



Essays on Public Health Insurance

Citation

Wettstein, Gal. 2016. Essays on Public Health Insurance. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:33493442>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Essays on Public Health Insurance

A dissertation presented

by

Gal Wettstein

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

April 2016

© 2016 Gal Wettstein

All rights reserved.

ESSAYS ON PUBLIC HEALTH INSURANCE

ABSTRACT

Over the last ten years there have been dramatic changes in the health insurance environment in the United States, spurred on by broad reforms in the public health insurance sector. In 2006 the Medicare Prescription Drug, Improvement and Modernization Act went into effect, providing broad access to prescription drug insurance for millions of elderly Americans. In 2014 the main provisions of the Patient Protection and Affordable Care Act began to be felt, dramatically changing health insurance markets, particularly for those seeking non-group coverage. These legislative changes both raise questions regarding how well the policy changes meet their goals, as well as offering new variation with the potential to answer questions of fundamental economic significance.

This dissertation addresses such important questions surrounding the effectiveness of public health insurance in meeting policymakers' goals, and the implications of public health insurance for private markets. In the three chapters of this dissertation I utilize the policy changes of Medicare Part D and the Affordable Care Act to provide quasi-experimental estimates of retirement lock, of the correlation of risk aversion and crowd-out of private insurance, and of the effectiveness of the individual health insurance mandate in expanding coverage.

The first part studies the implications of public drug insurance for labor markets. This part 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.

The second part of this dissertation considers the effect of public drug insurance on private drug coverage, with a focus on the correlation of crowd-out and risk aversion. I utilize Health and Retirement Survey data around the time of introduction of the Medicare Part D prescription drug insurance for the elderly in order to estimate crowd-out of private prescription drug insurance. I use individuals between the ages of 55 and 64, who are not eligible for the program, as a control group relative to individuals aged 65 to 75, who are eligible. I take a differences-in-differences approach to estimation by comparing outcomes before and after 2006, when Medicare Part D went into effect. I construct measures of risk aversion by exploiting unique questions eliciting risk preferences in the Health and Retirement Survey, as well as information on whether individuals have other kinds of insurance, or engage in risky behaviors. I find substantial differential crowd-out by risk aversion: every standard deviation increase in risk aversion was associated with about 5 percentage points less crowd-out, over a base crowd-out rate of 50%-60%. More risk averse individuals also saw greater reductions in out-of-pocket spending on prescription drugs due to Part D, particularly at high levels of spending: at the 85th percentile of spending an individual one

standard deviation more risk averse than the average experienced a decline of \$110/year due to Part D eligibility, above and beyond the gains for an averagely risk averse individual of \$382/year.

The third part of the dissertation estimates the effectiveness of the individual mandate in the Patient Protection and Affordable Care Act in expanding health insurance coverage. This paper studies the impact of the individual health insurance mandate in the Patient Protection and Affordable Care Act (PPACA) on health insurance coverage. This mandate went into effect in 2014, alongside various other elements of the PPACA. I focus on individuals ages 26-64 who are ineligible for the subsidies or Medicaid expansions included in the PPACA to isolate the effect of the mandate from these other components. To account for changes unrelated to the PPACA that occur over time and affect insurance coverage I utilize a control group of residents of Massachusetts who were already subject to mandated insurance following the 2006 health care reform in their state. Employing a differences-in-differences design applied to data from the American Community Survey, I find that the mandate caused an increase of 0.85 percentage points in health insurance coverage, or a 17% decline in the uninsurance rate. This increase was concentrated in coverage purchased directly by individuals, rather than acquired through an employer, and predominantly affected younger individuals. Both these observations are consistent with the mandate ameliorating adverse selection in the individual health insurance market.

TABLE OF CONTENTS

ABSTRACT	iii
ACKNOWLEDGMENTS	viii
PART I: RETIREMENT LOCK AND PRESCRIPTION DRUG INSURANCE: EVIDENCE FROM MEDICARE PART D	1
1.1 Introduction.....	1
1.2 Conceptual Framework of Retirement Lock.....	6
1.3 The Medicare Part D Program	12
1.4 The Health and Retirement Study Data and Empirical Strategy	14
1.5 Estimation of Prescription Drug Insurance Retirement Lock.....	21
1.6 Robustness Checks	48
1.7 Welfare Implications of Medicare Part D.....	57
1.8 Conclusions.....	67
PART II: HETEROGENEITY BY RISK AVERSION IN CROWD-OUT OF PRIVATE INSURANCE	70
2.1 Introduction.....	70
2.2 Data and Risk Aversion Indices.....	74
2.3 Institutional Details on Medicare Part D	79
2.4 Differences-in-Differences Estimation of Crowd-Out with Heterogeneity.....	81
2.5 Results.....	84
2.6 Conclusions.....	99
PART III: THE EFFECT OF INDIVIDUAL HEALTH INSURANCE MANDATES ON INSURANCE COVERAGE	101
3.1 Introduction.....	101
3.2 The PPACA’s Individual Mandate and Other Institutional Details.....	105
3.3 Data and Identification.....	109
3.4 Results.....	113
3.5 Conclusions.....	126
APPENDIX A TO PART I.....	130

APPENDIX B TO PART I.....	137
APPENDIX C TO PART I.....	139
APPENDIX D TO PART I.....	139
REFERENCES	144

ACKNOWLEDGMENTS

I am very grateful to my advisors David Cutler, Edward Glaeser, Nathaniel Hendren, and Lawrence Katz for their continuous advice and support; as well as to Alberto Alesina for his guidance through the years.

I would also like to express my deepest gratitude to my dear friends Itzik Fadlon, Mira Frick, Sangram Kadam, Marc Kaufmann, Danial Lashkari, Morgan McClellon, and Assaf Romm for profound and sincere interactions, and to my parents, my sister, and my partner for their love and constant encouragement. I could not have completed this dissertation without their emotional and physical support.

I am forever grateful to Dr. Alexandra Houck, Dr. David Fisher, and Michele Walsh, NP, as well as all the staff at Harvard University Health Services and the Dana-Farber Cancer Institute. I owe them my life.

I gratefully acknowledge generous financial support from the Boston College Center for Retirement Research.

Part I

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

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.

¹In 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).

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 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

³This 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.

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

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

⁴There 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).

⁵A 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).

⁶Two 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

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 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)

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.

⁷The 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.

⁸Other papers in this literature include Lichtenberg and Sun [2007], Kaestner et al. [2014], Abaluck et al. [2015], and Ayyagari and Shane [2015].

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:

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

⁹97% of the treatment group had some form of prescription drug coverage before becoming eligible for Medicare Part D; see table 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 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 : $\frac{dy_i}{dx_i} \equiv y'_i < 0, \frac{d^2y_i}{dx_i^2} \equiv y''_i > 0$.

$x_i(p)$ 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

¹⁰There 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).

¹¹Medicare 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.

given by:

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

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:

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

where c_{1i}^*, c_{0i}^* are the values of consumption after having optimally chosen the level of insurance conditional on labor choice. $\bar{v}(s)$ is the cutoff value of labor disutility above which individuals choose to stop working full-time. An individual with labor disutility $v_i = \bar{v}(s)$ 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

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:

¹²In 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 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).

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

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): $\frac{d\bar{v}(s)}{ds} = -E_Y[u'_{0i}(s) * \frac{dc_i}{ds}]$. Therefore the change in the actual share of individuals working full-time will be:

$$\frac{dG(\bar{v}(s))}{ds} = -g(\bar{v}(s)) * E_Y[u'_{i0}(s) * \frac{dc_i}{ds}] < 0. \quad (5)$$

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

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

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:

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

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:

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

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 $\frac{\frac{dG(\bar{v}(s))}{ds}/x_i}{\frac{dG(\bar{v})}{db}} = \frac{Cov(u'_{0i}(s), \frac{dc_i}{ds}) + E_Y[u'_{0i}(s)] * E_Y[\frac{dc_i}{ds}]}{x_i E_Y[u'_{0i}(s)]}$. 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 $\frac{\frac{dG(\bar{v}(s))}{ds}/x_i}{\frac{dG(\bar{v})}{db}}$ above 1:

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

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 $\frac{dG(\bar{v}(s))}{ds}$.

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:

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

For a small s the key quantity $s \frac{dG(\bar{v}(s))}{ds} \approx \frac{\Delta G(\bar{v}(s))}{\Delta s}$, 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

¹³The 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.

meant that every American over age 65 would have access to prescription drug insurance. 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.

¹⁴There 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.

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.

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

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

¹⁶The 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.

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 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:

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

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

¹⁷A 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.

¹⁸The 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.

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, 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

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

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.

²⁰Except 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.

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.

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 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.

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, 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).

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.

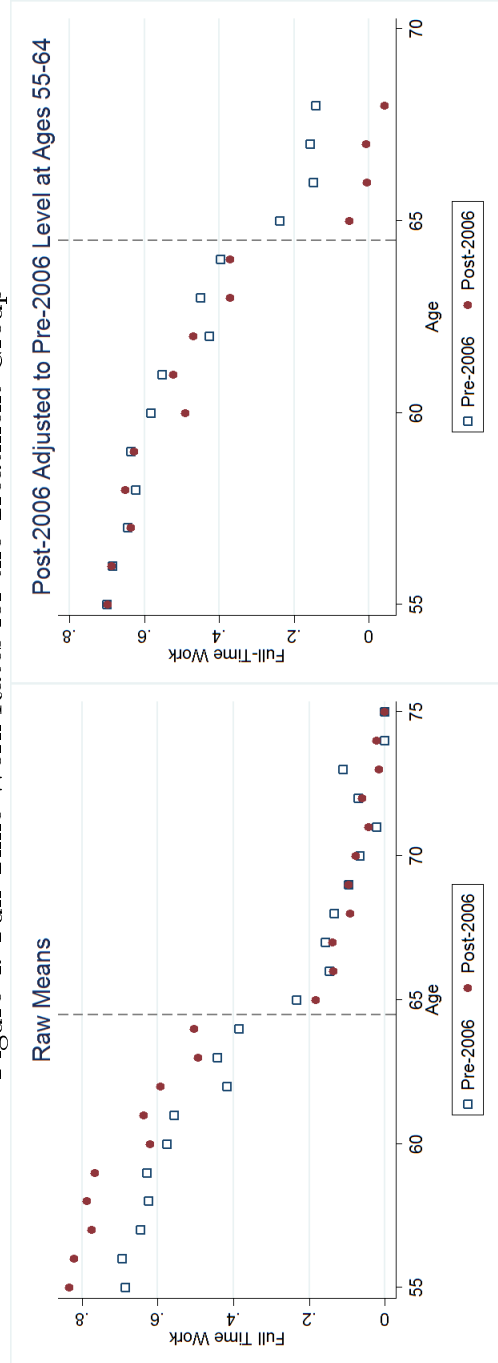
Of particular interest, however, is the noticeably larger decline in the full-time work rate at age 65 after 2006, relative to before 2006. This is a visual representation of the differences-in-differences estimation of the effects of Medicare Part D eligibility on full-time work. Also of note is the 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

²¹Very 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.


²²It 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.

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 55-64, 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.

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

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.

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

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

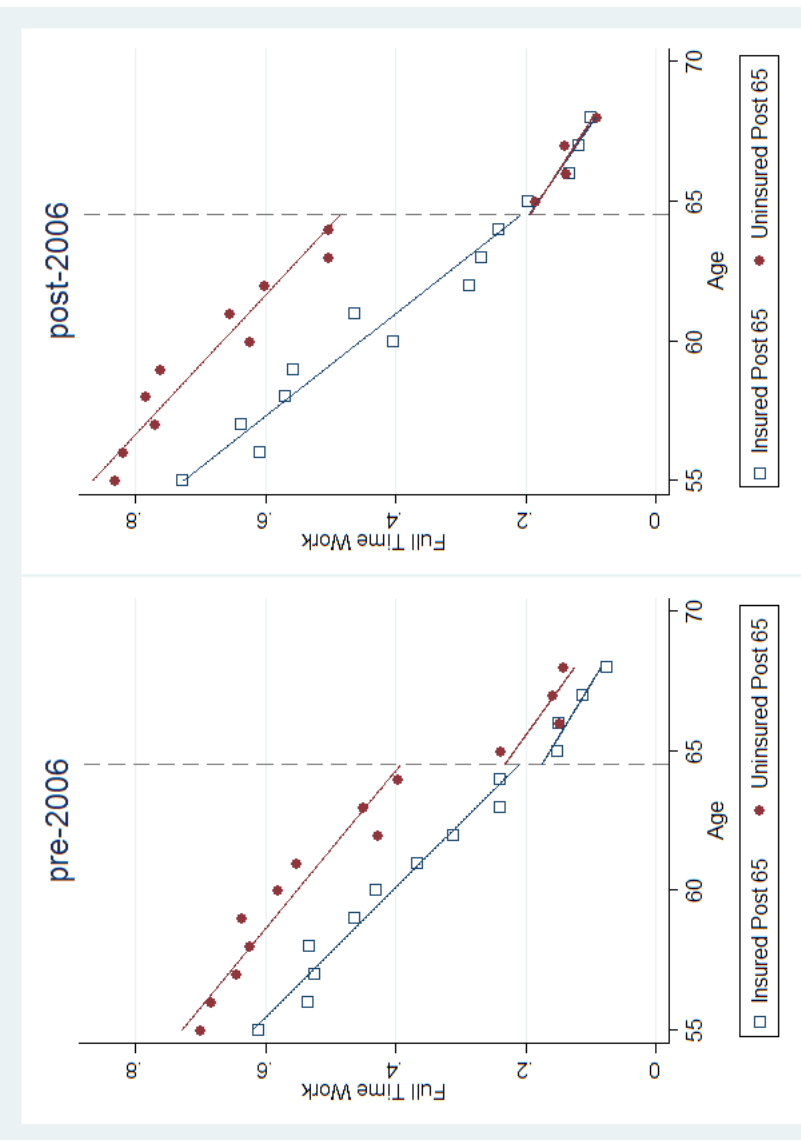


Notes: This figure shows the differences-in-differences of full-time work for the control group. The sample is individuals aged 55-75, in the years 2000 until 2010, who have retiree health insurance through their employer unlimited by age. The blue squares indicate rates of full-time work by age in the 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).

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

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.

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,

²³This 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.

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	
	No Demographic, Health Controls	Baseline	Baseline	Baseline	Baseline	Baseline
Post65*Post2006*Treated	-0.0845*** (0.0311)	-0.0836*** (0.0313)	0.0589* (0.0308)	-0.0247 (0.0338)		
Post65*Post2006	0.0234 (0.0216)	0.0199 (0.0217)	0.00157 (0.0218)	0.0215 (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.

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 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


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

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

Dependent Variable: Specification /Sub-Sample:	Full-Time Work		
	Full Sample	Only Ages 62-68	Excluding Medicaid and VA
Post65*Post2006*Treated	-0.0912*** (0.0290)	-0.0852*** (0.0318)	-0.0802* (0.0452)
Post65*Post2006	-0.0263 (0.0199)	0.0168 (0.0220)	0.0259 (0.0228)
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	No	Yes	Yes
Age and Year Dummies * Demographics	No	Yes	No
N	15,303	15,371	9,790
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.

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).

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

²⁵Results 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.

Table 5: Job Switches

Dependent Variable:	Part-Time Work	Job-Switching	Part-Time Work * Job-Switching
Post65*Post2006*Treated	0.0589* (0.0308)	0.0439** (0.0209)	0.0408** (0.0162)
Post65*Post2006	0.00157 (0.0218)	-0.00126 (0.0144)	-0.00481 (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.

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

²⁶In 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 \bar{v} , and thus also reduce labor supply.

Table 6: Effect on Annual Labor Earnings

Dependent Variable: Specification:	Annual Labor Earnings in 2010 Dollars			Wages
	No Individual Fixed Effects	Baseline	Enhanced Baseline	Baseline
Post65*Post2006*Treated	-4,138** (1,900)	-1,477 (1,915)	-4,269** (2,002)	-0.00949 (7)
Post65*Post2006	-348.2 (1,382)	1,557 (1,441)	1,959 (1,523)	4.834 (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.

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 similar heterogeneity by health

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

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

²⁸The 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).

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.

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 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.

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

Table 7: Heterogeneity by Health Status

Dependent Variable:	Full-Time Work		Part-Time Work	
	Sick	Healthy	Sick	Healthy
Sub-Sample:				
Post65*Post2006*Treated	-0.122*** (0.0374)	0.00536 (0.0652)	0.0990*** (0.0380)	-0.0113 (0.0586)
Post65*Post2006	0.0416 (0.0256)	-0.0269 (0.0466)	-0.0221 (0.0266)	0.0353 (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.

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 otherwise. The estimated effect of Part D is remarkably similar to results in the previous

Table 8: Heterogeneity by Spousal Health Status

Dependent Variable: Sub-Sample:	Full-Time		Part-Time	
	No Sick Spouse	Sick Spouse	No Sick Spouse	Sick Spouse
Post65*Post2006*Treated	-0.1704*** (0.0494)	-0.0135 (0.0437)	.0889* (0.0483)	0.0602 (0.0445)
Post65*Post2006	0.0436 (0.0326)	0.0184 (0.0312)	0.0086 (0.0326)	-0.00536 (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:	Hours/Week	Hours/Week (Conditional on Working)	More Than 35 Hours/Week
Post65*Post2006*Treated	-2.667** (1.349)	-4.914*** (1.655)	-0.0770** (0.0319)
Post65*Post2006	0.553 (0.947)	-0.922 (1.290)	0.0290 (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.

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.

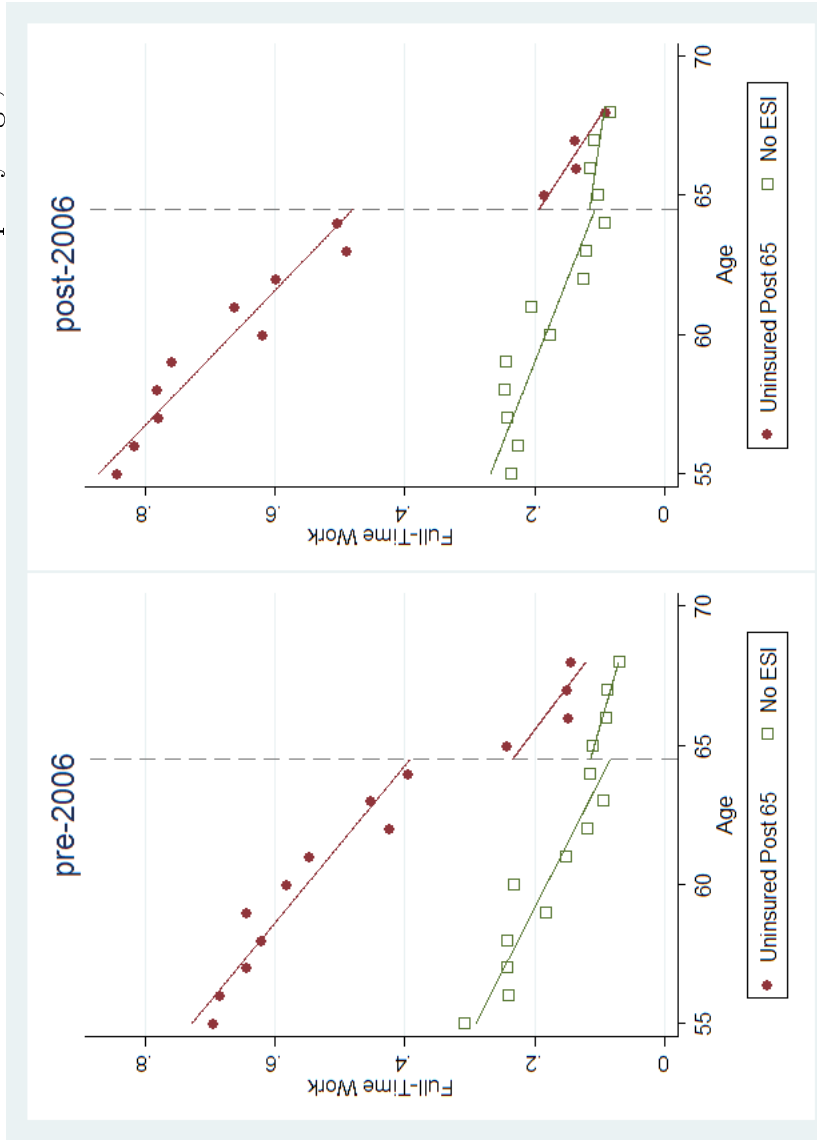
Those without any employer-sponsored insurance are less similar to the treatment group than those with retiree insurance to any age on virtually every observable, from gender distribution to income (see columns 1 and 3 of table 1). This second control group nevertheless allows me to test the robustness 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

³⁰Qualitatively 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.

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).

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

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** (0.0320)	0.0245 (0.0327)
Post65*Post2006	-0.00959 (0.0160)	0.0189 (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

Dependent Variable:	Only 2000-2006		Only 2000-2006 and 2010	
	Full-Time	Part-Time	Full-Time	Part-Time
Post65*Post2006*Treated	-0.0560 (0.0440)	0.0164 (0.0427)	-0.0928** (0.0392)	0.0433 (0.0370)
Post65*Post2006	0.0109 (0.0289)	0.0158 (0.0278)	0.0237 (0.0262)	0.00573 (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 Part 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.

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 2023³¹ (Medicare Board of Trustees, 2014). These benefits include the premiums enrollees pay themselves. I therefore subtract from these

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

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 $s \frac{dG(\bar{v}(s))}{ds} \approx \frac{\Delta G}{\Delta s} = 0.084$. It is now possible to calibrate R by comparing $s \frac{dG(\bar{v}(s))}{ds}$ to the effect on labor supply of increasing lifetime discounted Social Security wealth by b dollars, as in equation (10).

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 $b \frac{dG(\bar{v})}{db} = 0.011$, where $b = 10000$.

Rewriting R and plugging in the estimates yields:

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

R is estimated to be 2.03, substantially larger than 0.³³ In other words, the effect of a dollar

³²The 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].

³³Assuming there is only sampling error in my own estimate of $s \frac{dG(\bar{v}(s))}{ds}$, R is significantly larger than 0

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

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

at a 90% level of significance. Its 95% confidence interval is [-0.2,4.24].

³⁴Hurd 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 $R = 2.25$.

³⁵Samwick [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.

³⁶Van 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 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.

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

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

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

³⁸The estimation equation here is simplified in order to reduce computational complexity by excluding individual fixed effects and health controls and including only a quadratic in non-housing household wealth.

³⁹In 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.

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.

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:

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

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 $WTP =$

$\frac{s \frac{dG(\bar{v}(s))}{ds} / sx_i}{b \frac{dG(\bar{v})}{db} / b} = R + 1$. 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:

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

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

⁴¹Kleven 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.

where A signifies revenue per capita from sources other than income tax; $(1 - G(\bar{v}(s)))s$ is 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 $I(s) \equiv G(\bar{v}(s)) * I(1) + (1 - G(\bar{v}(s))) * I(0)$. The effect on the budget of offering another dollar of subsidy is therefore given by:

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

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:

$$(1 - G(\bar{v}(s))) = 0.65$$

$$s \frac{dG(\bar{v}(s))}{ds} = 0.084$$

$$\frac{sdI(s)}{ds} = 1,477$$

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:

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

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 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 $WTP * (1 - G(\bar{v}(s)))$. Combining this with the social costs estimated above gives:

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

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

⁴²This 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]).

⁴³In 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.

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 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.

⁴⁴Hendren [2013a] calculates the *MVPFs* 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.

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.

Part II

Heterogeneity by Risk Aversion in Crowd-Out of Private Insurance

1 Introduction

When assessing the benefits of a publicly provided good a crucial question is to what extent the public provision of the good crowds-out consumption of similar goods supplied through private markets. Typically, no social benefits can be gained when individuals merely replace a privately supplied good with a similar publicly supplied one.⁴⁵ Such considerations are central to the optimal design of provision of public goods in general, and of public insurance in particular. However, the average extent of crowd-out is not a sufficient statistic for net welfare gain from a public good when individuals are heterogeneous in their taste for the good, and this heterogeneity is correlated with the rate of crowd-out.

A natural index of individuals' valuation of insurance is their level of risk aversion. In this paper I estimate heterogeneity in crowd-out of private prescription drug insurance along the dimension of risk aversion. I find that there is less crowd-out for more risk averse individuals. While I leave a careful analysis of the welfare implications of this pattern aside, it suggests that the welfare gains from public drug insurance are larger than they would appear when assessing them using only the average crowd-out rate. This is because those for whom there is more net gain of insurance are those who most highly value that insurance.

In order to estimate crowd-out of private prescription drug insurance I use the 2006 introduction of Medicare Part D, which provides subsidies to prescription drug insurance

⁴⁵Although this is not always the case; see Gruber [1996], Greenberg [1997], Greenberg and Robins [2008], and Wettstein [2015].

for Americans over age 65. As a result of this change some individuals replaced the drug coverage they had before from a private source (e.g., their employer or a private Medigap policy) with the newly available public insurance. I employ a differences-in-differences design based on this policy change to estimate the change in overall net drug coverage associated with Medicare Part D. To the extent this falls short of a 1-to-1 increase in coverage due to take-up of Medicare Part D, that is my main measure of crowd-out.⁴⁶

The data I use in this analysis are from the Health and Retirement Study (HRS, Health and Retirement Study [2013]). These data are uniquely suited to estimating heterogeneity by risk aversion, as they include both direct measures of risk aversion (from questions designed to elicit risk preferences) and information about a number of behaviors which are conceptually associated with risk aversion (such as buying other kinds of insurance, or engaging in risky behaviors such as excessive drinking). Using these variables I construct indices of risk aversion, and use them to estimate how crowd-out varies with risk aversion.

My two main measures of risk aversion are a binary measure, which relies only on the risk aversion elicitation questions; and a continuous measure, defined by the principal component of the risk aversion category implied by those questions, as well as whether the individual has long-term care insurance and whether they engage in excessive drinking. Both these measures yield remarkably consistent estimates of the effect of risk aversion on crowd-out: they both imply that an increase of one standard deviation in risk aversion is associated with almost 5 percentage points less crowd-out, over a base crowd-out rate of 50%-60%.

The larger overall increase in coverage following take-up of public coverage among the highly risk averse also translates into greater reductions in out-of-pocket spending on prescription drugs. Public coverage reduces the probability of having out-of-pocket spending in the top 5% for the highly risk averse by 4.5 percentage points above and beyond the decline in this probability for the less risk averse. Furthermore, quantile regressions reveal that the

⁴⁶I find similar results when my measure of crowd-out is a decline in private coverage. The differences between these two approaches are discussed in detail in Cutler and Gruber [1996a].

more risk averse see larger declines in out-of-pocket spending due to eligibility for Medicare Part D at every part of the spending distribution. At the 85th percentile, for example, the average reduction in out-of-pocket spending due to eligibility for Part D was \$32 a month, or \$382 a year. However, for an individual one standard deviation more risk averse than the average that Part D-induced decline in spending was more than \$9 a month greater, leading to an annual reduction of \$492 a year.

Evidence from the number of health insurance plans individuals hold suggests the more risk averse increase their overall number of plans when they take up public coverage more than the less risk averse, perhaps as a means of supplementing the public insurance where it provides little protection (as in the Medicare Part D coverage gap). While public drug coverage leads to an increase of 0.1 in the mean number of health insurance plans held by low-risk aversion individuals, it increases the average number of plans held by the highly risk averse by 0.17. This is consistent with more risk averse individuals keeping more of their preexisting private coverage, or acquiring more new supplemental private coverage, alongside taking up new public coverage.

Furthermore, the highly risk averse had slightly lower levels of drug coverage in the pre-Part D period (about 1.7 percentage points less). This can be explained by higher participation in traditional Medicare at the expense of Medicare Advantage plans which covered drugs but also covered only limited health care provider networks. The evidence suggests that the risk averse want to avoid the risk of needing an out-of-network physician or hospital more than the risk of uninsured drug costs. This, along with the greater propensity of the highly risk averse to hold multiple plans, can explain the lower crowd-out rates among them.

This paper relates to a number of lines of previous research. First, it follows a long line of research dealing with crowd-out of private health insurance by public insurance. This literature tends to find substantial crowd-out of private health insurance by public alternatives. For example, Cutler and Gruber [1996a] find crowd-out of about 50% from

Medicaid expansions. Even more closely related to the current paper, Engelhardt and Gruber [2011] find 75% crowd-out of private prescription drug insurance by Medicare Part D.⁴⁷

Second, I rely on a broad literature dealing with risk aversion, its measurement, and its consistency across domains. In particular, constructing a measure of risk aversion utilizing behavior in different domains builds conceptually on Einav et al. [2012], who find evidence that there is a cross-domain general component of risk aversion.⁴⁸ Furthermore, the HRS questions which elicit risk aversion and their properties were examined in great detail in Barsky et al. [1997]. Implications of heterogeneity in risk aversion for health insurance have been studied, for example, in Cutler et al. [2008] and Fang et al. [2008]. While these papers generally find risky behaviors and risk tolerance can be associated with demand for insurance, the direction of that association can be very different across different insurance products.

In addition, the implications of heterogeneity in risk aversion for optimal social insurance are explored in Andrews and Miller [2013]. They modify the standard Baily-Chetty (Baily, 1978, Chetty, 2006) formula to account for heterogeneity in risk aversion and find that this may have important implications for welfare analysis. In the context on unemployment insurance they calibrate a model under different assumption on the distribution of risk preferences, and find that the covariance of drops in consumption at unemployment with risk aversion can change the estimate of the benefit from public unemployment insurance by 50%. The current paper brings together some insights from their work and the work on the impact of preexisting private insurance markets on welfare analysis in Chetty and Saez [2010].

Finally, this paper also contributes to the literature on Medicare Part D itself. An overview of early results on the structure and the effects of Part D is available in Duggan

⁴⁷A partial list of other papers in this literature includes Taylor et al. [1988], Bergstrom et al. [1986], Wolfe and Goddeeris [1991], Cutler and Gruber [1996b], Finkelstein [2004], Golosov and Tsyvinski [2007] and Chetty and Saez [2010].

⁴⁸Barseghyan et al. [2011] find evidence which qualifies the generality of risk preferences across domains, although it does not contradict a general component.

et al. [2008]. A great deal of research quantifies the effect of Medicare Part D on health expenditures and other outcomes: for example, Lichtenberg and Sun [2007] study the effect of Part D coverage on the utilization of prescription drugs.⁴⁹ This literature finds that Part D substantially increased prescription drug utilization among the elderly, while reducing their out-of-pocket expenses.

The rest of the paper proceeds as follows: section 2 describes the data and the construction of indices of risk aversion; section 3 provides institutional details on Medicare Part D; section 4 describes the empirical design; section 5 contains the results and a brief discussion of the possible mechanism; and section 6 concludes.

2 Data and Risk Aversion Indices

The data I use are primarily from the RAND version of the HRS (RAND HRS Data, 2014), supplemented as necessary from the raw HRS data. These data survey a random sample of non-institutionalized Americans over the age of 50 and their spouses, following up every two years. As the policy change I consider took place in 2006 and covers individuals over age 65, I limit the sample to years 1998-2010 (waves 4-10 of the HRS), and to ages 55-75.

The main dependent variable in the analysis is prescription drug insurance coverage, which takes the value of 1 (some drug insurance) or 0 (no drug insurance). In addition, some of the analysis will focus on out of pocket spending on drugs (in 2010 dollars/month). Another crucial variable is whether or not the individual has taken up public drug coverage. Descriptive statistics for these variables in the pre-2006 period can be found in table 12.

⁴⁹Other papers in this literature include Zhang et al. [2009], Blume-Kohout and Sood [2013], Lakdawalla et al. [2013], Kaestner et al. [2014], Abaluck et al. [2015], Ayyagari and Shane [2015], and Wettstein [2015].

Table 12: Descriptive Statistics for Sample Years 1998-2004, Ages 55-64

	Mean	Std. Dev.
Any Drug Coverage	0.86	0.35
Monthly Out-of-Pocket Drug Spending	\$76.41	490
Public Drug Coverage	0.06	0.23
Women	0.56	0.5
Age	64.55	5.7
Household Income	\$83,440	170637
Long-Term Care Coverage	0.09	0.29
Excessive Drinking	0.04	0.21
High Risk Aversion	0.65	0.48
Principal Component Risk Aversion	0	1

Notes: This table presents descriptive statistics for the control sample. The sample is restricted to ages 55-64 and years 1998-2004 (except for the statistics on age, high risk aversion and the principal component measure of risk aversion which are not restricted): before meeting the age criteria of Medicare Part D eligibility and only in the years before introduction of Medicare Part D in 2006. All monetary values are inflated to 2010 prices using the consumer price index. The first column shows the mean of the variable in that row; the second column shows the standard deviation.

Take-up of public drug coverage that is not associated with an increase in net coverage is how I measure crowd-out. To be precise, I will consider the effect of any public drug insurance, through Medicare Part D or otherwise, on net insurance coverage. The consideration of any public coverage rather than only Part D coverage is primarily due to the automatic transfer of individuals who had drug coverage through Medicaid to providers covered by Part D in 2006. Considering public coverage in general, rather than just Part D coverage, avoids the problem of these individuals being classified as newly going onto public coverage when, in fact, they are just switching between two sources of public coverage.

There remains a question of how to treat individuals who had coverage through Medicare Health Management Organizations (HMOs) before 2006. Many, though not all, such individuals received drug coverage through their HMO. It is not straightforward to classify such coverage as either private or public. On the one hand, these individuals were paying for supplemental coverage not included in Medicare, and so this could be considered private coverage. On the other hand, the prices these individuals paid were low due to cross-subsidization of risk, just as in Part D, and so such coverage resembles public coverage. On this question I follow the convention in Engelhardt and Gruber [2011] and classify individuals

covered by government HMOs as publicly covered. The main results are qualitatively robust to classifying government HMOs as private coverage.

Finally, the HRS allows individuals to claim more than one insurance plan. Employer plans tend to be more generous than public coverage, particularly with respect to drug coverage; and as a general rule employers with more than 20 employees are the primary insurer with respect to any additional Medicare coverage. Therefore, in cases where individuals claim multiple plans, of which one is an employer (private) plan and one is public I assume they have private coverage. The results are not very sensitive to this choice.

Measuring Risk Aversion

To estimate heterogeneity in crowd-out by risk aversion I construct two indices of risk aversion. The first relies on questions in the HRS of the following form: “*Suppose that you are the only income earner in the family. Your doctor recommends that you move because of allergies, and you have to choose between two possible jobs. The first would guarantee your current total family income for life. The second is possibly better paying, but the income is also less certain. There is a 50-50 chance the second job would double your total lifetime income and a 50-50 chance that it would cut it by $x\%$. Which job would you take – the first job or the second job?*” The potential loss of income, x , varies from 10% to 75%. Based on the answers to these question individuals can be divided into four groups by increasing risk aversion.⁵⁰

Some respondents to the survey in waves 1, 4, 5, 6, 7, and 8 were selected to answer this series of questions. To maximize the potential sample size I assume that risk aversion is largely stable over time.⁵¹ This allows me to impute risk aversion for many individuals when they are not asked these questions by carrying forward their answers from previous waves. In the regressions below I will also include a dummy variable for imputed risk aversion status.

⁵⁰Six groups are possible in some survey waves, but for consistency across waves I use the four-group partition.

⁵¹The R^2 of a regression of the raw risk aversion score only on its lag is 0.87.

The resulting 4-point “risk aversion score” forms the basis for the rest of the analysis around risk aversion heterogeneity. For an in-depth discussion of this variable and its properties see Barsky et al. [1997]. Briefly, they find that this measure is sensibly related to risky behaviors such as smoking and drinking, to having insurance, and to holding stocks rather than bonds. Nevertheless, they also find that risk aversion measured in this way generally explains only a small part of the variation in these variables.

For ease of interpretation I will rely on a dichotomous variable based on this risk aversion score: it takes a value of 1 for those in the highest category of risk aversion, and 0 otherwise; I will call this variable “high risk aversion”. Assuming constant relative risk aversion over the relevant income range, Barsky et al. [1997] calculate that the lower bound of the relative risk aversion parameter for people in the high risk aversion group is 3.76. Descriptive statistics for this variables are also in table 12. Note that about 65% of the sample falls in the highest risk aversion category group; and that the standard deviation of the high risk aversion variable is roughly 0.5, implying a movement from 0 to 1 is equivalent to an increase of around two standard deviations in risk aversion.

This is an imperfect measure of risk aversion for three main reasons. The first is due to the fact that the risk preference elicitation questions do not provide a perfect measure of risk aversion. The second is as a result of the fact that risk preferences in one domain may not be a perfect reflection of risk preferences in other domains (Barseghyan et al. [2011], Einav et al. [2012]). Finally, the third is because of the need to impute many of the values of risk aversion. As such this crucial explanatory variable is most likely subject to substantial measurement error; the results using this measure of risk aversion are therefore likely to be biased towards zero, and any estimated heterogeneity by risk aversion muted.

Alternative Risk Aversion Measure

A variation on this approach is to use additional variables which are conceptually related to risk aversion to augment the risk aversion score. While this does not solve the problem of measurement error in a straightforward way, it can at least provide added variation in

the measure of risk aversion and increase power. Furthermore, this approach provides a robustness check relative to using the risk aversion score alone. To implement this approach I use the first principal component of various sets of these additional variables and the 4-point risk aversion score.

The HRS provides many variables to choose from here. One set of variables which might be related to risk preferences is whether the individual is covered by other forms of insurance, besides health or prescription drug insurance. The HRS asks regarding long-term care insurance and life insurance. Of these long-term care insurance is more cleanly related to risk preferences regarding one's own consumption, and so I will use that variable.⁵²

The second set of variables which should be associated with risk preferences are risky behaviors such as smoking or excessive drinking.⁵³ Of the two, smoking is more likely to reflect experiences much earlier in life, particularly for the sample of those around age 65 in 2006 who likely started smoking before the risks of smoking were widely known. Therefore I use drinking behavior as an additional factor in calculating a risk aversion index. The HRS asks how many drinks the individual drinks in a day on which the individual drinks (set to 0 when the individual reports never drinking). I define heavy drinking as 1 when this variable is 4 or greater, and 0 otherwise.⁵⁴⁵⁵

To form the alternative risk aversion index I perform a principal component analysis of the three variables– the 4-point risk aversion score, possession of long-term care insurance,

⁵²Life insurance involves primarily individual preferences for leaving bequests, which might be quite different from risk aversion. Both these insurance choices clearly involve not only risk preferences but also risk. In the regressions below I will control for various health indicators. Furthermore, the choice of purchasing insurance is also likely to be associated with other personal characteristics such as income and wealth. In the regressions I control for income and wealth, as well as number of children.

⁵³Both of these can have health impacts which directly affect the need for prescription drug consumption, and therefore drug insurance. They can nevertheless provide a robustness check for the risk aversion measure based only on the risk preference questions.

⁵⁴Results are robust to perturbations of this cutoff.

⁵⁵The main results are qualitatively similar when using the smoking variable instead of, as well as in addition to, excessive drinking in the analysis which follows.

and excessive drinking— and take as an index the resulting first principal component. I then standardize this index so a 1 unit increase corresponds to an increase of 1 standard deviation in risk aversion. Descriptive statistics for the two additional variables, long-term care insurance and heavy drinking, and the standardized first principal component can be found in table 12. The signs of the weights on the three variables in the first principal component are sensible (0.45, 0.62, and -0.64 for the risk aversion score, long-term care insurance, and excessive drinking, respectively) and consistent with a higher value being associated with greater risk aversion.

3 Institutional Details on Medicare Part D

This section provides some institutional details regarding the Medicare Part D program: a change to traditional Medicare that 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).

For those over age 65 before 2006 private prescription drug insurance could be acquired through one of three essentially private options: an employer or union plan which covered drugs⁵⁶, a Medigap policy which supplemented traditional Medicare with drug coverage⁵⁷,

⁵⁶ In 2005 98% of employer plans also covered prescription drugs (Kaiser Family Foundation, 2014).

⁵⁷ Take-up of such plans was extremely low. 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).

or a stand-alone prescription drug plan.⁵⁸ In addition, public drug insurance was available to some low-income individuals through Medicaid, and to others through limited programs like Veterans Affairs. The largest somewhat public option for drug coverage was a Medicare Advantage managed care plan which covered drugs.⁵⁹ Overall, before 2006 a quarter of Medicare beneficiaries had no drug coverage whatsoever (Safran et al., 2005).

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 could have access to prescription drug insurance. 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. The program was highly effective in increasing coverage rates for those eligible, and by 2006 less than 10% of them lacked drug coverage (Engelhardt and Gruber, 2011).

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 individuals who were eligible for Medicaid and became eligible for Medicare Part D in 2006 were automatically enrolled in Part D plans. Roughly 7% of Americans over age 65 had drug coverage through Medicaid prior to 2006 (Safran et al., 2005). To avoid counting Medicaid coverage being switched to Medicare Part D coverage as crowd-out of

⁵⁸In practice such plans were almost completely unavailable (Pauly and Zeng, 2004).

⁵⁹Penetration of such plans was relatively low, between 9%-13% of Medicare beneficiaries in 2004 (Safran et al. [2005] and Mathematica Policy Research, 2008); and these plans often capped coverage for drugs (Pauly and Zeng, 2004).

private insurance, I take a general view of public insurance under any program crowding-out private coverage.

In sum, whereas before 2006 access to public prescription drug insurance was mostly restricted to those on Medicaid or in limited network Medicare Advantage plans, from 2006 onward everyone over age 65 had the option of purchasing subsidized prescription drug insurance. This new public insurance may have replaced some coverage acquired on the private market, primarily through employers. This sharp change in 2006 for individuals over age 65 forms the basis of my identification strategy, to which I turn in the next section.

4 Differences-in-Differences Estimation of Crowd-Out with Heterogeneity

The empirical strategy I use extends the approach taken in Engelhardt and Gruber [2011]: crowd-out is estimated using a differences-in-differences design where observations aged 55-64 provide a control group, and observations ages 65-75 are the treatment group, which is treated from year 2006 and onward. I then use this differences-in-differences to instrument for public coverage. The first stage of this estimation gives an estimate of the take-up rate of public drug insurance among those eligible for such coverage. In the second stage I regress the outcome of interest, primarily prescription drug insurance coverage, on the first stage estimated take-up of public insurance. In both stages I allow the treatment effect to vary by risk aversion by interacting the differences-in-differences with a measure of risk aversion.

The estimation equations are therefore:

$$\begin{aligned}
Public_{i,t,a} = & \beta_{1,1} * Post2006_{i,t} * Over65_{i,t,a} * RA_{i,t} + \beta_{1,2} * Post2006_{i,t} * Over65_{i,t,a} + \quad (17) \\
& \beta_{1,3} * Post2006_{i,t} + \beta_{1,4} * Over65_{i,t,a} + \beta_{1,5} * RA_{i,t} + \\
& \beta_{1,6} * RA_{i,t} * Post2006_{i,t} + \beta_{1,7} * RA_{i,t} * Over65_{i,t,a} + \\
& \alpha_{1,a} + \gamma_{1,t} + \delta_{1,a} * RA_{i,t} + \zeta_{1,t} * RA_{i,t} + \sum_{j=1}^k \theta_{1,j} X_{1,j,i,t,a} + \varepsilon_{1,i,t,a},
\end{aligned}$$

$$\begin{aligned}
Public_{i,t,a} * RA_{i,t} = & \beta_{2,1} * Post2006_{i,t} * Over65_{i,t,a} * RA_{i,t} + \beta_{2,2} * Post2006_{i,t} * Over65_{i,t,a} + \quad (18) \\
& \beta_{2,3} * Post2006_{i,t} + \beta_{2,4} * Over65_{i,t,a} + \beta_{2,5} * RA_{i,t} + \\
& \beta_{2,6} * RA_{i,t} * Post2006_{i,t} + \beta_{2,7} * RA_{i,t} * Over65_{i,t,a} + \\
& \alpha_{2,a} + \gamma_{2,t} + \delta_{2,a} * RA_{i,t} + \zeta_{2,t} * RA_{i,t} + \sum_{j=1}^k \theta_{2,j} X_{j,i,t,a} + \varepsilon_{2,i,t,a},
\end{aligned}$$

and

$$\begin{aligned}
Insured_{i,t,a} = & \beta_{3,1} * \widehat{Public}_{i,t,a} * RA_{i,t} + \beta_{3,2} * \widehat{Public}_{i,t,a} + \quad (19) \\
& \beta_{3,3} * Post2006_{i,t} + \beta_{3,4} * Over65_{i,t,a} + \beta_{3,5} * RA_{i,t} + \\
& \beta_{3,6} * RA_{i,t} * Post2006_{i,t} + \beta_{3,7} * RA_{i,t} * Over65_{i,t,a} + \\
& \alpha_{3,a} + \gamma_{3,t} + \delta_{3,a} * RA_{i,t} + \zeta_{3,t} * RA_{i,t} + \sum_{j=1}^k \theta_{3,j} X_{j,i,t,a} + \varepsilon_{3,i,t,a},
\end{aligned}$$

where *Public* is 1 if the individual has taken up public drug coverage, *Post2006* is 1 for observations in 2006 or later, *Over65* is 1 for observations who are 65 or older, *RA* is the measure of risk aversion, and *Insured* is 1 if the individual has prescription drug coverage.

$\widehat{Public}_{i,t,a}$ and $\widehat{Public}_{i,t,a} * RA_{i,t}$ are the estimated rates of take-up of public drug insurance and the estimated interaction of that take-up with the measure of risk aversion, as estimated in the first stage equations, 17 and 18, respectively. Furthermore, all specifications include age and year fixed effects, also interacted with the measure of risk aversion.

X is a vector of additional controls. These include gender, and full sets of dummies for being single, residence in each of the census divisions, years of education, race (white, African American, or other), religion (Protestant, Catholic, Jewish, None, or other), labor force status (full-time, part-time, unemployed, partially retired, retired, disabled, and not in the labor force), and fifth-order polynomials in non-housing household wealth and household income. Additional health controls are also included: 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, memory problems, or psychiatric conditions. Finally, a dummy is included for whether the risk aversion measure was imputed.⁶⁰⁶¹ All monetary variables are inflated to 2010 prices by the consumer price index. All standard errors are clustered at the individual level.⁶²

To the extent that the control group provides a counterfactual trend in insurance coverage in the absence of Part D, this estimation procedure reveals the causal effect of public drug insurance take-up on the rate of overall drug coverage. It also provides the correlation of this treatment effect with risk aversion. The resulting estimate is conceptually then simple to translate into an estimate of crowd-out; in the absence of any crowd-out the relation of public insurance take-up and coverage should be 1-to-1. Anything less than that reflects crowd-out.

⁶⁰Results are not sensitive to inclusion of this dummy.

⁶¹Individual fixed effects are generally not included because they absorb much of the variation in risk aversion. However, all main results remain similar in magnitude and sign if individual fixed effects are included.

⁶²Where possible, results are also robust to clustering at the household level.

Thus the estimate of crowd-out for a given level of risk aversion is $1 - \beta_{3,2} - \beta_{3,1} * RA$. If crowd-out declines with risk aversion, for example, it should be reflected in a positive value of $\beta_{3,1}$.

Whether or not the control group is actually suitable for this purpose is the next issue which must be addressed. For that I turn now to results.

5 Results

5.1 Overall Crowd-Out Estimates


Before turning to the main focus of heterogeneity of crowd-out by risk aversion, I first estimate a base average rate of crowd-out, ignoring any heterogeneity in the effect. This first step is useful in order to assess the plausibility of my identification strategy in a setting where visualization of the results is clear. In doing so I essentially replicate, on a different dataset, part of the analysis in Engelhardt and Gruber [2011]. This is helpful as a benchmark against which to scale the heterogeneity later on.

Graphical Evidence

A necessary condition for the control group, 55-64 year-olds, to provide a credible counterfactual trend for the treatment group of 65-75 year-olds is that in the pre-treatment period (years 1998-2004) the two groups move in parallel. For rates of public drug coverage this can be seen quite clearly in figure 13. In this figure the blue squares indicate the share of individuals holding public drug coverage at any year in the sample for individuals in the control group, while the treatment group rates are indicated by the red circles. Beyond the very similar trends between the two groups, there is a clear increase in public coverage for the treated group in 2006 of about 50 percentage points which is not mirrored by the control group (for whom there does not seem to be any substantial change in 2006). This increase in public coverage for the treatment group upon becoming eligible for Medicare Part D provides a visual counterpart to the first stage equation 17 (if heterogeneity by risk

aversion is neglected). This take-up rate is also consistent with the previous literature (e.g., in Engelhardt and Gruber [2011] the authors find public coverage in 2007 was about 70% for those over 65, and about 10% for those under 65).

Figure 13: Public Drug Insurance Coverage Rates



Notes: This figure shows the differences-in-differences of public prescription drug insurance coverage. The sample is individuals aged 55-75, in the years 1998 until 2010. The blue squares indicate rates of public prescription drug coverage for those aged 55-64 by year, while the red circles indicate public drug coverage for those aged 65-75. The dashed gray line differentiates between years before and after Medicare Part D.

Similarly parallel pre-trends for the treatment and control groups also hold for the main outcome of interest, any prescription drug insurance coverage. This can be seen in figure 14. On this outcome, as well, there is a dramatic increase for the treatment group in 2006, of about 15 percentage points, with no apparent change for the control group. This is the graphical version of the reduced form implied by equations 17-19 when all risk aversion terms are neglected.

Figure 14: Rate of Any Drug Coverage



Notes: This figure shows the differences-in-differences of prescription drug insurance coverage from any source. The sample is individuals aged 55-75, in the years 1998 until 2010. The blue squares indicate rates of prescription drug coverage for those aged 55-64 by year, while the red circles indicate drug coverage for those aged 65-75. The dashed gray line differentiates between years before and after Medicare Part D.

Regression Evidence

I now estimate regressions based on equations 17-19, still neglecting all terms involving risk aversion. Doing so shows that the visual results described above are not sensitive to adding controls, and provides more precise estimates along with their statistical significance. These results can be found in table 13. Column 1 shows the reduced form effect of Medicare Part D eligibility on coverage, an increase of nearly 17 percentage points. Column 2 shows the two-stage least squares estimate of the effect of public coverage on insurance coverage.⁶³ The estimate here implies that taking up public coverage increases net coverage by about 46 percentage points, implying a crowd-out rate of 54%. This is somewhat less crowd-out than in Engelhardt and Gruber [2011], where the authors estimated a 75% crowd-out; however it is of the same order of magnitude.

⁶³The first stage is not displayed but is highly significant. The F-statistic on the excluded instrument is greater than 4,000.

Table 13: Average Crowd-Out

	(1)	(2)
	Reduced Form	TOLS
<i>Post2006 * Over65</i>	0.168***	-
	(0.006)	-
<i>Public</i>	-	0.46***
	-	(0.018)
Controls	Yes	Yes
No. of Observations	66,284	66,284
No. of Clusters	19,066	19,066

Notes: This table presents estimates of the effects of public prescription drug coverage on overall prescription drug coverage. The sample is individuals ages 55-75 in the years 1998-2010. Column 1 shows the reduced form estimate of eligibility for Medicare Part D, being over age 65 in the years after 2006; column 2 shows the two-stage least squares estimate of the effect of having public coverage, instrumented by eligibility for Medicare Part D. All regressions include the following controls: age and year fixed effects, a full set of dummies for gender, being single, residence in each of the census divisions, years of education, race (white, African American, or other), religion (Protestant, Catholic, Jewish, None, or other), labor force status (full-time, part-time, unemployed, partially retired, retired, disabled, and not in the labor force), and fifth-order polynomials in non-housing household wealth and household income. Additional health controls are also included: 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, memory problems, or psychiatric conditions. Finally, a dummy is included for whether the risk aversion measure was imputed. 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.

5.2 Heterogeneity in Crowd-Out by Risk Aversion

I now turn to analyzing heterogeneity by risk aversion in crowd-out of private insurance by public insurance. As described in section 2, I will conduct the analysis in two parallel strands, using two measures of risk aversion. The first is a dummy for high risk aversion individuals; the second is the standardized first principal component of three variables: the raw 4-point risk aversion score from the survey questions which directly elicit risk preferences, a dummy for whether the individual has long-term care insurance, and a dummy for whether the individual habitually engages in heavy drinking (4 or more alcoholic drinks per day). Reassuringly, the results using both measures are very similar.⁶⁴

To begin, I estimate the reduced form equations of any insurance coverage, regressed on eligibility for Medicare Part D and its interaction with risk aversion (i.e., $Post2006 * Over65$, and $Post2006 * Over65 * RA$). Results for each of the risk aversion measures are in columns 1

⁶⁴Results using other measures (e.g., replacing drinking with smoking) are also generally similar.

and 3 of table 14. Column 1, for example, indicates that mere eligibility for Part D increases drug insurance coverage by about 15 percentage points for low risk-aversion individuals, while for high risk aversion individuals this increase is more than 3 percentage points larger (significant at the 5% level). Similar results are found using the principal component measure of risk aversion.

The reduced form estimates, however, do not account for possible differential take-up of Part D insurance by differentially risk averse individuals. For that I estimate the system of equations 17-19, now including all the terms involving risk aversion (I do this for each of the risk aversion measures). The results of these two-stage least squares estimations are in table 14, columns 2 and 4. In both specifications the base crowd-out rate is slightly less than 60%; however that rate declines with greater risk aversion. Using the binary measure of high risk aversion, those with high risk aversion have about 8 percentage points less crowd-out than those with low risk aversion (significant at the 5% level). This amounts to a 14% decrease in crowd-out for the highly risk averse. Using the principal component measure, every standard deviation of increased risk aversion reduces crowd-out by about 4.8 percentage points (significant at the 5% level).

It is worth noting here that the standard deviation of the extensive risk aversion measure is about 0.5. Thus a shift from low to high risk aversion corresponds to a two-standard deviation change. It is therefore reassuring that not only is the sign of the effect of risk aversion on crowd-out the same using both risk aversion measures, but also that the magnitude of the estimated effect is similar, roughly 4-5 percentage points per standard deviation.

Controlling for Other Dimensions of Heterogeneity

While the differences-in-differences design identifies the causal effect of public insurance coverage on overall coverage, the interaction of this treatment with risk aversion does not yield a causal effect. In particular, risk aversion is likely correlated with many other personal characteristics, and these characteristics might themselves modulate the effect of public insurance on net insurance coverage. Ideally one would like to have exogenous variation in

Table 14: Crowd-Out with Heterogeneity

Risk Aversion Measure	High Risk Aversion			Principal Component		
	(1)	(2)	(3)	(4)	(5)	(6)
	Reduced Form	Baseline	Augmented Baseline	Reduced Form	Baseline	Augmented Baseline
<i>Public</i>	-	0.413*** (0.031)	0.519*** (0.125)	-	0.448*** (0.021)	0.557*** (0.118)
<i>Public * RA</i>	-	0.081** (0.039)	0.075* (0.04)	-	0.048** (0.019)	0.054*** (0.019)
<i>Post2006 * Over65</i>	0.15*** (0.011)	-	-	0.162*** (0.008)	-	-
<i>Post2006 * Over65 * RA</i>	0.032** (0.014)	-	-	0.024*** (0.007)	-	-
<i>Public * Demographics</i>	-	-	Yes	-	-	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
No. of Observations	59,923	59,923	58,938	59,005	59,005	58,030
No. of Clusters	16,416	16,416	16,297	16,364	16,364	16,244

Notes: This table presents estimates of the effects of public prescription drug coverage on overall prescription drug coverage, allowing for heterogeneity in the effect by risk aversion. The sample is individuals ages 55-75 in the years 1998-2010. Columns 1 and 3 show the reduced form estimates of eligibility for Medicare Part D, being over age 65 in the years after 2006, interacted with measures of risk aversion: the extensive measure in column 1 and the principal component measure in column 3 (for precise definitions see text). Columns 2 and 4 show the two-stage least squares estimates of the effect of having public coverage, instrumented by eligibility for Medicare Part D. The estimated system of equations is detailed in equations 17-19. Columns 3 and 6 also present two-stage least squares estimates of the effect of public prescription drug coverage on net coverage allowing for heterogeneity by risk aversion and by additional demographic characteristics: gender, years of education, number of children, a dummy for being single, a dummy for veterans, household assets and household income. All regressions include the following controls: age and year fixed effects, also interacted with the measure of risk aversion, a full set of dummies for gender, being single, residence in each of the census divisions, years of education, race (white, African American, or other), religion (Protestant, Catholic, Jewish, None, or other), labor force status (full-time, part-time, unemployed, partially retired, retired, disabled, and not in the labor force), and fifth-order polynomials in non-housing household wealth and household income. Additional health controls are also included: 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, memory problems, or psychiatric conditions. Finally, a dummy is included for whether the risk aversion measure was imputed. 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.

risk aversion to be sure that the estimated effect is not driven by such omitted interactions of the treatment with other variables. It is, however, difficult to imagine what might produce such variation.

As a second best approach to dealing with this concern, I next estimate augmented versions of equations 17-19. The new first stage equations have as dependent variables the interaction of *Public* with other demographic variables, and on the right-hand side include third-order interactions of *Post2006*Over65* with these variables, as well as the second order interactions of *Post2006* and *Over65* with these same variables. The second stage equation then includes the estimated dependent variables from all the first stage equations, identified off of the interaction of eligibility for Medicare Part D with the demographic characteristics. The demographic characteristics I include here are: gender, years of education, number of children, a dummy for being single, a dummy for veterans, household assets and household income.

The results are in columns 3 and 6 (for the two measures of risk aversion) of table 14, and are very robust to inclusion of these additional controls. The effect of the extensive measure of risk aversion on crowd-out remains 8 percentage points (significant at 10%), while the effect of the principal component measure increases slightly to 5.4 percentage points (significant at 1%).⁶⁵ This stability suggests that the estimated effect of risk aversion on crowd-out is, in fact, due to risk aversion itself rather than to its correlation with other observable characteristics.

Intensive Margin

If the highly risk averse are actually getting more drug insurance coverage as a result of Part D, this should be reflected also in their out-of-pocket expenditures on drugs. Indeed, the more risk averse do decrease their drug expenditures more than the less risk averse when they take up public coverage. This can be seen in columns 1 and 4 of table 15. These

⁶⁵The change in the baseline crowd-out rate is not surprising, as this baseline now refers to individuals who have 0 in all the above demographic variables.

regressions have as their left-hand side variable a dummy for whether monthly out-of-pocket drug expenditures were high or not (defined as being in the 95th percentile of expenditures in the pre-2006 period, \$268/month).⁶⁶ It should be noted that the 95th percentile of spending here corresponds closely to where the Medicare Part D coverage gap is in the post-treatment period; \$268/month is annualized to \$3,216/year, precisely the beginning of the coverage gap in 2008 (although this varies slightly across years).

Looking at column 1 of table 15, for example, shows that while public drug coverage reduces the probability of high out-of-pocket drug expenses by about 3 percentage points for the less risk averse, the effect on the risk averse is more than twice as large, at 8 percentage points. The difference is statistically significant at the 5% level. The results using the principal component measure, in column 4, are less stark but have the same signs (albeit not statistically significant).

The division into high and low out-of-pocket expenditures is illustrative but the effect is more general than that. The scope for reducing costs is, naturally, increasing in the level of costs. Accordingly, at virtually every level of out-of-pocket prescription drug spending the effect of public drug insurance on spending is larger for the more risk averse. This can be seen in figure 15. This figure plots the reduced form quantile regression estimates of the treatment effect of Part D on monthly out-of-pocket spending interacted with the principal component measure of risk aversion (the 95% confidence interval is marked in dashed lines).⁶⁷ The estimated effect sizes therefore correspond to the reduction in out-of-pocket spending due to eligibility for Part D at every centile from the 40th to the 90th for a single standard

⁶⁶Qualitatively similar results are found using different cutoffs of out-of-pocket spending. Furthermore, the same signs are obtained using raw reported out-of-pocket spending, however those estimates have high variance due to the noisiness of the variable and its extreme skewness: fully a third of the sample report no drug spending whatsoever, while the 99th percentile is \$828/month. The large share of zeroes also makes a log-specification to account for the skewness unattractive.

⁶⁷This procedure is an elaboration on that used in Finkelstein and McKnight [2008] and Engelhardt and Gruber [2011]. They show treatment effects by centile using a reduced form differences-in-differences. Here I depict heterogeneity in the treatment effect by focusing instead on the treatment effect in a sub-sample in relation to the rest of the sample.

Table 15: Intensive Margin Changes


	High Risk Aversion			Principal Component		
	(1)	(2)	(3)	(4)	(5)	(6)
	High OOP	No. of Plans	No. of Plans	High OOP	No. of Plans	No. of Plans
<i>Public</i>	-0.033** (0.015)	-	0.101** (0.047)	-0.062*** (0.01)	-	0.128*** (0.03)
<i>Public * RA</i>	-0.047** (0.019)	-	0.07 (0.058)	-0.006 (0.01)	-	0.036 (0.031)
<i>RA</i>	0.009 (0.008)	0.034*** (0.007)	0.033 (0.021)	0.004 (0.005)	0.042*** (0.004)	0.066*** (0.013)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
No. of Observations	61,210	39,466	65,470	60,289	38,757	64,435
No. of Clusters	16,602	14,192	16,806	16,553	14,125	16,758

Notes: This table presents estimates of the effects of public prescription drug coverage on intensive margins of prescription drug insurance coverage, allowing for heterogeneity in the effect by risk aversion. The sample for columns 1, 3, 4, and 6 are individuals ages 55-75 in the years 1998-2010. For columns 2 and 5 the years are 1998-2004. Columns 1 and 4 show the two-stage least squares estimates of the effect of having public coverage, instrumented by eligibility for Medicare Part D interacted with the measure of risk aversion (the extensive measure in column 1 and the principal component measure in column 3, for precise definitions see text), on the probability of having high out-of-pocket drug expenditures, defined as monthly expenditures over \$268 (the 95th percentile of out-of-pocket drug spending in the years 1998-2004). Columns 2 and 5 show the correlations of risk aversion (the extensive measure in column 1 and the principal component measure in column 3) with number of health insurance plans held by the individual in the period 1998-2004. Columns 3 and 6 show the two-stage least squares estimates of the effect of having public coverage, instrumented by eligibility for Medicare Part D interacted with the measure of risk aversion (the extensive measure in column 1 and the principal component measure in column 3, for precise definitions see text), on the number of health insurance plans held by the individual. All regressions include the following controls: age and year fixed effects, a full set of dummies for gender, being single, residence in each of the census divisions, years of education, race (white, African American, or other), religion (Protestant, Catholic, Jewish, None, or other), labor force status (full-time, part-time, unemployed, partially retired, retired, disabled, and not in the labor force), and fifth-order polynomials in non-housing household wealth and household income. Additional health controls are also included: 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, memory problems, or psychiatric conditions. Finally, a dummy is included for whether the risk aversion measure was imputed. 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.

deviation of risk aversion, *above and beyond* the mean reduction due to Part D eligibility.⁶⁸ That the estimates are all negative shows that the reduction in spending is increasing with risk aversion (the difference is significant at the 5% level between the 57th and the 88th percentiles, and at the 10% level between the 48th percentile and the 90th).

⁶⁸The effect sizes below the 40th percentile are essentially 0, and those over the 90th have very large standard errors. I focus on the 40th-90th percentile range for clarity. I use the reduced form here, rather than TSLS, for ease of computation. Additionally, the set of controls used here is smaller, also for speed of computation, and includes only first order terms in household assets and income, and no interactions of age and year fixed effects with risk aversion.

Figure 15: Differential Treatment Effect on OOP Drug Spending for one Std. Dev. of Risk Aversion by Percentile of OOP Spending



Notes: This figure shows the estimated change at different percentiles of out-of-pocket spending on prescription drugs in that spending due to eligibility for Medicare Part D per one standard deviation in the principal component measure of risk aversion (for precise definition see text). At every percentile between the 40th and the 90th the point corresponds to the estimate of the coefficient from a quantile regression of that percentile on eligibility for Part D, interacted with the risk aversion measure. This gives an estimate for how much spending is reduced at that percentile for an individual one standard deviation more risk averse than the average due to becoming eligible for Part D, beyond what would be reduced for an individual of average risk aversion at that percentile of spending. Spending is average monthly spending in 2010 dollars. The controls in these quantile regressions are year and age fixed effects, household income and assets, number of children, and dummies for gender, race (white, African American, or other), religion (Protestant, Catholic, Jewish, none, or other), years of education, marital status, veteran status, census division. In addition, the regressions include a dummy for being over age 65, being observed in 2006 or later, the measure of risk aversion, and all second and third order interaction terms of these variables. The coefficients on the third order interaction are the plotted points. Standard errors are clustered at the individual level. The 95% confidence intervals of the estimated treatment effects are displayed in the dashed lines.

Discussion

What might explain the lower rates of crowd-out of private insurance with the introduction of Medicare Part D among the highly risk averse? I suggest two possible explanations: first, individuals may hold multiple insurance plans which cover drugs, and they may be a mix of public and private coverage. It seems plausible that individuals who hold multiple plans are disproportionately highly risk averse, perhaps keeping their private coverage in

order to supplement Part D where the latter provides little protection (i.e., in the coverage gap, or when purchasing specialty drugs).

The second possible explanation is that the highly risk averse may have simply had less drug coverage before Medicare Part D. This may be counterintuitive, but could arise from the institutional constraints on drug coverage before the reform. In particular, the primary source of public drug coverage before Part D was through Medicare Advantage HMOs which would have covered drugs but restricted provider networks. Highly risk averse individuals may have tended to be more averse to narrow provider networks than to uninsured drug risks. The introduction of Medicare Part D freed them from having to make this choice.

Recall that when an individual has both private and public drug insurance, the individual is classified as having private insurance. This accords with the intuitive idea of crowd-out: it measures to what extent individuals exchange private coverage for public coverage. If individuals keep their private coverage there is no crowding-out. The evidence presented thus far is consistent with highly risk averse individuals being more likely to keep both their preexisting private coverage and their new public coverage when taking up a Part D plan, or to supplement their public insurance with new private coverage. Private coverage tends to be more generous than Part D plans, and thus may be more attractive to the highly risk averse, even if keeping both plans involves higher premium payments.⁶⁹

Such a mechanism would require that the more risk averse be more likely to hold multiple insurance plans simultaneously. The data can only offer suggestive evidence on this point because it is not straightforward to infer which of an individual's health insurance plans is the primary insurer for prescription drugs, whether some classes of drugs are covered by different insurers (i.e., generics versus brand name drugs), and what the exact terms of the insurance are (whether a plan offers supplemental coverage in the Part D coverage gap, for example). Nevertheless, the limited evidence below accords with this interpretation of the

⁶⁹It is worth noting that many employer plans do not allow those they insure to have Medicare Part D plans in conjunction with the private coverage; however, some employer plans do allow this, as does Medicare Part D. This fact may account for some of the small effect sizes below.

main results.

To begin with, before Medicare Part D, more risk averse individuals were more likely to hold multiple health insurance plans.⁷⁰ This can be seen in columns 2 and 5 of table 15 (for the two respective measures of risk aversion). These regressions include only observations from before 2006. The results indicate that the highly risk averse had more insurance plans, on average, before Part D. Using the principal component measure of risk aversion, for example, shows that every standard deviation of risk aversion was associated with 0.04 more plans. The mean number of plans in the pre-Part D period was 0.71, thus every standard deviation of risk aversion was associated with about 6% more plans.


This greater number of plans in the pre-reform period for the more risk averse suggests they may indeed have a preference for a greater number of overall plans. Estimating the effect of Part D on their total number of health insurance plans with a specification like equations 17-19 gives results consistent with this (albeit statistically insignificant). These results are in columns 3 and 6 of table 15. For example, using the extensive risk aversion measure shows that for the less risk averse the overall number of insurance plans increases by about 0.1 due to Part D, however for the highly risk averse this increase is 70% greater, at 0.17. The results using the principal component measure of risk aversion are very similar. It thus seems as though the more risk averse may be supplementing their public coverage more than the less risk averse. Considering the results on out-of-pocket spending, and the sharpness of the reduction in spending around the region of the Part D coverage gap, it seems possible that for the very risk averse acquiring some protection from costs in that range may account for some of this effect.

Compounding this effect, it also seems to be the case that the highly risk averse had slightly less drug coverage than the less risk averse in the pre-2006 period. This can be seen in the upper panel of figure 16, which breaks down the type of insurance individuals aged 65-75 had in the years 1998-2004, separately for the low and high risk aversion groups (on

⁷⁰I top code the number of an individual's plans at 3, however results are not sensitive to this choice.

the left-hand and right-hand side panels, respectively). Unsurprisingly, a vast majority had Medicare coverage: about two-thirds of both risk aversion groups had traditional Medicare, and about another fifth were covered by a Medicare Advantage plan. Roughly 8% had Medicaid; another 3% had some form of private drug insurance, and less than 1% had no insurance whatsoever.

Figure 16: Insurance Composition for Ages 65-75 Before and After 2006, by Risk Aversion



Notes: This figure shows the shares of the sample having different kinds of health insurance in the years 1998-2004 (upper panel), and in the years 2006-2010 (lower panel), divided into the low and high risk aversion groups based on the extensive risk aversion measure (low risk aversion on the left, high risk aversion on the right; for precise definition see text). The shares set apart (No Drug Insurance and Traditional Medicare) are not insured against prescription drug expenses.

Recall that in the pre-reform period it is both the uninsured and those with traditional Medicare but no other source of drug coverage (such as an employer or retiree plan or a supplemental Medigap plan) who would have had no coverage for drugs (the roughly 20% set

apart from the rest). This latter group, those with traditional Medicare but no alternative drug coverage, is about 2.5 percentage points larger among the highly risk averse. This difference can be explained by their slightly lower participation in Medicare Advantage and in Medicaid. As a result, the highly risk averse did indeed have slightly lower rates of prescription drug insurance coverage in the pre-2006 period, helping to explain their lower rate of crowd-out. In a regression with the usual controls the high risk aversion group over age 65 has about 1.7 percentage points less drug coverage in the years 1998-2004 than the low risk aversion group, a difference which is significant at the 10% level.⁷¹

Accordingly, in the post-2006 period there was a slightly larger decline in the share of those with no prescription drug insurance among the high risk aversion group. For that group, the share uninsured declined by about 15.5 percentage points, compared with a decline of 14 percentage points in the low risk aversion group. While this does not directly measure the change in private coverage among the two groups, it illustrates that the increased public coverage reduced lack of coverage more for the more risk averse.

Furthermore, both groups experienced declines in Medicare Advantage coverage.⁷² The decline in Medicare Advantage in unconditional means is slightly larger for the less risk averse, as seen in 16. However, in regressions with the control group and other controls this sign is once again consistent, with a larger decline in public HMOs among the more risk averse, a difference which is marginally significant in some specifications.⁷³ All this is consistent with the more risk averse gaining more net coverage from the public coverage expansion at least in part because of a preference for broader provider networks. This preference led them to choose traditional Medicare (which did not cover drugs) in the pre-Part D period, rather

⁷¹Full results not displayed.

⁷²Membership in these plans has been on an upward trend since 2006, and this is evident in later years in the sample (i.e., 2008-2010) for both the high and low risk aversion groups

⁷³Results not displayed.

than limited-network plans which included drug coverage.⁷⁴

In sum, this section shows evidence that crowd-out of private prescription drug insurance by public insurance was decreasing with risk aversion. This finding is robust across specifications, definitions of crowd-out, and measures of risk aversion. While the estimated effect size is modest, it is probably a lower bound of the true magnitude given the inherent difficulty in measuring risk aversion. This finding can be explained by two observations: that the highly risk averse have a greater propensity to hold multiple insurance plans, and are correspondingly less likely to drop their private coverage when they take up public coverage; and that they had less drug coverage to begin with, perhaps because their risk aversion led them to prefer broader health care provider networks under traditional Medicare plans at the expense of drug insurance which was available through HMOs.

6 Conclusions

Crowd-out of privately supplied goods and services by public ones is a crucial input in assessing both the fiscal cost of public provision and its welfare implications. In the latter application it is often assumed that individuals are homogeneous in their risk preferences. A growing literature on heterogeneity in risk preferences and its impact on insurance markets is calling that assumption into question. When risk preferences differ across individuals average crowd-out rates are no longer sufficient for the welfare analysis of social insurance.

This paper estimated heterogeneity in crowd-out of private prescription drug insurance along the dimension of risk aversion. Using the quasi-experiment of the expansion of public drug insurance by addition of the Medicare Part D drug insurance program to traditional Medicare I quantify how much drug insurance coverage increased as a causal effect of the reform; and how much that increase varied by individuals' risk aversion.

⁷⁴Both risk aversion groups also experienced a large decline in traditional Medicare supplemented by some other drug coverage. While this is presumably mostly employer coverage, and thus mostly crowd-out, some of this coverage may have been from public sources such as through Veteran's Affairs. The regressions above control for veteran status.

I apply my methodology to HRS survey data. This data allows me to measure risk aversion in a number of ways. Using two measures of risk aversion based on these data I find that crowd-out declines with risk aversion. Consistently across the two measures and a number of different empirical specifications I find that insurance coverage increases more for the more risk averse. Every standard deviation of risk aversion is associated with a crowd-out rate roughly 5 percentage points lower, over a base crowd-out rate of 50%-60%. This lower crowd-out for the more risk averse also leads to greater declines in out-of-pocket drug spending for them, above and beyond what is experienced by the less risk averse.

In explaining these results two main mechanisms find suggestive support: that the more risk averse are more likely to hold more than one insurance plan, and thus are also more likely to keep holding a private drug plan even after taking up public coverage; and that the more risk averse had slightly lower rates of drug coverage before Part D, perhaps because they preferred the broader coverage of providers in traditional Medicare (which did not cover drugs) over the drug coverage coupled with limited networks in Medicare HMOs.

The pattern of results suggests that some welfare analysis using sufficient statistics formulas which do not account for heterogeneity in crowd-out by risk aversion may yield biased welfare estimates (this is not a universal problem; i.e., welfare estimates relying on estimation of demand and cost curves, as in Einav et al. [2010], are conceptually robust to this). In this case, the direction of the bias from ignoring heterogeneity would be to understate the welfare gains from the public provision of drug insurance: individuals gaining more net insurance from the program are those who most highly value insurance. Future work should quantify the welfare implications of this heterogeneity in crowd-out to assess how much of a bias its neglect might cause.

Part III

The Effect of Individual Health Insurance Mandates on Insurance Coverage

1 Introduction

Lack of health insurance is a significant barrier to access to health care.⁷⁵ In light of this, increasing health insurance coverage has long been a major policy goal. The Patient Protection and Affordable Care Act (PPACA) has as its primary aim such an increase in coverage, to which end it employs two major tools: subsidies to health insurance; and the individual mandate, a fine levied on uninsured individuals. While subsidies in general are a familiar policy tool, a mandate that individuals buy a product, like insurance, is nearly unprecedented. It is important to consider the effectiveness of the PPACA subsidies and individual mandate in isolation from one another, both to assess proposals to modify the law, and to learn about the pros and cons of such potentially substitutable policy tools for other contexts. Thus where previous literature has focused on subsidies to insurance, this paper seeks to identify the effect on insurance coverage of the PPACA's individual mandate.⁷⁶

Subsidies and mandates differ from each other in three important ways. The first is with respect to their fiscal costs: subsidies cost the government resources, while a mandate

⁷⁵Among the uninsured in 2014 27% reported going without needed care due to its cost, as opposed to 10% and 5% among those with public and private insurance, respectively (Majerol et al., 2015).

⁷⁶The literature on health insurance subsidies is very extensive. For example, see Ayyagari and Shane [2015], Baicker et al. [2014], Duggan et al. [2008], Engelhardt and Gruber [2011], Ericson and Kessler [2015], Garthwaite et al. [2014], Kaestner et al. [2014], Lichtenberg and Sun [2007], and Wettstein [2015].

costs nothing directly, and may even raise revenue. The second is with respect to who is induced to take up insurance coverage: there is limited evidence that the mandate is more effective at inducing younger and healthier individuals to purchase insurance than subsidies are (Hackmann et al., 2015). These advantages of the mandate notwithstanding, the third difference is on the political level, where the mandate has proven much less popular than the insurance subsidies in the PPACA, and may be more susceptible to change in the future due to political pressure. Thus it is a question of normative importance to what extent the PPACA (and other policies) should rely on subsidies versus mandates; and it is a question of practical importance what the positive implications of changes to the individual mandate in the PPACA might be.

To identify the effect of the mandate on coverage I exploit the quasi-experiment of the PPACA's introduction of the individual mandate in 2014 to estimate the effect of a fine for being uninsured. As a control group I use Massachusetts, which already had an individual mandate of its own due to the 2006 health care reform in that state. This control group provides a counterfactual trend in insurance coverage in the absence of the PPACA, to the extent that the Massachusetts trend mirrors the national trend. Overall, I find that the mandate substantially increased insurance coverage in the treatment group relative to the control group. Furthermore, I find suggestive evidence that this increased coverage contributed to amelioration of adverse selection in the individual health insurance market.

The data I use are from the American Community Survey (ACS), which is a sample of around 3.5 million households per year. These data are well suited for this analysis as they include variables on health insurance coverage, income, and geographic location at the state level. In particular, I use the 1-year estimates of the ACS for the years 2008-2014. I constrain the sample to individuals aged 26-64: those below age 26 are potentially covered by their parents' insurance as of 2010, while those over age 65 are entitled to Medicare.

To use the PPACA's introduction in order to estimate the effect of the mandate two main challenges must be overcome. The first is that the PPACA includes a bundle of many policy

changes beyond imposition of the individual mandate, and foremost among those things is the institution of health insurance subsidies.⁷⁷ The subsidies are given to all those earning between 138% and 400% of the Federal Poverty Line (FPL), on a sliding scale which decreases with income. At 400% of the FPL the subsidies are completely phased out. Therefore, to isolate the effect of the mandate I focus on households earning more than 400% of the FPL.

The use of Massachusetts residents to account for trends in coverage relies on the fact that Massachusetts introduced its own individual insurance mandate in 2006. While this mandate was not at the same level as the mandate under the PPACA, the differential increase in the penalty for lacking coverage provides variation for identification of the mandate's effect. The estimated effect of the mandate can be scaled by the differential penalty change between Massachusetts and the other states to get the effect of a \$1 penalty on insurance coverage. I find that the mandate increased coverage by 0.85 percentage points in the treatment states relative to Massachusetts; and that every \$1,000 of penalty increases insurance coverage by 0.73 percentage points, on average.

One major goal of the PPACA in general, and the individual mandate in particular, is to ameliorate adverse selection in the individual health insurance market. To be effective at this, the mandate must induce not only greater insurance coverage overall, but particularly induce younger and healthier individuals to purchase insurance. While the ACS has almost no information on the health of respondents, it does have information on their age. Analysis of heterogeneity in the treatment effect of the mandate by age shows that the increase in coverage was twice as large among younger individuals (below age 50) as it was among older individuals. Such a differential effect is consistent with the mandate being particularly effective at inducing the young to purchase health insurance, which can be considered a proxy for improving the overall risk pool in the individual market.

Similarly, additional evidence also points to the mandate reducing adverse selection in the

⁷⁷A number of additional provisions of the PPACA went into effect in 2014. For more detail on how they may impact the interpretation of the results refer to section 2 and the discussion in section 4.

individual market. Isolating the type of insurance individuals were induced to acquire by the mandate I show that the mandate increased coverage of directly purchased insurance much more than Employer-Sponsored Insurance (ESI). The latter already has less of an adverse selection problem than the former (see, for example, Hendren, 2013b). Therefore this, too, accords with the mandate reducing selection in the individual market, where there is more scope for such improvement.

There is a small but growing literature on the effects of insurance mandates. Chandra et al. [2011] examine the impact of the Massachusetts mandate on the mix of enrollees in non-group insurance plans. Their results are suggestive of the potential effects of the PPACA mandate, however they are mostly descriptive in the absence of a control group. Furthermore, as they note, the level of insurance subsidies may have non-trivial interactions with the effect of a mandate making direct inference from the Massachusetts experience to the federal law difficult. Closer in spirit to the current paper, Hackmann et al. [2015] use the introduction of the mandate in Massachusetts as a quasi-experiment with the rest of the US as a control group. Their focus is the effect of the mandate on adverse selection in the insurance market and they find that selection was ameliorated by the Massachusetts mandate.

Finally, Kolstad and Kowalski [forthcoming] use labor market responses to the Massachusetts mandate to compare the efficiency of mandates to that of a tax-based reform, extending the model in Summers [1989] to more closely mirror the actual policy. One of the empirical challenges they face in using the Massachusetts reform is that all the elements of that reform went into effect simultaneously: an individual mandate, insurance subsidies, and an employer mandate, among other provisions. The PPACA provides a clearer natural experiment to interpret because some of its similar provisions were phased in over time; in particular the employer mandate was not started concurrently with the individual mandate. This allows for a much more straightforward empirical model in estimating the effect of the individual mandate.

The papers above rely on the relatively small reform in Massachusetts alone. This makes generalization to the rest of the country difficult, particularly in the presence of much larger insurance subsidies under the PPACA than in the 2006 Massachusetts reform. This paper brings together the literature on individual insurance mandates with the new literature building on the quasi-experiments created by the PPACA to answer the question of how important the individual mandate is in achieving the aims of the PPACA: expanding health insurance coverage, and a resulting improvement in the risk pool participating in the individual insurance market.

The rest of the paper is organized as follows: section 2 provides institutional details regarding the PPACA's individual mandate; section 3 describes the data and the identification strategy in greater detail; section 4 describes the results; and section 5 concludes.

2 The PPACA's Individual Mandate and Other Institutional Details

The PPACA introduced numerous changes to the US health care system. Many of these changes have increasing health insurance coverage as their primary goal. These include regulation of the supply of insurance, establishment of state and federal exchanges to facilitate purchase of individual insurance, subsidies for plans purchased on those exchanges, and employer and individual mandates to offer and to acquire insurance, respectively. In this section I briefly describe the most important of these elements as they pertain to the choices I make in defining the sample and in my identification strategy, as well as regarding the interpretation of my results.

The focus in this paper is on the individual health insurance mandate. This provision of the PPACA was passed into law with the rest of the Act in 2010, but only went into effect in 2014. It generally requires all US citizens and legal residents to have minimal essential coverage, or pay a tax. The size of the tax depends on income and family size: in 2014 it was \$95 per adult and \$47.50 per child (up to \$285 per family) or 1% of family income,

whichever was greater. The tax increased in 2015 to \$325 per adult and \$162.50 per child (up to \$975 per family) or 2% of family income, whichever was greater; and will increase a final time in 2016 to \$695 per adult and \$347.50 per child (up to \$2,085 per family) or 2.5% of family income, whichever is greater. After 2016 the tax will increase by the cost of living. The tax is pro-rated by number of months in the year without coverage, although a gap in coverage of less than three months is exempt.

Some people are exempt from the mandate: those with religious objections to health care, those who are incarcerated, undocumented migrants, and members of Native American tribes. In addition, families below the tax filing threshold (\$10,150 for an individual or \$20,300 for a family, in 2014) are exempt; as are those who would have to pay more than 8% of their income for insurance after accounting for employer contributions and subsidies. Nevertheless, the mandate is otherwise very broad-based and for the vast majority the only way to avoid it is to have sufficient insurance coverage for at least nine months of the year.

Acquiring such coverage can be done in a variety of ways. For those eligible, Medicare and Medicaid provide the minimal necessary coverage. Eligibility for the former requires individuals only to be over age 65; the latter has complicated eligibility criteria which vary by state, generally including low income but in some states also covering only those who meet non-income criteria as well, such as pregnancy or disability. Other public sources of insurance which qualify are those for veterans, such as Tricare. Coverage for individuals under age 26 can be provided through their parents' plan, if they have one. However, for most individuals between age 26 and 64 the two predominant sources of insurance are ESI or plans purchased on the individual market.

The PPACA includes many other provisions besides the individual mandate which might impact insurance coverage. These other provisions may confound the effect of the mandate itself; isolating the effect of the mandate from them relies on the details of these additional changes. Foremost among these other changes are the insurance subsidies given to individuals who purchase insurance on the PPACA exchanges. Others include expansion of Medicaid;

guaranteed issue (the requirement that insurers provide insurance regardless of preexisting conditions and with severely regulated price discrimination); elimination of annual coverage limits; and requiring all creditable plans to cover essential health benefits. Furthermore, the mere construction of the state and federal health insurance exchanges may have made acquisition of insurance easier. It is worth noting that requirements for employers to offer group coverage were scheduled to begin in 2014 but were postponed until 2015 or later.

Subsidies for plans bought on the exchanges are distributed by, among other things, household income. The subsidies decrease as income increases, and go to zero at 400% of the FPL; this was about \$47,000 for an individual in 2014, or \$95,000 for a family of four. I isolate the effect of the individual mandate from that of the subsidies by restricting attention to individuals in households above 400% of the FPL, who are thus ineligible for subsidies. This strategy has an important limitation: if subsidies given to those below 400% of the FPL disproportionately induce the younger and healthier individuals in that group to acquire insurance on the exchanges the risk pool insurers face on the exchanges would improve. This, in turn, could lower prices for everyone using the exchanges, including those who are not themselves eligible for subsidies. Attributing the effect of the PPACA on coverage among individuals ineligible for subsidies to the individual mandate assumes that this spillover effect of the subsidies is negligible. Section 4 will provide some suggestive evidence that such a spillover is not the main driver of the results.

The other PPACA provision directly increasing coverage is the expansion of Medicaid to everyone under 138% of the FPL (in states that chose to accept this expansion). This element of the PPACA is also less relevant when interpreting the effect of the individual mandate once attention is focused on households making more than 400% of the FPL. For such relatively wealthy individuals Medicaid is not a major source of coverage regardless; virtually none of them are eligible for Medicaid either before or after its expansion in 2014.

The elimination of annual coverage limits, the required provision of essential health benefits, and guaranteed issue all have theoretically ambiguous effects on coverage. The first two

raise the value of available insurance plans, but presumably also raise their cost and their price. The latter expands the set of potential buyers of insurance to include those with pre-existing conditions; however it, too, is expected to raise the price of insurance for everyone else, potentially depressing demand among the healthy. In any case, it is hard to envision a mandate for consumers to buy insurance without accompanying it with a requirement that insurers offer insurance to everyone, making guaranteed issue an almost inseparable part of the individual mandate in the first place. In interpreting my results I assume that the effects of these provisions on coverage are small, on net. Section 4 also provides some evidence that guaranteed issue, at least, does not play a major role in driving the results.

The final large change initiated in 2014 under the PPACA is the activation of online exchanges where insurance plans can be bought and sold. This in itself may have made insurance easier to acquire, by reducing search costs. Two types of exchanges were established by the PPACA: the first, the American Health Benefits Exchanges, are online markets where individuals and families can purchase health insurance. The second type, the Small Business Health Options Program (SHOP) Exchanges are intended for small businesses with up to 100 employees. Some suggestive evidence that the effect of the existence of exchanges on coverage is small is offered in the discussion in section 4. I therefore assume that any confounding of the existence of exchanges and the individual mandate is negligible in attributing the effects I estimate to the mandate; however more careful analysis of this issue should be done in the future.

As explained in the next section, I use Massachusetts as a control group for the rest of the US. This relies on the fact that while the PPACA's individual mandate went into effect in 2014, Massachusetts underwent its own health care reform in 2006 – a reform which included an individual mandate similar to that in the PPACA. The penalty for individuals under the Massachusetts mandate for being uninsured was a fine equal to half the cost of the lowest annual premium which they could have bought. For individuals over age 27, with incomes above 400% of the FPL in 2013 this would have amounted to \$1,272 a year for an individual

(Massachusetts Department of Revenue, 2013).

As stated above, for individuals earning over 400% of the FPL under the PPACA the penalty in 2014 was 1% of income. The average income in my sample in 2014 is \$86,392 in Massachusetts and \$75,264 in the other states.⁷⁸ Thus the average penalty in Massachusetts actually declined from \$1,272 to \$864 while it increased in the other states from \$0 to \$753. Therefore the average change in penalty for the other states relative to the change in Massachusetts is \$1,160.⁷⁹ In discussing the effect of the mandate penalty in the results section I will show how the effects scale by this relative change in penalty, to get the effect on coverage of a \$1,000 fine.

3 Data and Identification

The data I use are from the one-year estimates of the American Community Survey. These estimates are based on a 2.5% sample of US households. As discussed in section 2, I exclude households below 400% of the FPL in order to focus on those who do not become newly eligible for subsidies under the PPACA; this also effectively excludes individuals who could newly qualify for Medicaid under its expansion. I also exclude those below age 26 or over age 65, as they could gain coverage through their parents or through Medicare, respectively.

I employ a cross-sectional differences-in-differences design with individuals living in Massachusetts as a control group. This provides the counterfactual change in insurance coverage when there is no new individual mandate introduced in 2014. Individuals in the rest of the United States make up the treatment group, with the treatment period being 2014. The state of Vermont had its own health care reform in 2006 which might bias results. Therefore

⁷⁸For the respective numbers in the years 2008-2013 see table 16.

⁷⁹This is an approximation of the relative change in penalty at the mean, as the penalty differs by household income and number of persons in the household, while I calculate these changes in penalty based on individual income.

Vermont residents are excluded from the following analysis.⁸⁰


The ACS is a repeated cross-section, so analyzing the effect of the mandate on particular individuals or households over time is not possible with this data. I include rich demographic controls in order to minimize potential omitted variable bias in such a design. These controls include a full set of gender, age, and detailed race, marital status, education, and employment dummies; as well as a sixth order polynomial in income.

Identification of the effect of the individual mandate on coverage relies on the fact that Massachusetts already had an individual mandate before the PPACA. The Massachusetts mandate has been in place since 2006; however questions in the ACS regarding health insurance coverage were only asked from 2008 onward. The data I use therefore span from 2008 until 2014.

The credibility of Massachusetts residents as a control group can be visually assessed by noting parallel trends in coverage rates between Massachusetts residents and those of the treatment states in the period before the PPACA individual mandate takes effect; i.e., in the years 2008-2013. Such parallel movement is apparent in figure 17; this figure plots the rate of insurance coverage each year, in Massachusetts (in blue squares) and the other states (red circles) separately. Such parallel movements suggest that had the individual mandate never gone into effect for the treatment group, the change in coverage for the treatment group between 2013 and 2014 should have been the same as in the control group.

⁸⁰None of the results is sensitive to including Vermont in the treatment states.

Figure 17: Health Insurance Coverage Over Time



Notes: This figure shows mean rates of health insurance coverage in the sample by year, separately for residents of Massachusetts and the rest of the United States, excluding Vermont. The former are in blue squares, the latter in red circles. The sample is individuals aged 26-64, in the years 2008 until 2014, who belong to families earning more than 400% of the Federal Poverty Line in the year they were sampled. The dashed gray line differentiates between years before the beginning of the Patient Protection and Affordable care Act's individual health insurance mandate, on the right, and 2014, when the mandate went into effect, on the left.

Following the sample restrictions above I am left with a sample of about 4.8 million individuals in 1.25 million households, of which Massachusetts residents make up 141,000 individuals in 37,500 households. Descriptive statistics for years 2008-2013 (before implementation of the individual mandate in the treatment states) for the sample, separated into Massachusetts and the other states, are in table 16. The treatment and control groups are quite similar, although Massachusetts residents are somewhat less likely to be members of a racial minority, and are somewhat more likely to have attended college, and have higher incomes. On all these measures the differences between the two groups are within one standard deviation of each other, except the share of whites.⁸¹ In terms of insurance, in the pre-treatment period Massachusetts residents have about 3.3 percentage points more insurance coverage, with higher levels of ESI, and lower levels of directly purchased insurance.

The regressions I estimate take the following form:

⁸¹Part of the reason why Massachusetts' and the other states' residents are so similar on observables is that the sample is already selected, particularly on income – everyone belongs to a household making over 400% of the FPL.

Table 16: Descriptive Statistics

	Massachusetts		Treatment States	
	Mean	Std. Deviation	Mean	Std. Deviation
Share Women	0.51	0.5	0.5	0.5
Age	46.7	10.6	47.4	10.5
Share White	0.9	0.3	0.84	0.36
Share College	0.81	0.39	0.78	0.41
Personal Income	78,184	80,881	70,398	71,944
Share with Any Insurance	0.98	0.13	0.95	0.22
Share with Employer Insurance	0.9	0.3	0.85	0.36
Share with Direct Insurance	0.09	0.29	0.11	0.32
No. of Individuals	141,457	-	4,697,505	-

Notes: This table presents descriptive statistics for the two experimental groups separately, Massachusetts and the rest of the United States, excluding Vermont. The years of the sample for the purposes of these statistics are restricted to the pre-treatment period, 2008-2013, except for the number of individuals which spans the entire sample period, 2008-2014. The sample is individuals aged 26-64 who belong to families earning more than 400% of the Federal Poverty Line in the year they were sampled. The first two columns display the means of the variables and their standard deviations for Massachusetts, respectively. The third and fourth columns show the means and standard deviations of the rest of the sample.

$$y_{i,t,s} = \beta_1 * Post_{i,t} * Treat_{i,s} + \beta_2 * Post_{i,t} + \beta_3 * Treat_{i,s} + \alpha_s + \gamma_t + \sum_{j=1}^k \theta_j X_{j,i,t,s} + \varepsilon_{i,t,s}, \quad (20)$$

where $Post_{i,t}$ is 1 for observations in 2014; $Treat_{i,s}$ is 1 for residents of states other than Massachusetts; α_s are state dummies (excluding Massachusetts, which is the control group, and using Wyoming as a reference group); and γ_t (excluding 2014, which is the only treatment period, and using 2008 as a reference group) are year dummies. X contains the other controls detailed above. All observations are weighted by their household sampling weights and standard errors are clustered by household.⁸²

4 Results

This section describes the estimated effect of the PPACA individual mandate on insurance coverage. The analysis reveals that the mandate causes an increase in coverage, which is mostly concentrated in insurance bought on the individual market, as opposed to group coverage acquired through an employer. Furthermore, the increase in coverage is greater among younger individuals, consistent with the mandate reducing adverse selection by improving the risk pool in the individual market. A brief discussion section analyzes whether some alternative interpretations of the mechanism behind the estimated effects are plausible.

To begin, figure 17 shows the effect of the introduction of the individual mandate on health insurance coverage. As described in section 3, this figure plots insurance coverage rates by year and by treatment state (Massachusetts or otherwise). In addition to the parallel trends in coverage between the treatment and control groups, the treatment effect can be observed in the sharply higher increase in coverage in 2014 in the treatment group relative to the control group. This visually demonstrates the effect of the mandate in increasing coverage: the introduction of the individual mandate in the treatment states increases the

⁸²Results are robust to clustering at the state level or to using unclustered robust standard errors.

coverage rate by about 1 percentage point, with no substantial change in the control group.

This estimate is robust to adding controls and is statistically significant. In order to see this, consider column 1 of table 17. This column has as its dependent variable health insurance coverage from any source, with the specification detailed in equation 20. The estimated effect of the mandate on coverage is an increase of 0.85 percentage points, representing a decline of about 17% in the uninsured rate among residents of the treatment states.⁸³ This effect is highly statistically significant.⁸⁴

This change in the uninsured rate is also large in terms of economic significance. One way to see this is to consider the overall change in the uninsurance rate in the US between 2013 and 2014. This change encompasses the entire effect of the PPACA on coverage, not just that of the individual mandate (e.g., it includes the effect of the subsidies, the Medicaid expansion, etc.). Among non-elderly adults, the number of uninsured declined by 7.8 million between 2013 and 2014, a change of 4.2 percentage points (Majerol et al., 2015).⁸⁵ A back-of-the-envelope calculation thus implies that the 0.85 percentage point decline in the uninsurance rate due to the individual mandate alone makes up more than 20% of the total decline between 2013 and 2014.⁸⁶

A second approach to clarifying the magnitude of this effect is to scale it by the dollar amount of the penalty. As stated in section 2, the penalty in Massachusetts for lack of insurance before 2014 was different from that imposed by the PPACA. The approximate

⁸³The uninsured rate in the treatment group in 2008-2013 was 5%.

⁸⁴Including household fixed effects has little impact, either on the point estimate or on the statistical significance of this estimate.

⁸⁵This comparison does not have a control group, however, and is therefore simply the change over time of coverage, with no counterfactual estimate for what would have happened to coverage in the absence of the PPACA.

⁸⁶This assumes that the effect of the mandate on households making less than 400% of the FPL is the same as its effect on households in the sample, who make more than 400% of the FPL. Because the mandate penalty is proportional to income, this is likely an upper bound for the share of the overall reduction in the uninsured rate due to the mandate alone.

Table 17: Health Insurance Coverage

	(1)	(2)	(3)	(4)
	Any Insurance	Employer Insurance	Directly Purchased Insurance	States with Guaranteed Issue
Treatment Effect	0.00848*** (0.00168)	0.00368 (0.0034)	0.00981*** (0.00295)	0.00904*** (0.00208)
Year Fixed Effects	Yes	Yes	Yes	Yes
Age Fixed Effects	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes
Observations	4,838,962	4,838,962	4,838,962	692,101
Clusters	1,247,806	1,247,806	1,247,806	187,922

Notes: This table presents the results of estimating equation (1) with the stated dependent variable at the head of columns 1, 2, and 3; and with any insurance as the dependent variable in column 4. Column 4 restricts the sample to be residents of states that had guaranteed issue throughout the sample period: Maine, New York, Vermont, and Washington as treatment states, and Massachusetts as the control state. The sample is individuals aged 26-64, in the years 2008 until 2014, who belong to families earning more than 400% of the Federal Poverty Line in the year they were sampled. All the regressions include dummy variables for being a resident of a treatment state, and for being observed in 2014, as well as the interaction of those two variables. The coefficient on the latter is the displayed estimate of the treatment effect of the individual mandate on the outcome variables, in the first row. Its associated standard error is in parentheses in the second row. Furthermore, all the regressions include year, age, and state fixed effects; as well as a full set of dummies for gender, race, marital status, education, and employment status, and a sixth order polynomial in individual income. Standard errors are clustered by household, and observations are weighted by their sampling weights. The final two rows indicate the number of individuals included in each regression and the number of households, respectively. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

dollar difference in penalty in the treatment states relative to Massachusetts was on average \$1,160. As a result, my estimate implies that a \$1,160 increase in the penalty for lack of insurance causes a 0.85 percentage point increase in insurance coverage, or an increase of 0.73 percentage points for every \$1000 of penalty.

Accounting for the Effect of Guaranteed Issue

As described in section 2, one of the PPACA provisions that went into effect nationally in 2014 (and existed in Massachusetts since 2006) was guaranteed issue, the requirement to sell insurance to anyone with only very limited price discrimination between observably different consumers. This provision can be considered an inseparable part of the individual mandate: it would be perverse to require individuals to buy insurance while allowing insurers not to sell insurance to potential customers. However, it is important to see whether the increased coverage in treatment states in 2014 is due to the mandate itself, or due to guaranteed issue.

To do this it is helpful to note that there were a number of states that had guaranteed issue in the years 2008-2013, aside from Massachusetts itself. These states were Maine, New York, Vermont, and Washington. Column 4 of table 17 shows results for the sample limited to residents of these states that had guaranteed issue for the entire sample period.⁸⁷ The estimate of the causal effect of the mandate alone, with no change in the status of guaranteed issue, is very similar to that from the full sample, at 0.9 percentage points. Per a \$1,000 increase in the mandate penalty this translates into an increase in coverage of 0.75 percentage points, almost exactly like the 0.73 percentage points using the whole sample.⁸⁸

The estimated percentage change in the uninsurance rate is also very similar for the states with guaranteed issue to what it was for the entire sample. The uninsurance rate for

⁸⁷This regression includes Vermont, however results are virtually identical if Vermont is excluded here, as well.

⁸⁸Average income in my sample of residents of the treatment states with guaranteed issue in 2014 was \$80,200, so a penalty of 1% of income was on average \$802, compared to a decline in the mandate penalty in Massachusetts of, on average in my sample, \$408, as described in section 2. Thus the average individual penalty increased by \$1,210 in the treatment states with guaranteed issue relative to Massachusetts.

the treatment states with guaranteed issue in the years 2008-2013 was 4.7 percentage points in my sample, and this number declined by 19% due to the mandate, compared with 17% for the whole sample. These comparisons all serve to provide evidence that the increased coverage in 2014 in treatment states relative to Massachusetts does not stem from newly guaranteed issue, but rather from the mandate itself.

Source of New Insurance

The individual mandate incentivizes acquisition of insurance by any means. One of the main goals of this element of the PPACA, however, is to improve the risk profile of consumers of insurance on the individual market, reducing adverse selection and lowering costs for insurers and thus premiums for consumers. If the mandate is effective at this, it would make plans purchased directly on the individual market disproportionately more attractive than those acquired on the group market (such plans are not significantly affected by the PPACA in 2014): while demand for insurance is increased on both markets, the cost of supplying insurance should be mostly impacted on the individual market.

The ACS asks respondents whether they are covered by an employer (current or former) or union plan, their own or a family member's. I define these individuals as having employer insurance. Furthermore, the ACS also asks whether individuals are covered by a plan purchased directly, for themselves or by a family member. I define these individuals as having direct insurance.⁸⁹

Figure 18 displays graphically the change in coverage in these two types of insurance over the sample period, separately for the control and treatment groups; the left-hand panel plots ESI rates while the right-hand panel does the same for direct insurance. In both cases the parallel trends between the treatment and control groups hold in the pre-treatment period. However, while there is no apparent break in these trends for ESI in 2014, in direct insurance

⁸⁹Note that these two categories, of ESI and directly purchased insurance, are neither exhaustive nor exclusive: there are potentially other sources of insurance (e.g., public insurance) and individuals may be covered both by an employer and by a directly purchased plan. What follows is therefore not a decomposition of the increase in insurance into these two categories.

a general downward trend in both Massachusetts and in the other states is reversed in 2014 only in the treatment states.⁹⁰ This suggests that the effect of the mandate was to increase insurance predominantly by increasing direct coverage, rather than ESI.

Figure 18: Source of Health Insurance Coverage



Notes: This figure shows mean rates of health insurance coverage by source of coverage in the sample by year, separately for residents of Massachusetts and the rest of the United States, excluding Vermont. The former are in blue squares, the latter in red circles. The sample is individuals aged 26-64, in the years 2008 until 2014, who belong to families earning more than 400% of the Federal Poverty Line in the year they were sampled. The dashed gray line differentiates between years before the beginning of the Patient Protection and Affordable care Act's individual health insurance mandate, on the right, and 2014, when the mandate went into effect, on the left. The left-hand panel shows health insurance acquired through an employer or a union of the individual or one of their family members. The right-hand panel shows health insurance acquired directly by an individual or a family member.

Columns 2 and 3 of table 17 show regression estimates of the effect of the individual mandate on these two categories of insurance coverage, using the specification in equation 20. These results accord with the graphical evidence above: column 2 shows a relatively small (and statistically insignificant) increase of 0.37 percentage points in employer coverage due to the mandate. In contrast, column 3 shows a much larger (and highly statistically significant) increase in direct insurance, of nearly 1 percentage point. This pattern is consistent with the individual mandate reducing adverse selection in the individual market, as intended.

However, it is worth noting that such a pattern of a larger treatment effect on direct insurance than on employer insurance is also consistent with other mechanisms, albeit less direct ones. For example, it is possible that the PPACA subsidies for households making less than 400% of the FPL create spillover effects on my sample of individuals in households

⁹⁰In fact, it seems that there might have been a decline of direct insurance in Massachusetts relative to the trend in 2014. This could be consistent with the penalty Massachusetts residents faced in 2014 declining, on average, as described in section 2.

making more than that cutoff. The subsidies may induce healthier individuals in poorer households to enter the individual market, lowering insurance costs and premiums for all individual market consumers, including those who are not themselves eligible for subsidies. A second alternative interpretation is that the increase in direct insurance above beyond the increase in employer insurance may be due to the newly established online insurance exchanges which serve individuals purchasing insurance directly, but not those selecting group plans.

Heterogeneity in the Treatment Effect

While the source of new insurance is consistent with reduced adverse selection, a more direct test of whether the individual mandate encourages healthier individuals to buy insurance is to test for a larger treatment effect on healthier individuals relative to sicker ones. This would also help in distinguishing between the various possible mechanisms behind the overall increase in coverage. The ACS does not contain information on health status; however, it does include the age of respondents. I use age as a proxy for health, assuming younger individuals have lower health care costs for insurers to cover.

The simplest approach to testing for heterogeneity in the treatment effect by age is to allow for a linear interaction of age and treatment. This adds an interaction of age and $Post_{i,t} * Treat_{i,s}$ to equation 20 (as well as an interaction of age with $Post_{i,t}$ and $Treat_{i,s}$ separately). The results of this specification can be seen in column 1 of table 18. The estimate of the interaction of age and treatment is negative (and highly statistically significant) at -0.00035, implying a larger increase in coverage due to the individual mandate among younger individuals than older. The results of this specification state that for 26 year-olds, the youngest group in the sample, the individual mandate is associated with an increase of 1.55 percentage points in coverage. In contrast, the implied increase in coverage for the eldest

Table 18: Heterogeneity by Age

	(1)	(2)	(3)
	Linear Age	Below Age 49	Over Age 50
Treatment Effect	0.0246*** (0.00259)	0.01018*** (0.00256)	0.00576*** (0.00184)
Treat*Post*Age	-0.00035*** (0.000004)	-	-
Year Fixed Effects	Yes	Yes	Yes
Age Fixed Effects	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes
Observations	4,838,962	2,536,190	2,302,772
Clusters	1,247,806	997,352	968,216

Notes: This table presents the results of estimating equation (1) with any insurance as the dependent variable. The sample is individuals aged 26-64, in the years 2008 until 2014, who belong to families earning more than 400% of the Federal Poverty Line in the year they were sampled. All the regressions include dummy variables for being a resident of a treatment state, and for being observed in 2014, as well as the interaction of those two variables. The coefficient on the latter is the displayed estimate of the treatment effect of the individual mandate on insurance coverage, in the first row. Its associated standard error is in parentheses in the second row. Furthermore, all the regressions include year, age, and state fixed effects; as well as a full set of dummies for gender, race, marital status, education, and employment status, and a sixth order polynomial in individual income. Column 1 includes an interaction of age with the dummy for treatment state, the dummy for year 2014, and a third order interaction with the product of those two variables. The coefficient on this third order interaction is given in the third row, and its standard error is given in parentheses in the fourth row. Columns 2 and 3 are restricted to observations aged 49 or less, and 50 or more, respectively. Standard errors are clustered by household, and observations are weighted by their sampling weights. The final two rows indicate the number of individuals included in each regression and the number of households, respectively. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

individuals in the sample, 64 year-olds, is only 0.2 percentage points.⁹¹

Other descriptive dimensions of heterogeneity are noteworthy. The first is by gender: columns 1 and 2 of 19 show the results of estimating 20 separately for women and for men, respectively. The effect of the mandate on women is an increase in coverage of 1 percentage point. For men the point estimate is smaller, an increase of only 0.6 percentage points. This difference is not statistically significant.

A third dimension of heterogeneity is by education. There are numerous ways to parse the data by education; columns 3 and 4 of 19 do so by whether the individual has no college education or any college education, respectively. The effect of the mandate on those with no college education seems descriptively larger than for those with any college education. The former experience an increase of about 1 percentage point in coverage, while the latter increase coverage only by three quarters of one percentage point. This difference is not statistically significant. Only 25% of the sample have no college education, leading to relative imprecision in their estimate. Nevertheless, the same qualitative results hold in other specifications of education, such as by including a linear interaction of years of education and treatment status; i.e., the slope of this interaction is negative but not significant at conventional levels.⁹²

Finally, table 20 shows the result of estimating equation 20 with all three dimensions of heterogeneity included. Thus this regression contains three third order interactions of age, gender, and a dummy for any college education, each with the product of *Post* and *Treat*,

⁹¹Qualitatively similar results are obtained from a quadratic interaction of age and treatment status. Similar results are also obtained using non-parametric approaches. E.g., the results for the above- and

below-median age (age 49) subsamples are in column 2 and 3 of table 18, respectively. For the young the individual mandate is estimated to increase coverage by about 1 percentage point; for the old the estimated increase in coverage is only just over half a percentage point. Dividing the sample by age quartiles (not displayed, full results available upon request) shows that the point estimate of the treatment effect is largest in the youngest quartile (those strictly below age 40), at 0.94 percentage points, and decreases in the second and third quartiles (ages 40-49 and 50-55, respectively) with treatment effects of 0.87 and 0.35 percentage points. Finally the estimated treatment effect increases somewhat to 0.59 percentage points for the oldest quartile (ages 56-64).

⁹²The results are not displayed, but are available upon request.

Table 19: Other Dimensions of Heterogeneity

	(1)	(2)	(3)	(4)
	Women	Men	No College	Any College
Treatment Effect	0.01007*** (0.00185)	0.00635*** (0.00233)	0.01018* (0.00567)	0.00747*** (0.00152)
Year Fixed Effects	Yes	Yes	Yes	Yes
Age Fixed Effects	Yes	Yes	Yes	Yes
State Fixed Effects	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes
Observations	2,433,461	2,405,501	1,049,948	3,789,014
Clusters	1,171,269	1,164,444	633,288	1,188,050

Notes: This table presents the results of estimating equation (1) with any insurance as the dependent variable. The sample is individuals aged 26-64, in the years 2008 until 2014, who belong to families earning more than 400% of the Federal Poverty Line in the year they were sampled. All the regressions include dummy variables for being a resident of a treatment state, and for being observed in 2014, as well as the interaction of those two variables. The coefficient on the latter is the displayed estimate of the treatment effect of the individual mandate on insurance coverage, in the first row. Its associated standard error is in parentheses in the second row. Furthermore, all the regressions include year, age, and state fixed effects; as well as a full set of dummies for gender, race, marital status, education, and employment status, and a sixth order polynomial in individual income. Columns 1 and 2 are restricted to women and to men, respectively. Columns 3 and 4 are restricted to individuals with no college or with any college education, respectively. Standard errors are clustered by household, and observations are weighted by their sampling weights. The final two rows indicate the number of individuals included in each regression and the number of households, respectively. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

as well as the main treatment effect of $Treat * Post$.⁹³ In this regression age is still highly negatively correlated with the magnitude of the effect of the mandate, even when controlling for these other dimensions of heterogeneity.

Discussion

There are two main provisions of the PPACA which went into effect simultaneously with the individual mandate which might yet account for the increased coverage estimated by the differences-in-differences estimator: the first is the insurance subsidies which are distributed to households below 400% of the FPL, and the second is establishment of online insurance exchanges. The subsidies could drive the estimated effect by inducing healthy individuals in relatively poor households excluded from the sample to buy insurance, causing a decline in insurance premiums for those in the sample leading them to take up coverage through such a general equilibrium spillover. The availability of insurance exchanges may reduce search costs for insurance, which could induce consumers to take up insurance. Both these mechanisms compete with the individual mandate in explaining the observed coverage increase.

Some of the patterns in the data observed above can help in arguing in favor of the individual mandate being the main driver of the increase in coverage. As detailed below, the interpretation of a spillover effect from the subsidies is less plausible in light of the heterogeneity in the effect by age. Attribution of the entire effect to the online exchanges is hard to reconcile with the effect being much larger for direct insurance than for ESI.

The spillover interpretation would require that insurance premiums for the young be disproportionately reduced by an improvement in the risk pool driven by increased coverage among the (relatively) poor.⁹⁴ This is unlikely: the cost to insurers of insuring the young was

⁹³For reasons of power, due to there only being one treatment year, the regression does not include the second order interactions of age, gender, and any college with *Post* and with *Treat* separately. Results are qualitatively similar when including these controls, however they lose from their statistical significance. Results are also similar, and highly significant, when including only years 2013 and 2014, in which case there should not be time trends by definition.

⁹⁴The spillover interpretation could also be justified by the young being more sensitive to changes in insurance premiums even if the actual change for them is the same, or smaller, than for the elderly. However,

Table 20: Heterogeneity by Age, Gender and College Education

	Any Insurance
Treatment Effect	0.03358*** (0.00313)
Treat*Age	-0.00037*** (0.00004)
Treat*Woman	-0.00385*** (0.00063)
Treat*College	-0.00739*** (0.00131)
Year Fixed Effects	Yes
Age Fixed Effects	Yes
State Fixed Effects	Yes
Other Controls	Yes
Observations	4,838,962
Clusters	1,247,806

Notes: This table presents the results of estimating equation (1) with any insurance as the dependent variable. The sample is individuals aged 26-64, in the years 2008 until 2014, who belong to families earning more than 400% of the Federal Poverty Line in the year they were sampled. Included are dummy variables for being a resident of a treatment state, and for being observed in 2014, as well as the interaction of those two variables. The coefficient on the latter is the displayed estimate of the treatment effect of the individual mandate on insurance coverage, in the first row. Further included are the third-order interactions of treatment state, year 2014 and age, gender, and a dummy for any college education. The coefficients on these three interactions are in rows 3, 5, and 7, respectively. Standard errors for each coefficient are in parentheses below the coefficient. The regression also includes year, age, and state fixed effects; as well as a full set of dummies for gender, race, marital status, education, and employment status, and a sixth order polynomial in individual income. Standard errors are clustered by household, and observations are weighted by their sampling weights. The final two rows indicate the number of individuals included and the number of households, respectively. (***) indicates significance at the 1% level; (**) indicates significance at the 5% level; (*) indicates significance at the 10% level.

already relatively low before the PPACA. Furthermore, age is observable to insurers so it is to be expected that the price of insurance for the young would have been low before 2014, and would weakly increase in 2014 once price discrimination on age and previous health status was limited. As a result, while reduced selection in the individual market could be expected to lower premiums for the elderly, the effect on the premiums of the young is ambiguous at best.

The heterogeneity in the estimate by age is therefore incongruent with the spillover interpretation in two ways: first, the spillover mechanism would imply that the level of the effect on the young should be relatively small, whereas in fact it seems to be relatively large, with the largest effects observed among the youngest individuals. Second, it would imply that the gradient of the effect by age should be positive, while it is in fact negative. It is therefore unlikely that the strong differences-in-differences estimate for the young could be a result of decreased adverse selection as a spillover effect from subsidized households. Rather, the strong impact on the young relative to the elderly is likely the direct result of the individual mandate, which itself reduces selection by age even within the set of individuals who are ineligible for the subsidies.

The other element of the PPACA which went into effect in 2014 and might be responsible for the change in coverage is the establishment of online insurance exchanges in the treatment states. Alongside the individual exchanges which were established in 2014, the PPACA also established health insurance exchanges for small businesses which were activated in 2014. Large employers almost all offered health benefits before the PPACA as well as in 2014.⁹⁵ Thus if an increase in employer sponsored insurance were to occur due to establishment of

it is not clear why this would be the case; all the regressions control very flexibly for income and other variables which might matter here such as gender or education.

⁹⁵In 2014 98% of firms with more than 200 employees offered health insurance to at least some workers, a rate not substantially different from prior years (this share has remained above 97% since 1999, Kaiser Family Foundation, 2014).

exchanges, it would be expected to take place among small employers in any case.⁹⁶

Both individuals and small employers gained access to online exchanges in 2014. The fact that the effect on directly purchased insurance is substantially larger than on employer sponsored insurance is suggestive that it is not merely the establishment of such exchanges which drives the entire increase in coverage, although it may account for some of the effect. If the entire effect were driven by availability of exchanges, there should have been similar coverage increases in ESI and in individual plans, whereas in fact there was a much larger effect on individual plans.

In sum, this section documents statistically and economically significant increases in coverage caused by the PPACA's individual mandate. Furthermore, it provides suggestive evidence that this increase in coverage ameliorates adverse selection in the individual insurance market: the effect on directly purchased insurance is much larger than on employer sponsored coverage; and the increased coverage is concentrated among younger individuals. In addition, it seems most in line with the evidence that the effect estimated by the differences-in-differences estimator is, in fact, due to the individual mandate itself rather than spillover effects of subsidies going to individuals outside the sample, or due to the mere establishment of online insurance exchanges.

5 Conclusions

The mandate in the PPACA that individuals purchase health insurance on the private market is an almost unprecedented policy. The only similar preexisting policy was a mandate that individuals have insurance in the state of Massachusetts. This paper uses Massachusetts as a control for the rest of the US around the time of the implementation of the PPACA's individual mandate to estimate the effect of such a policy on health insurance coverage. This is done in a differences-in-differences design, where the treatment group, residents of states

⁹⁶Among employers with less than 200 employees the rate of offering health benefits to some employees was only 54% in 2014; this was actually a decline from the level in 2013, of 57% (Kaiser Family Foundation, 2014).

other than Massachusetts, are treated in 2014 relative to a reference period of the years 2008-2013.

I focus on individuals who, by virtue of their relatively high household income, are ineligible for the insurance subsidies and the expansion of Medicaid the PPACA introduced simultaneously with the individual mandate. This leaves the individual mandate as the only large health care policy change in 2014 for the resulting sample, making attribution of reduced form differences-in-differences estimates to the mandate more straightforward. The Massachusetts control group then allows me to account for any other change which was not explicitly tied to the PPACA between 2014 and the previous years. Graphical evidence of parallel trends in insurance coverage between the control and treatment groups is consistent with Massachusetts residents being a good control group in this setting.

The estimates point to the individual mandate causing an increase of 0.85 percentage points in health insurance coverage, or a 17% decline in the uninsured rate. Furthermore, when estimating the effect separately for ESI and for directly purchased insurance, the latter is much more responsive to the mandate. This is consistent with the mandate reducing adverse selection in the individual market, as it was intended to do by policymakers. Similarly, heterogeneity in the treatment effect by age shows substantially greater impact of the mandate among younger individuals, also consistent with the mandate's effectiveness in reducing selection in insurance coverage.

One threat to the interpretation of the reduced form estimate as being a result of the mandate rather than other elements of the PPACA is the possibility that the PPACA subsidies induce individuals who are healthy and eligible for them (i.e., below 400% of the FPL) to enter the market, improving the risk pool for those over 400% of the FPL and driving down their premiums. However, the young being the "low-risk" group in the population a priori yet nevertheless displaying the largest increase in coverage in 2014 is not consistent with this spillover mechanism being the main driver of the results. The heterogeneity in the treatment effect is therefore also consistent with the interpretation that the individual

mandate itself, rather than some general equilibrium effect due to the overall change in the composition of consumers of individual insurance, is behind the estimated results.

The other main threat to the mandate interpretation is that the estimated increase in coverage in the treatment states relative to Massachusetts in 2014 was due to establishment of online insurance exchanges. The fact that small employer exchanges were also introduced in 2014 but that there was only a very small increase in ESI is not consistent with the individual exchanges contributing very much to the increase in coverage; they may, however, have contributed some of the effect and further work should be done to isolate these two potential mechanisms.

The data required to estimate the early effects of the individual mandate is new. However, the mandate itself is still evolving; in 2015 the size of the penalty increased from 1% of income to 2%, and in 2016 it will increase again to 2.5% of income. Each of these changes promises new variation to reestimate the effect of the mandate on insurance coverage, with new advantages and disadvantages to each such quasi-experiment. The 2015 increase in the penalty will not coincide with introduction of subsidies or online exchanges; however it will have to be separated from the effects of new employer mandates for large employers (over 100 employees). The 2016 increase will occur together with an employer mandate for smaller employers (50-99 employees). Consideration of all these sources of variation in the aggregate promises to help address some of the questions left open in the current analysis.

The PPACA is a complex policy change with many moving parts. Its use of an individual mandate to purchase insurance was novel, and such mandates may be considered in other contexts. In addition, the mandate is one of the least popular elements of the PPACA, and numerous proposals have been made to change or eliminate it. Due to these two considerations it is important to know how important the mandate was in achieving the goals of the PPACA, an increase in insurance coverage being chief among them, and thus how effective such a mandate might be as a component of other policies. This paper answers those questions, concluding that the mandate accounts for roughly 20% of the decline in the

uninsurance rate due to the main components of the PPACA. This underscores the importance of the individual mandate in the PPACA, and the potential for similar policies in the future.

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 analysis⁹⁷ 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

⁹⁷See below for construction of the sample for the robustness check using an alternative control group in section 6.2.

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 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 21). 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

Table 21: 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 (0.0428)	-0.0893** (0.0413)	-	0.1202*** (0.0375)	0.0949** (0.0384)
Post65*Post2006	-0.0805** (0.0351)	-0.0491* (0.0252)	0.014 (0.0234)	0.0933*** (0.0315)	0.0094 (0.0251)	-0.0106 (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 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.

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 21 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 21 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 21 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 22.

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

Table 22: 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, 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.

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 .

$$U_i = u_i(c(l_i), s(l_i)) - v_i * l_i \tag{21}$$

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:

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

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:

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

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:

$$\bar{v}_0 - \bar{v}_1 = u_i(c_0, 1) - u_i(c_0, 0) \equiv \Delta u \quad (22)$$

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:

$$\Delta G(v) \equiv G(\bar{v}_0) - G(\bar{v}_1) = \int_{\bar{v}_0 - \Delta u}^{\bar{v}_0} G(v) dv \quad (23)$$

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:

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

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, $\tilde{v} < \bar{v}_0$, 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:

$$\bar{v}_0 - \tilde{v} = u_i(\tilde{c}_0, 0) - u_i(c_0, 0) \equiv \widetilde{\Delta}u \quad (24)$$

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

$$\Delta\tilde{G}(v) \equiv G(\bar{v}_0) - G(\tilde{v}) = \int_{\bar{v}_0 - \widetilde{\Delta}u}^{\bar{v}_0} G(v)dv \quad (25)$$

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 (22) and equation (24). (2) follows immediately from the definitions of equation (23) and equation (25).

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.
- Isaiah Andrews and Conrad Miller. Optimal social insurance with heterogeneity. Technical report, MIT Working Paper. 8, 2013.
- 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.
- Levon Barseghyan, Jeffrey Prince, and Joshua C. Teitelbaum. Are risk preferences stable across contexts? evidence from insurance data. *American Economic Review*, 101(2):591–631, 2011.
- Robert B. Barsky, F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro. Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study*. *The Quarterly Journal of Economics*, 112(2):537–579, 1997.
- Theodore Bergstrom, Lawrence Blume, and Hal Varian. On the private provision of public goods. *Journal of Public Economics*, 29(1):25–49, 1986.
- 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.
- Margaret E. Blume-Kohout and Neeraj Sood. Market size and innovation: Effects of medicare part d on pharmaceutical research and development. *Journal of Public Economics*, 97:327–336, 2013.
- Gary Burtless. Social security, unanticipated benefit increases, and the timing of retirement. *The Review of Economic Studies*, 53(5):781–805, 1986.

- Amitabh Chandra, Jonathan Gruber, and Robin McKnight. The importance of the individual mandate - evidence from massachusetts. *New England Journal of Medicine*, 364(4): 293–295, 2011.
- 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 and Emmanuel Saez. Optimal taxation and social insurance with endogenous private insurance. *American Economic Journal: Economic Policy*, 2(2):85–114, 2010.
- 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.
- David M. Cutler and Jonathan Gruber. Does public insurance crowd out private insurance? *The Quarterly Journal of Economics*, 111(2):391–430, 1996a.
- David M. Cutler and Jonathan Gruber. The effect of medicaid expansions on public insurance, private insurance, and redistribution. *The American Economic Review*, 86(2): 378–383, 1996b.
- David M. Cutler, Amy Finkelstein, and Kathleen McGarry. Preference heterogeneity and insurance markets: Explaining a puzzle of insurance. *American Economic Review*, 98(2): 157–62, 2008.
- 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.
- Liran Einav, Amy Finkelstein, and Mark R. Cullen. Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics*, 125(3):877, 2010.
- Liran Einav, Amy Finkelstein, Iuliana Pascu, and Mark R. Cullen. How general are risk preferences? choices under uncertainty in different domains. *The American Economic Review*, 102(6):2606–2638, 2012.

- 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.
- Keith Marzilli Ericson and Judd B Kessler. The articulation of government policy: Health insurance mandates versus taxes. *Journal of Economic Behavior & Organization*, 2015.
- 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.
- Hanming Fang, Michael P. Keane, and Dan Silverman. Sources of advantageous selection: Evidence from the medigap insurance market. *Journal of Political Economy*, 116(2):303–350, 2008.
- 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.
- Amy Finkelstein. The interaction of partial public insurance programs and residual private insurance markets: evidence from the us medicare program. *Journal of Health Economics*, 23(1):1–24, 2004.
- Amy Finkelstein and Robin McKnight. What did medicare do? the initial impact of medicare on mortality and out of pocket medical spending. *Journal of public economics*, 92(7):1644–1668, 2008.
- 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.
- Mikhail Golosov and Aleh Tsyvinski. Optimal taxation with endogenous insurance markets. *The Quarterly Journal of Economics*, 122(2):487–534, 2007.
- 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.
- Martin B. Hackmann, Jonathan T. Kolstad, and Amanda E. Kowalski. Adverse selection and an individual mandate: When theory meets practice. *The American Economic Review*, 105(3):1030–1066, 2015.
- 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.

Jonathan T. Kolstad and Amanda E. Kowalski. Mandate-based health reform and the labor market: Evidence from the massachusetts health insurance reform. *Journal of Health Economics*, forthcoming.

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.

Darius Lakdawalla, Neeraj Sood, and Qian Gu. Pharmaceutical advertising and medicare part d. *Journal of health economics*, 32(6):1356–1367, 2013.

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.

Melissa Majerol, Vann Newkirk, and Rachel Garfield. The uninsured: A primer. Technical report, The Kaiser Family Foundation, November 2015.

Massachusetts Department of Revenue. Technical information releases 13-1: Individual mandate penalties for tax year 2013, 2013. URL <http://www.mass.gov/dor/businesses/help-and-resources/legal-library/tirs/tirs-by-years>

Mathematica Policy Research. Tracking medicare and prescription drug plans monthly report, 2001-2007. Technical report, 2008.

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.
- Dana Gelb Safran, Patricia Neuman, Cathy Schoen, Michelle S. Kitchman, Ira B. Wilson, Barbara Cooper, Angela Li, Hong Chang, and William H. Rogers. Prescription drug coverage and seniors: Findings from a 2003 national survey. *Health Affairs*, 24:W5, 2005.
- 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.
- Lawrence H. Summers. Some simple economics of mandated benefits. *The American Economic Review*, pages 177–183, 1989.
- 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.
- Amy K. Taylor, Pamela Farley Short, and Constance M. Horgan. Medigap insurance: friend or foe in reducing medicare deficits? *Health Insurance: Public or Private, Pacific Institute for Public Policy Research, San Francisco, CA*, 1988.

- 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.
- Gal Wettstein. Retirement lock and prescription drug insurance: Evidence from medicare part d. Working paper, 2015.
- John R. Wolfe and John H. Goddeeris. Adverse selection, moral hazard, and wealth effects in the medigap insurance market. *Journal of Health Economics*, 10(4):433–459, 1991.
- Yuting Zhang, Julie M. Donohue, Judith R. Lave, Gerald O’Donnell, and Joseph P. Newhouse. The effect of medicare part d on drug and medical spending. *New England Journal of Medicine*, 361(1):52–61, 2009.