2016

Optimal Health Insurance and the Distortionary Effects of the Tax Subsidy

David Powell, *Rand Corporation*

Optimal Health Insurance and the Distortionary Effects of the Tax Subsidy*

David Powell†

September 2016

Abstract

The tax exclusion of health insurance premiums represents the largest source of tax expenditures in the United States while reducing the after-tax price of insurance for the majority of households. This paper develops a model of optimal health insurance in the presence of a tax-deductible premium as well as considering the implications of the Affordable Care Act's "Cadillac tax." While there is a long literature discussing the possible consequences of subsidizing health insurance through the tax code, we have little evidence about how this tax subsidy distorts the optimal cost-sharing schedule for a household. This paper provides theoretical and empirical evidence about how the tax subsidy encourages coverage at low levels of medical care expenditures as well as how these distortions change throughout the medical care distribution. I use federal legislative tax schedule changes to generate variation in the generosity of the tax subsidy and find large effects of the subsidy on generosity at the bottom of the medical care distribution. This responsiveness decreases before forming a U-shape pattern, which is consistent with the predictions of the model. I also estimate the effect of the tax subsidy on annual medical expenditures and private insurer payments, finding large effects for both outcomes. I estimate that the tax subsidy annually leads to deadweight loss of $34-$44 billion.

Keywords: Tax Elasticity, Generous Health Care Coverage, Cost-Sharing, Tax Subsidy, Cadillac Tax, Deadweight Loss
JEL classification: H24, H31, I13

*This project was funded through AHRQ grant R03HS023628. I am grateful for comments from participants at the Southern California Conference for Applied Microeconomics, the Annual Conference on Taxation, and the Conference of the American Society of Health Economists, especially from my discussants Jeffrey Clemens, Michael Geruso, and Mark Pauly. I also had helpful discussions with Abby Alpert, Jim Hosek, Mireille Jacobson, and Paco Martorell.

†RAND, dpowell@rand.org.

1 Introduction

The tax exclusion for employer-sponsored health insurance represents the largest source of tax expenditures in the United States while subsidizing health insurance premium payments for the majority of households. Employee and employer contributions to purchase health insurance can be paid in pre-tax dollars, reducing the relative price of insurance and encouraging demand for more generous health coverage. The Congressional Budget Office projects that the tax exclusion of premiums for employment-based insurance will cost the United States Treasury over $250 billion in 2016 (Congressional Budget Office (2016a)). A major component of the Patient Protection and Affordable Care Act (ACA) is an excise tax on high-premium health insurance plans, the "Cadillac tax," which is designed to counter the distortionary effects of the tax exclusion. The Cadillac tax was initially scheduled for implementation in 2018 but was delayed until 2020, and there is growing political sentiment to fully repeal the tax.

Concerns about the distortionary effects of the tax subsidy are almost as old as the literature concerning the welfare loss associated with moral hazard in health insurance. Feldstein (1973) argues that households in the United States are overinsured and that overly-generous health insurance reduces welfare by several billion dollars per year while suggesting that a major source of this overinsurance is due to the favorable tax treatment received by health insurance. Feldstein and Friedman (1977) blames much of the "health care crisis" on the preferential treatment of health insurance payments in the United States tax code. Pauly (1986) describes how "[t]he tax treatment of health insurance has warped the choice process for such insurance, resulting in the purchase of insurance that, at the market level, is both excessive in quantity and distorted in form."

In this paper, I develop a model of optimal health insurance in the presence of a tax-deductible premium while providing complementary empirical evidence about the relationship between the health insurance tax exclusion and the non-linear cost-sharing schedule in private health insurance plans. The tax subsidy distorts health insurance generosity by reducing cost-sharing, but this relationship is not uniform throughout the medical care distribution. The empirical analysis uses variation in the generosity of the subsidy to study an array of outcomes related to health insurance generosity and medical care consumption. The magnitude of the tax subsidy is directly related to the marginal tax rate, which provides variation useful for the empirical analysis of this paper. Because individuals facing higher marginal tax rates are different on observable and unobservable dimensions compared to people with lower tax rates, I use federal legislative tax schedule changes over the 1999-2012 time period to relate variation in the tax subsidy to changes in health insurance outcomes. Using data from the Medical Expenditure Panel Survey (MEPS), I am able to provide representative estimates of the distortionary effects of the tax subsidy.

Much of the early literature on the distortionary effects of the tax subsidy relied on time series variation in marginal tax rates (e.g., Phelps (1986); Long and Scott (1982); Vroman and Anderson (1984); Turner (1987)), cross-sectional variation in marginal tax rates (e.g., Phelps (1973); Woodbury (1983); Holmer (1984)), or simulations (e.g., Feldstein and Friedman (1977)). More recent work has used natural experiments to estimate the effects of the tax exclusion on a variety of outcomes. Gruber and Poterba (1994) studies the effect of the tax subsidy on the probability that an individual is insured, using the introduction of a tax subsidy for self-employed individuals as part of the Tax Reform Act of 1986. They find that the tax subsidy increases health insurance coverage. Finkelstein (2002) studies employer-provided health insurance coverage in Canada. In 1993, the Quebec government removed the deduction of employer contribution to health benefits from taxable income while other provinces did not. Finkelstein (2002) estimates an elasticity of health insurance coverage with respect to the tax price of -0.5.

Gruber and Lettau (2004) examines the effect of the tax subsidy on the provision of health insurance and firm-level costs using federal tax schedule changes as well as cross-state variation in income tax progressivity. They find that higher subsidies increase the amount spent by firms on health costs. Using a microsimulation model, Gruber and Levitt (2000) estimates that the tax subsidy decreases the fraction of people that are uninsured but at a substantial cost. Cogan et al. (2011) estimates the effect of the tax subsidy on total medical care expenditure consumption by exploiting the discontinuity at the Social Security maximum taxable earnings for identification. The paper employs a regression discontinuity design and estimates a large effect on total medical expenses for the population with labor earnings close to the taxable earnings maximum.

The papers above can be summarized as studying how the tax subsidy affects broad measures of health insurance coverage rates or proxies for generosity (such as average costs or consumption). A primary motivation of this paper is to provide theoretical and empirical evidence which looks within the "black box" to understand *how* the tax subsidy affects costs and medical consumption. More expensive health insurance should provide more generous coverage, but we have little insight about how the tax subsidy directly impacts cost-sharing within private insurance plans. In addition, this paper is the first to rely solely on federal legislative tax changes and the differential effects that these changes have across households for identification while also accounting for income effects. This approach permits estimation of compensated elasticities for a nationally-representative population.

There is widespread interest in understanding the demand for generous health insurance plans. Insurance protects households from large financial losses. Health insurance plans, however, are frequently designed to provide compensation at even low levels of medical expenditures. Such plans decrease the marginal cost of care to consumers, increasing incentives to consume additional care and raising premiums with little benefit in terms of consumption smoothing. Intuitively, we know that the tax subsidy should promote the purchase of more generous health insurance, but we have little evidence about what "more generous" implies. The framework introduced in this paper is well-suited to model the mechanisms through which incentives in the tax code can alter the optimal cost-sharing structure of health insurance plans. I provide theoretical results about the distortionary impact of the tax subsidy on cost-sharing and then complement these results with an empirical strategy which looks at how exogenous changes in the magnitude of the tax subsidy affect cost-sharing throughout the medical expenditure distribution.

In the 2012 MEPS, households paid on average only 39% of the total cost of their first claim in the year and 77% of households paid less than the full cost of the first claim.[1] Selecting on first claims less than \$100 in total costs, households still only paid 54% of the cost of the claim and 62% of households did not pay the full amount of their care. Thus, households are "insured" even at low levels of annual medical expenditures. An important component of this paper is to understand and quantify the role of the tax subsidy in explaining generous cost-sharing arrangements at relatively small amounts of medical care consumption, including a focus on coverage generosity of "first claims."

This paper makes several contributions. First, it is rare in this literature to exploit tax schedule changes for identification. I implement a strategy similar to the one applied in the elasticity of taxable income literature (Gruber and Saez (2002); Saez et al. (2012)) which exploits policy-driven changes in tax rates while separately identifying substitution and income effects. Second, this paper formalizes the theoretical relationship between the

[1] I select on households without any public insurance and with at least one employed adult to create these statistics.

tax subsidy on health insurance generosity in a plan with a nonlinear cost-sharing schedule. This model is extended to include the Cadillac tax and predict its implication on health insurance generosity. The model also provides insight on the appropriate summary statistics required for proper welfare analysis. Third, I provide empirical support for the model while studying an array of outcomes that the model suggests should be impacted by the tax exclusion of health insurance, including measures of first-claim coverage, cost-sharing throughout the medical care distribution, total insurer payments, and total medical care expenditures. Fourth, I study distributional effects. While the literature provides evidence of mean effects on total costs or expenditures, there is no work regarding which parts of the medical care distribution are most affected. Finally, I provide estimates of the annual deadweight loss resulting from the tax exclusion of health insurance premiums.

The theoretical model predicts that the relationship between the marginal tax rate and the cost-sharing schedule should vary throughout the medical expenditure and exhibit a U-shape pattern. It also predicts large effects at the very bottom of the distribution. Both of these predictions are supported by the empirical analysis. I estimate high responsiveness of the magnitude of first-claim coverage to the generosity of the tax subsidy. At higher levels of expenditures, this cost-sharing relationship follows the predicted U-shape pattern. There is little statistical evidence of any effect on cost-sharing above $2150 in annual household medical expenditures. I also estimate that the tax subsidy has large effects on annual household medical expenditures and private insurer payments. These effects are also not uniform as the lower parts of the distribution are most responsiveness to changes in the generosity of the subsidy. The estimates imply that elimination of the tax subsidy would reduce national medical expenditures by $92-$107 billion and total insurer payments by $68-$88 billion. The resulting annual deadweight loss due to the tax subsidy is estimated at $34-$44 billion.

The next section provides additional background about the tax subsidy and the literature studying its effects on health care in the United States. Section 3 develops a model to understand how the tax subsidy distorts cost-sharing schedules. The empirical strategy of the paper is discussed in Section 4. Results are provided in Section 5 while Section 6 concludes.

2 Background

2.1 Tax Exclusion

Because premium payments are paid before income and payroll taxes are levied, each additional dollar paid for health insurance only decreases after-tax household income by $\$(1-\tau)$ where τ is the marginal tax rate. The existence of the tax subsidy for employer-sponsored health insurance likely explains why the majority of households obtain coverage in this manner – in 2014, 55.4% of the United States population obtained health insurance coverage through their employer-based insurance (Smith and Medalia (2015)). The tax subsidy was established as part of the Internal Revenue Code of 1954, and there were several calls to repeal or replace the subsidy in the years prior to the passage of the ACA. President George W. Bush's 2008 budget proposed to replace the tax exclusion with a fixed $7,500 deductible for individuals with health insurance. Presidential candidate John McCain proposed a similar policy during his 2008 campaign, except with refundable tax credits of $2,500 for single people and $5,000 for families.[2] As part of the ACA, the Cadillac tax aims to counter some of the distortionary effects of the tax subsidy by penalizing generous health insurance plans. While the Cadillac tax is levied on insurers, the incidence of this tax is assumed to fall on consumers.[3]

Recent estimates of the distortionary effects of the tax subsidy are rare. Cogan et al. (2011) studies the relationship between changes in tax subsidy generosity and medical care expenditures for 1996-2005, estimating an elasticity of -0.7. The identification strategy focuses on the population close to the Social Security earnings maximum, a relatively high-earning population. Cogan et al. (2011) also provides evidence on the responsiveness of premiums using 1996-1998 data and a measure of predicted (based on firm and insurance policy characteristics) premiums. Gruber and Lettau (2004) studies employer costs for the 1983-1995 time period, estimating an elasticity of -0.7 for firm-level spending conditional on offering and an elasticity of insurance offering of -0.25. The empirical strategy used below resembles the Gruber and Lettau (2004) approach. Gruber and Lettau (2004) uses federal and state legislative tax changes *and* cross-sectional variation within each state-year due to the progressivity of the state tax code. I rely solely on federal legislative tax changes for identification and do not select on households close to the Social Security earnings

[2] http://www.nytimes.com/2008/05/01/us/politics/01mccain.html?_r=0 (accessed August 11, 2016)

[3] See page 13 of Congressional Budget Office (2016b).

maximum.

Gruber (2011) discusses the costs and benefits of the tax exclusion, noting that subsidizing employer-sponsored health insurance may reduce distortions operating through adverse selection by promoting risk-pooling. With the introduction of the individual mandate as part of the ACA, it is likely that the gains associated with risk-pooling at the firm-level have been reduced. The welfare analysis of this paper will consider the distortionary nature of the tax subsidy but does not consider any welfare gains due to risk-pooling.

The tax code subsidizes health care consumption through other mechanisms as well and this paper provides evidence more broadly about the consequences of incentivizing health care consumption through the tax code. Health savings accounts (HSAs) are available to individuals enrolled in high deductible health plans (HDHPs) and the contributions are not subject to federal taxes. Health reimbursement arrangements (HRAs) also act as tax-preferred savings to use for out-of-pocket medical expenses. Medical flexible spending accounts (FSAs) are tax-advantaged, though not tied to HDHP enrollment. The ACA introduced subsidized coverage available on the health insurance exchanges for families making up to 400% of the federal poverty level. Finally, medical costs that exceed 10% of a household's adjusted gross income can be deducted. The focus of this paper is on the tax exclusion, which is – by far – the largest subsidy source in the tax code.

2.2 Health Insurance Generosity

There is substantial interest in understanding why health insurance often provides coverage at low levels of annual medical care expenditures. Coverage at these levels provides little value in terms of consumption smoothing but comes with a high moral hazard cost. Health insurance also provides access to care (Nyman (1999)), though this access motivate also cannot explain generous health insurance cost-sharing arrangements since it is more relevant to very expensive treatments, not low-cost services.

There is little work explaining the determinants of generous health insurance at low levels of medical care. Baicker et al. (2015) discusses reasons that plans could encourage high value care through reduced co-payments for specific services. The authors assume a single co-payment and do not model how cost-sharing should vary throughout the medical care distribution, but this "behavioral hazard" could explain generous coverage at any point of the distribution. Other psychological factors, such as viewing generous health insurance as a commitment device, are also possible explanations.

The tax subsidy offers one straightforward explanation as it encourages households to purchase coverage for low levels of medical expenditures even if this generosity results in higher premiums. This paper quantifies the role of the tax subsidy in explaining generous health insurance coverage, including analysis which specifically analyzes generosity for the first claim incurred by a household.

2.3 Cadillac Tax

The Cadillac tax is one of the primary mechanisms in the ACA to slow the growth rate of health costs in the United States (Gruber (2010)). It is also an important revenue source and is projected to raise tax revenues by $70 billion by 2025 (Congressional Budget Office (2016a)). In 2020, the most generous health insurance plans will be subject to an excise tax. Plans with employee premiums in excess of $10,800 for individual coverage and $29,100 for family coverage will be taxed by 40% of the amount above these thresholds.[4] The purpose of this tax is to gradually, though imperfectly, phase out the tax subsidy for health insurance in the United States tax code as health insurance premiums rise over time (outpacing inflation). Initially, the Cadillac tax was scheduled for implementation in 2018, but was delayed in December 2015 with calls to repeal it entirely. In a 2014 survey of 317 employers, 88% of employers favored repeal of the Cadillac tax.[5] Given its initial 2018 implementation, Herring and Lentz (2011) predicted that 16% of plans would be affected in 2018 and 75% affected by 2028. There is evidence that employers reduced the generosity of their plans in anticipation of implementation (Piotrowski (2013), Abelson (2013), Mrkvicka et al. (2013)).

3 Theory

In this section, I develop a guiding framework to determine the optimal cost-sharing relationship for a household. The framework permits a nonlinear cost-sharing schedule. Other studies have modeled cost-sharing in health insurance, often assuming a linear schedule. Feldstein and Friedman (1977) models the relationship between the tax subsidy and a health insurance

[4]These thresholds are projected (Congressional Budget Office (2016b)).

[5]http://aon.mediaroom.com/2014-10-16-New-Aon-Hewitt-Survey-Shows-Majority-of-Companies-Taking-Immediate-Steps-to-Minimize-Exposure-to-Excise-Tax (accessed August 12, 2016)

plan characterized entirely by the coinsurance rate. This model imposes linear cost-sharing, converts the value of medical care into a dollar value, and then assumes constant absolute risk aversion to translate this value of medical care plus the protection of financial loss into preferences for coinsurance rates. Arrow (1976) models the welfare effects of changes in linear coinsurance rates, recognizing the tradeoff between protection from financial risk and moral hazard. Zeckhauser (1970) also assumes a linear schedule to model the tradeoffs between moral hazard and risk protection.

I build on Mirrlees (1971) which models the optimal income tax schedule, applying the same principles to model the optimal nonlinear health insurance plan, where I define "optimal" as the plan structure which maximizes the household's expected utility. Spence and Zeckhauser (1971) also notes the appropriateness of modeling optimal nonlinear insurance contracts in the same manner as optimal income tax schedules. Blomqvist (1997) applies this approach to model health insurance and discusses a similar model as the one developed in this paper, though it does not explicitly incorporate the tax subsidy into the theoretical framework.

It is helpful to first discuss some intuition for how the tax subsidy can impact the structure of cost-sharing. Assume that the household knows with certainty that it will consume exactly \$100 in medical care in the year. Without the tax subsidy, the household is indifferent between (a) paying a \$100 premium for full insurance and (b) not purchasing insurance and paying \$100 at the time of care. With a tax subsidy, the household prefers paying the insurer a \$100 premium because the after-tax price to the household is equal to \$100$(1-\tau)$. The model will add uncertainty and moral hazard, but the general result is that each household has some optimal relationship between low cost-sharing (with high premiums) and low premiums (with more cost-sharing). The tax subsidy distorts this optimal relationship.

3.1 Model

Each household has a distribution of possible health types for the upcoming year, known to the household and insurer. The health insurance plan provides payments based on the level of medical expenditures consumed since the insurer does not observe the realized health type of the household. The household maximizes its expected utility and must pay a tax-deductible premium equal to the expected value of payments received from the insurer.

Once a plan is designed, households will maximize their utility subject to the incentives inherent in the plan. Low coinsurance rates promote additional medical care consumption since households value medical care. This moral hazard necessitates a series of incentive compatibility constraints.

3.1.1 Household Risk and Utility

The household designs its optimal health insurance plan with a risk-neutral insurer that charges a premium equal to its expected payouts to the household. The household is risk-averse and has a distribution of possible health types for the year, $\theta \sim F\,[0, \infty)$, where higher θ refers to a sicker health state. This distribution is known to the insurer. Once a health insurance contract is agreed upon, the household learns its health type for the year, but the insurer only observes medical expenditures m. The household's out-of-pocket medical payments are represented by $s(m)$ and I refer to $\frac{\partial s}{\partial m}$ as the coinsurance rate. The household maximizes expected utility represented by

$$\int_0^\infty U(c, h; \theta) f(\theta)\, d\theta. \tag{1}$$

The household values consumption (c) and health (h). I define health as a function of health type and medical care such that $h \equiv m - \theta$.

3.1.2 Insurer

The household pays a premium n to the insurer which is equal to the expected value of the payments received from the insurer:

$$n = \int_0^\infty \left[m(\theta) - s\left(m(\theta)\right)\right] f(\theta)\, d\theta. \tag{2}$$

The insurer acts as a means of transferring wealth between different possible states. The expected value of payments is a function of the health risk distribution of the household and the structure of the plan. A generous plan will promote high levels of medical expenditures for a given θ, resulting in a high premium.

3.1.3 State-Specific Utility Maximization

After the household agrees upon an insurance plan and learns its health type, it maximizes utility subject to a state-specific budget constraint. The household solves the following problem:

$$\max_{c,m} U(c, m-\theta)$$
$$\text{subject to} \quad c + s(m) \leq (w-n)(1-\tau),$$

where w represents earnings, which are exogenous in this model. The tax rate is τ and also taken as given. I assume $\tau \geq 0$ and, for notational simplicity, I assume a linear tax rate.[6] n and the health insurance plan (represented by $s(m)$) are exogenous at this stage. $(w-n)(1-\tau)$ represents the household's after-tax income. I also assume $U_c, U_h > 0; U_{cc}, U_{hh} < 0; U_{cc}U_{hh} - U_{ch}^2 \geq 0$. Utility maximization implies:

$$\frac{\partial s}{\partial m} = \frac{U_h}{U_c}. \tag{3}$$

3.1.4 Incentive Compatibility Constraints

The insurer does not observe the household's health status so all payments must be based on m. Because the contract is based on m, the household actually demands that higher medical expenditures require higher out-of-pocket costs. To commit to a specific level of medical care consumption, the chosen health insurance plan must meet a series of incentive compatibility constraints. These constraints can be represented by

$$U(c(\theta), m(\theta)-\theta) \geq U(c(\tilde{\theta}), m(\tilde{\theta})-\theta) \quad \text{for all} \quad \tilde{\theta}, \theta.$$

These constraints ensure that a household with health risk θ will consume $m(\theta)$ in medical care. One implication of these constraints is that higher medical expenditures must be associated with lower consumption (i.e., $\frac{\partial s}{\partial m} > 0$).

[6] The assumption of a linear tax has little effect on the conclusions. The tax reduction due to the tax subsidy is proportional to the marginal tax rate. An exception is if the premium drops the household down to a lower tax bracket. In this case, little is lost by defining τ as the weighted average of the marginal tax rates above and below the kink.

3.1.5 Plan Budget Constraint

The budget constraint can be written as

$$\int_0^\infty [c(\theta) + s(m(\theta))]\, f(\theta)\, d\theta \leq (w - n)(1 - \tau).$$

Plugging in for n using equation (2), the budget constraint is

$$\int_0^\infty [c(\theta) + (1 - \tau)m(\theta) + \tau s(m(\theta))]\, f(\theta)\, d\theta \leq w(1 - \tau).$$

The premium is tax deductible (i.e., n is multiplied by $(1 - \tau)$), implying that all medical expenditures not paid out-of-pocket shift the budget constraint outward by τ. With the tax subsidy, more medical care consumption allows for more "total consumption" $(c + m)$.

3.1.6 Non-Negative Out-of-Pocket Payments

This constraint requires non-negative out-of-pocket payments by the household:

$$s(m(\theta)) \geq 0$$

Without this constraint, a household would purchase a plan with its full earnings. The plan would return a large fraction of this payment back to the consumer, now tax-free. This constraint prevents this scenario and states that the insurer can only pay the individual up to the total amount of medical expenditures. This constraint can be replaced with $s(m(0)) \geq 0$. The IC constraint, then, ensures $s(m(\theta)) \geq 0$ for $\theta > 0$.

3.2 Optimal Coinsurance Rates

For ease of interpretation, I assume $U_{ch} = 0$. Define $A(\theta) \equiv \int_\theta^\infty \frac{1}{U_c(t)} f(t)dt$, representing the expected value of the inverse of $U_c(t)$ for $t \in [\theta, \infty)$; $A(0)$ is the expected value of $\frac{1}{U_c}$.

Proposition 3.1 *Assuming $U_{ch} = 0$, the optimal coinsurance rate at θ is defined by*

$$1 - \frac{\partial s}{\partial m} = \underbrace{\frac{1}{1-\tau}}_{(a)} \underbrace{\frac{U_{hh}}{f(\theta)}}_{(b)} \underbrace{\{A(\theta) - [1 - F(\theta)]\, A(0)\}}_{(c)} \tag{4}$$

The left-hand side variable can be interpreted as the marginal insurance payment, the amount that the insurer pays on the marginal dollar of medical care consumption. I've divided the right-hand side into three separate elements:

(a) The marginal tax rate has a direct effect on the coinsurance rate. This term is positive assuming that $\tau < 1$.

(b) This term is negative by $U_{hh} < 0$. More probable health states (high $f(\theta)$) are associated with higher coinsurance rates to reduce moral hazard in the most likely health states. Using the household first-order condition, it is possible to replace U_{hh} with $\frac{U_c}{\nu}$ where ν represents the responsiveness of medical care to the out-of-pocket price (in state θ). This implies that coinsurance rates increase when the household is more responsive to the out-of-pocket price of medical care.

(c) Due to concavity of the utility function, this term is non-positive. $A(\theta)$ is the expected value of $\frac{1}{U_c}$ over $[\theta, \infty)$. Since $\frac{1}{U_c}$ is decreasing in θ, $A(\theta)$ must be smaller than $[1 - F(\theta)]\, A(0)$ for $\theta \in (0, \infty)$. Finally, it should be noted that this term is equal to 0 for $\theta = 0$ and converges to 0 as $\theta \to \infty$.

Thus, (a) is positive, (b) is negative, and (c) is non-positive. Consequently, $1 - \frac{\partial s}{\partial m}$ is non-negative: $\frac{\partial s}{\partial m} \leq 1$. This conclusion is consistent with our intuition. The appendix includes a proof of Proposition 3.1.

3.3 Impact of Tax Subsidy

This paper is interested in the impact of the tax subsidy on coinsurance rates. The model shows that an increase in τ leads to a decrease in the coinsurance rate:

$$\frac{\partial(1 - \frac{\partial s}{\partial m})}{\partial \left(\frac{1}{1-\tau}\right)} = \frac{U_{hh}}{f(\theta)} \{A(\theta) - [1 - F(\theta)]\, A(0)\} \geq 0. \tag{5}$$

This result matches our intuition. When tax rates increase, the price of health insurance coverage relative to other goods decreases. In response, the individual (holding after-tax income constant) demands more generous coverage, represented here by lower coinsurance rates.

This tax distortion is not uniform throughout the health distribution. While we have little theoretical guidance about how $\frac{U_{hh}}{f(\theta)}$ varies over the distribution, the (c) term provides some insight about how cost-sharing preferences vary over the medical expenditure distribution. The (c) term in equation (4) is equal to 0 at both $\theta = 0$ and as $\theta \to \infty$ and, consequently, there is no tax subsidy effect at the top or the bottom of the distribution.

Taking the derivative of $\{A(\theta) - [1 - F(\theta)]\, A(0)\}$ with respect to θ, we arrive at $f(\theta)\left\{-\frac{1}{U_c(\theta)} + A(0)\right\}$. At low θ, this term is negative. The term becomes less negative as θ increases and eventually positive. Consequently, $\{A(\theta) - [1 - F(\theta)]\, A(0)\}$ is always non-positive, decreases until some $\tilde{\theta}$, and increases thereafter. This term suggests that the cost-sharing schedule should be U-shaped. Which of these terms dominates is an empirical question, motivating the analysis below.

3.4 Extension to the Cadillac Tax

The Cadillac tax will implement a tax on plans with premiums above a threshold at a rate of 40%. Define t as the Cadillac tax rate and n_0 as the threshold. While health plans pay the tax, I will assume that workers bear the burden of the tax through lower wages, though it would be straightforward to modify the model if this incidence assumption is incorrect. The lower wage earnings due to the Cadillac tax reduce total income tax burden. I assume that the Cadillac tax is binding in this section (i.e., $n > n_0$). With the Cadillac tax, I replace the budget constraint with

$$\int_0^\infty [c(\theta) + s(m(\theta))]\, f(\theta)\, d\theta \leq (w - n - t(n - n_0))(1 - \tau).$$

Plugging in for n as before, the new budget constraint is

$$\int_0^\infty \left[c(\theta) + (1-\tau)(1+t)m(\theta) + [1 - (1-\tau)(1+t)]\, s(m(\theta))\right] f(\theta)\, d\theta \leq w(1-\tau) + tn_0(1-\tau).$$

The main result can be modified to include the Cadillac tax:

$$1 - \frac{\partial s}{\partial m} = \frac{1}{(1-\tau)(1+t)} \frac{U_{hh}}{f(\theta)} \left\{ A(\theta) - [1 - F(\theta)] \, A(0) \right\}. \tag{6}$$

The preceding discussions about the tax subsidy can be applied to the Cadillac tax as the Cadillac tax operates in the opposing manner. In my sample, $1 - \tau$ is equal to 0.76 (see Section 4). Consequently, a Cadillac tax of 40% implies that $(1 - \tau)(1 + t)$ is equal to 1.06, approximately canceling out the distortionary effects of the tax subsidy.

3.5 First Claim Coverage

If $m(0) > 0$, then households know with certainty that they will consume positive medical expenditures. In this section, I consider the cost-sharing relationship for medical expenditures below $m(0)$. The probability that annual medical expenditures are less than $m(0)$ is, by definition, zero and the insurer payment is equal to the medical expenditures (due to the restriction imposed in Section 3.1.6). We should expect full coverage of health expenditures incurred with certainty. For coverage below $m(0)$, the payments are inframarginal with no moral hazard cost to the household. This matches our intuition that households will pre-pay for certain health expenditures in pre-tax dollars. Consequently, we should expect first-claim coverage rates to be especially responsive to the generosity of the tax subsidy.

3.6 Welfare Analysis

The tax subsidy distorts demand for health insurance plan generosity by reducing the after-tax cost of the premium to households. The model shows that this distortion acts by increasing insurer payments throughout the medical care distribution. For each health state θ, there is an associated deadweight loss, DWL(θ). This deadweight loss is due to the additional insurer payments for that given health state. Integrating over all health states provides the total deadweight loss. The additional insurer payments results from two factors. First, the tax subsidy increases insurer payments for a fixed consumption of medical care. Households demand more payments for a given level of care because these payments are cheaper to purchase. Second, the tax subsidy encourages the consumption of additional care by distorting cost-sharing and a fraction of this care is paid by the insurer, further increasing insurer payments for that health state.

Deadweight loss in state θ can be represented by the marked region in Figure 1. I define deadweight loss by the following equations. I use $s(m(\theta), p)$ to signify that the chosen amount of insurer payments depends on the price, represented here by p:

$$\text{DWL}(\theta) = \int_{1-\tau}^{1} s(m(\theta), p)dp \tag{7}$$

$$\text{DWL} = \int_{0}^{\infty} \text{DWL}(\theta) f(\theta) d\theta \tag{8}$$

Total deadweight loss is the sum of the deadweight loss in every health state. The sufficient statistic for welfare analysis is the distortion in insurer payments. Total household medical care spending should also increase due to the tax subsidy because of more generous cost-sharing arrangements, but welfare losses only occur to the extent that this spending increases insurer payments. The household fully internalizes cost-sharing payments because they are paid in after-tax dollars.

3.7 Discussion of Model

The framework considered above is quite general. In practice, health insurance plans are more than simple cost-sharing agreements. Nyman (1999) argues that health insurance increases access to care since plans make some types of expensive care financially accessible. More generally, health insurance plans can negotiate the cost of care and consumers value the lower prices that they receive through their insurer. The model can incorporate these price reductions as well. Instead of defining $h = m - \theta$, it is possible to use $h = \alpha m - \theta$ where α is a measure of the efficiency of medical expenditures. If an insurer negotiates prices down (or makes expensive care more affordable even before cost-sharing is considered), the medical expenditures have larger effects on health. I have put little structure on the effect of health on utility and this efficient medical care is similar to placing different value on medical care.

The model does not include adverse selection, a major source of distortion in health insurance markets. In practice, adverse selection may make this optimal plan impossible to find in the market, but we are still interested in defining how those preferences map into the household's optimal plan. Moreover, the model is household-specific and does not consider any externalities or general equilibrium effects of health insurance coverage on other households.

While the model presumes that each household can design its optimal insurance plan, this is not the case in practice as individuals are offered a small set of options. From a policy perspective, we are interested in how the tax subsidy affects actual cost-sharing, not the individual's optimal cost-sharing schedule. The model informs us about what households would like to choose, but the empirical analysis provides complementary information about the practical impact of these preferences.

4 Data and Empirical Strategy

Variation in the marginal tax rate acts as variation in the magnitude of the tax subsidy for health insurance. I use an approach similar to the identification strategy found in the elasticity of taxable income literature. The empirical strategy of this paper relies on an instrumental variable strategy which exploits federal legislative tax schedule changes in the spirit of Auten and Carroll (1999), Gruber and Saez (2002), and Saez et al. (2012), though the implementation will be most comparable to the approach found in Burns and Ziliak (2016) for repeated cross-sections. Federal tax schedule changes alter the marginal tax rates and tax liabilities of households holding household behavior constant. These tax code changes are often not uniform and, instead, affect some households significantly while leaving others relatively unaffected.

I take advantage of changes in federal tax policy during the study period that altered tax-based incentives for reasons unrelated to household behavioral changes. During the sample period, there were two key federal (non-payroll) tax reforms: 1) the Economic Growth and Tax Relief Reconciliation Act (EGTRRA) of 2001, which reduced tax rates for nearly every tax bracket with especially large changes for those with low income; 2) the Jobs and Growth Tax Relief Reconciliation Act (JGTRRA) of 2003, which also reduced tax rates, primarily focusing on relatively higher income households. Figure 2 shows changes in the federal marginal tax rate across the study period. For married couples filing jointly, reductions in the marginal tax rate over this time period ranged from 0 to 46 percent, depending on the household's adjusted gross income.

In addition, there were two temporary payroll tax changes. In 2009 and 2010, the

Making Work Pay Tax Credit provided a tax credit up to $400 for single heads of households and up to $800 for married couples filing jointly. The credit increased after-tax labor earnings by 6.2 percent up to the maximum and was phased out at higher levels of income. In 2011, the payroll tax holiday replaced the Making Work Pay Tax Credit. The "Temporary Employee Payroll Tax Cut" reduced the employee portion of the payroll tax from 6.2 percent to 4.2 percent on earnings up to the taxable earnings maximum ($106,800 in 2011). This tax holiday was extended to 2012.

Before discussing how I exploit these tax code changes for identification, I first describe the data sources that I will use for my analysis.

4.1 Data

The Medical Expenditure Panel Survey (MEPS) contains an extensive set of demographic, health, and medical care variables. I focus on the years 1999-2012 because 1999 is the first year that the MEPS defines "health insurance eligibility units," the individuals in the household that would typically be eligible for coverage under the same private health insurance plan. These eligibility units, which I will refer to as "households," will be the units of observations for the cost-sharing analysis. Furthermore, 2012 is the final year in which the claims data include dates, which are needed because I will use the temporal order of claims to construct measures of first-claim coverage and the cost-sharing schedule.

The MEPS collects and reports several important sets of variables for the analysis of this paper. First, the MEPS includes the Household Component Full-Year files which have detailed income data. The detailed income information is useful to derive marginal tax rates and tax liability measures. I use NBER's TAXSIM program (Feenberg and Coutts (1993)), which takes information on different types of income, number of dependents, and filing status. Given this information, TAXSIM provides federal and FICA tax statistics for each household.[7]

Second, the Full-Year files also include data on individual-level annual medical expenditures and sources of payment such as private insurer payments and out-of-pocket payments. These variables are useful to study the effect of the tax subsidy on annual consumption of medical care.

[7] The MEPS does not include state identifiers so I cannot generate state tax rates. The identification strategy relies on changes in federal taxes so missing the state component is unlikely to be problematic.

Third, the MEPS Household Component Event files include claims data with information about the date of claim, total costs, and sources of payments. This information is useful to construct the observed household cost-sharing relationship and how it varies throughout the medical care distribution. To study the cost-sharing schedule, I use the household as the primary unit of analysis given that individuals can purchase plans for their entire family (or the tax subsidy may impact behaviors that affect the ability and preferences to cover family members). I order each household's claims by date and construct a measure of insurer payments for a given level of medical expenditures. This approach "traces out" the cost-sharing relationship for households that consume at least that level of medical expenditures. The advantages and limitations of this approach are discussed below.

When studying the relationship between the tax subsidy and annual medical care expenditures or annual private insurer payments, I perform the analysis at the individual-level, which permits estimation of models which account for individual characteristics. In all analyses, I exclude people ages 65 and over and select on households in which the head of the household is between 25 and 60. I exclude heads of households with any public insurance (i.e., Medicaid). I also limit the sample to households with at least one employed person. I test the sensitivity of the results to these selection criteria in Section 5.4.2. These selections leave me with 213,456 individuals in 89,419 households. Summary statistics for my sample are provided in Table 1.

The average share of the first household claim paid by a private insurer is 59.5%.[8] I also show the share of expenditures paid by private insurers throughout the entire medical care distribution in Appendix Figure A.1. For the first $50 of annual household medical expenditures, insurers pay – on average – 62% of this amount. The share steadily increases throughout the distribution, reaching 84% at $5,000 in medical expenditures.

In the sample, average medical expenditures are $2,300 (in 2012 dollars). This variable is highly-skewed as the median is only $539. Mean private insurer payments are $1,576; the median value is $208.

4.2 Empirical Strategy

I perform two sets of complementary analyses. In one set, I study how the magnitude of the tax subsidy affects annual individual medical expenditures and private insurer payments.

[8] This metric varies slightly from a similar statistic provided in the introduction due to sample restrictions.

In the other, I study the effect of the tax subsidy on cost-sharing at the household-level throughout the medical expenditure distribution. I first describe the model used to estimate the relationship between the tax subsidy and annual individual medical spending outcomes. The cost-sharing analysis builds on this model.

4.2.1 Total Medical Expenditures and Insurer Payments

I study the empirical relationship between outcomes related to health insurance generosity and the marginal net-of-tax rate. I model the conditional (on all exogenous variables, represented by Z) expectation of medical expenditures as

$$E\left[M_{it}|Z_{it}\right] = \exp\left[\gamma_t + X_{it}'\delta + \beta_1 \ln\left(1-\tau_{it}\right) + \beta_2 \ln y_{it}\right], \tag{9}$$

where M_{it} is an outcome measure such as medical expenditures or private insurer payments for individual i in year t. The specification includes covariates (X_{it}) and year fixed effects (γ_t). I model outcomes as a function of the log of the marginal net-of-tax rate ($1-\tau_{it}$) and a measure of the log of after tax income (y_{it}) discussed in more detail below.[9] These tax variables are consistent with the specification estimated in Gruber and Saez (2002). I define the marginal tax rate as the federal tax rate plus the employee portion of payroll taxes. I use an exponential specification and estimate the above specification using instrument variables Poisson regression for three reasons. First, medical care outcomes are highly-skewed and estimates of linear specifications are sensitive to extreme values. Second, a nontrivial fraction of observations have zero medical expenditures (or private insurer payments) so using $\ln M$ in a linear specification would be inappropriate. Finally, the literature has frequently noted that we are interested in the conditional expectation of medical expenditures, not the expectation of the log of medical expenditures (Manning and Mullahy (2001); Manning et al. (2005)). By Jensen's inequality, we know that $\ln E[M] \neq E[\ln M]$. Santos Silva and Tenreyro (2006) makes a similar point and further shows that a log specification assumes that the error term is multiplicative while Poisson estimation relaxes this restriction and permits multiplicative

[9] Gruber and Lettau (2004) also includes the employer portion of payroll taxes in the denominator when defining the marginal tax rate. Since there were no legislative changes to the employer portion of payroll taxes over this time period, there is little loss in excluding these terms.

and additive error terms.[10]

The covariates are based on characteristics of the heads of the households. They include indicators based on: marital status, 5 year age groups, education,[11] and family size.[12] The first three sets of fixed effects are interacted with gender indicators. Including these covariates in the model is critical to the empirical strategy as outlined below.

Since households may appear in the data twice and because some outcomes are studied at the individual-level, I adjust all standard errors for clustering at the household-level.

4.2.2 Cost-Sharing

To study the tax subsidy's impact on the household cost-sharing schedule, I estimate equation (9) with outcome variable equal to $1 - \frac{s}{m}$, the insurer share of payments for the household's first m medical expenditures in the year.[13] This metric is the average payment share while the solution from the model represented in equation (4) pertains to the marginal payment share. The average share is straightforward to calculate in the data and since I will estimate the cost-sharing relationship at several values of m, the movements in the average effects reflect movements in the marginal effects. I construct these measures using the MEPS claims data, merging the following sets of claims together: inpatient care, outpatient visits, office-based medical provider visits, and emergency room visits.

For this analysis, equation (9) is estimated at different values of m, permitting the parameters of interest to vary based on the level of medical expenditures consumed by the

[10]This literature also highlights that while it is commonly believe that Poisson regression assumes that the conditional mean equals the variance, this restriction is *not* enforced to estimate the parameters. Moreover, similar techniques, such as negative binomial regression, do enforce restrictions related to the relationship between the conditional mean and variance. Thus, other models may be more efficient when these restrictions on the mean and variance hold, but Poisson regression estimates are robust to misspecification of this relationship. See Wooldridge (2010) for a further discussion. Nichols (2010) makes similar points.

[11]I use 4 education groups: less than high school degree, high school degree, some college, and college degree.

[12]I include indicators for family sizes of 1, 2, 3, 4, and 5+.

[13]I only observe payments by source for the entire claim and a claim will not necessarily place the household at precisely m in medical expenditures. In these cases, I impute the distribution of payments for the final claim (which puts the household at m or above in expenditures for the year) using the average for that claim. For example, when analyzing the cost-sharing relationship for the first $100 of medical claim, I may observe households in which the first claim totals $60 and the second claim totals $50. In this case, I multiply the insurer payments for the second claim by 0.8 and add to the insurer payments of the first claim to arrive at the total insurer payments for the first $100 of care.

household up to that point in the year. It is necessary to select the sample on households with at least m in annual medical care expenditures, and I will provide several tests to determine whether this selection criterion is problematic. Ideally, I would observe the cost-sharing arrangements for each household for the entire medical care distribution regardless of how much medical care they actually consume in a year. Unfortunately, there are no representative data sets with both insurance plan information and adequate income and earnings data information to generate tax rates for each household.[14] Furthermore, note that even given detailed insurance plan information, it is difficult to trace out the cost-sharing schedule since most plans include varying prices for different services even at fixed levels of medical expenditures. Using an empirical, realized measure is advantageous in this respect because it is the observed relationship between household medical care consumption and the amount paid by the insurer without guessing what services (and their associated prices) would have been consumed.

There are two primary concerns with this approach. First, the approach requires selecting on individuals with at least m medical expenditures. Since the tax subsidy encourages consumption of more medical care, this selection may not be random. Note, however, that it is not obvious that this is problematic. The strategy provides consistent estimates even if households observed with at least m in medical expenditures are different from those that are not and if the tax subsidy increases the probability that a household will meet that threshold. For example, an increase in the generosity of the subsidy may decrease cost-sharing for a set of households, some of which consume more medical care in response. The household which consume more care accurately trace out the cost-sharing schedule in a representative manner. Alternative mechanisms, however, may cause more systematic selection. To test for the empirical importance of this concern, I show results for samples in which each household has a high probability of consuming m medical expenditures in the year such that selection is less likely to be an issue. I find similar estimates even when I use these selected samples.

Second, the sample for a given m includes households that consume m in medical expenditures and households that consume more than m. I will show results in which I also control for annual medical care expenditures to account for heterogeneity across households. The results are robust to this control (and, though not shown, more flexible functions of annual medical expenditures).

[14]The MEPS provides a link between the household and insurer survey for a select group of years (1996-1999). However, the linked sample is not random and the years avoid any major tax schedule changes.

The cost-sharing analysis does not include subsidies resulting from HSAs, FSAs, or HRAs since we do not observe the tax deductions. In general, since tax-advantaged accounts disproportionately benefit households with high marginal tax rates (the same households benefitting primarily from the tax subsidy), the use of these accounts likely biases the analysis against finding effects. During the time period of the analysis, however, the use of these accounts was relatively rare. In 2012, the use of HSAs and HRAs (combined) peaked at 10.6% among employed individuals[15] and was significantly lower for most of the sample time period. FSAs are also not common (14% in 2010).[16]

4.3 Instrumental Variables

Tax rates and demand for medical care are potentially correlated for reasons unrelated to the tax subsidy. Better health may be associated with less medical care consumption and higher earnings. This underlying correlation motivates the use of an instrumental strategy given the mechanical relationship between earnings and taxes. I create a set of instruments to isolate plausibly exogenous variation in the above tax variables.

The elasticity of taxable income literature relies on variation resulting from federal legislative tax changes and the differential effects that they have on households based on initial household income and characteristics. In this paper, I generate tax instruments by using the differential effect of tax policy changes on households based on observable household characteristics. This approach is most similar to the instrumental variable strategy used in Burns and Ziliak (2016).

The instrumental variable strategy involves pooling all observations in the data and calculating the tax variables, for each observation, for each year from 1999 through 2012. I then estimate the relationship between the simulated tax variables and the covariates defined by X_{it} in equation (9) in each year and predict values for each tax variable based only on covariates. Finally, each observation is assigned the predicted values based on their covariates and the year – these are the "simulated instruments." Consequently, two households with the same covariates X will be assigned different values of the instrument *only* because they face different tax schedules due to federal legislative changes.

[15] In 2012, 19% of covered workers enrolled in an HRA or HSA (Exhibit 8.5 in Claxton et al. (2015)) and 56% of all workers are covered (Exhibit 3.1).

[16] In 2010, 39% of workers had access to an FSA and 37% participated (Mulvey (2012)).

The method for constructing the instruments can be summarized by the following steps. I discuss construction of the simulated log of the marginal net-of-tax rate. The same procedure is used for the after-tax income variable.

1. Holding real income and household characteristics constant, I simulate the log of the marginal net-of-tax rate for every observation in the sample assuming that they were subject to the year s tax code. For example, I take the entire sample and calculate the log of the marginal net-of-tax rate under the 1999 tax code. I represent this variable by $\ln(1-\tau_{it}^{1999})$, where τ_{it}^{s} is the tax rate that household i in year t would have faced under the year s tax schedule. I implement this step for all years 1999-2012.

2. For each s and using the full pooled sample, I regress $\ln(1-\tau_{it}^{s})$ on X_{it}, where X_{it} is the same vector of covariates in equation (9). Let δ_s represent the coefficients on X_{it} for the regression for year s. These estimated coefficients parameterize the relationship between covariates and the predicted log of the marginal net-of-tax rate.

3. Using these coefficients, I predict the log of marginal net-of-tax rate for each observation in the sample in each year: i.e., $\ln(\widehat{1-\tau_{it}^{s}}) = X_{it}'\widehat{\delta}_s$.

4. I assign the predicted log of marginal net-of-tax rate based on t. Observations in year t are assigned $\ln(\widehat{1-\tau_{it}^{t}})$.

The sample and covariates are held constant across regressions in Step 2. Moreover, the inputs used to generate $\ln(1-\tau_{it}^{s})$ are identical across years with the slight caveat that I adjust all income measures for inflation to hold real income constant. Consequently, the *only* reason that $\ln(1-\tau_{it}^{s})$ varies for $s \neq s'$ is because the tax code has changed. As a result, $\widehat{\delta}_s$ (and, consequently, the instruments) only varies from one year to another due to tax code changes.

Restricting the instruments to vary based only on covariates (X_{it}) makes it straightforward to account for the *exact* function that generated the instruments in the specification. I control for X_{it}, accounting for the independent effects of household characteristics on medical care outcomes, and I control for the tax policy changes (i.e., year fixed effects).

Identification originates solely from the interaction of X_{it} and the tax policy changes.[17]

The tax literature following Gruber and Saez (2002) predicts household-level tax changes due to tax code changes for fixed initial household income and characteristics. One concern with this approach is that the instruments, which are functions of initial income, may be predictive of mean reversion and secular income trends (Weber (2014); Burns and Ziliak (2016)). This concern is less important in this context since the dependent variable is a medical care outcome that is not directly used to predict tax changes. However, predicting the instruments based on covariates further alleviates concerns of trends and mean reversion.

To test the robustness of the results to differential shocks and trends, I will also show results in which I include controls related to the mean value of the tax instruments for the person (based on covariates) across all years (i.e., $\frac{1}{T}\sum_s \ln(\widehat{1-\tau_{it}^s})$) interacted with year dummies. This permits individuals predicted to have high tax rates throughout the sample to have different year-to-year shocks. These controls reduce concerns that individuals with high tax rates experienced faster growth in medical care expenditures over the sample period for reasons unrelated to the tax subsidy. For these models, I include both the marginal net-of-tax rate and after-tax income instrument averages interacted with year dummies as controls.

4.4 Income Effects

The identification strategy introduced in Gruber and Saez (2002) uses the nonlinearities in the tax schedule to separately identify substitution and income effects, permitting the parameter associated with the log of the marginal net-of-tax rate to be interpreted as a compensated elasticity. Holding initial household income and characteristics fixed, households experience different changes in their marginal tax rates when the tax schedule changes. Moreover, households experiencing the same marginal tax rates changes may experience different shocks to after-tax income such that the two instruments have independent variation.

[17]Note that it is possible to use more information about each household to construct the instruments (and control for that information in the analysis) such as the characteristics of the other members of the household. I only use information from the head of the household as a simplification, but this choice should have no effect on the consistency of the estimates. Additional information would potentially increase power. I reserve the individual-level characteristics for a robustness test. Given the empirical strategy, the inclusion of these variables in the model should not qualitatively affect the results despite their obvious explanatory power and the consistency of the results serves as a useful test.

The specification in Gruber and Saez (2002) includes observed after-tax income. As discussed in Powell and Shan (2012), this is not the appropriate structural relationship because the observed after-tax income includes the behavioral response to the budget constraint shift.[18] The income effect is the behavioral response to the shift in the budget constraint, which is best measured by the change in after-tax income holding pre-tax income constant. This is equivalent to the instrument as generated using the method in the previous section. I use this measure as the after-tax income variable in all specifications.

5 Results

I begin by providing some graphical evidence of the relationship between the tax subsidy and annual medical care spending outcomes. I follow this preliminary evidence by studying the relationship between the tax subsidy and cost-sharing with an initial focus on health insurance generosity for the first claim. The final analysis in this section studies annual medical care expenditures and private insurer payments.

5.1 Graphical Analysis

To illustrate the approach of this paper, I create three categories based on the predicted change in the marginal net-of-tax rate (using the method described in Section 4.3) between 1999 and 2012. The legislative tax schedule changes over this time period affected some groups more than others. I divide the sample into groups with (1) a predicted decrease in the marginal net-of-tax rate; (2) a predicted increase between 0% and 10%; and (3) a predicted increase greater than 10%. I show the mean change in medical expenditures over the same 1999-2012 time period. There is a monotonic relationship between the change in the tax subsidy and growth in medical expenditures, as shown in Figure 3a. The group that experienced a decrease in the predicted marginal net-of-tax rate consumed, on average, $729 (in 2012 dollars) more in medical care in 2012 relative to 1999. The group with the largest increase in the predicted marginal net-of-tax rate measure consumed only $376 more. This difference is consistent with the distortionary nature of the tax subsidy as those

[18]Imagine a household that perfectly "income targets." When tax liability increases, the household will respond by earning additional income until there is no observed change in after-tax income. Although they are very responsive to changes in income, this would not be observable in a specification which models behavior as a function of after-tax income (which has not changed).

experiencing the largest increase in generosity due to the subsidy increased spending more. Figure 3b shows the corresponding results for private insurer payments with a similar pattern. The increase in private insurer payments between 1999 and 2012 ranges from \$625 to \$262 per person. Larger subsidy increases are associated with larger increases in insurer payments.

5.2 First-Claim Coverage

There is little economic research explaining the prevalence of coverage at low levels of annual medical expenditures in private health insurance. I study this margin by using the claims data in the MEPS and selecting on each household's first claim in the year. The outcome variable is $1 - \frac{s}{m}$ where $\frac{s}{m}$ is the fraction of the total claim paid by the consumer such that $1 - \frac{s}{m}$ is the share paid by the insurer. I perform the analysis, first, using the full sample. I also present estimates using households with only one or two members. The responsiveness to the subsidy may vary by family size if the mapping from health status to household-level medical expenditures is different between small and large households.

Table 2 presents the estimates. The Panel A, Column (1) estimate implies that a 10% decrease in the tax subsidy decreases the magnitude of first-claim coverage by 9.5%. The predicted values presented in Table 2 are based on the 2012 sample. These estimates are generated by predicting outcomes when setting $\ln(1 - \tau) \equiv 0$ for the 2012 data, which is equivalent to eliminating the tax exclusion, and then reporting the change from observed values in 2012. The Column (1) estimate predicts that elimination of the tax subsidy would reduce the share of the first claim paid by the insurer by over 14 cents per dollar. The magnitude is even larger for small families with an estimated elasticity of -2.1 (Panel B, Column (1)), predicting that the subsidy is increasing private insurer payment shares for the first claim by 28 cents per dollar.

A primary concern of this approach is that we only observe the outcome variable for households with at least one claim in the year, excluding 18% of households in the data. The tax subsidy (through first-claim coverage generosity) may itself affect the probability of having at least one claim and, in fact, the results presented below in Section 5.4 suggest that there is such a relationship.

To test for the importance of selection, I replicate the analysis while selecting on households that are predicted – based on covariates and the instruments – to have a high probability of at least one claim. For households with very high predicted probabilities, selection becomes less of a potential issue. I estimate a probit model with the covariates and instruments as the explanatory variables and generate predicted probabilities for each observation in the sample. In Column (2) of Table 2, I select on households with at least an 80% chance (for the full sample) of having at least one claim. On average, this sample has over a 90% probability of having at least one claim. The magnitudes of the estimates increase using this smaller sample. In Column (3), I select on households such that, on average, there is over a 95% chance of having a claim (with a minimum of 93%). The estimates become noisier due to the smaller sample size, but they are similar to the main estimates, suggesting the selection is not driving the results. While the subsidy encourages more spending and increases the probability of having any medical care consumption in a year, there is no evidence that this selection is systematic and, in fact, the pattern suggests that any selection effect is biasing *against* finding an effect.

Next, I control for differences in household log of (end-of-year) medical care expenditures. These estimates are presented in Column (4). Since total expenditures are also a function of the subsidy, the inclusion of this control should attenuate the effect. There is some evidence of attenuation, but the predicted changes in cost-sharing are still large at 12 cents per dollar and statistically significant from zero at the 5% level. The robustness to controlling for differences in annual medical care is also suggestive that selection is not driving the results. For small families, the results are also consistent when adding this control.

Finally, the sample of first claims includes large claims in which households may value risk protection from losses of that magnitude. In the final column, I select on first claims that were $250 or less and replicate the above analysis. The estimates are generally consistent with the previous estimates, suggesting that the results are not driven by large first claims. Instead, the evidence suggests that the tax subsidy drives a significant proportion of generous coverage observed at low levels of medical care expenditures.

5.3 Cost-Sharing Schedule

The model presented in this paper makes predictions about the relationship between the tax subsidy and the cost-sharing schedule. First, a higher net-of-tax rate should lead to reduced per-dollar insurer payments. Second, this relationship should vary throughout the medical expenditure distribution. Specifically, the model provides theoretical reasons to expect a

U-shape relationship. I test the model empirically using the MEPS claims data, ordering the claims by date, and constructing the total insurer payments for a given level of medical expenditures for each household. This process traces out the empirical cost-sharing relationship between household medical expenditures and private insurer payments. For a given level of medical expenditures (m), the outcome variable is the fraction of those expenditures paid by private health insurance. With this outcome variable, I estimate equation (9). The sample is households that consumed at least m in medical expenditures and the outcome variable is $1 - \frac{s}{m}$. This approach permits β_1 (and all other parameters) to vary based on m. I estimate the model for each \$50 increment up to \$5000.[19]

The estimates are presented in Figure 4.[20] Near the bottom of the distribution, we observe a "first-claim" effect discussed in the previous section. For low medical care expenditures which households can predict that they will reach, the tax subsidy encourages full coverage of these costs. The estimates decrease in magnitude until \$250 as the "first claim" effect should diminish at higher levels of care (i.e., households are less likely to have certain medical expenditures at larger amounts). The insurer payment elasticity at \$250 of medical care expenditures is -0.71. The magnitude of this estimate generally increases until reaching a peak of -1.17 at \$1300 of medical care. The magnitude then steadily decreases, tracing out a U-shape relationship. At \$2150 of medical care expenditures, the elasticity is no longer statistically significant from zero at the 10% level and the point estimates reach approximately zero at around \$3500. At high levels of medical care, there is little effect of the tax subsidy on cost-sharing, as hypothesized.

Following the robustness tests implemented in Section 5.2, I replicate the above analysis while testing for the importance of selection. The sample for each estimate consists only of households which actually consumed at least m medical expenditures in the year. In Appendix Figure A.2, I control for the log of annual medical expenditures consumed by the household. The estimates are generally similar, suggesting that differences in medical care across households are not driving the results. In Appendix Figure A.3, I select on households with a high probability of consuming m in medical expenditures for the year. Given that I estimate the model for high levels of medical expenditures, I set this probability at 0.5 and only estimate the model up to \$1000. If systematic selection into the sample is

[19] I have estimated the model above \$5000 as well. The tax elasticity estimates are generally close to zero and not statistically significant from zero.

[20] The corresponding after-tax income elasticities are not shown and are generally not statistically significant from zero.

driving the results, we should observe large differences between the main estimates and the estimates selecting on the high probability households. The estimates are generally similar in Appendix Figure A.3, suggesting that selection is not biasing the estimates.

Finally, Appendix Figure A.4 presents the same results but selecting on families with only one or two members. While noisier, the general pattern of the results is similar, though the magnitudes are often larger for this subset. Overall, I do not find evidence that aggregating all family types together is responsible for the main estimates presented in Figure 4. The results in this section are strongly supportive of the hypothesis that the tax subsidy makes plans more generous and that this causal effect is not uniform throughout the entire cost-sharing schedule.

5.4 Annual Medical Expenditures and Private Insurer Payments

The tax subsidy encourages the purchase of more generous health insurance which potentially leads to higher insurer payments and consumption of more medical care. While the literature has provided some estimates of the relationship between the tax subsidy and total individual (or household) medical care spending, it has not used legislative tax changes to identify this relationship.

5.4.1 Distributional Estimates

Given the implications of the model, it is likely that the effect of the tax subsidy on medical care expenditures is heterogeneous throughout the entire distribution of medical care spending. I model the probability that individual-level annual medical spending is smaller than some threshold q as a function of time fixed effects, covariates, and the tax variables. I estimate

$$\mathbf{1}\left(M_{it} < q\right) = \gamma_t(q) + X_{it}'\delta(q) + \beta_1(q)\ln\left(1 - \tau_{it}\right) + \beta_2(q)\ln y_{it} + \epsilon_{it}(q), \qquad (10)$$

where q indexes the parameters and residual. I estimate equation (10) for $q = \$1, \$100, \$200, \$300, \ldots, \$10{,}000$ using 2SLS. The advantage of this approach is that it permits the effect of the tax subsidy on medical care to vary at different points of the distribution. This analysis is conducted at the individual-level. I expect the estimates of

$\beta_1(q)$ to be positive as higher values of the marginal net-of-tax rate imply a less generous subsidy and should increase the probability of lower values of medical care spending.

The estimates are shown in Figure 5a. I find that the tax subsidy substantially shifts the cumulative distribution function at the bottom of the medical care spending distribution. A 10% increase in the marginal net-of-tax rate increases the probability of having no medical care spending by 4.2 percentage points, statistically significant from zero at the 10% level. This effect increases monotonically up to the $600 threshold and then steadily declines to around zero. The estimates are generally statistically significant from zero at the 5% level until the $3800 medical care spending threshold. At higher points in the distribution, there is little evidence of any effect.

Overall, we observe an interesting pattern. The tax subsidy appears to encourage consumption of more care at lower parts of the distribution with little effect at the top of the distribution. This result is consistent with more elastic behavior among relatively healthy individuals while individuals consuming significant amounts of health care are less responsive to changes in the price of care. It is also consistent with the conclusions of the model and the empirical cost-sharing estimates that the tax subsidy causes large cost-sharing changes in some parts of the distribution and smaller changes elsewhere.

The corresponding effects of the log of after-tax income are estimated jointly and presented in Figure 5b. There is little statistical evidence that additional income due to legislative budget constraint shifts alters the medical care spending cumulative distribution function (CDF).

Figure 6a presents the relationship between the tax subsidy and the cumulative distribution function of private insurer payments. As discussed before, studying distortions in private insurer payments is important for measuring the welfare loss associated with subsidizing health insurance premiums. The estimates in Figure 6a reveal a similar pattern as those observed for total medical care spending, though the magnitudes are larger for private insurer payments. These larger magnitudes reflect that the tax exclusion more directly subsidizes insurer payments. As before, the estimates monotonically increase from the bottom of the distribution to a modest level of annual insurer payments ($300) and then gradually decrease to zero. All estimates are statistically different from zero at the 5% level until the $3000 threshold. I observe little statistical evidence of any effect on the CDF above that point. The pattern is consistent with predictions from the model concerning differential effects of the subsidy throughout the medical care distribution.

Figure 6b presents the corresponding income effects. As before, there is little statistical evidence that tax-driven shifts in the budget constraints affect demand for health insurance plan generosity.

To test the robustness of these estimates, I control for individual-level characteristics as well and present the figures in the Appendix. Appendix Figure A.5 shows the change in the CDF for total expenditures; Appendix Figure A.6 shows the change in the CDF for insurer payments. There is little effect on the pattern of estimates for either outcome when these additional controls are included in the model. Furthermore, Appendix Figures A.7 and A.8 replicate the results for families with only 1 or 2 people. We might expect that the responsiveness at different points in the distribution could vary based on family size. For a given level of medical expenditures, larger families likely have different cost-sharing demands than smaller families, impacting the relationship between the marginal net-of-tax rate and the medical care distributions. While noisier due to the smaller sample, the results presented in Appendix Figures A.7 and A.8 are similar to the main results.

5.4.2 Mean Estimates

To summarize the relationship between the tax subsidy and medical care spending outcomes, I estimate equation (9) using instrumental variables Poisson regression.[21] The estimates are presented in Table 3. The outcome variable is total individual annual medical expenditures. The Column (1) estimate implies that a 10% increase in the marginal net-of-tax rate leads to a 20% reduction in medical care spending, holding after-tax income constant. Based on this result, the elimination of the tax subsidy would reduce per capita medical care spending by \$758 and total national medical care spending by \$107 billion. As before, these estimates are generated by predicting medical expenditures when setting $\ln(1-\tau) \equiv 0$ for the 2012 data. All predictions are partial equilibrium estimates of household-level responses and exclude spillovers which may occur if the tax exclusion were eliminated.

The model shown in Column (1) includes the household-level covariates necessary for identification due to the instrumental variable strategy. In Column (2), I add individual-level controls including indicators for race, educational attainment, age, and sex. The estimate is

[21]Because of the highly-skewed nature of the outcome variables, I drop individuals in the top 1% of the medical care expenditure distribution. Given that the results in the previous section found little evidence that the tax subsidy has any effect at the top of the distribution, this selection should not affect the consistency of the estimates and serves only to reduce the variance.

relatively unchanged by the inclusion of these controls. Column (3) adds the mean value of the instruments for each person interacted with year dummies. The estimates, again, are not meaningfully affected, suggesting that any differential trends or shocks based on household characteristics are not correlated with the tax instruments and biasing the estimates. Finally, in Column (4), I exclude prescription drug expenditures. The estimates are similar, though the "total change" metric decreases because of the smaller base.

Panel B of Table 3 includes the corresponding estimates for the private insurer payments outcome variable. The pattern of estimates is similar, though the magnitudes are larger. As discussed in the previous section, the larger magnitudes are consistent with the tax exclusion having a more direct impact on health insurance plan generosity. The health insurance premium tax exclusion encourages the purchase of insurance which provides more inframarginal and marginal (which encourage additional consumption) payments to consumers. While I estimate larger proportional effects for insurer payments, the estimates of the total change in private insurer payments if the exclusion were eliminated are smaller than the equivalent estimates for the change in medical care expenditures because of the smaller base. The estimates imply that the elasticity is between -2.1 and -1.7, and that the tax exclusion increases annual national private insurer payments by $65 to $88 billion.

While these estimates are large, they are quite similar to estimates found in the literature. Using the Column (1), Panel A estimate in Table 3, the results imply that elimination of the tax subsidy would decrease total spending by 32.2% for the employed population without public insurance. Cogan et al. (2011) reports a 32.5% decrease for the population close to the Social Security earnings maximum while Gruber (2002) estimates that elimination of the subsidy would decrease spending by 35.4% (among those offered insurance).

The Column (1), Panel B estimate implies that elimination of the tax exclusion would decrease insurer payments by 37.8%. The estimates presented in Gruber and Lettau (2004) suggest a 45% decrease. While Gruber and Lettau (2004) estimates a smaller spending elasticity, this elasticity is estimated conditional on offering insurance, which they find is itself responsive to the tax subsidy. When both the extensive and intensive margins are combined to summarize the total effect of the tax exclusion, the estimates are relatively consistent with those found in this paper.

The above analysis selected on households with at least one employed adult. If tax policy affects the decision to work, then the tax instruments may be correlated with changes in the sample. In Appendix Table A.2, I replicate Table 3 but keep the families with no employment in the sample. Regardless of the household's tax rate, I set $\ln(1-\tau)=0$ for this group since they receive no subsidy. Overall, the results are similar (in general, the magnitudes are larger) and the selection criteria used for this paper do not appear to be driving the conclusions.

5.5 Welfare Estimates

Estimates of the deadweight loss estimates associated with the tax subsidy are rare in the literature. Using a simulation approach, Feldstein and Friedman (1977) estimates a welfare loss of \$8 billion in 1970 dollars, equivalent to \$47.3 billion in 2012 dollars. Calculating the welfare consequences of the health insurance tax exclusion follows from estimates of the relationship between the tax subsidy and private insurer payments. The tax exclusion subsidizes the premium which, in the model, is directly related to expected payouts from the insurer. Thus, deadweight loss is a function of the change in insurer payments as illustrated earlier in Figure 1. Because the approach of this paper conditions on and separately identifies the budget constraint shift due to legislative tax schedule changes, I interpret the elasticities in the previous section as compensated elasticities, which are necessary for proper welfare analysis. I use $\frac{1}{2} \times \Delta$Insurer Payments to estimate deadweight loss. Consequently, I estimate an annual deadweight loss between \$34 billion and \$44 billion (using Columns (1)-(3) in Table 3B). This estimate does not incorporate any externalities of health insurance, include possible general equilibrium effects, or consider the risk-pooling benefits that potentially arise from the existence of the tax subsidy.

6 Discussion and Conclusion

The tax exclusion of health insurance premiums is costly and distorts demand for health care. In this paper, I provide theoretical and complementary empirical evidence concerning how the tax subsidy affects the cost-sharing schedule in employer-sponsored health insurance plans. Using variation in the subsidy due to legislative changes in the tax schedule, I estimate large effects on first-claim coverage generosity and, similarly, cost-sharing at low levels of annual medical care expenditures. The estimates suggest that elimination of the tax subsidy would reduce first-claim coverage by about 14 cents per dollar of care consumed.

The model also predicts that the subsidy's effect on cost-sharing should express a U-shape and I find empirical support for this prediction. Responsiveness peaks at around \$1300 of care before gradually converging to zero; there is little evidence of any effects above \$2,150.

I also estimate large effects on annual private insurer payments and medical care expenditures. I, again, find that these effects are not uniform throughout the distribution. There are larger effects closer to the bottom of the distribution for both outcomes. The mean estimates imply that that a 10% increase in the marginal net-of-tax rate (i.e., decrease in the tax subsidy) reduces individual medical expenditures by 20% and insurer payments by 27%. The effect on insurer payments is expected to be larger given that the subsidy more directly distorts insurer payments. While these estimates are large, they are consistent with estimates found in the literature, which often use different sources of variation and select on different populations.

The analysis of this paper studies household-level behavior in response to changes in the magnitude of the tax subsidy. While previous work has studied firm-level decisions, the estimates of this paper incorporate a host of possible mechanisms which could affect medical care consumption and health insurance plan generosity such as selection into firms with generous health insurance offerings which may be missed by studying outcomes at the firm-level. This paper is also the first to identify solely off of legislative tax schedule changes, and I jointly estimate income effects to permit estimation of compensated elasticities. My estimates imply that the tax exclusion is responsible for \$34-\$44 billion of annual deadweight loss.

Figures

Figure 1: Deadweight Loss in Health State θ

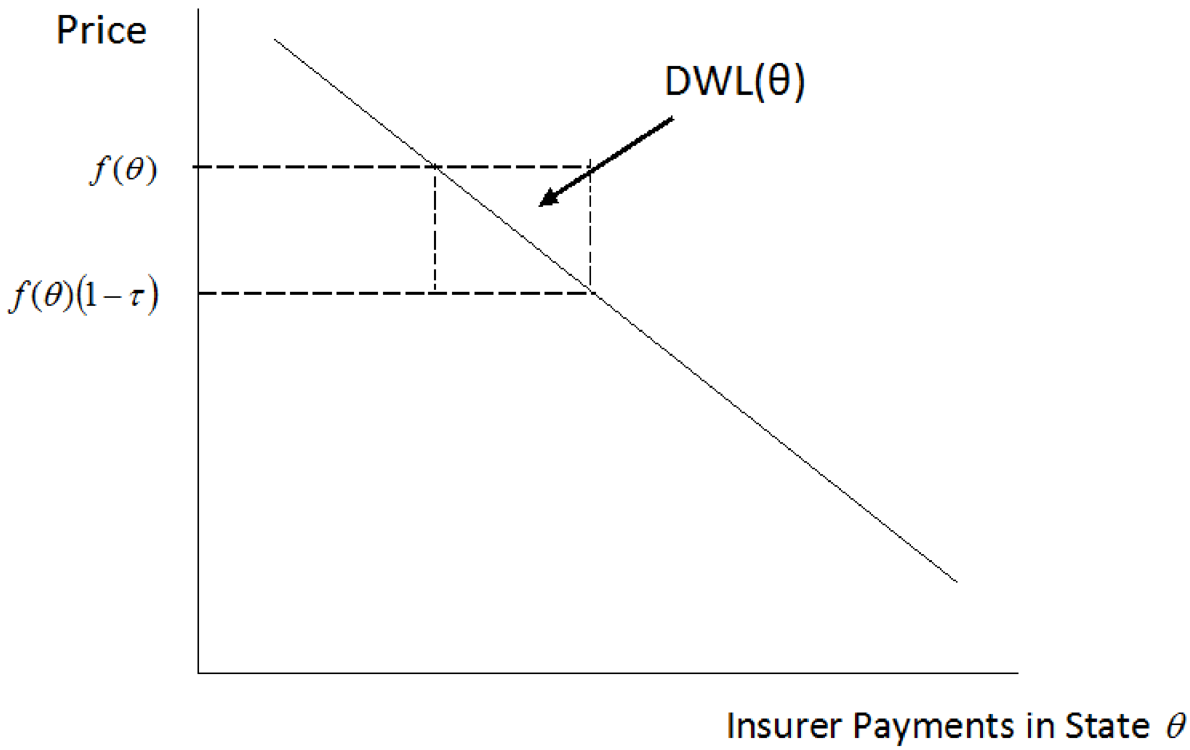

Notes: This figure shows a compensated demand curve for private insurer payments in health state θ, defined in the model.

Figure 2: Federal Marginal Tax Rates, Pre- and Post-Tax Reforms

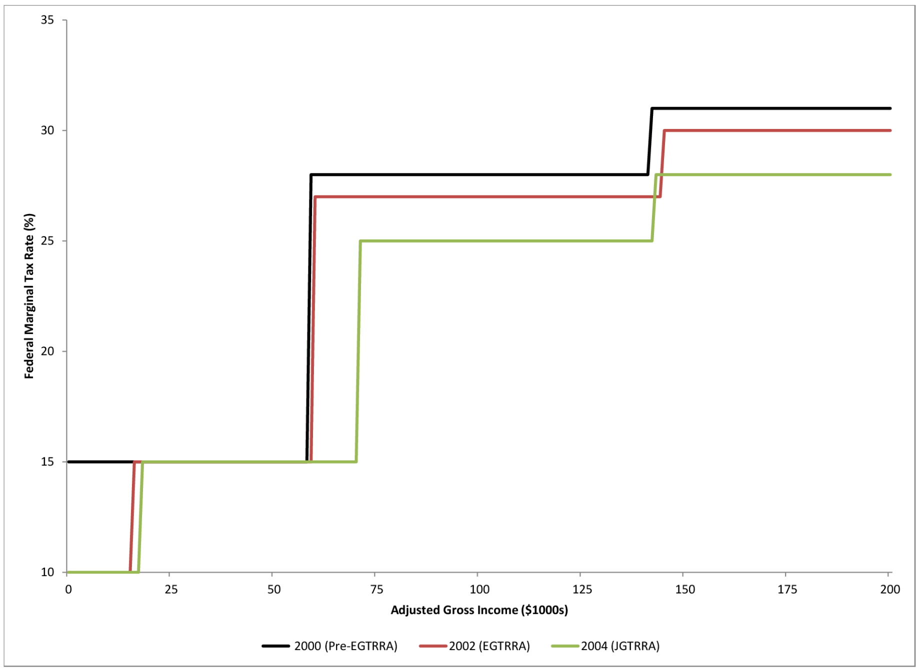

Notes: Marginal tax rates are for married couples filing jointly. Income is in constant 2013 dollars. Source: "U.S. Federal Individual Income Tax Rates History, 1862-2013 (Nominal and Inflation Adjusted Brackets)," Tax Foundation.

Figure does not include the Making Work Pay Tax Credit, enacted for 2009 and 2010, in this figure since the credit size was dependent on labor earnings, not income. In 2009 and 2010, the Making Work Pay Tax Credit provided a tax credit up to $400 for single heads of households and up to $800 for married couples filing jointly. The credit increased labor earnings by 6.2 percent up to the maximum and was phased out at higher levels of income. The above figure also does not include the 2011 and 2012 payroll tax holiday which reduced the payroll tax on individual earnings from 6.2 to 4.2 percentage points. Again, this tax was dependent on earnings, not income.

Figure 3a: Change in Medical Expenditures, 1999-2012

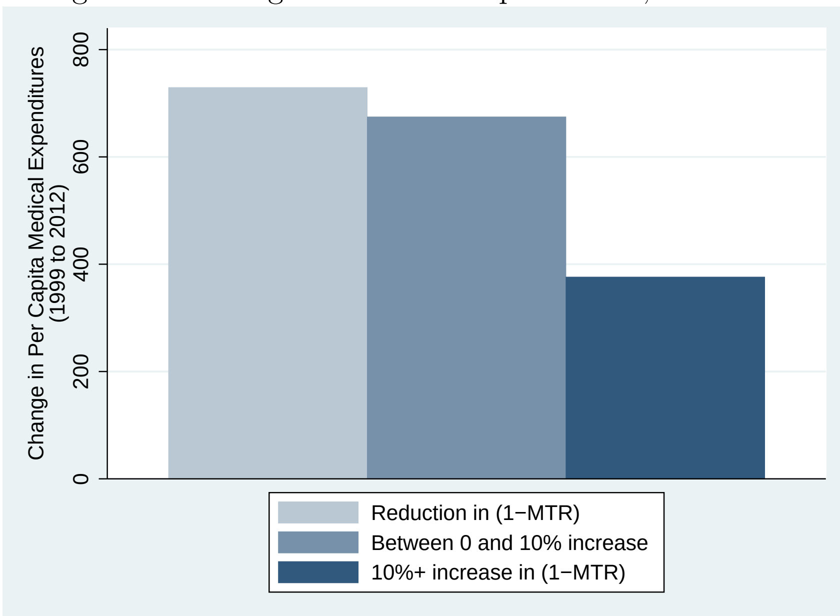

Figure 3b: Change in Insurer Payments, 1999-2012

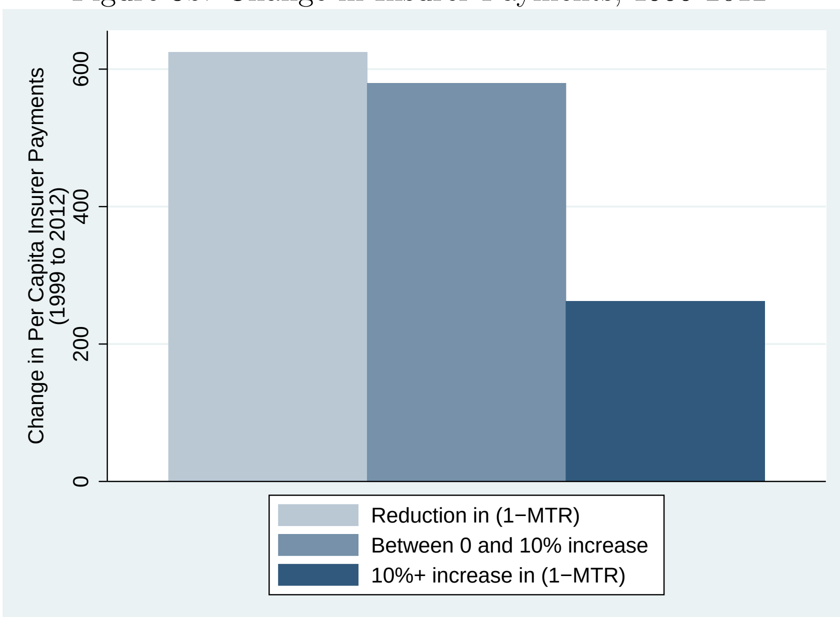

Notes: Using the instruments generated by the method described in Section 4.3, I group each 1999 and 2012 observation based on the predicted change in the marginal net-of-tax rate over that time period. I then calculate the mean change in the observed medical expenditures (insurer payments) between 1999 and 2012 for each group. This method parallels the main empirical strategy of the paper.

Figure 4: Cost-Sharing Marginal Net-of-Tax Rate Elasticity by Total Medical Expenditures

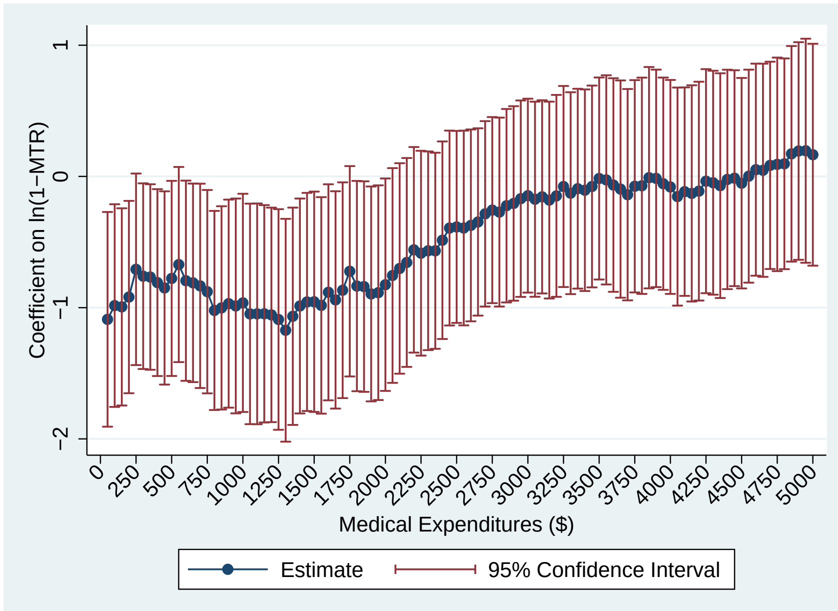

Notes: Outcome is share of payments made by private insurers. Each point represents an estimation of equation (9) for the given level of medical expenditures (in 2012 dollars). Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, year fixed effects, and predicted after-tax income. The covariates listed are determined by the head-of-the-household. Estimation uses instrumental variables Poisson regression. Confidence intervals adjusted for within-household clustering.

Figure 5a: Effect of Log of Marginal Net-of-Tax Rate on Medical Expenditure CDF

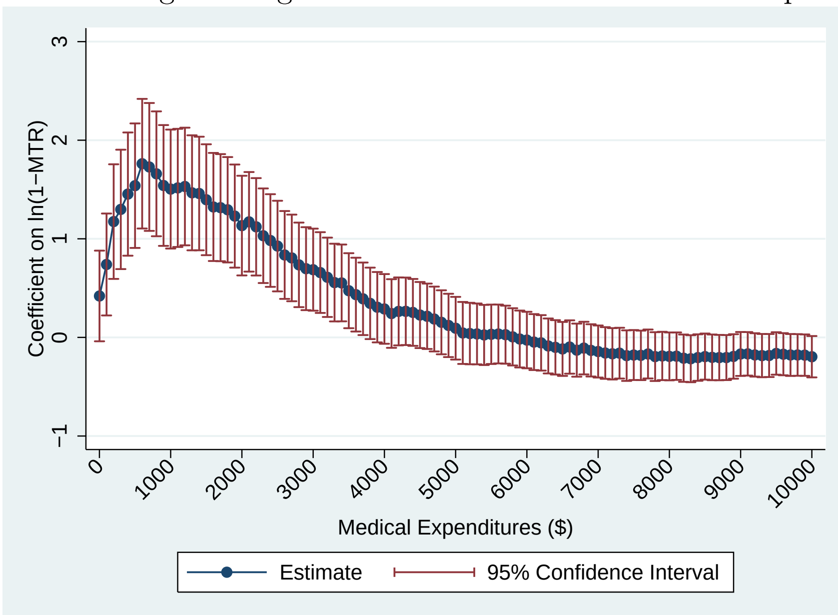

Figure 5b: Effect of Log of After-Tax Income on Medical Expenditure CDF

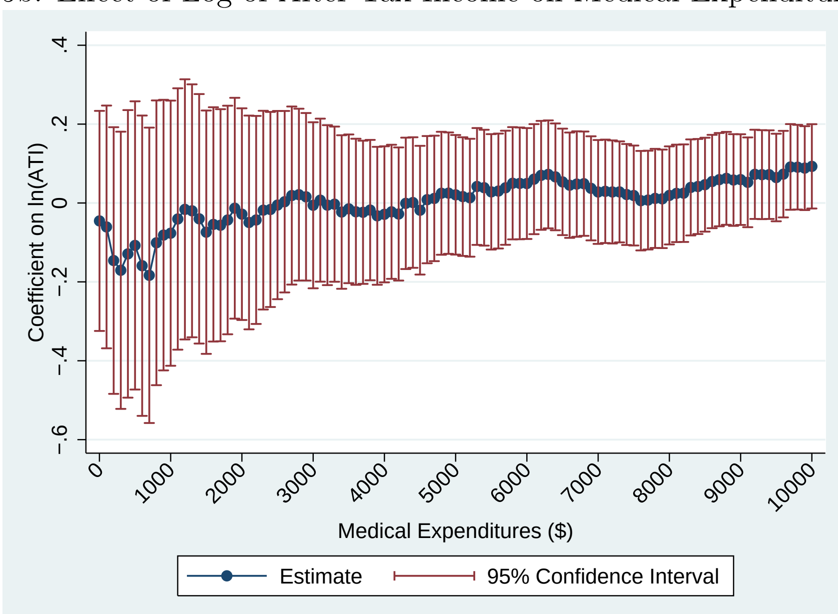

Notes: Outcome is indicator equal to 1 if annual medical expenditures less than threshold. Each point represents an estimation of equation (10) for the given level of medical expenditures (in 2012 dollars) – the corresponding points on the two figures are estimated jointly. Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, and year fixed effects. The covariates listed are determined by the head-of-the-household. Estimation uses 2SLS. Confidence intervals adjusted for within-household clustering.

Figure 6a: Effect of Log of Marginal Net-of-Tax Rate on Private Insurer Payments CDF

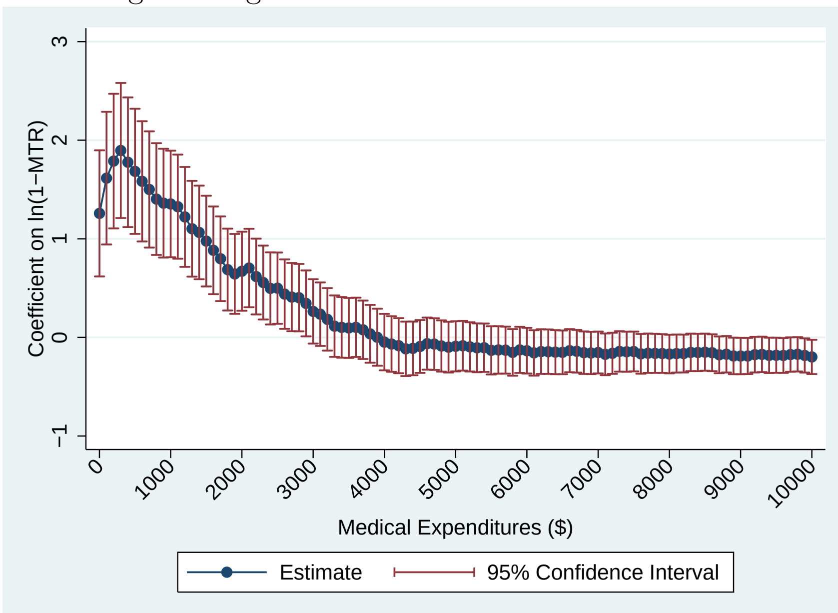

Figure 6b: Effect of Log of After-Tax Income on Private Insurer Payments CDF

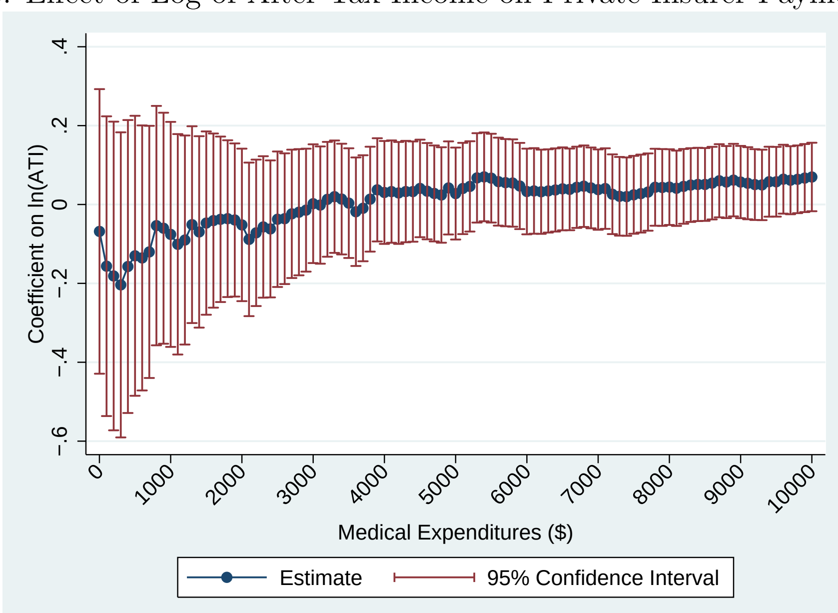

Notes: Outcome is indicator equal to 1 if annual private insurer payments less than threshold. Each point represents an estimation of equation (10) for the given level of private insurer payments (in 2012 dollars) – the corresponding points on the two figures are estimated jointly. Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, and year fixed effects. The covariates listed are determined by the head-of-the-household. Estimation uses 2SLS. Confidence intervals adjusted for within-household clustering.

Tables

Table 1: Summary Statistics

	Head of Household (N=89,419)	
	Mean	Standard Deviation
Marginal Tax Rate	24.03	11.77
After-Tax Income	$55,675.33	$42,279.19
Married (%)	61%	49%
Family Size	2.65	1.48
12 Years of Education (%)	28%	45%
13-15 Years of Education (%)	23%	42%
16+ Years of Education (%)	29%	45%
Share of First Claim Paid by Private Insurer (%)	59.5%	39.3%
	Individual (N=213,456)	
	Mean	Standard Deviation
Age	31.8	17.3
Medical Expenditures	$2,299.68	$7,261.52
Insurer Payments	$1,576.28	$6,412.75

Notes: All dollar values expressed in 2012 dollars. Sample excludes individuals ages 65+ (other sample restrictions are discussed in text). The marginal tax rate is defined as the federal marginal tax rate plus the employee portion of Federal Insurance Contributions Act (FICA) rate.

Table 2: First-Claim Coverage

	(1) Full Sample	(2) High Probability of Positive Expenditures	(3) Very High Probability of Positive Expenditures	(4) Control for log(Expenditures)	(5) Small Claims Only
Panel A: All Family Sizes					
ln(1-MTR)	-0.948**	-1.303***	-1.239*	-0.808*	-0.998*
	(0.446)	(0.494)	(0.739)	(0.439)	(0.543)
ln(ATI)	0.098	0.056	0.089	0.073	0.202**
	(0.084)	(0.088)	(0.116)	(0.086)	(0.102)
Change in Insurer Share if No Tax Subsidy	-0.142**	-0.192***	-0.206**	-0.123**	-0.136**
	(0.060)	(0.062)	(0.105)	(0.061)	(0.066)
N	73,417	55,550	19,720	73,417	53,215
Panel B: Family Size ≤ 2					
	(1) Full Sample	(2) High Probability of Positive Expenditures	(3) Very High Probability of Positive Expenditures	(4) Control for log(Expenditures)	(5) Small Claims Only
ln(1-MTR)	-2.112***	-2.395***	-2.350*	-1.979***	-2.399***
	(0.660)	(0.766)	(1.354)	(0.650)	(0.860)
ln(ATI)	0.049	0.298	0.157	0.095	0.124
	(0.285)	(0.257)	(0.352)	(0.280)	(0.325)
Change in Insurer Share if No Tax Subsidy	-0.278***	-0.307***	-0.329**	-0.265***	-0.282***
	(0.066)	(0.071)	(0.135)	(0.066)	(0.073)
N	36,196	27,164	10,435	36,196	27,028

Notes: Significance levels represented by * 10%, ** 5%, *** 1%. Standard errors are adjusted for within-household clustering. Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, and year fixed effects. The covariates listed are determined by the head-of-the-household. "Small claims" defined as $250 or less in 2012 dollars. "Change in Insurer Share if No Tax Subsidy" is predicted based on 2012 sample.

Table 3: Tax Elasticities for Medical Expenditures and Insurer Payments

Panel A: Annual Expenditures				
	(1)	(2)	(3)	(4)
ln(1-MTR)	-2.044***	-1.856**	-1.716**	-2.121**
	(0.823)	(0.816)	(0.891)	(0.927)
ln(ATI)	-0.126	-0.135	1.966	1.124
	(0.586)	(0.596)	(1.424)	(1.501)
Per Capita Change	-758.16***	-701.08***	-654.21**	-619.44***
with No Tax Subsidy in 2012	(246.85)	(255.36)	(291.88)	(226.15)
Total Change with No Tax Subsidy	-107***	-99***	-92**	-87***
in 2012 (in billions)	(35)	(36)	(41)	(32)
Family Characteristic Controls	Yes	Yes	Yes	Yes
Individual Characteristic Controls	No	Yes	Yes	Yes
Tax Instruments x Year Controls	No	No	Yes	Yes
Include Prescription Drugs?	Yes	Yes	Yes	No
Panel B: Annual Insurer Payments				
	(1)	(2)	(3)	(4)
ln(1-MTR)	-2.743***	-2.447***	-1.966*	-2.433**
	(0.913)	(0.914)	(1.036)	(1.057)
ln(ATI)	-0.039	-0.072	2.942	2.166
	(0.686)	(0.698)	(1.838)	(1.939)
Per Capita Change	-625.76***	-575.84***	-483.66**	-460.12***
with No Tax Subsidy in 2012	(160.32)	(172.38)	(220.09)	(165.76)
Total Change with No Tax Subsidy	-88***	-81***	-68**	-65***
in 2012 (in billions)	(23)	(24)	(31)	(23)
Family Characteristic Controls	Yes	Yes	Yes	Yes
Individual Characteristic Controls	No	Yes	Yes	Yes
Tax Instruments x Year Controls	No	No	Yes	Yes
Include Prescription Drugs?	Yes	Yes	Yes	No

Notes: Significance levels represented by * 10%, ** 5%, *** 1%. Standard errors are adjusted for within-household clustering. Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, and year fixed effects. The covariates listed are determined by the head-of-the-household. The same covariates are included at the individual-level in Columns (2)-(4). Columns (3) and (4) include additional controls based on the mean level of the instruments for each individual interacted with year indicators. The "Total Change" metrics use the MEPS household-level weights to scale up to national estimates.

A Appendix

A.1 Preliminary Condition

First, it is necessary to introduce a monotonicity condition similar to the Spence-Mirrlees condition. The assumption of the lemma below is met under the $U_{ch} =$ assumption. The lemma states that unhealthy households spend more on medical care and less on general consumption. The assumption driving this result holds trivially for $U_{ch} \geq 0$ since $U_{hh} < 0$. This condition is important in interpreting the results because it states that unhealthy households receive less non-medical consumption.

Lemma A.1 *If* $U_{ch} \geq \frac{U_c U_{hh}}{U_h}$, *then* $\frac{\partial c}{\partial \theta'} \leq 0$ *and* $\frac{\partial m}{\partial \theta'} \geq 0$.

Proof:

Define $V(\theta, \theta') \equiv U(c(\theta'), m(\theta') - \theta)$.

Incentive compatibility implies that the following conditions hold:

1. $\frac{\partial V}{\partial \theta'}(\theta, \theta) = 0$.
2. $\frac{\partial^2 V}{\partial \theta'^2}(\theta, \theta) \leq 0$.

Differentiating the first condition gives: $\frac{\partial^2 V}{\partial \theta'^2} + \frac{\partial^2 V}{\partial \theta \partial \theta'} = 0$ such that the second condition can be replaced by $\frac{\partial^2 V}{\partial \theta \partial \theta'} \geq 0$.

The first condition implies

$$\frac{\partial V}{\partial \theta'} = U_c \frac{\partial c}{\partial \theta'} + U_h \frac{\partial m}{\partial \theta'} = 0$$

This gives us the useful condition:

$$\frac{\partial c}{\partial \theta'} = -\frac{U_h}{U_c} \frac{\partial m}{\partial \theta'} \tag{11}$$

The second condition can be rewritten:

$$\begin{aligned} \frac{\partial^2 V}{\partial \theta \partial \theta'} &= -U_{ch} \frac{\partial c}{\partial \theta'} - U_{hh} \frac{\partial m}{\partial \theta'} \\ &= \frac{\partial m}{\partial \theta'} \left[\frac{U_{ch} U_h}{U_c} - U_{hh} \right] \geq 0 \end{aligned}$$

This last condition implies that $\frac{\partial m}{\partial \theta'} \geq 0$ whenever $\frac{U_{ch}U_h}{U_c} - U_{hh} \geq 0$. Equation (11) then implies that $\frac{\partial c}{\partial \theta'} \leq 0$.

A.2 Main Result

Defining $V(\theta) \equiv U(c(\theta), m(\theta) - \theta)$, I use Pontryagin's maximization principle where $V(\theta)$ is the state variable and $m(\theta)$ is the control variable. Using the envelope theorem, I can replace the incentive compatibility constraints above with

$$V'(\theta) = -U_h$$

The final household maximization problem can be written as

$$\max_{V,m} \int_0^\infty V(\theta) f(\theta)\, d\theta$$

subject to

$$\begin{aligned} &\text{(IC)} & V'(\theta) = -U_h \\ &\text{(BC)} & \textstyle\int_0^\infty \left[c(\theta) + (1-\tau)m(\theta) + \tau s(m(\theta))\right] f(\theta)\, d\theta \leq w(1-\tau) \\ &\text{(Non-Negative OOP)} & s(m(0)) \geq 0 \end{aligned}$$

The Hamiltonian is

$$H = V(\theta)f(\theta) + \lambda \left[w(1-\tau) - c(\theta) - (1-\tau)m(\theta) - \tau s(m(\theta))\right] f(\theta) - \mu(\theta) U_h(c(\theta), m(\theta) - \theta).$$

The optimal health insurance plan is defined by the following equations:

$$\frac{\partial H}{\partial m} = 0 \tag{12}$$

$$\frac{\partial H}{\partial V} = -\mu'(\theta) \tag{13}$$

$$\mu(0) = \lim_{t \to \infty} \mu(t) = 0 \tag{14}$$

Proposition 3.1 The optimal coinsurance rate at θ is defined by

$$1 - \frac{\partial s}{\partial m} = \frac{1}{1-\tau} \frac{U_{hh}}{f(\theta)} \left\{ A(\theta) - [1 - F(\theta)]\, A(0) \right\}$$

Proof:

Equations (12) and (13) lead to

$$\lambda \left[\frac{U_h}{U_c} - (1-\tau) - \tau(\frac{\partial s}{\partial m}) \right] f(\theta) = \mu(\theta) U_{hh}, \tag{15}$$

$$f(\theta) - \lambda \left[\frac{1}{U_c} \right] f(\theta) = -\mu'(\theta). \tag{16}$$

Deriving $\mu(\theta)$ from (16) and using $\mu(0) = 0$,

$$\mu(\theta) = 1 - F(\theta) - \lambda \int_{\theta}^{\infty} \frac{1}{U_c(t)} f(t)\, dt,$$

$$\mu(0) = 1 - \lambda \int_{0}^{\infty} \frac{1}{U_c(t)} f(t)\, dt = 0.$$

These conditions give us

$$\lambda = \frac{1}{A(0)}$$

$$\mu(\theta) = 1 - F(\theta) - \frac{A(\theta)}{A(0)}$$

I replace $\mu(\theta)$ and λ in equation (15). I also use the individual FOC to replace $\frac{U_h}{U_c}$ with $\frac{\partial s}{\partial m}$ and arrive at a formula for the optimal coinsurance rate

$$1 - \frac{\partial s}{\partial m} = \frac{1}{1-\tau} \frac{U_{hh}}{f(\theta)} \Big\{ A(\theta) - [1 - F(\theta)]\, A(0) \Big\}$$

B Additional Figures

Figure A.1: Share of Total Payments Paid by Insurer for Different Levels of Medical Expenditures

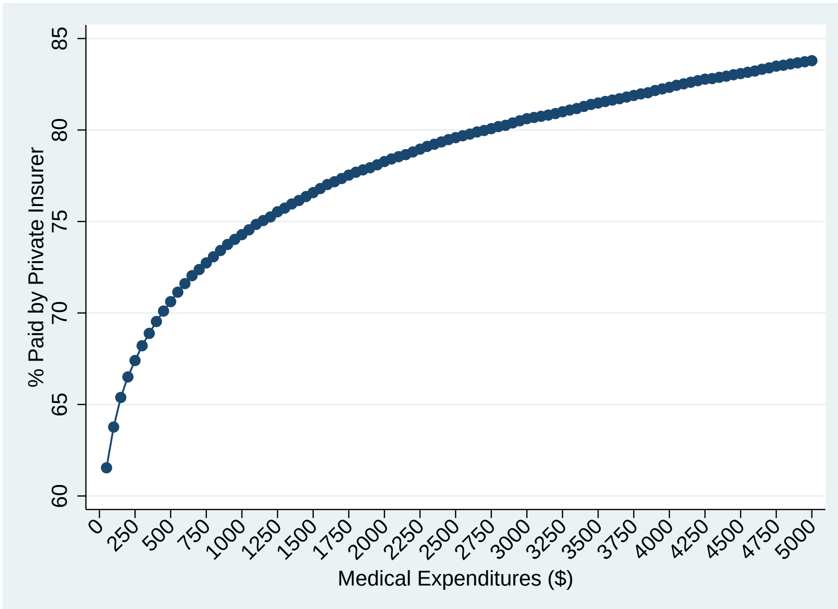

Notes: Each point represents the percentage of the first $X in medical expenditures paid by private insurers, where X is defined on the x-axis.

Figure A.2: Cost-Sharing Marginal Net-of-Tax Rate Elasticity by Total Medical Expenditures: Controlling for log(Total Expenditures)

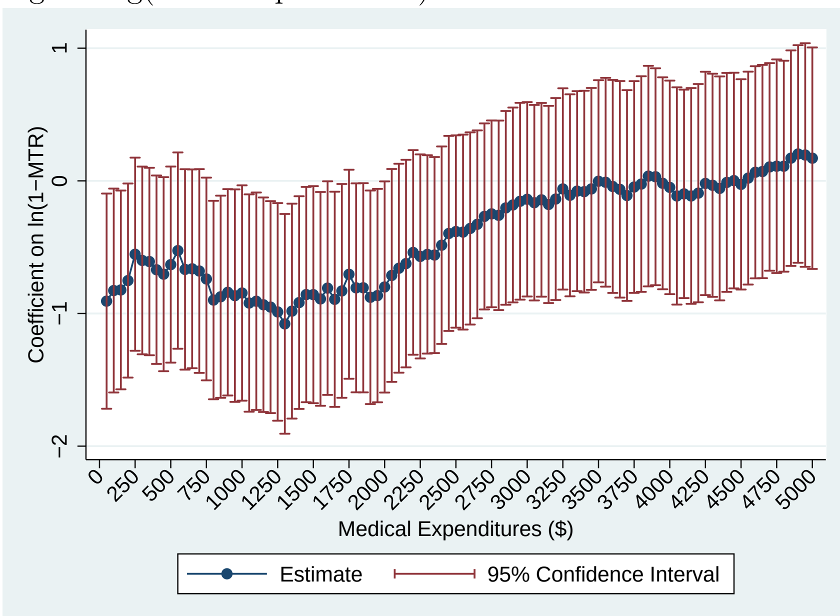

Notes: Outcome is share of payments made by private insurers. Each point represents an estimation of equation (9) for the given level of medical expenditures (in 2012 dollars). Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, year fixed effects, predicted after-tax income, and the log of total annual household medical expenditures. The covariates listed are determined by the head-of-the-household. Estimation uses instrumental variables Poisson regression. Confidence intervals adjusted for within-household clustering.

Figure A.3: Cost-Sharing Marginal Net-of-Tax Rate Elasticity by Total Medical Expenditures: High Probability Households

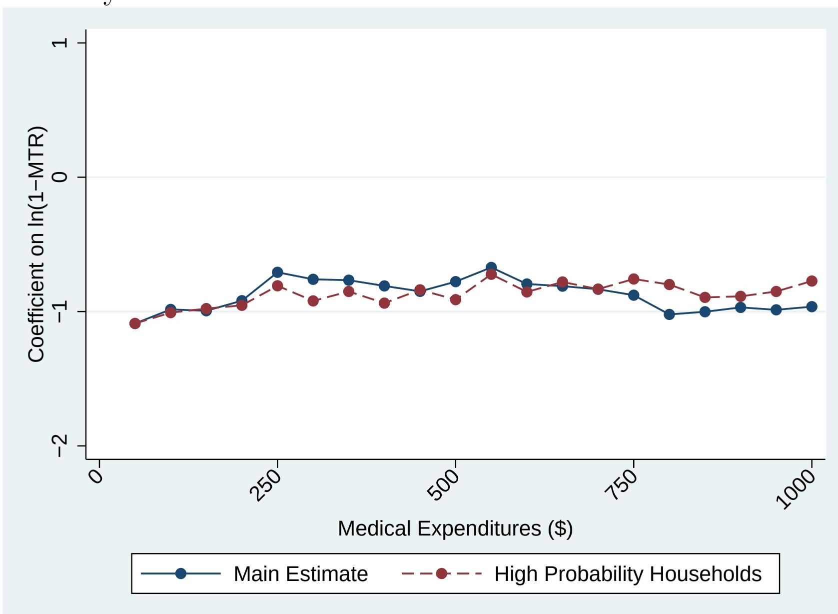

Notes: Outcome is share of payments made by private insurers. Each point represents an estimation of equation (9) for the given level of medical expenditures (in 2012 dollars). Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, year fixed effects, predicted after-tax income, and the log of total annual household medical expenditures. The covariates listed are determined by the head-of-the-household. Estimation uses instrumental variables Poisson regression. Confidence intervals suppressed for legibility. "High Probability Household" sample selected on households with at least a predicted 50% chance of incurring the given level of medical expenditures. y-axis selected to match previous figures.

Figure A.4: Cost-Sharing Marginal Net-of-Tax Rate Elasticity by Total Medical Expenditures: Family Size ≤ 2

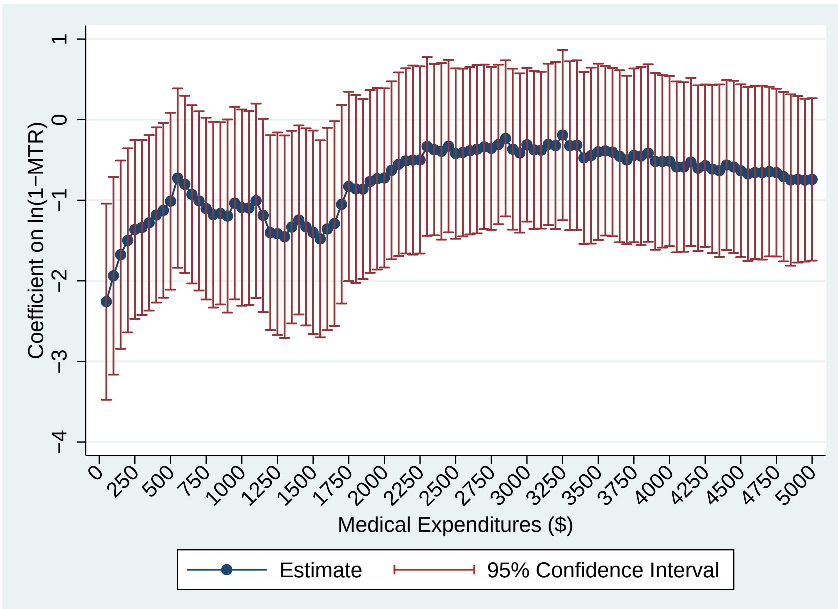

Notes: Outcome is share of payments made by private insurers. Each point represents an estimation of equation (9) for the given level of medical expenditures (in 2012 dollars). Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, year fixed effects, and predicted after-tax income. The covariates listed are determined by the head-of-the-household. Estimation uses instrumental variables Poisson regression. Confidence intervals adjusted for within-household clustering.

Figure A.5: Effect of Log of Marginal Net-of-Tax Rate on Medical Expenditure CDF: Additional Controls

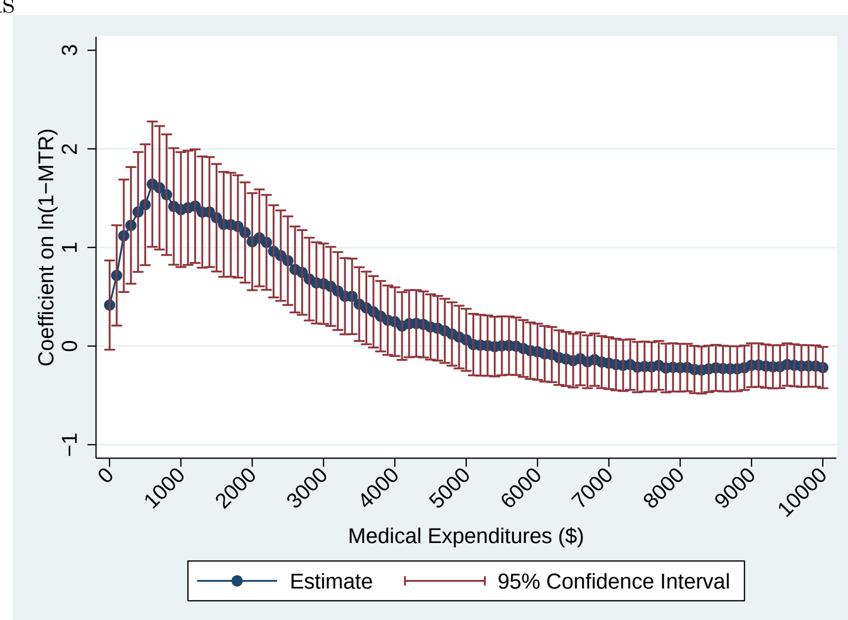

Notes: Outcome is indicator equal to 1 if annual medical expenditures less than threshold. Each point represents an estimation of equation (10) for the given level of medical expenditures (in 2012 dollars) – the corresponding points on the two figures are estimated jointly. Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, and year fixed effects. The covariates listed are determined by the head-of-the-household. In this analysis, I also include the same covariates at the individual-level. Estimation uses 2SLS. Confidence intervals adjusted for within-household clustering.

Figure A.6: Effect of Log of Marginal Net-of-Tax Rate on Private Insurer Payments CDF: Additional Controls

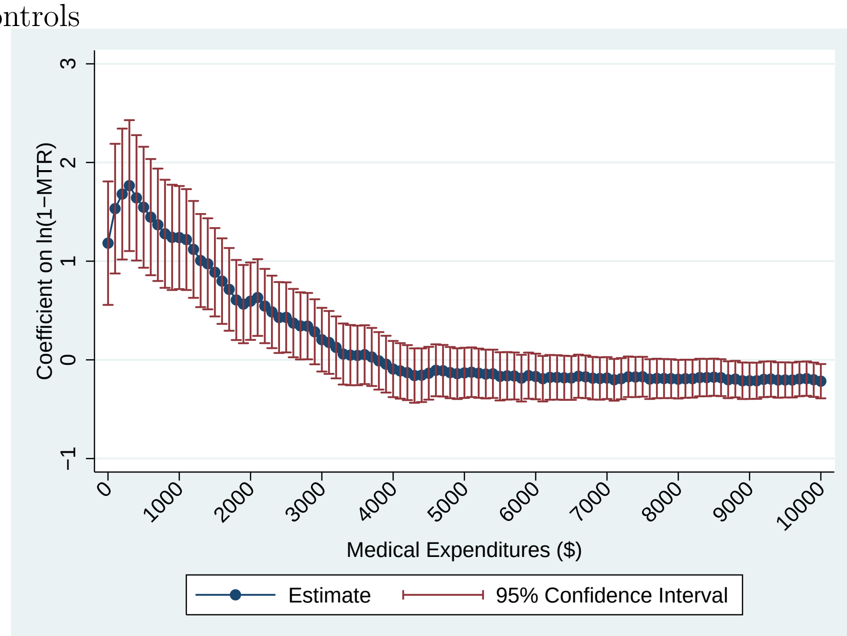

Notes: Outcome is indicator equal to 1 if private insurer payments less than threshold. Each point represents an estimation of equation (10) for the given level of private insurer payments (in 2012 dollars) – the corresponding points on the two figures are estimated jointly. Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, and year fixed effects. The covariates listed are determined by the head-of-the-household. In this analysis, I also include the same covariates at the individual-level. Estimation uses 2SLS. Confidence intervals adjusted for within-household clustering.

Figure A.7: Effect of Log of Marginal Net-of-Tax Rate on Medical Expenditure CDF: Family Size ≤ 2

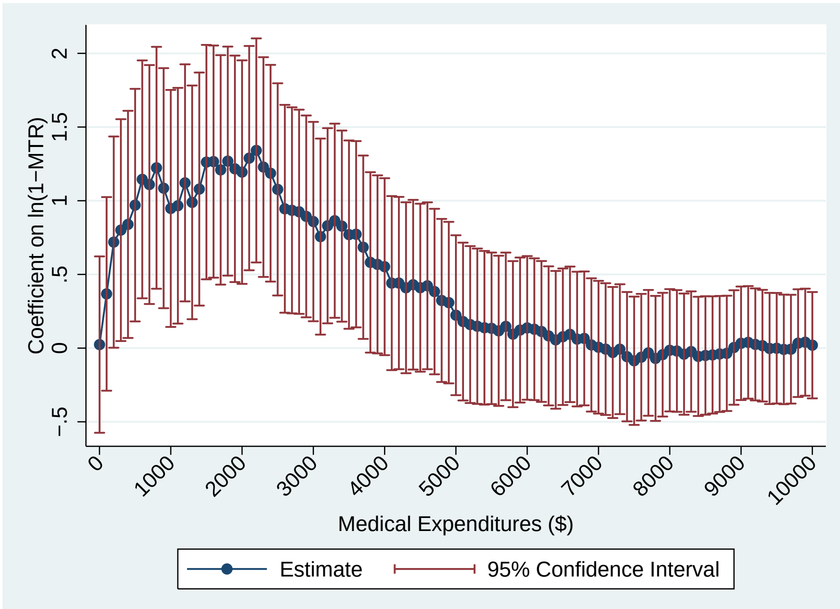

Notes: Outcome is indicator equal to 1 if annual medical expenditures less than threshold. Each point represents an estimation of equation (10) for the given level of medical expenditures (in 2012 dollars) – the corresponding points on the two figures are estimated jointly. Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, and year fixed effects. The covariates listed are determined by the head-of-the-household. Estimation uses 2SLS. Confidence intervals adjusted for within-household clustering.

Figure A.8: Effect of Log of Marginal Net-of-Tax Rate on Private Insurer Payments CDF: Family Size ≤ 2

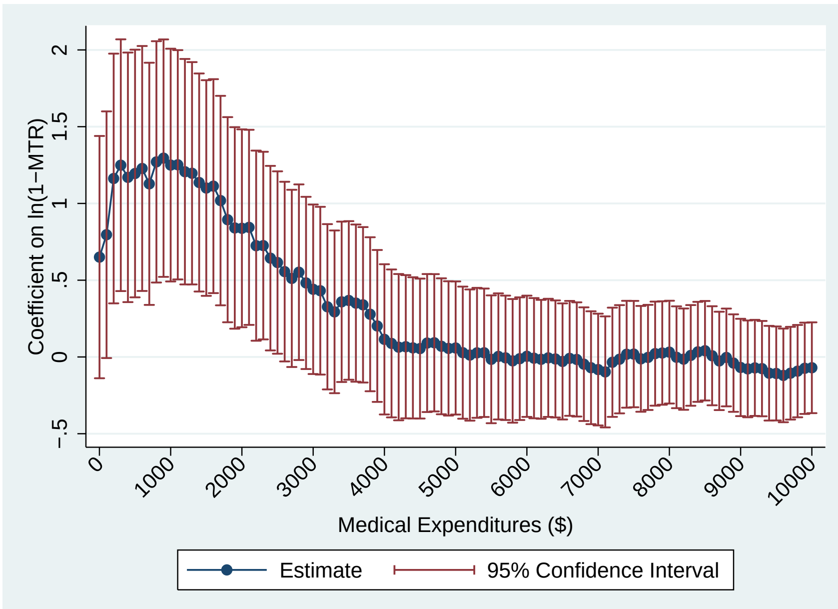

Notes: Outcome is indicator equal to 1 if private insurer payments less than threshold. Each point represents an estimation of equation (10) for the given level of private insurer payments (in 2012 dollars) – the corresponding points on the two figures are estimated jointly. Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, and year fixed effects. The covariates listed are determined by the head-of-the-household. Estimation uses 2SLS. Confidence intervals adjusted for within-household clustering.

C Additional Tables

Table A.1: First Stage

Outcome:	ln(1-MTR)		
	(1)	(2)	(3)
$\ln(1 - \widehat{\text{MTR}})$	1.067***	1.073***	1.221***
	(0.131)	(0.131)	(0.144)
ln(ATI)	0.196***	0.198***	-0.047
	(0.073)	(0.073)	(0.169)
Family Characteristic Controls	Yes	Yes	Yes
Individual Characteristic Controls	No	Yes	Yes
Tax Instruments x Year Controls	No	No	Yes
First Stage F-Statistic	66.22	67.36	71.69

Notes: Significance levels represented by * 10%, ** 5%, *** 1%. Standard errors are adjusted for within-household clustering. Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, and year fixed effects. The covariates listed are determined by the head-of-the-household. The same covariates are included at the individual-level in Columns (2)-(3). Column (3) also includes additional controls based on the mean level of the instruments for each individual interacted with year indicators.

Table A.2: Tax Elasticities for Medical Expenditures and Insurer Payments, Including Families with No Employment

Panel A: Annual Expenditures				
	(1)	(2)	(3)	(4)
ln(1-MTR)	-2.190***	-1.968**	-1.934**	-2.329**
	(0.676)	(0.714)	(0.757)	(0.812)
ln(ATI)	-0.138	-0.166	1.566	0.818
	(0.672)	(0.678)	(1.003)	(1.422)
Per Capita Change	-780.73***	-715.07***	-700.61**	-646.26***
with No Tax Subsidy in 2012	(183.83)	(204.74)	(204.99)	(174.30)
Total Change with No Tax Subsidy	-114***	-104***	-102***	-94***
in 2012 (in billions)	(27)	(30)	(30)	(25)
Family Characteristic Controls	Yes	Yes	Yes	Yes
Individual Characteristic Controls	No	Yes	Yes	Yes
Tax Instruments x Year Controls	No	No	Yes	Yes
Include Prescription Drugs?	Yes	Yes	Yes	No
Panel B: Annual Insurer Payments				
	(1)	(2)	(3)	(4)
ln(1-MTR)	-2.994***	-2.684***	-2.320**	-2.770***
	(0.788)	(0.832)	(0.931)	(0.909)
ln(ATI)	-0.156	-0.201	2.932*	2.155
	(0.847)	(0.872)	(1.514)	(1.889)
Per Capita Change	-633.85***	-585.67***	-520.93**	-479.90***
with No Tax Subsidy in 2012	(121.81)	(137.89)	(158.67)	(120.34)
Total Change with No Tax Subsidy	-92***	-85***	-76**	-70***
in 2012 (in billions)	(18)	(20)	(23)	(18)
Family Characteristic Controls	Yes	Yes	Yes	Yes
Individual Characteristic Controls	No	Yes	Yes	Yes
Tax Instruments x Year Controls	No	No	Yes	Yes
Include Prescription Drugs?	Yes	Yes	Yes	No

Notes: Significance levels represented by * 10%, ** 5%, *** 1%. Standard errors are adjusted for within-household clustering. Other variables included but not displayed: age group-gender interactions, marital status-gender interactions, education-gender interactions, family size indicators, and year fixed effects. The covariates listed are determined by the head-of-the-household. The same covariates are included at the individual-level in Columns (2)-(4). Columns (3) and (4) include additional controls based on the mean level of the instruments for each individual interacted with year indicators. The "Total Change" metrics use the MEPS household-level weights to scale up to national estimates. Households with no workers are assigned a marginal net-of-tax rate equal to one.

References

Abelson, Reed, "High-End Health Plans Scale Back to Avoid Cadillac Tax," *New York Times*, May 2013.

Arrow, Kenneth, "Welfare analysis of changes in health coinsurance rates," in "The role of health insurance in the health services sector," NBER, 1976, pp. 1–34.

Auten, Gerald and Robert Carroll, "The effect of income taxes on household income," *Review of Economics and Statistics*, 1999, *81* (4), 681–693.

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein, "Behavioral Hazard in Health Insurance," *The Quarterly Journal of Economics*, 2015, *1623*, 1667.

Blomqvist, Åke, "Optimal non-linear health insurance," *Journal of Health Economics*, 1997, *16* (3), 303–321.

Burns, Sarah K and James P Ziliak, "Identifying the elasticity of taxable income," *The Economic Journal*, 2016.

Claxton, Gary, Matthew Rae, Michelle Long, Nirmita Panchal, Anthony Damico, Kevin Kenward, and Heidi Whitmore, "Employer Health Benefits: 2015 Annual Survey," Technical Report, The Kaiser Family Foundation and Health Research & Educational Trust 2015.

Cogan, John F., R. Glenn Hubbard, and Daniel P. Kessler, "The Effect of Tax Preferences on Health Spending," *The National Tax Journal*, 2011, *64* (3), 795.

Congressional Budget Office, "Federal Subsides for Health Insurance Coverage for People Under Age 65: 2016 to 2026," March 2016.

_ , "Private Health Insurance Premiums and Federal Policy," February 2016.

Feenberg, Daniel and Elisabeth Coutts, "An introduction to the TAXSIM model," *Journal of Policy Analysis and Management*, 1993, *12* (1), 189–194.

Feldstein, Martin and Bernard Friedman, "Tax subsidies, the rational demand for insurance and the health care crisis," *Journal of Public Economics*, 1977, *7* (2), 155–178.

Feldstein, Martin S, "The Welfare Loss of Excess Health Insurance," *The Journal of Political Economy*, 1973, *81* (2 Part 1), 251–280.

Finkelstein, Amy, "The effect of tax subsidies to employer-provided supplementary health insurance: evidence from Canada," *Journal of Public Economics*, June 2002, *84* (3), 305–339.

Gruber, Jon and Emmanuel Saez, "The elasticity of taxable income: evidence and implications," *Journal of Public Economics*, April 2002, *84* (1), 1–32.

Gruber, Jonathan, "Taxes and health insurance," in "Tax Policy and the Economy, Volume 16," MIT Press, 2002, pp. 37–66.

_ , "The cost implications of health care reform," *New England Journal of Medicine*, 2010, *362* (22), 2050–2051.

_ , "The Tax Exclusion For Employer-Sponsored Health Insurance," *National Tax Journal*, 2011, *64* (2), 511–30.

_ **and James Poterba**, "Tax Incentives and the Decision to Purchase Health Insurance: Evidence from the Self-Employed," *The Quarterly Journal of Economics*, August 1994, *109* (3), 701–33.

_ **and Larry Levitt**, "Tax Subsidies for Health Insurance: Evaluating the Costs and Benefits," *Health Affairs*, January/February 2000, *19* (1), 72–85.

_ **and Michael Lettau**, "How Elastic is the Firm's Demand for Health Insurance," *Journal of Public Economics*, 2004, *88*, 1273–1293.

Herring, Bradley and Lisa Korin Lentz, "What Can We Expect from the Cadillac Tax in 2018 and Beyond?," *Inquiry*, 2011, *48* (4), 322–337.

Holmer, Martin, "Tax policy and the demand for health insurance," *Journal of Health Economics*, 1984, *3* (3), 203–221.

Long, James E and Frank A Scott, "The income tax and nonwage compensation," *The Review of Economics and Statistics*, 1982, pp. 211–219.

Manning, Willard G and John Mullahy, "Estimating log models: to transform or not to transform?," *Journal of Health Economics*, 2001, *20* (4), 461–494.

_ **, Anirban Basu, and John Mullahy**, "Generalized modeling approaches to risk adjustment of skewed outcomes data," *Journal of Health Economics*, 2005, *24* (3), 465–488.

Mirrlees, James A, "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, April 1971, *38* (114), 175–208.

Mrkvicka, Neil, Justin Held, Julie Stich, and Kelli Kolsrud, "2013 Employer-Sponsored Health Care: ACA's Impact," 2013.

Mulvey, Janemarie, "Health Care Flexible Spending Accounts," Technical Report, Congressional Research Service 2012.

Nichols, Austin, "Regression for nonnegative skewed dependent variables," in "BOS10 Stata Conference," Vol. 2 Stata Users Group 2010, pp. 15–16.

Nyman, John A, "The value of health insurance: the access motive," *Journal of Health Economics*, 1999, *18* (2), 141–152.

Pauly, Mark V, "Taxation, health insurance, and market failure in the medical economy," *Journal of Economic Literature*, 1986, *24* (2), 629–675.

Phelps, Charles, "Demand for health insurance: a theoretical and empirical investigation," 1973.

Phelps, Charles E, "Large-scale tax reform: The example of employer-paid health insurance premiums," *University of Rochester Center for Economic Research Working Paper*, 1986, (37).

Piotrowski, Julie, "Health Policy Brief: Excise Tax on 'Cadillac' Plans," *Health Affairs*, September 2013.

Powell, David and Hui Shan, "Income taxes, compensating differentials, and occupational choice: How taxes distort the wage-amenity decision," *American Economic Journal: Economic Policy*, 2012, *4* (1), 224–247.

Saez, Emmanuel, Joel Slemrod, and Seth H Giertz, "The elasticity of taxable income with respect to marginal tax rates: A critical review," *Journal of Economic Literature*, 2012, pp. 3–50.

Silva, JMC Santos and Silvana Tenreyro, "The log of gravity," *The Review of Economics and Statistics*, 2006, *88* (4), 641–658.

Smith, Jessica C. and Carla Medalia, *Health Insurance Coverage in the United States: 2014*, Washington D.C.: U.S. Government Printing Office, 2015.

Spence, Michael and Richard Zeckhauser, "Insurance, information, and individual action," *The American Economic Review*, 1971, *61* (2), 380–387.

Turner, Robert W, "Are taxes responsible for the growth in fringe benefits?," *National Tax Journal*, 1987, pp. 205–220.

Vroman, Susan and Gerard Anderson, "The effect of income taxation on the demand for employer-provided health insurance," *Applied Economics*, 1984, *16* (1), 33–43.

Weber, Caroline E, "Toward obtaining a consistent estimate of the elasticity of taxable income using difference-in-differences," *Journal of Public Economics*, 2014, *117*, 90–103.

Woodbury, Stephen A, "Substitution between wage and nonwage benefits," *The American Economic Review*, 1983, *73* (1), 166–182.

Wooldridge, Jeffrey M, *Econometric analysis of cross section and panel data*, MIT press, 2010.

Zeckhauser, Richard, "Medical insurance: A case study of the tradeoff between risk spreading and appropriate incentives," *Journal of Economic Theory*, 1970, *2* (1), 10–26.