'_{0i}(s) * \frac{dc_i}{ds}]/x_i$: the expected marginal utility of consumption times as many dollars of consumption as she expects to receive in consumption from a single dollar of subsidy.

The value of a dollar paid for such an increase is precisely the marginal utility from a dollar of consumption, $E_y[u'_{0i}(s)]$. The ratio of these two quantities is her willingness to pay for a one dollar increase in subsidy, which by equation (7) is also exactly equal to the ratio of labor supply changes due to a dollar of subsidy versus a dollar of retirement income:

individual fixed effects and health controls and including only a quadratic in non-housing household wealth.

[39] In Engelhardt and Gruber [2011] the sample was not constrained to individuals who could have had coverage if they worked; thus it is to be expected that there should be greater crowd-out in my setting, and a correspondingly smaller impact of Part D on out-of-pocket spending.

$$WTP \equiv \frac{E_Y[u'_{0i}(s) * \frac{dc_i}{ds}]/x_i}{E_Y[u'_{0i}(s)]} = \frac{\frac{dG(\overline{v}(s))}{ds}/x_i}{\frac{dG(\overline{v})}{db}}. \tag{12}$$

Recalling equation (8), it follows that if the insurance market is constrained efficient the $WTP = 1$. Thus the extent of retirement lock distortion exactly identifies the willingness to pay more than \$1 per dollar of insurance, mirroring the willingness to work for a dollar of insurance above and beyond willingness to work for income.

Calibration of Willingness to Pay The ratio in equation (12), as in Section 7.1, can also be calculated directly from the observed labor market effects of a subsidy to prescription drug insurance, noting that $WTP = \frac{s\frac{dG(\overline{v}(s))}{ds}/sx_i}{b\frac{dG(\overline{v})}{db}/b} = R + 1$. Therefore we get:

$$WTP = 3.03. \tag{13}$$

This implies that retirees are willing to pay \$3 in return for a \$1 increase in the subsidy to their prescription drug insurance.[40] Individuals were not able to optimize their choice of insurance, and thus the subsidy is valued at more than one dollar per dollar. This is expressed in the labor market by oversupply of labor: as individuals value a lower cost of insurance more than they value income, they are willing to work even when income does not fully compensate for their labor disutility in return for a lower price of insurance. The excess WTP above 1 quantifies how much individuals who have employer-sponsored drug insurance conditional on working are willing to pay to move to an environment in which they could have drug insurance without working.

7.2.2 Marginal Value of Public Funds in Medicare Part D

The willingness to pay above accounts for the private gains from Medicare Part D. Its large magnitude indicates a large scope for welfare gains from the Medicare Part D subsidy. However, the retirement of individuals who would have otherwise continued working full-time imposes a fiscal externality on the government budget due to tax revenue which is lost. This lost revenue is socially costly but is not accounted for in the individual's decision to retire. The following accounts for the cost to the government of increasing the subsidy.[41]

Define the government budget per capita as:

$$B \equiv A - (1 - G(\overline{v}(s)))sx + \tau_a * I(s), \tag{14}$$

where A signifies revenue per capita from sources other than income tax; $(1 - G(\overline{v}(s)))s$ is

[40] The 95% confidence interval is [0.8,5.24].

[41] Kleven and Kreiner [2006] show that in cases where there are multiple margins of response, such as intensive and extensive labor, the elasticity of taxable income is no longer a sufficient statistic for deadweight loss due to a change in government policy, as in Feldstein [1999]. The labor supply response to Medicare Part D is precisely such a case. Hendren [2013a] shows that the impact of a policy change on the government budget, rather than on the tax base, is a sufficient statistic for deadweight loss even in such cases.

the average subsidy to the prescription drug insurance of those not working full-time per unit of insurance; x is the average quantity of insurance they purchase; τ_a is the average income tax rate; and $I(s)$ is average income, so that $I(s) \equiv G(\overline{v}(s)) * I(1) + (1 - G(\overline{v}(s))) * I(0)$. The effect on the budget of offering another dollar of subsidy is therefore given by:

$$\frac{1}{x}\frac{dB}{ds} = -(1 - G(\overline{v}(s))) + s\frac{dG(\overline{v}(s))}{ds} - (1 - G(\overline{v}(s)))\frac{s}{x}\frac{dx}{ds} + \frac{\tau_a}{sx} * \frac{sdI(s)}{ds} \tag{15}$$

The first term is the mechanical cost of the subsidy, the additional dollar given to all those who were already retired; the second term states that the entire subsidy must now be given to individuals who choose to retire due to the change in subsidy; the third term indicates that the entire subsidy must be given to additional units of insurance that retirees are induced to purchase due to the lower price of insurance; the final term captures the reduction in income tax revenues due to individuals' behavioral responses to the subsidy, their lower rate of work. These last three terms together make up the fiscal externality.

Calibrating the Social Cost of Medicare Part D All the terms in equation (15) were estimated in Section 5, with the exception of the elasticity of demand for insurance with respect to the subsidy, $\frac{s}{x}\frac{dx}{ds}$. This latter term is estimated in Appendix C using the same differences-in-differences research design as the main specification of Section 5, with prescription drug insurance coverage as the outcome variable. The result of that estimation is that $\frac{s}{x}\frac{dx}{ds} = 0.15$.

The other quantities used in the calibration are, based on the results from Section 5:
$(1 - G(\overline{v}(s))) = 0.65$
$s\frac{dG(\overline{v}(s))}{ds} = 0.084$
$\frac{sdI(s)}{ds} = 1,477$

an average income tax rate of $\tau_a = 0.28$ (using 2006 rates for federal and average state income taxes, Tax Policy Center, 2014)

and $sx = 1,588$.[42] Plugging these numbers into equation (15) and normalizing by the share of the population receiving the subsidy gives:

$$\frac{1}{x}\frac{dB}{ds}/(1 - G(\overline{v}(s))) = 1.68.$$

I.e., every dollar spent subsidizing the prescription drug insurance of retirees costs the government an extra 68 cents due to the behavioral responses to the subsidy: increased retirement, increased demand for insurance, and lower income tax revenue.

Calibrating the Marginal Value of Public Funds Following Hendren [2013a] we can get the marginal value of public funds ($MVPF$) spent on the subsidy to prescription drugs

[42] This is different than the number used in Section 7.1 because it is the subsidy for one year, rather than discounted over the lifetime, to keep it in the same units as the change in annual labor income due to Part D estimated in Section 5. The value $sx = 1588$ is the net subsidy per capita in 2010 (Medicare Board of Trustees [2014]).

of retirees by integrating the WTP for one dollar of subsidy over the entire population, and accounting for the whole cost of providing a dollar of subsidy, the sum of the mechanical cost and the fiscal externality.

The WTP estimated above is the average willingness to pay among retirees. The willingness to pay of full-time workers for a subsidy they do not benefit from is 0.[43] Therefore the average willingness to pay in the population is $WTP * (1 - G(\overline{v}(s)))$. Combining this with the social costs estimated above gives:

$$MVPF \equiv \frac{(1 - G(\overline{v}(s))) * WTP}{-\frac{1}{x}\frac{dB}{ds}} = 1.80. \tag{16}$$

Equation 16 gives the ratio of the social benefit from an additional dollar of subsidy to its full social cost, the sum of the mechanical dollar spent and the fiscal externality associated with the additional subsidy. All of these are denoted in terms of the welfare gain from an additional dollar of income to retirees.

Note that this calculation does not account for the cost of raising funds. The question the $MVPF$ answers is how to spend funds already raised by the government. With such funds in hand, the $MVPF$ of various policies can be compared and the funds allocated where they provide the highest social return. Such alternative policies could include not raising such funds to begin with.[44]

8 Conclusions

Medicare Part D was the largest expansion of a public health insurance program in forty years at the time of its implementation. While it was primarily considered a safety net for uninsured elderly faced with high prescription drug costs, it also had the effect of aiding individuals who were already insured through their employers who would have liked to retire but for the loss of their coverage.

This paper provides clear evidence of retirement lock stemming from employer-sponsored prescription drug insurance. It does so by focusing on individuals who had employer sponsored retiree health insurance but only till Medicare eligibility at age 65. At that age before 2006 such individuals would have had to remain in (typically full-time) work in order to maintain their drug coverage. After 2006 drug coverage was no longer contingent upon work.

Estimates based on this sharp change in 2006 at age 65 show that individuals indeed reduced their labor supply substantially, decreasing their full-time work rate by about 8.4 percentage points, with no significant effect for a control group of individuals with retiree

[43]In the static model in Section 2 an individual with low disutility of labor is assumed to have a lifelong low disutility of labor. In a richer dynamic model individuals would have a willingness to pay for the subsidy that would vary in each period.

[44]Hendren [2013a] calculates the $MVPF$s of various policies, among them reducing the top marginal tax rate; the $MVPF$ of that policy has a broad range, from 1.33 to 2.

health insurance to any age. 70% of this reduction occurs on the intensive margin, moving from full-time to part-time work. The remaining 30% consist of individuals moving from full-time work directly into full retirement.

The large labor response to the Part D subsidy is not in itself evidence of any distortion. To test for such distortion I compare the estimated labor supply effect of a dollar of drug insurance subsidy to the effect from previous literature of an additional dollar of Social Security. I find that the magnitude of labor supply decline in response to a dollar of subsidy is equivalent to the decline which would be expected from three dollars of additional Social Security benefits. This demonstrates that individuals work for employer insurance above and beyond what they are willing to work for income, implying a distortion in labor supply.

In addition to documenting a large retirement lock effect, I suggest a method of quantifying the welfare gains from relaxing retirement lock. Using the labor supply responses to the policy change of Medicare Part D's introduction I estimate the willingness to pay of retirees for a subsidy to prescription drug insurance. This estimate directly mirrors the extent to which labor responds more to the subsidy than to retirement income, and thus I estimate that individuals are willing to pay $3 per every $1 of subsidy. The valuation of the subsidy at greater than a dollar per dollar can be thought of as the value to these individuals of the existence of an individual prescription drug insurance market.

Accounting for the fiscal externality of the subsidy allows me to estimate the social cost of providing the Part D subsidy. I estimate, again based on the labor responses to Part D, that the fiscal cost of a dollar of subsidy is $1.68. Combining this cost with the estimate of the willingness to pay for the subsidy allows me to calibrate the marginal value of public funds in Medicare Part D to be $1.80 per dollar, or a net social gain of 80 cents per dollar. This welfare improvement through subsidizing the prescription drug insurance of the elderly can be compared to other policies in order to inform policymakers in deciding how to allocate funds across programs.

It is important to note that this welfare analysis includes benefits which have been implicitly assumed to be 0 in previous analyses of Medicare Part D. The estimated welfare gains accrue to individuals who had access to private prescription drug insurance before Medicare Part D so long as they worked. The benefits for these individuals arise largely from relaxation of retirement lock rather than from additional insurance. These large gains serve to provide a scale of the cost to society implicit in the existence of retirement lock. Such potential gains should also be taken into account when assessing other public programs which allow more flexibility in labor supply.

Tables and Figures

Table 1: Descriptive Statistics by Experimental Group at Ages 55-64, Years 2000-2004

	Treatment	Main Control	Alternative Control
Share women	0.513	0.495	0.639
	(0.5)	(0.5)	(0.48)
Age	62.4	62.28	62.8
	(3.86)	(3.85)	(3.79)
Years of Education	13.07	13.51	11.11
	(2.67)	(2.63)	(3.46)
Non-Housing Household Assets	350,119	405,274	214,341
	(1,293,232)	(2,039,182)	(874,601)
Share with Prescription Drug Insurance	0.969	0.985	0.591
	(0.173)	(0.122)	(0.492)
Share with Public Prescription Drug Insurance	0.003	0.005	0.204
	(0.051)	(0.071)	(0.403)
Out-of-Pocket Spending on Drugs/Month	71.79	55.48	94.79
	(258)	(200)	(1049)
Share Working Full-Time	0.554	0.4	0.186
	(0.5)	(0.49)	(0.389)
Share Working Part-Time	0.143	0.158	0.161
	(0.35)	(0.364)	(0.368)
Annual Labor Earnings	32,930	28,104	6,374
	(31,404)	(32,931)	(14,945)
Number of Individuals	3,717	4,048	5,773

Notes: This table presents descriptive statistics for the three experimental groups in the analysis: column 1 shows the treatment group of individuals with retiree health insurance until age 65; column 2 shows the main control group of individuals with retiree insurance past age 65; column 3 shows the alternative control group of individuals with no employer sponsored insurance. The sample is restricted to ages 55-64 (except for the statistics on age and number of individuals) and years 2000-2004: before meeting the age criteria of Medicare Part D eligibility and only in the years before introduction of Medicare Part D in 2006. For the row of age the sample is ages 55-68, years 2000-2004. All monetary values are inflated to 2010 prices using the consumer price index. Annual labor earnings are top-coded at $100,000. The number of individuals is the number of unique individuals included in the baseline specification of Equation (6) in the complete sample, within each experimental group; i.e., all individuals aged 55-68, in the years 2000-2010 in each of the experimental groups. Note that there are individuals who may appear in more than one group at different survey waves (e.g., if they move from a job which does not offer any employer-sponsored insurance to one which offers retiree insurance). Each row besides the last presents the mean of the variable listed in that row for the three experimental groups, with standard deviations in parentheses.

Table 2: Differences-in-Differences Estimates of the Effect of Medicare Part D Eligibility on Full-Time Work

Dependent Variable:	Full-Time Work
Post65*Post2006	-0.0703**
	(0.0305)
Age and Year Dummies	Yes
Demographic and Health Controls	Yes
Individual Fixed Effects	Yes
N	6,850
Number of Clusters	3,717

Notes: This table presents the differences-in-differences estimates of the effect of Medicare Part D eligibility on full-time work. The sample is restricted to individuals in the "treatment group"– those having employer-sponsored retiree health insurance only until age 65. The dependent variable is an indicator of full-time work. The first row provides the differences-in-differences estimate of Medicare Part D eligibility on full-time work. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. All monetary variables are inflated to 2010 prices by the consumer price index. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 3: Triple Differences Estimates of the Effect of Medicare Part D Eligibility on Labor

Dependent Variable:	Full-Time Work		Part-Time Work	Any Work
Specification:	No Demographic, Health Controls	Baseline	Baseline	Baseline
Post65*Post2006*Treated	-0.0845***	-0.0836***	0.0589*	-0.0247
	(0.0311)	(0.0313)	(0.0308)	(0.0338)
Post65*Post2006	0.0234	0.0199	0.00157	0.0215
	(0.0216)	(0.0217)	(0.0218)	(0.0234)
Age and Year Dummies	Yes	Yes	Yes	Yes
Age and Year Dummies * Treated	Yes	Yes	Yes	Yes
Demographic and Health Controls	No	Yes	Yes	Yes
Individual Fixed Effects	Yes	Yes	Yes	Yes
N	15,828	15,382	15,382	15,382
Number of Clusters	6,819	6,516	6,516	6,516

Notes: This table presents the triple differences estimates of the effect of Medicare Part D eligibility on full-time work, part-time work and any work. The dependent variable in the first two columns is an indicator for full-time work. In column 3 the dependent variable is part-time work; in column 4 it is any work. The controls included in each specification are indicated in the table. The first row provides the triple-differences estimates of Medicare Part D eligibility on full-time work for individuals with employer-sponsored retiree health insurance only until age 65. The third row provides the estimates of the effect of Medicare part D eligibility on the dependent variable for the control group of individuals with retiree health insurance unlimited by age. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 4: Triple Differences Estimates of the Effect of Medicare Part D Eligibility on Full-Time Work– Robustness Checks

Dependent Variable:	Full-Time Work			
Specification/Sub-Sample:	Full Sample		Only Ages 62-68	Excluding Medicaid and VA
Post65*Post2006*Treated	-0.0912***	-0.0852***	-0.0802*	-0.0938***
	(0.0290)	(0.0318)	(0.0452)	(0.0329)
Post65*Post2006	-0.0263	0.0168	0.00405	0.0259
	(0.0199)	(0.0220)	(0.0286)	(0.0228)
Age and Year Dummies	Yes	Yes	Yes	Yes
Age and Year Dummies * Treated	Yes	Yes	Yes	Yes
Demographic and Health Controls	Yes	Yes	Yes	Yes
Individual Fixed Effects	No	Yes	Yes	Yes
Age and Year Dummies * Demographics	No	Yes	No	No
N	15,303	15,371	9,790	14,345
Number of Clusters	6,479	6,509	4,785	6,226

Notes: This table presents robustness checks for the triple differences estimates of the effect of Medicare Part D eligibility on full-time work. The dependent variable is an indicator for part-time work. The controls included in each specification are indicated in the table. The first two columns are estimated on the full sample. Column 3 is estimated only on a sample of 62-68 year olds. Column 4 is estimated on a sample excluding individuals on Medicaid or veteran's insurance. The first row provides the triple-differences estimates of Medicare Part D eligibility on full-time work for individuals with employer-sponsored retiree health insurance only until age 65. The third row provides the estimates of the effect of Medicare part D eligibility on the dependent variable for the control group of individuals with retiree health insurance unlimited by age. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. The first column includes additional demographic controls: gender, a full set of dummies for years of education, veteran status, and dummies for race (white, African American or other) and religion (Protestant, Catholic, Jewish, None or other). The demographic variables interacted with age and year are gender, a full set of dummies for years of education and a quadratic in non-housing household wealth. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 5: Job Switches

Dependent Variable:	Part-Time Work	Job-Switching	Part-Time Work * Job-Switching
Post65*Post2006*Treated	0.0589*	0.0439**	0.0408**
	(0.0308)	(0.0209)	(0.0162)
Post65*Post2006	0.00157	-0.00126	-0.00481
	(0.0218)	(0.0144)	(0.0112)
Age and Year Dummies	Yes	Yes	Yes
Age and Year Dummies * Treated	Yes	Yes	Yes
Demographic and Health Controls	Yes	Yes	Yes
Individual Fixed Effects	Yes	Yes	Yes
N	15,382	15,382	15,382
Number of Clusters	6,516	6,516	6,516

Notes: This table presents estimates of the effect of Medicare Part D eligibility on job switching, and decomposition of the shift from full-time to part-time work into movements involving a change in employer and those reflecting only a reduction of work intensity with the same employer. The dependent variable for each column appears in the column's heading. The controls included in each specification are indicated in the table. The first row provides the triple-differences estimates of Medicare Part D eligibility on the dependent variable for individuals with employer-sponsored retiree health insurance only until age 65. The third row provides the estimates of the effect of Medicare part D eligibility on the dependent variable for the control group of individuals with retiree health insurance unlimited by age. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 6: Effect on Annual Labor Earnings

Dependent Variable:	Annual Labor Earnings in 2010 Dollars			Wages
Specification:	No Individual Fixed Effects	Baseline	Enhanced Baseline, Excluding Medicaid and VA	Baseline
Post65*Post2006*Treated	-4,138**	-1,477	-4,269**	-0.00949
	(1,900)	(1,915)	(2,002)	(7)
Post65*Post2006	-348.2	1,557	1,959	4.834
	(1,382)	(1,441)	(1,523)	(7.137)
Age and Year Dummies	Yes	Yes	Yes	Yes
Age and Year Dummies * Treated	Yes	Yes	Yes	Yes
Demographic and Health Controls	Yes	Yes	Yes	Yes
Individual Fixed Effects	No	Yes	Yes	Yes
Age and Year Dummies * Demographics	No	No	Yes	No
N	15,000	15,076	14,051	6,694
Number of Clusters	6,428	6,465	6,167	3,688

Notes: This table presents estimates of the effect of Medicare Part D eligibility on annual labor earnings and wages. This is measured in dollars inflated to 2010 prices by the consumer price index, and top coded at $100,000, the 95th percentile of earnings for full-time workers. The third column excludes individuals who are on Medicaid or have veterans' health insurance. The dependent variable of the column 4 is wages, defined as: $w_{i,t,a} = \frac{AnnualLaborEarnings_{i,t,a}}{UsualWeeklyHours_{i,t,a}*52}$. The controls included in each specification are indicated in the table. The first row provides the triple-differences estimates of Medicare Part D eligibility on annual labor earnings for individuals with employer-sponsored retiree health insurance only until age 65. The third row provides the estimates of the effect of Medicare part D eligibility on annual labor earnings for the control group of individuals with retiree health insurance unlimited by age. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. The first column includes additional demographic controls: gender, a full set of dummies for years of education, veteran status, and dummies for race (white, African American or other) and religion (Protestant, Catholic, Jewish, None or other). The demographic variables interacted with age and year are gender, a full set of dummies for years of education and a quadratic in non-housing household wealth. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 7: Heterogeneity by Health Status

Dependent Variable:	Full-Time Work		Part-Time Work	
Sub-Sample:	Sick	Healthy	Sick	Healthy
Post65*Post2006*Treated	-0.122***	0.00536	0.0990***	-0.0113
	(0.0374)	(0.0652)	(0.0380)	(0.0586)
Post65*Post2006	0.0416	-0.0269	-0.0221	0.0353
	(0.0256)	(0.0466)	(0.0266)	(0.0428)
Age and Year Dummies	Yes	Yes	Yes	Yes
Age and Year Dummies * Treated	Yes	Yes	Yes	Yes
Demographic and Health Controls	Yes	Yes	Yes	Yes
Individual Fixed Effects	Yes	Yes	Yes	Yes
N	10,733	4,649	10,733	4,649
Number of Clusters	4,856	2,320	4,856	2,320

Notes: This table presents heterogeneity of the effect of Medicare Part D eligibility on full-time and part-time work by health status. The dependent variable of the first two columns is full-time work, and for the latter two columns part-time work. The sub-sample of each column is detailed in the column's heading, where "sick" and "healthy" groups are defined in the text. The controls included in each specification are indicated in the table. The first row provides the triple-differences estimates of Medicare Part D eligibility the dependent variable for individuals with employer-sponsored retiree health insurance only until age 65. The third row provides the estimates of the effect of Medicare part D eligibility on the dependent variable for the control group of individuals with retiree health insurance unlimited by age. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; and body-mass index. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 8: Heterogeneity by Spousal Health Status

Dependent Variable:	Full-Time		Part-Time	
Sub-Sample:	No Sick Spouse	Sick Spouse	No Sick Spouse	Sick Spouse
Post65*Post2006*Treated	-0.1704***	-0.0135	.0889*	0.0602
	(0.0494)	(0.0437)	(0.0483)	(0.0445
Post65*Post2006	0.0436	0.0184	0.0086	-0.00536
	(0.0326)	(0.0312)	(0.0326)	(0.0322)
Age and Year Dummies	Yes	Yes	Yes	Yes
Age and Year Dummies * Treated	Yes	Yes	Yes	Yes
Demographic and Health Controls	Yes	Yes	Yes	Yes
Individual Fixed Effects	Yes	Yes	Yes	Yes
N	7,268	8,114	7,268	8,114
Number of Clusters	3,613	3,708	3,613	3,708

Notes: This table presents heterogeneity of the effect of Medicare Part D eligibility on full-time and part-time work by spousal health status. The dependent variable of the first two columns is full-time work, and for the latter two columns part-time work. The sub-sample of each column is detailed in the column's heading, where "sick spouse" and "no sick spouse" groups are defined in the text. The controls included in each specification are indicated in the table. The first row provides the triple-differences estimates of Medicare Part D eligibility on the dependent variable for individuals with employer-sponsored retiree health insurance only until age 65. The third row provides the estimates of the effect of Medicare part D eligibility on the dependent variable for the control group of individuals with retiree health insurance unlimited by age. Demographic controls include a dummy for being single (omitted in columns 2 and 4), a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 9: Alternative Definitions of Labor Supply

Dependent Variable:	Average Hours/Week	Average Hours/Week (Conditional on Working)	More Than 35 Hours/Week
Post65*Post2006*Treated	-2.667**	-4.914***	-0.0770**
	(1.349)	(1.655)	(0.0319)
Post65*Post2006	0.553	-0.922	0.0290
	(0.947)	(1.290)	(0.0224)
Age and Year Dummies	Yes	Yes	Yes
Age and Year Dummies * Treated	Yes	Yes	Yes
Demographic and Health Controls	Yes	Yes	Yes
Individual Fixed Effects	Yes	Yes	Yes
N	15,076	7,511	15,076
Number of Clusters	6,465	4,038	6,465

Notes: This table presents estimates of the effect of Medicare Part D eligibility on various measures of labor supply. The dependent variable of each column appears in its heading. Individuals reporting more than 70 hours of work in a typical week are omitted. In the second column only individuals reporting strictly positive hours are included. The controls included in each specification are indicated in the table. The first row provides the triple-differences estimates of Medicare Part D eligibility on the dependent variable for individuals with employer-sponsored retiree health insurance only until age 65. The third row provides the estimates of the effect of Medicare part D eligibility on the dependent variable for the control group of individuals with retiree health insurance unlimited by age. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 10: Triple-Differences Using Alternate Control Group of Individuals with No Employer Sponsored Insurance

Dependent Variable:	Full-Time Work	Part-Time Work
Post65*Post2006*Treated	-0.0668**	0.0245
	(0.0320)	(0.0327)
Post65*Post2006	-0.00959	0.0189
	(0.0160)	(0.0186)
Age and Year Dummies	Yes	Yes
Age and Year Dummies * Treated	Yes	Yes
Demographic and Health Controls	Yes	Yes
Individual Fixed Effects	Yes	Yes
N	19,224	19,224
Number of Clusters	8,913	8,913

Notes: This table presents estimates of the effect of Medicare Part D eligibility on full-time and part-time work relative to a control group of individuals who had no employer-sponsored insurance. The dependent variable of each column appears in its heading. The controls included in each specification are indicated in the table. The first row provides the triple-differences estimates of Medicare Part D eligibility on the dependent variable for individuals with employer-sponsored retiree health insurance only until age 65. The third row provides the estimates of the effect of Medicare part D eligibility on the dependent variable for the control group of individuals with no employer-sponsored insurance. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 11: Excluding the Great Recession

	Only 2000-2006		Only 2000-2006 and 2010	
Dependent Variable:	Full-Time	Part-Time	Full-Time	Part-Time
Post65*Post2006*Treated	-0.0560	0.0164	-0.0928**	0.0433
	(0.0440)	(0.0427)	(0.0392)	(0.0370)
Post65*Post2006	0.0109	0.0158	0.0237	0.00573
	(0.0289)	(0.0278)	(0.0262)	(0.0253)
Age and Year Dummies	Yes	Yes	Yes	Yes
Age and Year Dummies * Treated	Yes	Yes	Yes	Yes
Demographic and Health Controls	Yes	Yes	Yes	Yes
Individual Fixed Effects	Yes	Yes	Yes	Yes
N	11,646	11,646	13,039	13,039
Number of Clusters	5,741	5,741	6,201	6,201

Notes: This table presents estimates of the effect of Medicare Part D eligibility on full-time and part-time work excluding the years of the Great Recession. The sub-sample of years after 2006 included in the estimation for each column is in the column's heading, with only 2006 in the post Medicare Part D period being included in the first two columns and 2006 and 2010 comprising the post Pat D period in the third and fourth columns. The dependent variable of each column appears in its heading. The controls included in each specification are indicated in the table. The first row provides the triple-differences estimates of Medicare Part D eligibility on the dependent variable for individuals with employer-sponsored retiree health insurance only until age 65. The third row provides the estimates of the effect of Medicare part D eligibility on the dependent variable for the control group of individuals with retiree health insurance unlimited by age. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 12: Main Results Using only Pre-Age 65 Values to Determine Experimental Group

Dependent Variable:	Full-Time Work			Part-Time Work		
	(1)	(2)	(3)	(4)	(5)	(6)
Post65*Post2006*Treated	-	-0.0482	-0.0893**	-	0.1202***	0.0949**
	-	(0.0428)	(0.0413)	-	(0.0375)	(0.0384)
Post65*Post2006	-0.0805**	-0.0491*	0.014	0.0933***	0.0094	-0.0106
	(0.0351)	(0.0252)	(0.0234)	(0.0315)	(0.0251)	(0.0236)
Age and Year Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Age and Year Dummies * Treated	Yes	Yes	Yes	Yes	Yes	Yes
Demographic and Health Controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual Fixed Effects	Yes	No	Yes	Yes	No	Yes
N	4,956	11,920	11,982	4,956	11,920	11,982
Number of Clusters	2,782	4,906	4,934	2,782	4,906	4,934

Notes: This table presents the effects of Medicare Part D eligibility on the main outcomes of full-time and part-time work when the sample includes only individuals who can unambiguously be assigned to treatment and control groups based on answers they gave. No inference of retiree insurance status is made from observed insurance status after age 65. The dependent variable of the first three columns is full-time work, and for the latter three columns part-time work. Columns (1) and (4) show the differences-in-differences estimates of the effect of Medicare Part D eligibility on the treatment group with no control group. The sample for these columns is restricted to the treatment group: individuals with employer-sponsored retiree health insurance only until age 65. Columns (2) and (5) show the same in a triple-differences design with a control group of individuals who have retiree health insurance past age 65, with no individual fixed effects. Columns (3) and (6) show the same including individual fixed effects, the baseline specification. The controls included in each specification are indicated in the table. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. Columns (2) and (5) include additional demographic controls: gender, a full set of dummies for years of education, veteran status, and dummies for race (white, African American or other) and religion (Protestant, Catholic, Jewish, None or other). Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Table 13: Estimation of Demand for Insurance Elasticity with Respect to Medicare Part D

Dependent Variable:	Full-Time Work
Post65*Post2006	0.13***
	(0.022)
Age and Year Dummies	Yes
Demographic and Health Controls	Yes
Individual Fixed Effects	Yes
N	6,557
Number of Clusters	3,628

Notes: This table presents the effects of Medicare Part D eligibility on prescription drug insurance coverage. The sample is restricted to the treatment group: individuals with employer-sponsored retiree health insurance only until age 65. The controls included are indicated in the table. Demographic controls include a dummy for being single, a set of dummies for each of the census divisions and a fifth-order polynomial in non-housing household wealth. Health controls include a set of dummies for self-reported health on a scale of 1-5; body-mass index; and a set of dummies for having any of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. Robust standard errors clustered at the level of the individual are in parentheses. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

Figure 1: Distribution of Occupations for Treatment and Control Groups, before and after 2006

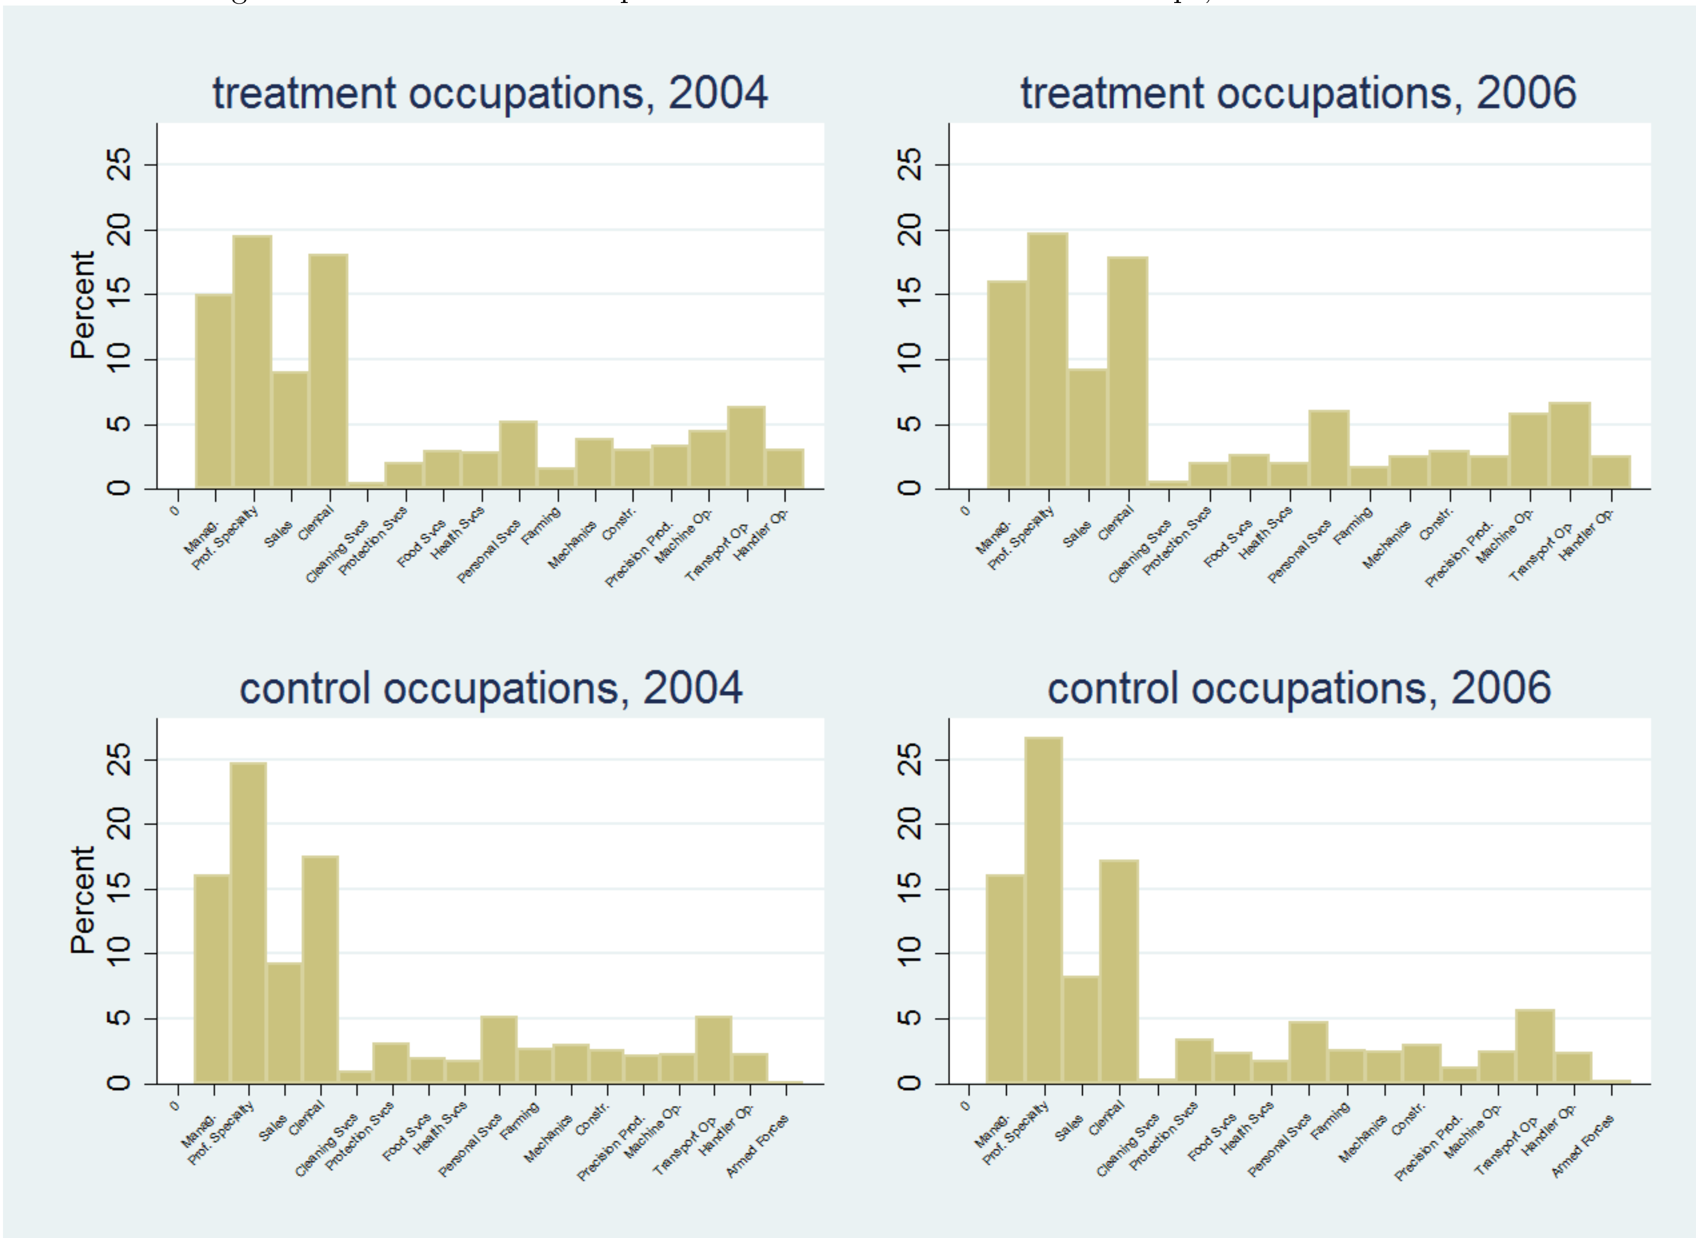

Notes: This figure represents the share of the relevant population in each of the occupations listed along the x-axis. The relevant population in each panel is: treatment group in 2004, treatment group in 2006, control group in 2004 and control group in 2006 for the upper left, upper right, lower left and lower right panels, respectively. Individuals who are no longer working are excluded.

Figure 2: Distribution of Industries for Treatment and Control Groups, before and after 2006

Notes: This figure represents the share of the relevant population in each of the industries listed along the x-axis. The relevant population in each panel is: treatment group in 2004, treatment group in 2006, control group in 2004 and control group in 2006 for the upper left, upper right, lower left and lower right panels, respectively. Individuals who are no longer working are excluded.

Figure 3: Rate of Public Prescription Drug Insurance

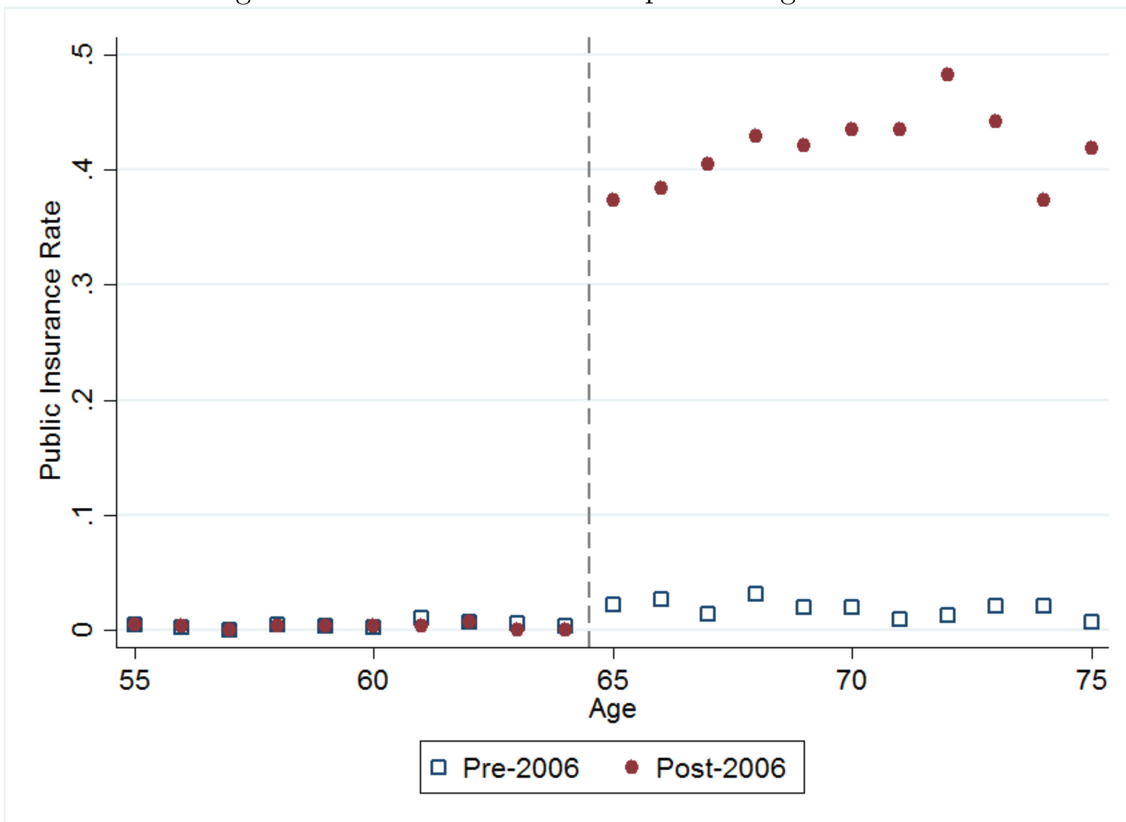

Notes: This figure shows rates of public prescription drug coverage. The sample is individuals aged 55-75, in the years 2000 until 2010, who have retiree health insurance through their employer only until age 65. The blue squares indicate coverage rates by age in the years 2000-2004, while the red circles indicate coverage rates by age for years 2006-2010. The dashed gray line differentiates between ages eligible for Medicare Part D, on the right, and those ineligible, on the left (in the post-2006 period).

Figure 4: Full-Time Work Rates for the Treatment Group

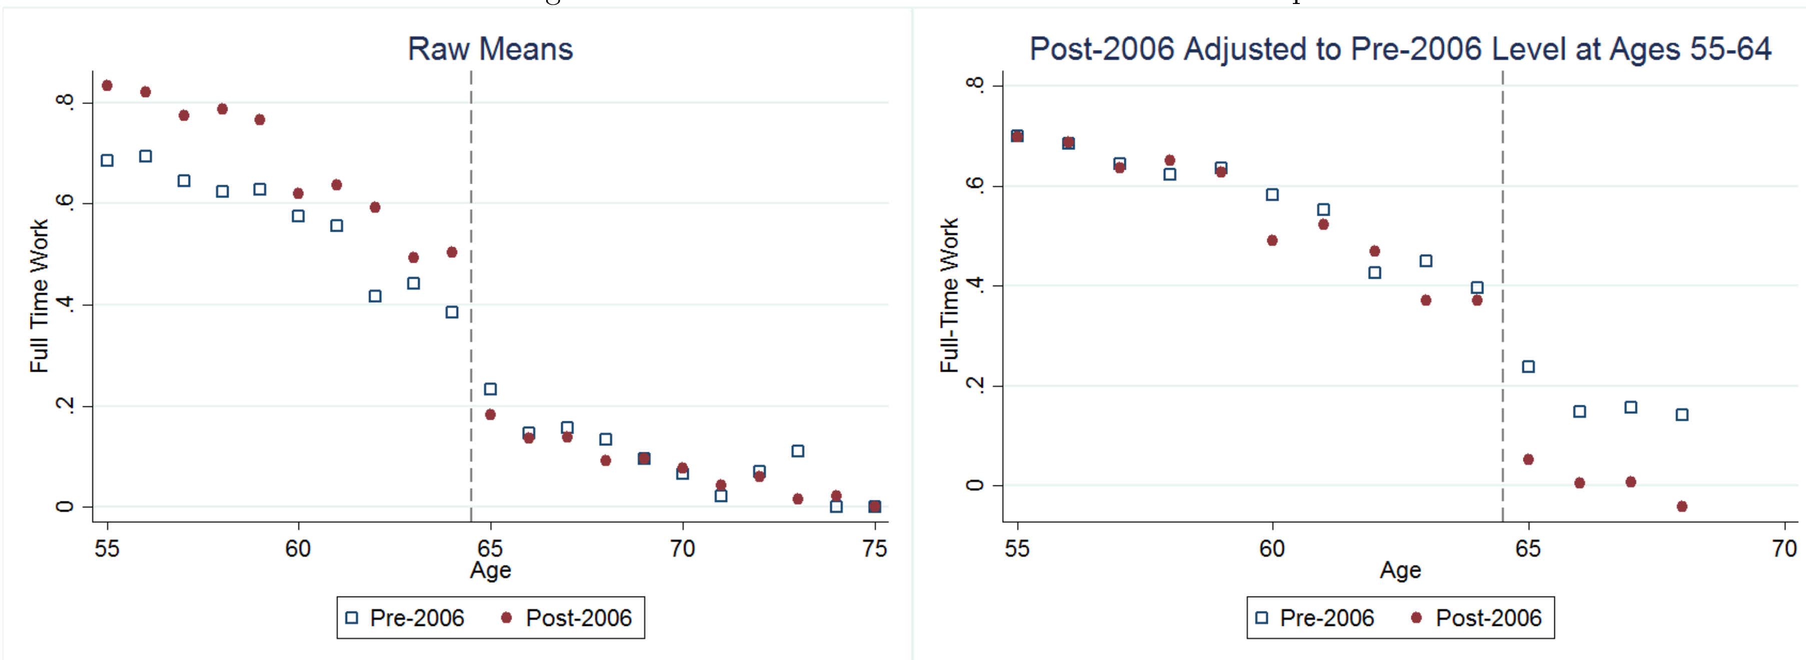

Notes: This figure shows the differences-in-differences of full-time work for the treatment group. On the left panel are the raw means of full-time work. The sample is individuals aged 55-75, in the years 2000 until 2010, who have retiree health insurance through their employer only until age 65. The blue squares indicate rates of full-time work by age in the years 2000-2004, while the red circles indicate full-time work rates by age for years 2006-2010. The dashed gray line differentiates between ages eligible for Medicare Part D, on the right, and those ineligible, on the left (in the post-2006 period). On the right panel is the same data, for ages 55-68, with the means of the post-2006 observations adjusted to match the pre-2006 observations at ages 55-64; i.e., after the first difference of the differences-in-differences.

Figure 5: Labor Force Participation Rate for Individuals Aged 55-64, Years 1985-2010

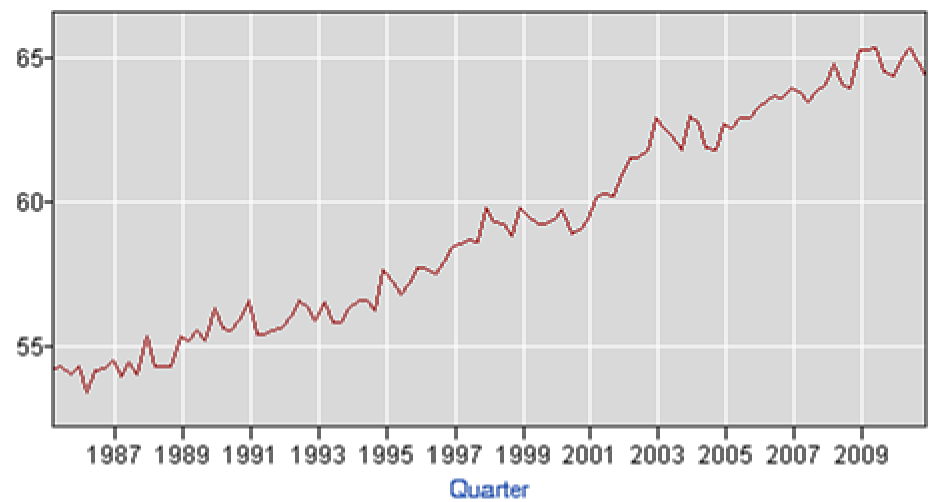

Notes: This figure represents the labor force participation rate from the Current Population Survey for individuals aged 55-64, from 1985 until 2010 at a quarterly frequency.

Figure 6: Full-Time Work Rates for the Control Group

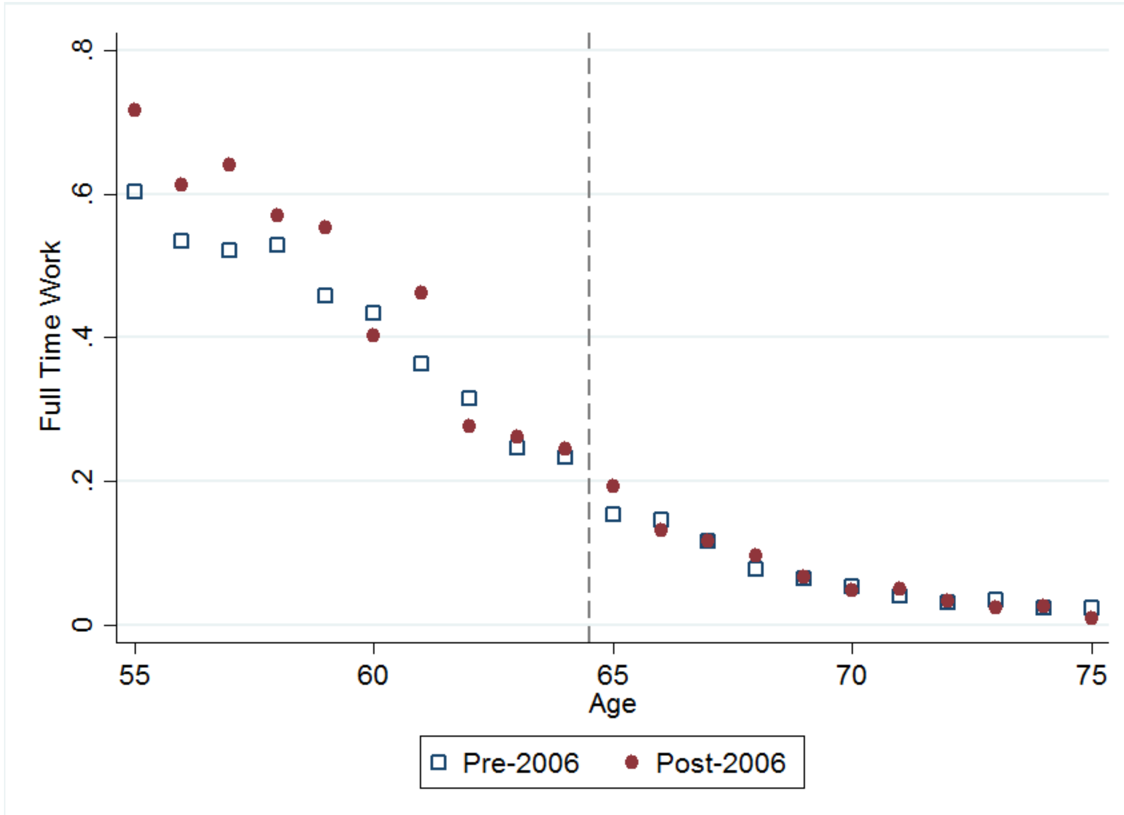

Notes: This figure shows the differences-in-differences of full-time work for the control group. The sample is individuals aged 55-75, in the years 2000 until 2010, who have retiree health insurance through their employer unlimited by age. The blue squares indicate rates of full-time work by age in the years 2000-2004, while the red circles indicate full-time work rates by age for years 2006-2010. The dashed gray line differentiates between ages eligible for Medicare Part D, on the right, and those ineligible, on the left (in the post-2006 period).

Figure 7: Triple Differences– Full-Time Work Rates in the Treatment and Control Groups by Age, before and after 2006

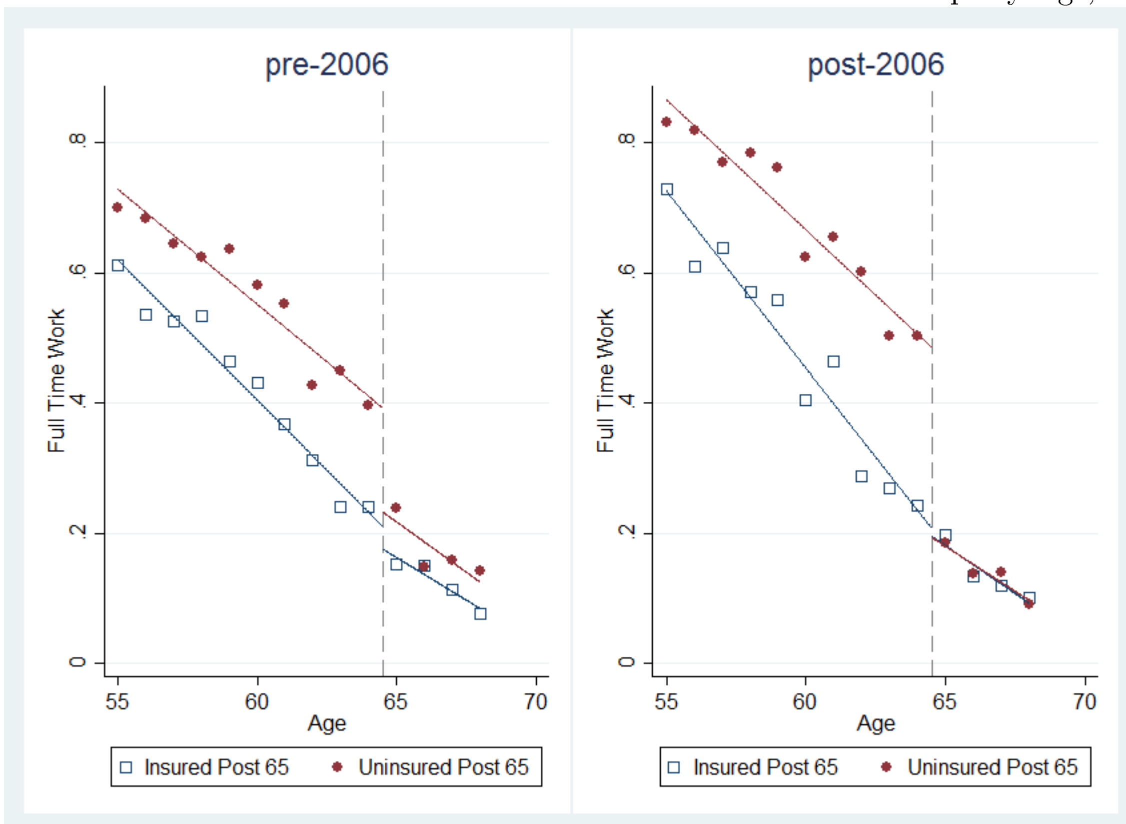

Notes: This figure shows the triple differences of full-time work. The sample is individuals aged 55-68, in the years 2000 until 2010. The blue squares depict the rates of full-time work by age for the control group of individuals who have retiree health insurance through their employer unlimited by age. The red circles depict full-time work rates by age for the treatment group of individuals who have retiree health insurance through their employer only until age 65. The panel on the left consists of observations in the years 2000-2004, before Medicare Part D; the panel on the right consists of observations from the years 2006-2010, after the introduction of Medicare Part D. The dashed gray line differentiates between ages eligible for Medicare Part D, on the right, and those ineligible, on the left (in the post-2006 period).

Figure 8: Triple Differences– Full-Time Work Rates in the Treatment and Control Groups by Age and by Year

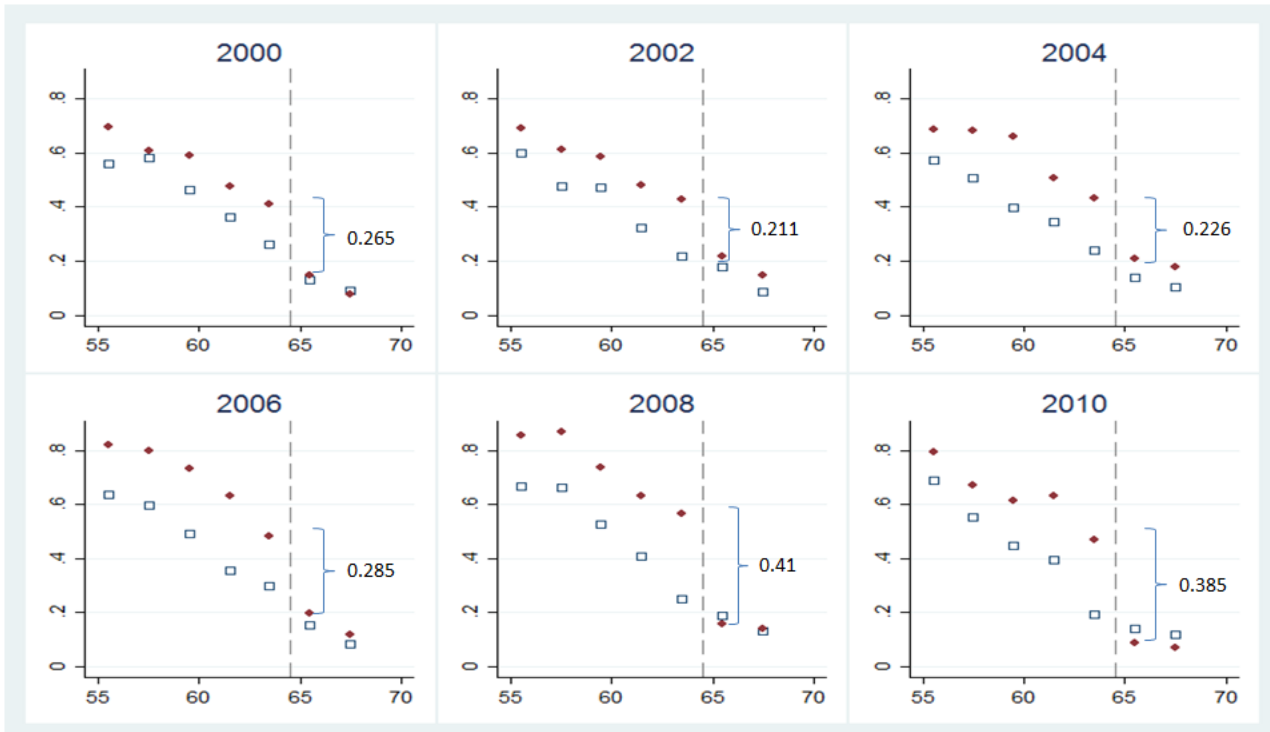

Notes: This figure shows the triple differences of full-time work, on a year-by-year level. The sample is individuals aged 55-68, in the years 2000 until 2010. The blue squares depict the rates of full-time work by every two consecutive ages for the control group of individuals who have retiree health insurance through their employer unlimited by age. The red circles depict full-time work rates by every two consecutive ages for the treatment group of individuals who have retiree health insurance through their employer only until age 65. Each panel in the top row represents observations from the years 2000-2004, before Medicare Part D; each panel on the bottom row consists of observations from the years 2006-2010, after the introduction of Medicare Part D. The dashed gray line differentiates between ages eligible for Medicare Part D, on the right, and those ineligible, on the left (in the post-2006 period). The brackets indicate the difference in full-time work rates for the treated group between ages 63-64 and 65-66 in every survey wave.

Figure 9: Part-Time Work Rates for the Treatment Group

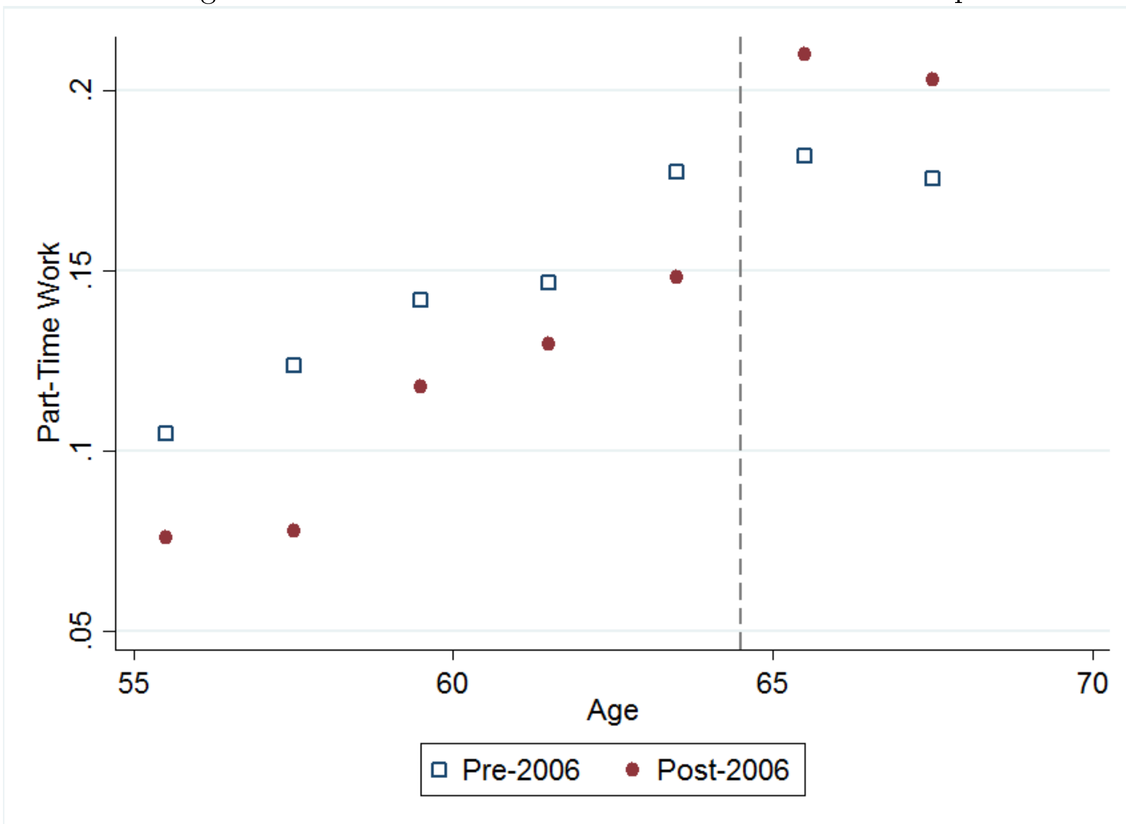

Notes: This figure shows the differences-in-differences of part-time work for the treatment group. The sample is individuals aged 55-68, in the years 2000 until 2010, who have retiree health insurance through their employer only until age 65. The blue squares indicate rates of full-time work for every two consecutive ages in the years 2000-2004, while the red circles indicate full-time work rates for every two consecutive ages for years 2006-2010. The dashed gray line differentiates between ages eligible for Medicare Part D, on the right, and those ineligible, on the left (in the post-2006 period).

Figure 10: Full-Time Work Rates for the Treatment Group, by Health Status

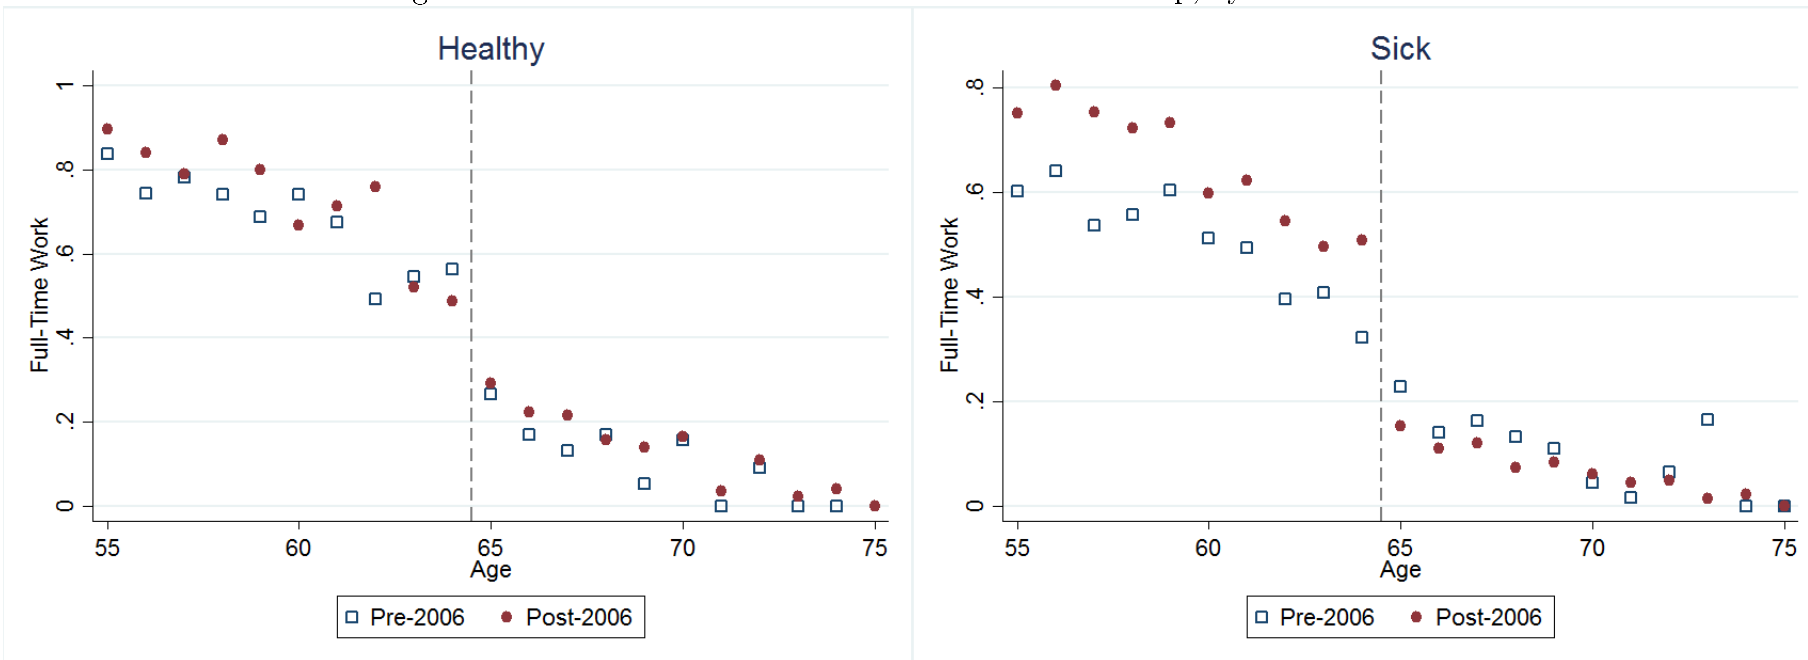

Notes: These figures show the differences-in-differences of full-time work for the treatment group, broken down by individual health status. The sample is individuals aged 55-75, in the years 2000 until 2010, who have retiree health insurance through their employer only until age 65. The blue squares indicate rates of full-time work by age in the years 2000-2004, while the red circles indicate full-time work rates by age for years 2006-2010. The dashed gray line differentiates between ages eligible for Medicare Part D, on the right, and those ineligible, on the left (in the post-2006 period). The sample is divided into "sick" and "healthy" groups, with the sick group including any individual who, at the time of the survey, had one of the following physician-diagnosed conditions: cancer, lung disease, heart disease, stroke, arthritis or psychiatric conditions. Individuals were classified as healthy otherwise. The left panel includes only healthy individuals, while the right panel includes those who were sick.

Figure 11: Full-Time Work Rates in the Treatment and Alternative Control Groups by Age, before and after 2006

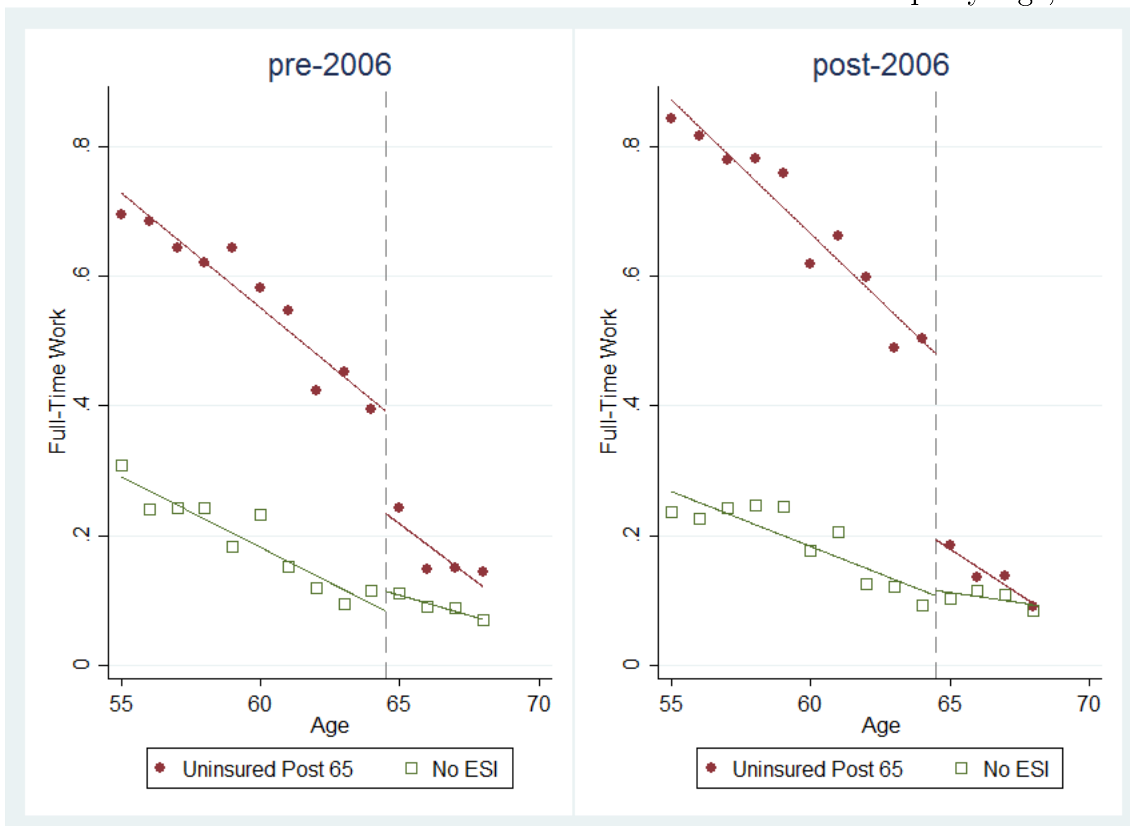

Notes: This figure shows the triple differences of full-time work using an alternate control group off individuals who had no employer-sponsored insurance (ESI). The sample is individuals aged 55-68, in the years 2000 until 2010. The green squares depict the rates of full-time work by age for the control group of individuals who have no ESI whatsoever. The red circles depict full-time work rates by age for the treatment group of individuals who have retiree health insurance through their employer only until age 65. The panel on the left consists of observations in the years 2000-2004, before Medicare Part D; the panel on the right consists of observations from the years 2006-2010, after the introduction of Medicare Part D. The dashed gray line differentiates between ages eligible for Medicare Part D, on the right, and those ineligible, on the left (in the post-2006 period).

Figure 12: Reductions in Monthly Drug Out-of-Pocket Spending by Percentile

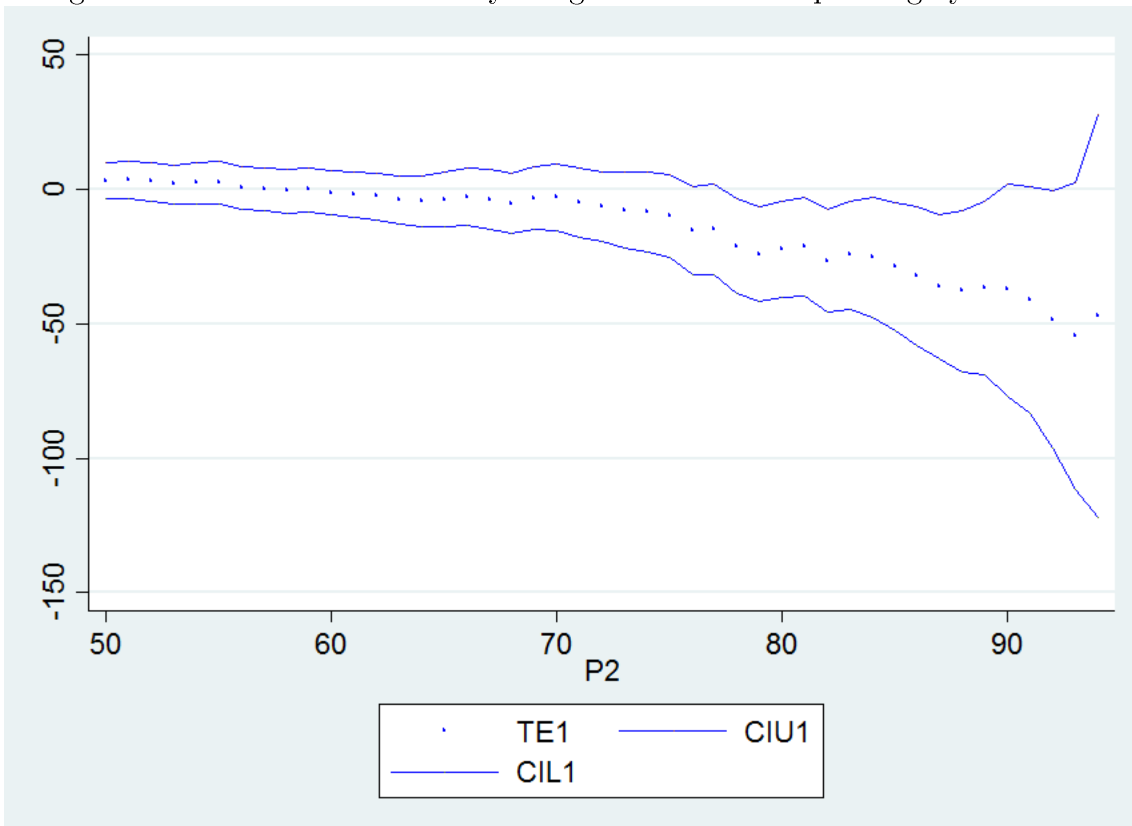

Notes: This figure shows reductions in monthly out-of-pocket (OOP) expenditures on prescription drugs as a result of eligibility for Medicare Part D at different points in the distribution of OOP spending. OOP spending is measured in 2010 dollars. Results are shown for every percentile between the 50th and the 95th. The points represent the triple differences estimates of the change in OOP spending for the treated group of individuals who have retiree health insurance from their employer until age 65; the control group is individuals who have retiree health insurance from their employer unlimited by age. The solid lines are the 95% confidence intervals for the point estimates.

A Data Appendix

Defining Treatment and Control Groups Based on Employer-Sponsored Retiree Health Insurance

The RAND version of the HRS data contains information regarding whether or not individuals are offered retiree health insurance from their employer. Questions on whether a respondent has retiree health insurance of any sort (limited to age 65 or not) are asked from wave 3 and onward (1996 and later). For waves 5 and onward (interviews conducted in 2000 and later), this is taken from a question asking individuals under age 65 whether they have employer-sponsored plans that offer retiree health insurance (their own or their spouse's). If they reply that they do, a follow-up question asks if this coverage would extend past age 65.

The main variable I use to determine coverage is based on these questions from wave 5 and later, and provides a summary of the information regarding all employer-sponsored plans the individual reports (up to three different plans). These questions are asked only of individuals who have employer sponsored insurance while working and are under age 65. The possible values this variable takes are: "not covered in retirement"; "covered [in retirement] just to age 65"; "covered [in retirement] to age 65, don't know over"; "covered [in retirement] to and over age 65"; and a number of possible missing values: "age is 65 or older"; "don't know"; "source missing, question"; "missing"; "no respondent employer provided insurance"; "refused to answer"; "question not asked"; and "spouse is non response".

To be included in the sample for the main analysis[45] an individual must be either in the treatment group (covered in retirement only until age 65) or in the control group (covered in retirement to and over age 65). If the individual cannot be definitively allocated into one of these groups she will not be included in the sample (e.g., if she gave an answer of "covered to age 65, don't know over", but see strategies below for inferring insurance status in the absence of clear answers).

The questions about retiree coverage are only asked of respondents below age 65. In order to allocate observations with age above 65 to the treatment and control groups I employ two complementary strategies: 1) use of lagged values of the same individual from before age 65 to ascertain what manner of retiree coverage, if any, she would have after age 65; 2) infer from current employment and insurance statuses (after age 65) what manner of retiree coverage, if any, she has after age 65. I will now detail for each of these approaches what information is required and what assumptions are made.

1) Use of Lagged Values from before Age 65. This approach is relatively straightforward. If a respondent is interviewed both under and over age 65 at different survey waves then from answers given regarding retiree health insurance offered by employers when asked at waves when she was younger than 65, it can be inferred what retiree insurance she will

[45] See below for construction of the sample for the robustness check using an alternative control group in section 6.2.

have when over 65. For example, a respondent replying at age 64 that she will be covered by her employer plan in retirement only until age 65 will be allocated to the treatment group (of individuals with retiree insurance only until age 65) in all waves, including later waves when she is not asked this question because she is 66 or older.

The assumption made in this approach is that employers cannot change the terms of their retiree insurance plans for employees or retirees already covered by those plans, when they are over 65 years of age. This assumption is not completely innocuous: for example, employers who face financial distress such as bankruptcy may change the terms of their retiree health plans. It is assumed that such cases are relatively rare.

Misallocation of observations to the treatment and control groups due to this assumption should generally operate in the direction of allocation of a treated individual to the control group: an individual was promised retiree coverage for life but at some point the employer decided not to honor that promise and the individual becomes de facto only covered until age 65 (if the abrogation of the promise occurs before age 65 then respondents' answers to the HRS question regarding retiree health coverage should reflect this and no error in allocation would be made). In this case the identification concerns raised by violations of this assumption pertain to the hypothesis that the control group is, in fact, untreated. To the extent that no significant effect on the control group was found, this should not be of grave concern. Moreover, if any bias is implied by this regarding the effect of Medicare Part D eligibility on the treated group it is to bias that effect towards 0.

This concern is further allayed by use of the alternative control group in section 6.2. The finding of a significant effect on the treated group relative to this alternative control and the null effect of Medicare Part D eligibility on individuals with no employer sponsored insurance whatsoever provides further evidence of the mechanism of retirement lock irrespective of the assumption made here.

2) Inference of Experimental Group from post Age 65 Employment and Insurance Statuses. This approach is a little more complex, though the idea is simple: consider individuals who reported that they have retiree insurance but do not know if it is limited to age 65 or not, or who have missing values for the question on retiree insurance for any reason. If over time they retire it can be inferred whether or not their retiree coverage extends past age 65 by observing whether they are covered by an employer plan when they retire and are over age 65.

This is especially useful for individuals who were 61-64 in 1996 or 63-64 in 1998: such individuals were asked if they had retiree insurance but were not asked if it was limited to age 65 or not. In future waves with more detailed questions they were not probed further because they were already over age 65, and thus not asked questions regarding retiree insurance. There were 684 such individuals in 1996, and 295 such individuals in 1998.

The main difficulty in putting this approach into practice is that it will not reveal the retiree insurance status of respondents over age 65 who still work. This difficulty can be partially circumvented by observing the same individual over time until she is retired. If her employer plan continues to cover her in retirement then it can be inferred that she was covered by a plan that would cover her in retirement even when she was employed. As she is over 65, this places her unambiguously in the control group of individuals with retiree health insurance past age 65.

If, however, the respondent is observed retired and over age 65 without insurance, then it is not immediately clear if she would have had retiree insurance only until age 65 (and thus belong in the treatment group) or whether she had no retiree insurance at all (and thus should not be included in the sample). To deal with this ambiguity we must refer again to lagged responses of the same individual from before age 65. If at those ages the individual at some point replied she had retiree health insurance then she can be included in the treatment group. Otherwise she is assumed not to have had retiree insurance at all, and thus is excluded from the sample.

Concretely, the approach I take is to consider for each respondent the first period after age 65 in which she is retired and check whether or not she has retiree insurance at that point. If she does, I assign her to the control group in all previous periods as well. If she does not I check whether before age 65 she claimed she would have some form of retiree insurance should she retire. If she did she is assigned to the treatment group in all periods. If she did not she is excluded from the sample.

This approach substantially increases the size of the sample, salvaging many observations with missing values or unknown age limits for retiree insurance. For example, in the baseline specification it increases the number of individuals observed from 4,934 using just strategy (1) to 6,515 using both (see table 3 and table 12). However, it implies some selection of workers out of the sample. Specifically, individuals who continue working throughout the period they are observed in the HRS cannot reveal their retiree insurance status in this way.

It is not clear that this selection should be different across the treatment and control groups and its overall magnitude is small as the vast majority of individuals do, in fact, retire by the later ages considered (I check for retirement among individuals as old as 75-76 in 2010, covering to these ages even the youngest individuals in 1996-1998 who would not be asked about their post-65 retiree insurance status). Nevertheless, in order to be sure that this selection is not biasing the results to any great extent, Appendix B replicates the main results of the paper using a sample where treatment and control groups are constructed only based on strategy (1). This leads to a smaller sample and thus larger standard errors but the qualitative results remain quite robust.

A constraint that the HRS survey, even coupled with these two procedures, places on the sample is that in the early years of the sample (2000 and 2002) individuals of particularly advanced ages cannot have their retiree insurance status identified if they do not actually have retiree insurance after age 65: for example, an individual aged 69 in 2000 would not have been asked regarding retiree health insurance in wave 5 (2000) or waves 3 and 4 (1996-1998) because she would have been over age 65 in all those survey waves, and retiree health insurance was not inquired about in previous waves. If she is retired and insured by an employer past age 65 she can be placed in the control group- but if she is uninsured it is impossible to tell whether it is because her insurance was limited to age 65 or because she had no retiree insurance at all. It is for this reason that the sample for the entire analysis is based on ages 55-68: respondents over 68 who should properly belong in the treatment group would be unrepresented in most of the pre-Medicare Part D period.

Definition of Alternative Control Group for Section 6.2 and Descriptive Statistics

Section 6.2 examines the robustness of the central results to use of a different control group. The treatment group in all analyses is the same (individuals with retiree insurance from their employer only until age 65); however while the main analysis is done with a control group of individuals who have retiree health insurance from their employers past age 65, section 6.2 uses a control group of individuals who have no employer-sponsored insurance whatsoever. Construction of this latter group is straightforward: it includes only individuals who have no employer sponsored insurance. This includes insurance from a current or previous employer or union, of one's own or of one's spouse. All respondents are asked this question and there are few missing values (an average of 210 missing values out of about 20,000 observations each wave).

Table 1 provides a comparison of descriptive statistics for the three experimental groups: the treatment group (individuals with retiree insurance only until age 65), the main control group (individuals with retiree insurance past age 65), and the alternative control group (individuals with no employer-sponsored insurance).

In general the demographic characteristics of the treatment group and the main control group are quite similar, although the control group is slightly wealthier and more educated on average. While the two groups differ substantially in levels with respect to full-time work rates, the identification strategy I employ requires parallel trends between these two groups, rather than identical levels in their outcomes. The assumption of parallel trends in the absence of the treatment can be assessed visually in the pre-2006, pre-age 65 trends of the treatment and main control group, as depicted in figures 7 and 8.

The alternative control group is much less similar to the treatment group in both demographic characteristics and in levels of the outcome variables than the main control group is. As is to be expected, the alternative control group is less educated, less wealthy, has lower income and has a much higher share of women than the treatment group. They are also less likely to have prescription drug insurance coverage, and more likely to have public prescription drug insurance coverage before introduction of Medicare Part D. Furthermore, their rate of full-time work is lower pre-treatment, as are their average annual labor earnings. However, here too, the identifying assumption is one of parallel trends in the absence of treatment rather than identical levels. This assumption can be assessed by examining the pre-trends in figure 11. The very different baseline characteristics of the alternative control group and the treatment group are a motivating factor in the choice of individuals with retiree health insurance past age 65 as the main control group for the analysis, rather than individuals with no employer-sponsored insurance at all.

B Main Results with Experimental Groups Defined only by Lagged Values from before Age 65

As discussed in Appendix A, construction of the sample requires knowledge of the employer-sponsored insurance status of respondents after retirement. If they are insured in retirement but only until age 65 they are in the treatment group; if they are insured in retirement past age 65 they are in the control group; if they are neither then they are not included in the sample (except for the sample in section 6.2, see the main text and Appendix A for details).

While the HRS contains all the necessary information for construction of these groups for individuals below age 65, at age 65 and over questions regarding retiree health insurance are not asked. It is therefore necessary to infer retiree insurance status for observations aged 65 or over. This is done by two strategies detailed in Appendix A. The first uses answers given by individuals interviewed when they were younger than 65 to infer their retiree insurance status after age 65. The second fills in the gaps due to missing or ambiguous answers by inferring from the observed retiree insurance status after age 65 for a given individual what that individual's employer offered retirees.

This second method admits into the sample individuals who are observed retired and over age 65 at some point during the sample period. For them it is possible to see if they are insured in retirement past age 65, and thus infer that when they were not retired they were plausibly nevertheless offered retiree coverage past age 65 should they retire. However, this method cannot admit into the sample individuals who are never observed retired, and thus selects out of the sample by construction some individuals who keep working throughout the sample period. The implications of this selection are discussed in Appendix A. This appendix aims to demonstrate robustness of the main results to using a sample constructed using only the first method, which does not involve any possible selection on work status.

Table 12 replicates the main results of the paper using this smaller but less potentially selected sample. The table shows the effect of Medicare Part D eligibility on rates of full-time and part-time work. The estimation method is the same as that used in Section 5: differences-in-differences, comparing mean outcomes for individuals just over age 65 to those just under age 65 after 2006, and subtracting from that the same difference between those just over and just under age 65 before 2006. This comparison (including the controls in the baseline equation) is displayed in columns (1) and (4) for full-time work and part-time work, respectively, with the sample restricted to those whose retirement lock was relaxed by Part D's introduction– individuals with retiree health insurance from their employers only until age 65. Columns (2), (3), (5), and (6) estimate the effect of Part D eligibility using triple differences, with the same treated group as above, but now also differencing out the differences-in-differences for a control group of individuals who have employer-sponsored retiree health insurance unlimited by age. Columns (2) and (5) do this without individual fixed effects (instead including richer demographic controls). Columns (3) and (6) do this using the baseline specification.

The results in table 12 are remarkably similar to those in tables 2, and 3. In all specifications, with both samples, the rate of full-time work declines for the treated group by around 8 percentage points (except column (2) where the estimated effect is 5 percentage points). The effect on part-time work is also of similar magnitude using this smaller sample. For the main analysis the estimated effect was an increase of around 6 percentage points, while with this smaller sample it is between 9 and 12 percentage points. Furthermore, in all specifications in both samples the control group has no significant effect, as expected (except for a marginally significant effect in column (2)). The main difference between the results in the main analysis and here are the standard errors. Unsurprisingly, the standard errors in table 12 are somewhat larger. This is due to the smaller sample used here.

C Estimation of Elasticity of Insurance Demand

In this appendix I estimate the response of insurance coverage to introduction of Medicare Part D for use in the calibration of Part D's costs in Section 7. Estimation is based on the basic differences-in-differences design described in Section 4. The dependent variable is a dummy for prescription drug insurance coverage. Results are in table 13.

The prescription drug insurance coverage rate increases by 13 percentage points upon Part D eligibility for the treatment group. The baseline insurance coverage rate for this group is 0.887. Thus the elasticity of coverage is $0.13/0.887 = 0.15$. This is a proxy for the parameter required in 15, assuming that everyone who buys insurance buys the average quantity of insurance.

D Extensive Policy Change Model

The policy change considered in Section 2 is a marginal increase in subsidy for prescription drug insurance for retirees. However, the introduction of Medicare Part D was not an incremental increase in a subsidy but large change in such a subsidy, from no subsidy at all to around \$1,800 worth of subsidy per capita a year. In addition, Part D is more than just a subsidy; for example, it involved the creation of online "markets" to compare and select different plans. Individuals may value these miscellaneous changes apart from their valuation of dollars of subsidy.

The model in Section 2 also assumes some structure on preexisting insurance markets, and how individuals interact with them. This provides intuition regarding what might drive a valuation of the Part D subsidy above and beyond valuation of simple income. However, the fact of such excess valuation is not dependent on the specifics of the modeling assumptions made, but rather can be inferred from the estimation of labor responses to Part D irrespective of such assumptions. In this appendix I present a simple variation of the model in Section 2 which allows for a discrete policy change, which is not necessarily denoted in dollars. Furthermore, I will impose no structure here on the insurance markets mechanisms underlying individuals' valuations of the subsidy.

This analysis is based on finding the equivalent variation of the Part D policy change, with the only difference from the typical equivalent variation analysis being that it is primarily measured in labor supply responses, which are then put into dollar terms. Such an approach is closely related to that described in Hendren [2013a].

Setup Individuals have preferences over two goods, consumption, c, and some policy, $s \in \{0,1\}$, as well as a disutility from labor, v_i.

$$U_i = u_i(c(l_i), s(l_i)) - v_i * l_i \tag{17}$$

Labor is once again modeled as an extensive margin decision, where $l_i = 1$ if individual i works full-time, and $l_i = 0$ if not. Consumption is assumed to be equal to income in this static model and so is larger when working full-time than when not, $c(1) \equiv c_1 > c_0 \equiv c(0)$. Labor disutility is distributed according to a cumulative distribution function $G(v)$, with a probability density function of $g(v)$.

Furthermore, the policy is dependent on labor supply. Before the policy change $s(l_i) = 0$ for all i, and for any l. After the policy change retirees enjoy the policy while full-time workers do not: $s(1) = 0, s(0) = 1$.

Optimal Labor Choice Before the policy change individual i works full-time if and only if:

$u_i(c_1, 0) - u_i(c_0, 0) \equiv \overline{v}_0 > v_i$

In other words, i works full-time only if the utility from the added consumption of full-time work minus her labor disutility is larger than the utility of consumption from less than full-time work. This defines a labor disutility cutoff below which individuals work full-time and above which they do not.

Similarly, after policy change individual i works full-time if and only if:

$u_i(c_1, 0) - u_i(c_0, 1) \equiv \overline{v}_1 > v_i$

Here, too, there is a labor disutility cutoff below which individuals work full-time and above which they do not. This cutoff is now lower because the utility in the non-working state is higher due to the policy.

Analysis of the Policy Change Define the change in utility when not working full-time due to s as:

$$\overline{v}_0 - \overline{v}_1 = u_i(c_0, 1) - u_i(c_0, 0) \equiv \Delta u \tag{18}$$

This change in utility is precisely equal to the change in the labor disutility cutoff. Therefore the policy change will lead to a decline in labor supply associated with a decline in the cutoff labor disutility of full-time work. The change in labor supply associated with s is therefore:

$$\Delta G(v) \equiv G(\overline{v}_0) - G(\overline{v}_1) = \int_{\overline{v}_0 - \Delta u}^{\overline{v}_0} G(v) dv \tag{19}$$

Equivalent Variation Calculation Consider a different policy change, which increases retirement consumption, c_0, to $\tilde{c}_0 = c_0 + \Delta c$:

As before, prior to the policy change individual i works full-time if and only if:

$u_i(c_1, 0) - u_i(c_0, 0) \equiv \overline{v}_0 > v_i$

After the policy change individual i works full-time if and only if:

$u_i(c_1, 0) - u_i(\tilde{c}_0, 0) \equiv \tilde{v} > v_i$

Where $\tilde{v}$ is the labor disutility cutoff when retirement consumption has been increased by Δc. As before, $\tilde{v} < \overline{v}_0$, this time due to the added utility of additional consumption in retirement.

Define the change in utility when not working full-time due to Δc as:

$$\overline{v}_0 - \tilde{v} = u_i(\tilde{c}_0, 0) - u_i(c_0, 0) \equiv \widetilde{\Delta u} \tag{20}$$

As above, this decline in the labor disutility cutoff leads to a decline in the share of the population working full-time:

$$\Delta \tilde{G}(v) \equiv G(\overline{v}_0) - G(\tilde{v}) = \int_{\overline{v}_0 - \widetilde{\Delta u}}^{\overline{v}_0} G(v) dv \tag{21}$$

Lemma. If Δc is such that $\widetilde{\Delta u} = \Delta u$ then: 1) Individuals value Δc precisely as much as they value the policy s; 2) $\widetilde{\Delta u} = \Delta u$ if and only if $\Delta \tilde{G}(v) = \Delta G(v)$.

Proof. (1) follows immediately from the definitions in equation (18) and equation (20). (2) follows immediately from the definitions of equation (19) and equation (21).

This lemma shows that if we choose Δc such that $\Delta \tilde{G}(v) = \Delta G(v)$ then we will have found the equivalent variation of s such that individuals value s as much as they value Δc.

Calibration Section 5 estimated precisely that $\Delta G(v) = 0.0836$. As described in Section 7, Gelber et al. [2015] found that \$10,000 increase in Social Security leads to a decline in participation of 0.011 among 64-66 year-olds. Thus the equivalent variation of Medicare Part D is \$76,000; i.e., Part D is valued as another \$76,000 of lifetime discounted (annually at 3%) social security wealth.

A further assumption regarding the nature of Part D can get us to the same willingness to pay calculation in Section 7.2. If we assume the sum of the policy change implicit in Medicare Part D is the additional subsidy to prescription drug insurance, the same calibration used in Section 7 can get us that the monetary value of Part D is $25,000. Therefore willingness to pay for one dollar of the subsidy can be calibrated as $\frac{76{,}000}{25{,}000} \approx 3$, as in Section 7.2.

References

Jason Abaluck and Jonathan Gruber. Heterogeneity in choice inconsistencies among the elderly: Evidence from prescription drug plan choice. *American Economic Review*, 101 (3):377–81, 2011.

Jason Abaluck, Jonathan Gruber, and Ashley Swanson. Prescription drug use under medicare part d: A linear model of nonlinear budget sets. Working Paper 20976, National Bureau of Economic Research, February 2015.

America's Health Insurance Plans. A survey of medigap enrollment trends, 2006, 2006. America's Health Insurance Plans Center for Policy Research.

Padmaja Ayyagari and Dan M. Shane. Does prescription drug coverage improve mental health? evidence from medicare part d. *Journal of Health Economics*, 41:46–58, 2015.

Katherine Baicker, Amy Finkelstein, Jae Song, and Sarah Taubman. The impact of medicaid on labor market activity and program participation: Evidence from the oregon health insurance experiment. *The American Economic Review*, 104(5):322–328, 2014.

Martin Neil Baily. Some aspects of optimal unemployment insurance. *Journal of Public Economics*, 10(3):379–402, 1978.

David M. Blau and Donna B. Gilleskie. Retiree health insurance and the labor force behavior of older men in the 1990s. *Review of Economics and Statistics*, 83(1):64–80, 2001.

David M. Blau and Donna B. Gilleskie. Health insurance and retirement of married couples. *Journal of Applied Econometrics*, 21(7):935–953, 2006.

Gary Burtless. Social security, unanticipated benefit increases, and the timing of retirement. *The Review of Economic Studies*, 53(5):781–805, 1986.

Raj Chetty. A general formula for the optimal level of social insurance. *Journal of Public Economics*, 90(10):1879–1901, 2006.

Raj Chetty. Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy*, 116(2):173–234, 2008.

Raj Chetty. Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annu. Rev. Econ.*, 1(1):451–488, 2009.

Raj Chetty, John N. Friedman, Tore Olsen, and Luigi Pistaferri. Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records. *The Quarterly Journal of Economics*, 126(2):749, 2011.

Courtney Coile and Jonathan Gruber. Future social security entitlements and the retirement decision. *The review of Economics and Statistics*, 89(2):234–246, 2007.

Dora L. Costa. *The Evolution of Retirement: An American Economic History, 1880-1990*. University of Chicago Press, 1998.

David Cutler. Market failure in small group health insurance. Working Paper 4879, National Bureau of Economic Research, October 1994.

Mark Duggan, Patrick Healy, and Fiona Scott Morton. Providing prescription drug coverage to the elderly: America's experiment with medicare part d. *The Journal of Economic Perspectives*, pages 69–92, 2008.

Michael W. L. Elsby, Bart Hobijn, Ayşegül Şahin, Lawrence F. Katz, and Robert Shimer. The labor market in the great recession [with comments and discussion]. *Brookings Papers on Economic Activity*, pages 1–69, 2010.

Gary V. Engelhardt and Jonathan Gruber. Medicare part d and the financial protection of the elderly. *American Economic Journal: Economic Policy*, 3(4):77–102, 2011.

Itzik Fadlon and Torben Heien Nielsen. Household responses to severe health shocks and the design of social insurance. Working Paper 21352, National Bureau of Economic Research, July 2015.

Martin Feldstein. Tax avoidance and the deadweight loss of the income tax. *Review of Economics and Statistics*, 81(4):674–680, 1999.

Martin Feldstein and Jeffrey B. Liebman. Social security. *Handbook of Public Economics*, 4: 2245–2324, 2002.

Eric French and John Bailey Jones. The effects of health insurance and self-insurance on retirement behavior. *Econometrica*, pages 693–732, 2011.

Craig Garthwaite, Tal Gross, and Matthew J. Notowidigdo. Public health insurance, labor supply, and employment lock. *The Quarterly Journal of Economics*, 129(2):653–696, 2014.

Alexander Gelber, Adam Isen, and Jae Song. The effect of social security benefits on earnings: Novel evidence from the social security notch. work in progress, 2015.

David H. Greenberg. The leisure bias in cost-benefit analyses of employment and training programs. *Journal of Human Resources*, pages 413–439, 1997.

David H. Greenberg and Philip K. Robins. Incorporating nonmarket time into benefit–cost analyses of social programs: An application to the self-sufficiency project. *Journal of Public Economics*, 92(3):766–794, 2008.

Jonathan Gruber. The incidence of mandated maternity benefits. *The American Economic Review*, pages 622–641, 1994.

Jonathan Gruber. Disability insurance benefits and labor supply. Working Paper 5866, National Bureau of Economic Research, December 1996.

Jonathan Gruber and Brigitte C. Madrian. Health-insurance availability and the retirement decision. *The American Economic Review*, pages 938–948, 1995.

Jonathan Gruber and Brigitte C. Madrian. Health insurance, labor supply, and job mobility. *Health Policy and the Uninsured*, page 97, 2004.

Alan L. Gustman and Thomas L. Steinmeier. The 1983 social security reforms and labor supply adjustments of older individuals in the long run. *Journal of Labor Economics*, pages 237–253, 1985.

Alan L. Gustman and Thomas L. Steinmeier. Employer-provided health insurance and retirement behavior. *Industrial & Labor Relations Review*, 48(1):124–140, 1994.

Alan L. Gustman and Thomas L. Steinmeier. Effects of social security policies on benefit claiming, retirement and saving. *Journal of Public Economics*, 129:51 – 62, 2015. ISSN 0047-2727.

Martin B. Hackmann, Jonathan T. Kolstad, and Amanda E. Kowalski. Health reform, health insurance, and selection: Estimating selection into health insurance using the massachusetts health reform. *American Economic Review*, 102(3):498–501, 2012.

Health and Retirement Study. "public use data", 2013. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740). Ann Arbor, MI, (2013).

Nathaniel Hendren. The policy elasticity. Working Paper 19177, National Bureau of Economic Research, June 2013a.

Nathaniel Hendren. Private information and insurance rejections. *Econometrica*, 81(5): 1713–1762, 2013b.

Michael D. Hurd and Michael J. Boskin. The effect of social security on retirement in the early 1970s. *The Quarterly Journal of Economics*, 99(4):767–790, 1984.

Robert Kaestner, Cuiping Long, and G. Caleb Alexander. Effects of prescription drug insurance on hospitalization and mortality: Evidence from medicare part d. Working Paper 19948, National Bureau of Economic Research, February 2014.

Kaiser Family Foundation. Employer health benefits, 2014 annual survey, 2014. Kaiser Family Foundation and Health Research and Education Trust.

Kanika Kapur. The impact of health on job mobility: A measure of job lock. *Industrial & Labor Relations Review*, 51(2):282–298, 1998.

Lynn A. Karoly and Jeannette A. Rogowski. The effect of access to post-retirement health insurance on the decision to retire early. *Industrial & Labor Relations Review*, 48(1): 103–123, 1994.

Henrik Jacobsen Kleven and Claus Thustrup Kreiner. The marginal cost of public funds: Hours of work versus labor force participation. *Journal of Public Economics*, 90(10): 1955–1973, 2006.

Alan Krueger and Jorn-Steffen Pischke. The effect of social security on labor supply: a cohort analysis of the notch generation. *Journal of Labor Economics*, 10(4):412–437, 1992.

Alan B. Krueger and Bruce D. Meyer. Labor supply effects of social insurance. *Handbook of Public Economics*, 4:2327–2392, 2002.

Frank R. Lichtenberg and Shawn X. Sun. The impact of medicare part d on prescription drug use by the elderly. *Health Affairs*, 26(6):1735–1744, 2007.

Robin L. Lumsdaine, James H. Stock, and David A. Wise. Pension plan provisions and retirement: men and women, medicare, and models. In *Studies in the Economics of Aging*, pages 183–222. University of Chicago Press, 1994.

Brigitte C. Madrian and Nancy Beaulieu. Does medicare eligibility affect retirement? In *Inquiries in the Economics of Aging*, pages 109–131. University of Chicago Press, 1998.

Brigitte C. Madrian, Gary Burtless, and Jonathan Gruber. The effect of health insurance on retirement. *Brookings Papers on Economic Activity*, pages 181–252, 1994.

Medicare Board of Trustees. 2014 annual report of the board of trustees of the federal hospital insurance and federal supplementary medical insurance trust funds, 2014. Medicare Board of Trustees.

Steven Nyce, Sylvester J. Schieber, John B. Shoven, Sita Nataraj Slavov, and David A. Wise. Does retiree health insurance encourage early retirement? *Journal of Public Economics*, 104:40–51, 2013.

Mark V. Pauly and Yuhui Zeng. Adverse selection and the challenges to stand-alone prescription drug insurance. In *Forum for Health Economics & Policy*, volume 7, 2004.

Franco Peracchi and Finis Welch. Trends in labor force transitions of older men and women. *Journal of Labor Economics*, 12(2):210, 1994.

RAND HRS Data. Version N., 2014. Produced by the RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration. Santa Monica, CA (September 2014).

Jeannette Rogowski and Lynn Karoly. Health insurance and retirement behavior: evidence from the health and retirement survey. *Journal of Health Economics*, 19(4):529–539, 2000.

Christopher J. Ruhm. Bridge jobs and partial retirement. *Journal of Labor Economics*, pages 482–501, 1990.

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

John P. Rust. Behavior of male workers at the end of the life cycle: An empirical analysis of states and controls. In *Issues in the Economics of Aging*, pages 317–382. University of Chicago Press, 1990, 1990.

Andrew A. Samwick. New evidence on pensions, social security, and the timing of retirement. *Journal of Public Economics*, 70(2):207–236, 1998.

Robert Shimer and Iván Werning. Reservation wages and unemployment insurance. *The Quarterly Journal of Economics*, pages 1145–1185, 2007.

John B. Shoven and Sita Nataraj Slavov. The role of retiree health insurance in the early retirement of public sector employees. *Journal of Health Economics*, 38:99–108, 2014.

Tax Policy Center. Historical average federal tax rates for all households, 2014. URL `http://www.taxpolicycenter.org/taxfacts/displayafact.cfm?Docid=456`. Tax Policy Center.

Wilbert Van der Klaauw and Kenneth I. Wolpin. Social security and the retirement and savings behavior of low-income households. *Journal of Econometrics*, 145(1):21–42, 2008.