

Optimal Health Insurance and the Distortionary Effects of the Tax Subsidy

David Powell*

December 2, 2012

Abstract

This paper introduces a model of optimal health insurance. This model provides theoretical guidance of the relationship between household preferences, cost-sharing, and premiums. I apply this model to understand how the income tax subsidy distorts optimal cost-sharing in health insurance. Typically, insurance protects individuals from financial risk. Health insurance plans, however, are frequently designed to provide coverage at non-catastrophic levels of financial loss. The presence of a health insurance subsidy in the United States tax code, which enables individuals to pay premiums in pre-tax dollars, encourages the purchase of more generous health insurance plans. Little is known about how the tax subsidy affects preferences for the structure of cost-sharing in private plans. This model illustrates how the tax subsidy can distort the optimal cost-sharing schedule. The model is tested empirically using claims data in the Medical Expenditure Panel Survey and a regression discontinuity strategy that uses discrete changes in the marginal tax rate at the Social Security taxable maximum for identification.

Keywords: Optimal Health Insurance, Income Taxes, Cost Sharing

JEL classification: H24, H31, I13

*RAND, dpowell@rand.org.

1 Introduction

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. In this paper, I introduce a model of optimal health insurance which shows the relationship between premiums, cost-sharing, and medical care consumption. The literature provides little guidance concerning the optimal design of insurance plans, and this model should be useful broadly in the health economics field. I use this framework to illustrate the mechanisms through which incentives in the tax code can alter the optimal cost-sharing structure of health insurance plans. The income tax subsidy affects the structure of cost-sharing in health insurance plans, directly impacting medical care utilization, medical costs, and premiums. Intuitively, we know that the tax subsidy should promote the purchase of more generous health insurance, but the literature offers little guidance about what “more generous” implies. 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.

I introduce a model that is similar to the Mirrlees [1971] model of optimal income taxation. To my knowledge, this is the first comprehensive model defining a household’s optimal health insurance plan. It includes moral hazard, uncertainty, and household preferences. The result is an equation of an optimal health insurance plan. The model allows a household to design a health insurance plan defined by the relationship between medical expenditures and out-of-pocket payments for the entire possible distribution of medical costs. The household is responsible for paying a premium equal to the insurer’s expected payouts. The insurer-household contract is based solely on the observable variable of medical expenditures. Households value medical care in any possible “health state” but relatively more in unhealthy states. The household values protection from large financial losses and insurer payouts. However, lower coinsurance rates increase the premium. To commit to consumption of lower medical care, the household may actually demand higher coinsurance rates. Moral hazard is modeled as a series of incentive compatibility constraints. Coinsurance rates must be high enough that households find it utility-maximizing to consume less medical care. It is straightforward in this framework to include the tax subsidy and understand how subsidizing the premium changes demand for cost-sharing throughout the entire distribution of possible

medical expenditures.

Many possible reasons have been given for the demand for generous health coverage beyond protection from large financial losses. The presence of a health insurance subsidy in the United States tax code, which enables individuals to make premium payments in pre-tax dollars, encourages people to purchase more generous health plans than they would in the absence of such a subsidy. The tax subsidy creates a direct relationship between marginal tax rates and the effective price of insurance. Higher marginal tax rates lower the relative cost of health insurance. A family facing a high marginal tax rate may be willing to purchase an expensive health plan which begins payments at the first dollar of medical care consumption. Instead of only protecting themselves against large annual expenditures, households may want to “insure” against even low levels of medical expenditures.

Labor earnings are subject to several types of taxes, including federal, state, and Federal Insurance Contribution Act (FICA) taxes. The employee and employer contributions to purchase health insurance can be paid in pre-tax dollars. Consequently, for every \$1 spent on health insurance, the individual’s after-tax income decreases by only $\$(1 - \tau)$, the marginal net-of-tax rate. Because the premium is related to expected medical costs, the subsidy allows households to essentially pre-pay for medical expenditures in pre-tax dollars. According to a 2007 report by the Congressional Budget Office, the tax subsidy cost the United States Treasury over \$200 billion (note 1 in Auerbach and Hagen [2007]).

The tax subsidy could potentially be changing health insurance plans in a way which encourages high costs with relatively little impact on the overall wellbeing of households. The tax subsidy distorts household preferences for insurance coverage beyond simple protection from large financial losses. Without the tax subsidy, a household should be indifferent between (a) paying an insurer a certain amount and subsequently receiving medical care equal to that exact amount and (b) paying that exact amount for the medical care itself. The tax subsidy distorts this indifference by making premium payments to insurers cheaper in relative terms. Health insurance purchases are, of course, more complicated than this one decision since medical care consumption for the year is unknown at the time of insurance purchase and because the design of the plan can itself affect consumption of care. However, there is still an optimal tradeoff between high payments to the insurer (premiums) and high payments when actually receiving care (high cost-sharing), and the tax subsidy distorts this tradeoff.

The Affordable Care Act (ACA) makes many important changes to the United States

health care system, including new regulations of health insurance markets. These regulations seek to change health insurance plans by mandating coverage of preventive care or enforcing minimum loss ratios. The structure of health insurance plans and types of plans that we see in the market, however, will potentially experience some of its greatest change due to regulations on so-called “Cadillac” health insurance plans. In 2018, the most generous health insurance plans will be subject to taxes. The purpose of this tax is to gradually (as health insurance premiums rise over time) though imperfectly phase out the tax subsidy for health insurance in the United States tax code. We have little theoretical guidance in the literature or empirical estimates of how this Cadillac tax will change preferences for cost-sharing in health insurance.

Since health insurance premiums are a function of expected insurer payments, the tax subsidy should lead to higher premiums. The model illustrates the relationship between the tax subsidy and the premium. There is a direct tradeoff between more generous plans and lower premiums. Plans with low coinsurance rates have higher premiums for two important reasons. First, for any given level of medical care consumption, the insurer must pay more. A risk neutral insurer increases the premium in response. Second, low coinsurance rates encourage additional medical care consumption. Since the insurer is responsible for a fraction of these expenditures, the premium must rise even more to compensate the insurer ex-ante. According to a recent survey, health insurance premiums increased 9% between 2010 and 2011 (Claxton et al. [2011]). The model in this paper illustrates how the tax subsidy can increase premiums.

I make some simple calculations using the 2009 Medical Expenditure Panel Survey for suggestive evidence about the generosity of health insurance in the United States. Using claims records in the Household Component Event files (data details explained later), I look at each person’s *first* medical claim in the year, excluding claims with any public component such as Medicaid, Veterans Administration, etc. The comparison is out-of-pocket payments versus private insurance payments. Even for the first claim, the average percentage of the claim that is covered by private health insurance is 62%. The large magnitude may be potentially due to the inclusion of very large claims. However, even when I select on first claims which are less than or equal to \$100, private insurance covers 53% of the medical costs. While there are many caveats and possible explanations to these numbers, they illustrate that private insurance plans can be very generous at even low levels of medical expenditures.

This paper aims to understand the distortion caused by the tax subsidy in two ways. First, I introduce a new model of health insurance plan design which allows individuals to design an “optimal” cost-sharing structure. The model shows how a tax subsidy can distort the structure of this optimal plan. Second, I test the conclusions of this model empirically using claims data in the Medical Expenditure Panel Survey (MEPS), comparing the out-of-pocket expenditures relationship with total expenditures for households with different marginal tax rates. Because the tax rate is potentially correlated with many characteristics that could also affect health insurance generosity preferences, I employ a regression discontinuity approach that exploits the change in the marginal tax rate at the Social Security taxable earnings maximum.

The next section discuss the related literature. Section 3 introduces and details the model. The empirical component of the paper is provided in section 4. Section 5 concludes.

2 Literature

The model in this paper builds on a foundation found in the income tax literature. Mirrlees [1971] derives formulae for optimal income taxation. In this model, there is a distribution of ability throughout the population, which is unobservable to the social planner or government. The government can only observe and tax labor supply. To prevent high-ability types from pretending that they are low-ability types by providing less labor supply, the government must satisfy incentive compatibility constraints. The Mirrlees [1971] model maximizes total utility in the society subject to a social budget constraint and incentive compatibility constraints for each ability type in the economy. The solution provides an optimal non-linear income taxation schedule. The resulting equations, however, are hard to interpret without assumptions on the utility function. Diamond [1998] assumes a separable utility function that is linear in consumption and provides some intuitive theoretical results. Saez [2001] extends the Mirrlees result by translating it into elasticities.

The model presented in this paper is similar to the one found in Mirrlees [1971]. Instead of a distribution of people with different ability levels, the model pertains to a household with a distribution of possible health types for the upcoming year. The health insurance plan provides payments based on the level of medical expenditures consumed since the insurer does not observe the actual health type of the household. Thus, the insurer and

household can only contract on medical expenditures. The household must pay a premium equal to the expected value of payments received from the insurer. This provides an incentive for the household to commit to lower medical consumption through higher coinsurance rates. Finally, a major departure from the Mirrlees model is that the budget constraint is not fixed. I include the tax subsidy as a reduction in the effective price of the premium to the household. This reduction shifts the household budget constraint, but the magnitude of the budget constraint shift is proportional to the premium chosen by the household. I solve this model and show that the coinsurance rate is dependent on the subsidy. Furthermore, the effect of the subsidy on the coinsurance rate varies throughout the medical expenditure distribution.

A small literature studies the relationship between the tax subsidy and health insurance coverage. Gruber and Poterba [1994] study the effect of the tax subsidy on the probability that an individual is insured. The Tax Reform Act of 1986 (TRA86) introduced a tax subsidy for self-employed individuals, allowing the authors to use a differences-in-differences approach to estimate the effect of this reform on coverage. 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 the other provinces did not. Finkelstein [2002] uses this shock to estimate the elasticity of health insurance coverage and finds a large effect.

Gruber and Lettau [2004] look at the effect of the tax subsidy on the provision of health insurance, measured by the firm-level total costs. They find that higher subsidies increase the amount spent by firms on health costs. Using a microsimulation model, Gruber and Levitt [2000] estimate that the tax subsidy decreases the fraction of people that are uninsured but at a substantial cost.

Cogan et al. [2011] uses the MEPS to estimate the effect of the tax subsidy on total medical care expenditure consumption. They use the discontinuity at the Social Security maximum taxable earnings for identification. People directly below the maximum face a higher marginal tax rate and, consequently, larger tax subsidy than people above the maximum. The paper employs a regression discontinuity design and estimates a very large effect. A higher tax rate leads households to consume significantly more medical care.

While the literature has recognized that the tax subsidy meaningfully affects health insurance coverage rates, there is little work on how the tax subsidy changes the structure

of health insurance plans. The papers above can be summarized as studying how the tax subsidy affects broad measures of health insurance status or proxies for generosity. If the tax subsidy causes people or firms to spend more on health insurance, then the purchased health insurance plans is more generous. A primary motivation of this paper is to look within the “black box” to begin to understand how the tax subsidy affects costs and medical consumption. More expensive health insurance should provide more generous coverage, but we have little theoretical insight or empirical evidence about how the tax subsidy directly impacts cost-sharing within private insurance plans. Intuitively, we can infer that a more generous plan resulting from the tax subsidy should have less cost-sharing, but we do not have a framework to understand at what levels of medical expenditures that cost-sharing is affected the most.

This paper studies this point in two ways. First, I model the structure of health insurance. This model is new to the health economics literature and provides a useful framework for understanding optimal health insurance design. The framework allows me to model cost-sharing as a function of the marginal tax rate. Second, I empirically compare the out-of-pocket expenditures in the MEPS throughout the medical expenditure distribution for households with different tax rates. I use a regression discontinuity design in the spirit of the empirical strategy found in Cogan et al. [2011]. The theoretical and empirical evidence in this paper shows how the tax subsidy affects cost-sharing, which directly impacts medical care consumption and premiums.

3 Theory

3.1 Motivating Example

Since I will include the tax subsidy in the model, it is helpful to discuss some intuition for how the tax subsidy can impact the structure of cost-sharing before delving into the full model. As an example, assume that the household knows with certainty that it will consume exactly \$100 in medical care in the year. Also assume that there is no tax subsidy for health insurance. In this case, the household is indifferent between (a) paying a \$100 premium for full insurance and (b) paying \$0 for no insurance.

Now add the tax subsidy. With a tax subsidy, the household prefers paying the

insurer a \$100 premium because the after-tax price to the household is equal to $\$100(1 - \tau)$.

In practice, households rarely predict their exact annual medical care. Furthermore, health insurance generosity can induce households to consume additional care through moral hazard. Both of these factors obfuscate the conclusions of the above example. However, the broader point offers some insight into the possible effects of the tax subsidy. Each household has some optimal relationship between low cost-sharing (with high premiums) and low premiums (with more cost-sharing). High cost-sharing plans offer households insurance against medical risk. Low cost-sharing plans reduce moral hazard and, consequently, premiums. Without the tax subsidy, a household should be indifferent between payments to the insurance company and payments at the time of receipt of medical care. The tax subsidy distorts this optimal relationship.

The literature has recognized that the tax subsidy encourages purchasing more generous insurance. More generous insurance raises premiums for two reasons. First, the insurance company pays more for a given level of total medical costs. Second, generous coverage promotes additional medical care consumption, further raising insurer payouts. A risk-neutral insurer will charge higher premiums for more generous plans to account for both of these factors. To keep its premiums down, a risk-averse household will likely forgo generous coverage for low levels of medical expenditures. The household wants insurance from large financial risks and protection from significant annual medical costs, but is less concerned with “healthy states” where they require less medical care. In healthy states, the household wants to commit to lower medical care consumption. They can do so through high coinsurance rates at low levels of expenditures. I model “moral hazard” as a series of incentive compatibility constraints. The premium is equal to the expected insurer payouts. The household, then, must commit to a certain level of medical care consumption in each possible health state. The incentive compatibility constraints are this commitment device and are operationalized through the coinsurance rates.

The tax subsidy distorts the optimal cost-sharing structure. As in the example, there is an incentive to pre-pay for medical care by using pre-tax dollars. The higher the marginal tax rate, the greater of an incentive that the household has to purchase insurance with generosity at modest levels of care.

3.2 Model

I present a model for the household's optimal structure of cost-sharing in health insurance. The literature provides little guidance about the optimal structure of cost-sharing in health insurance. In this paper, I build on the foundation of Mirrlees [1971] which models optimal income tax structure. I apply the same principles to model the optimal health insurance plan, where I define "optimal" as the plan structure which maximizes the household's expected utility. The household maximizes its expected utility subject to several constraints. Each household has a distribution of possible health risk for the upcoming year. The health insurance plan provides payments based on the level of medical expenditures consumed since the insurer does not observe the actual health type of the household. The household must pay a premium equal to the expected value of payments received from the insurer. Thus, the insurer and household can only contract on medical expenditures. This provides an incentive for the household to commit to lower medical consumption through higher coinsurance rates.

Once a plan is designed and agreed upon, 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 requires a series of incentive compatibility constraints. This ensures that for any given health type, there is a level of consumption and medical care that maximizes utility in that state. To understand why these constraints are necessary, note that a household cannot design a plan with free medical care but promise the insurer that they will only consume modest amounts of medical care in healthy states. Insurers do not observe the realized health state so households could simply "pretend" to be in a poor health state. Instead, households must design a plan which forces them to commit to lower levels of medical care in those healthy state. They do this through higher coinsurance rates which ensure that high levels of medical care are not utility-maximizing in those states. Otherwise, the insurer predicts high expected level of expenditures and demands a high premium.

I include the tax subsidy, modeled as a reduction in the premium, directly into the framework. The optimal health insurance plan will correspond to the $\tau = 0$ case. Premium payments are effectively discounted by the marginal tax rate when compared to medical payments made during the year and other consumption. Consequently, higher premiums actually shift the budget constraint outward. I will now discuss details of the model.

3.2.1 Household Risk and Utility

This model allows a household to design its optimal health insurance plan with a risk-neutral insurer that charges a premium equal to its expected payouts to the household. In this model, a risk-averse household has a distribution of possible health types for the year, $\theta \sim F [0, \infty)$, where higher θ refers to a sicker household. This distribution is known to a risk-neutral insurer. Once a health insurance contract is agreed upon, the household learns its health type for the year, but the insurer does not observe the type. The insurer can only observe and contract on medical expenditures, m . The insurer provides payments to the household based on 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

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

The household values consumption (c) and health (h). I define health as a function of health type and medical care. In practice, I model this as $h \equiv m - \theta$ so that the utility function can be represented as $U(c, m - \theta)$. Under an assumption that $U_{hh} < 0$, this parametrization implies that high θ (sicker) households receive more utility from additional medical care at any given level of care.

3.2.2 Insurer

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

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

The insurer acts as a means of transferring wealth between different possible states. Note that 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.

3.2.3 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:

$$\max_{c,m} U(c, m - \theta) \quad (3)$$

$$\text{subject to } c + s(m) \leq (w - n)(1 - \tau) \quad (4)$$

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

The FOC of this maximization problem is:

$$\frac{\partial s}{\partial m} = \frac{U_h}{U_c} \quad (5)$$

This result is useful in the solution of the final model.

3.2.4 Incentive Compatibility Constraints

If the insurer could observe and contract on θ , then the household would purchase full insurance for every possible θ due to concavity of the utility function. Because the insurer cannot observe θ , the household has an incentive to “pretend” to be sicker than it is in each given state. As stated before, the household cannot design a plan with low premiums and low co-insurance rates while simply promising the insurer that it will not consume high levels of care in healthy states. This is where moral hazard impacts the model.

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. If additional medical care were free, the household would consume more in each state, raising the premium. Consequently, a higher $\frac{\partial s}{\partial m}$ acts as a commitment device to keep the premium down. There is a tradeoff -

the individual wants to insure against poor health states while committing to consume lower medical expenditures in good health states so as to pay a smaller premium.

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

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

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, or $\frac{\partial s}{\partial m} > 0$.

These constraints are necessary because a household cannot simply promise the insurer that it will consumer specific amounts of medical care for a given, unobservable health state. Instead, the health insurance plan must be designed in a way that the household has no incentive to deviate in any state from consuming $m(\theta)$. This allows the insurer to know the exact possible distribution of medical expenditures and, thus, the expected value of payments.

3.2.5 Budget Constraint

The household is also subject to a budget constraint for the year. This budget constraint differs from equation (4) because it holds before the household learns its risk type for the year. 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 can be rewritten as

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

The premium is tax deductible, which is why n is multiplied by $(1 - \tau)$. When combined with equation (2), this implies that all possible medical expenditures not paid out-of-pocket shift the budget constraint outward by τ . This equation distinguishes the setup

from the Mirrlees model because the budget constraint is not fixed but, instead, dependent on the structure of the health insurance plan. With the tax subsidy, more medical care consumption allows for more “total consumption” ($c + m$).

3.2.6 Non-Negative Out-of-Pocket Payments

Because the budget constraint is not fixed, another constraint is critical. This constraint provides some intuition about how the tax subsidy can encourage demand for generous insurance. 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. This 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. In practice, this constraint can be replaced with $s(m(0)) \geq 0$. The IC constraint ensures $s(m(\theta)) \geq 0$ for $\theta > 0$. Consequently, I ignore this constraint when forming the Hamiltonian.

3.2.7 Final Model

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

$$V'(\theta) = -U_h \tag{6}$$

The final household maximization problem can be written as

$$\max_{V, m} \int_0^\infty V(\theta) f(\theta) d\theta \tag{7}$$

subject to

$$(IC) \quad V'(\theta) = -U_h \quad (8)$$

$$(BC) \quad \int_0^\infty [c(\theta) + (1 - \tau)m(\theta) + \tau s(m(\theta))] f(\theta) d\theta \leq w(1 - \tau) \quad (9)$$

$$(Non-Negative OOP) \quad s(m(0)) \geq 0 \quad (10)$$

Before proceeding to the solution, I discuss an important property.

Lemma 3.1 *If $U_{ch} \geq \frac{U_c U_{hh}}{U_h}$, then $\frac{\partial c}{\partial \theta'} \leq 0$ and $\frac{\partial m}{\partial \theta'} \geq 0$.*

This lemma is a monotonicity condition similar to the Spence-Mirrlees condition. 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 I am assuming that $U_{hh} < 0$. A proof of this lemma can be found in the appendix. This condition is important in interpreting the results because it states that unhealthy households receive less non-medical consumption.

The Hamiltonian is

$$H = V(\theta)f(\theta) + \lambda [w(1 - \tau) - c(\theta) - (1 - \tau)m(\theta) - \tau s(m(\theta))] f(\theta) + \mu(\theta)U_h(c(\theta), m(\theta) - \theta). \quad (11)$$

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

$$\frac{\partial H}{\partial m} = 0 \quad (12)$$

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

$$\mu(0) = \lim_{t \rightarrow \infty} \mu(t) = 0 \quad (14)$$

For ease of interpretation, I assume that $U_{ch} = 0$ for the rest of this paper. The resulting utility function still has many desirable properties. For notational simplicity, define $A(\theta) \equiv \int_\theta^\infty \frac{1}{U_c(t)} f(t) dt$. Note that $A(0)$ is the expected value of $\frac{1}{U_c}$. Let $\nu = \frac{\partial m}{\partial p}$. ν is how responsive the household is to the price of medical care at $m(\theta)$.

Proposition 3.1 *The optimal coinsurance rate at θ is defined by*

$$1 - \frac{\partial s}{\partial m} = \underbrace{\frac{1}{1-\tau}}_{(a)} \underbrace{\frac{U_c}{f(\theta)} \frac{1}{\nu}}_{(b)} \underbrace{\{A(\theta) - [1 - F(\theta)] A(0)\}}_{(c)} \quad (15)$$

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, as hypothesized, has a direct effect on the coinsurance rate. This term is positive assuming that $\tau < 1$.
- (b) I used the household first-order condition to replace U_{hh} with $\frac{U_c}{\nu}$. This term is negative by $U_{hh} < 0$. Note that more probable health states (high $f(\theta)$) are associated with higher coinsurance rates to reduce moral hazard. Similarly, coinsurance rates increase when the household is more responsive to the out-of-pocket price of medical care. Finally, coinsurance rates are lower for larger U_c . Given that $U_{cc} < 0$ and Lemma 3.1 states that $\frac{\partial c}{\partial \theta'} \leq 0$, everything else equal, coinsurance rates are decreasing in θ .
- (c) By Lemma 3.1 and $U_{cc} < 0$, this term is non-positive. $A(\theta)$ is the expected value of $\frac{1}{U_c}$ over $[\theta, \infty)$. Note that $\frac{1}{U_c}$ is decreasing in θ so $A(\theta)$ 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 asymptotes to 0 as $\theta \rightarrow \infty$.

To sign equation (24), I note that (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 so I push this model further. The model suggest that an increase in τ leads to a decrease in the coinsurance rate:

$$\frac{\partial(1 - \frac{\partial s}{\partial m})}{\partial(\frac{1}{1-\tau})} = \frac{U_c}{f(\theta)} \frac{1}{\nu} \{A(\theta) - [1 - F(\theta)] A(0)\} \geq 0. \quad (16)$$

This result also matches our intuition. When tax rates increase, the price of health

insurance coverage relative to other goods decreases. In response, the individual demands more generous coverage, represented here by lower coinsurance rates.

We are also interested in understanding how the coinsurance response to the tax subsidy responds differentially throughout the distribution of possible medical care consumption. Again, note that the (c) term in equation (24) is equal to 0 at both $\theta = 0$ and as θ approaches ∞ , implying that there should be no effect at the top or the bottom of the distribution. Next, note that we have little theoretical guidance about how $\frac{1}{f(\theta)\nu}$ varies over the distribution. Instead, I focus on the terms that provide more theoretical insight.

As θ increases, we know that U_c increases, implying that the coinsurance rate should be decreasing. The $\{A(\theta) - [1 - F(\theta)] A(0)\}$ term is more complicated. Taking the derivative with respect to θ , we arrive at $-\frac{1}{U_c(\theta)} + f(\theta)E\left[\frac{1}{U_c(\theta)}\right]$. At low θ , this term is negative. The term becomes less negative as θ increases and eventually positive. Consequently, $\{A(\theta) - [1 - F(\theta)] A(\underline{\theta})\}$ is always non-positive, decreases until some $\tilde{\theta}$, and increases thereafter. It approaches 0 as $\theta \rightarrow \infty$. This term suggests that the coinsurance rate should be U-shaped. Which of these terms dominates is an empirical question.

3.3.1 Relationship with Motivating Example

Interpreting equation (24) in light of the motivating example offered in section 3.1 is also helpful. Note that the example can be modeled as pertaining to circumstances where a household has positive medical expenditures when $\theta = 0$. In other words, the households always has a certain level of medical expenditures, even in its healthiest possible state. Equation (24) does not explicitly tell us the optimal coinsurance rates for these medical expenditures. However, section 3.1 explains that we should expect to see a large response to the tax subsidy at these levels.

4 Data and Empirical Strategy

4.1 Data

The Medical Expenditure Panel Survey (MEPS) contains an extensive set of demographic, health, and medical variables for the years 1996-2009. The MEPS collects and reports two

especially important sets of variables. First, the MEPS includes the Household Component Full-Year files which have detailed income data, including wage earnings. Wage earnings are necessary to determine whether the individual is above or below the Social Security taxable earnings maximum. The detailed income information is also useful to derive marginal tax rates. I use NBER's Taxsim program (Feenberg and Coutts [1993]), which takes information on different forms of income, number of dependents, and filing status. It provides federal marginal taxes and the marginal FICA tax rates for each household.

Second, the MEPS Household Component Event files have claims data including the breakdown of payments. This allows me to construct total medical expenditures and out-of-pocket payments. My expenditures and out-of-pocket payments variables include the following medical services consumed within a year: inpatient services, outpatient services, emergency room visits, and office-based medical provider visits. I exclude all claims with payments made by a person or organization other than the family or a private insurance company. This selection excludes claims with payments by Medicaid, Medicare, Worker's Compensation, etc.

The model above assumes that a household can design an optimal insurance plan. In practice, of course, this is not the case. To empirically test the model, I consider individuals offered insurance at their current job. While they may not have designed the plan, they can make an optimal choice given their options. Consequently, I select on individuals that report being offered a choice of insurance plans at their current main job. The underlying experiment is to compare the choices that people make, assuming that this approximates their optimal insurance plan. Note that I have avoided modeling the firm-level decision about what types of plan to offer. My empirical strategy will compare people who look similar, except for a discrete change in the marginal tax rate, and study the structure of the plan that they choose. While they are not offered every possible health insurance structure as in the model, I believe that the estimates below are valuable for two reasons. First, while some chosen plans might be "too generous" at certain points of the health risk distribution for an individual, some chosen plans might not be generous enough at those same point in the distribution. The individuals will choose the optimal plan from a given set of choices and while the coverage might deviate from the plan that they would design, it should not do so systematically.

Second, from a policy perspective, we are interested in how the tax subsidy affects actual cost-sharing, not the individual's optimal cost-sharing structure. Thus, the results

below are more policy relevant than estimating changes in optimal plan structure. The model informs us about what households would like to choose, but the estimates themselves tell us the practical impact of these preferences.

I select on individuals making between 60% and 140% of the Social Security taxable earnings maximum in that year. The limit changes annually, steadily growing from \$62,700 in 1996 to \$106,800 in 2009. While this selection may appear to sacrifice representativeness for identification, this population is one that is most likely to purchase very generous plans. Policy has recently been concerned with reducing so-called “Cadillac plans,” making this population especially interesting.

I do not include the person’s family medical expenditures because it is unknown whether the employer offered insurance which covers the entire family. While this is a limitation of the data, the individual’s medical expenditures and out-of-pocket payments should provide evidence of the model. Finally, I exclude people ages 65 and over.

These selections leave me with 9,268 individuals with positive medical expenditures. I merge each person to their claims data. Summary statistics for my sample are provided in Table I. Given the selection based on earnings, the sample has high earnings. The marginal tax rate reported includes both the employer and employee portion of the FICA tax since both parts are relevant to the effective price of the offered insurance plans.

Table I: Summary Statistics

Variable	Mean	Standard Deviation
Total Medical Expenditures	\$3,609.71	\$11,496.94
Out-of-Pocket Payments	\$267.15	\$709.64
Wage Earnings	\$79,262.56	\$17,085.27
Marginal Tax Rate (Federal + FICA)	40.49	5.06
Fraction Below FICA Maximum	83.56%	37.07%
Age	44.63	9.65
Male	57.21%	49.48%

All dollar values in 2008 dollars.

4.2 Empirical Strategy

The underlying experiment is to find people with different marginal tax rates and observe differences in the structure of their cost-sharing. Households with higher marginal tax rates

face a reduced effective price for a given insurance plan. Consequently, I can estimate the impact of the tax subsidy by using variation in tax rates since the subsidy impacts households differentially based on their tax rate. Tax rates, however, are not random and, generally, increase with income. I employ a strategy similar to the one found in Cogan et al. [2011] which uses the tax discontinuity at the Social Security taxable earnings maximum. At this maximum, the marginal tax rate on labor earnings drops by 12.4 percentage points (employer plus employee contribution) during the years in my sample. Cogan et al. [2011] use a regression discontinuity design to estimate the effect of tax rate on total medical expenditures, arguing that the health risk of households on either side of the discontinuity is the same, conditional on earnings. They find that families directly above the maximum consume less medical care than families directly below the maximum.

The implicit reason for this relationship is that the tax subsidy is encouraging the purchase of more generous plans, but we know very little about the mechanisms causing these higher expenditures. This paper looks inside this “black box” and estimates the effect of the tax subsidy on the structure of cost-sharing in private health insurance plans.

Importantly, the MEPS allows me to observe the relationship between total expenditures and the out-of-pocket payments of the household, and I can do so for households with varying marginal tax rates. In the MEPS, I do not observe the structure of the chosen plan or any of the offered plans. However, it is not clear that knowing details about the plans themselves would be very helpful for two reasons. First, a household may not believe that it has a meaningful chance of consuming certain levels of medical care. For example, a family that knows that it will consume substantial amounts of medical care likely care about their out-of-pocket costs at high levels of medical care consumption with little thought about the out-of-pockets costs for the first \$50 of total expenditures. By using the actual distribution of expenditures, I account for the parts of the plans that families actually use, properly weighting households by their actual consumption of care. Second, health insurance plans are complicated and it can be difficult to predict out-of-pockets payments for a given amount of medical expenditures. I use actual payments and expenditures so that such predictions are not necessary.

The strategy, then, is to compare out-of-pockets payments for a given level of medical care consumption for households facing different marginal net-of-tax rates. Notice that this comparison provides information about the response of the *average* cost of medical care, not the marginal cost. While the model highlights that the marginal cost is the important

parameter, studying the average cost is more practical in the data. Furthermore, I am estimating the response of the average cost throughout the expenditure distribution. The response of the marginal and average costs are directly related and follow the same trajectory. A decreasing marginal price implies that we should observe a decreasing average price. Little is lost by discussing the average cost response and the insights of the model remain the same. Since equation (16) specifies $(1 - \frac{\partial s}{\partial m})$, I use $\ln(1 - \frac{s}{m})$ as my outcome variable.

This outcome metric is “censored” at both $-\infty$ and 0 since a household cannot pay more than it consumes and cannot pay less than zero. My outcome of interest, then, is

$$y_i \equiv \max \left[\ln \left(1 - \frac{\text{OOP}_i}{m_i} \right), C \right] \quad (17)$$

where OOP_i represents the out-of-pocket payments of household i and m_i remains the total medical expenditures for the household. Because of the censoring, I use a median regression approach since the median is robust to censoring at the tails. Given a sample of households with medical expenditures m , I could model the median in the following manner:

$$Q_y(0.5|\tau, L) = \gamma_t + \beta_m \ln(1 - \tau) + f(L) \quad (18)$$

where L is labor earnings as a percentage of the FICA maximum in that year and τ is the marginal tax rate. I use $\ln(1 - \tau)$, which is commonly used in the tax literature (see Gruber and Saez [2002] for one example). The model states that the relevant variable is $\frac{1}{1-\tau}$ so it is important to note that $\ln(\frac{1}{1-\tau}) = -\ln(1 - \tau)$. Thus, the model suggests that an increase in the marginal net-of-tax rate should decrease the average insurer payment. I could estimate the above specification with an IV strategy where I use the regression discontinuity design previously discussed, comparing people directly above and directly above the FICA maximum.

In practice, I do not observe people at each point in the medical expenditure distribution. Instead, I utilize a local IV quantile estimation strategy. For each level of medical expenditures, m , I create weights for each observations based on its distance from m and the FICA maximum ($= 100$ since I express all labor earnings as a percentage relative to the FICA maximum):

$$w_i = d(L_i - 100, m_i - m)$$

The final specification of interest controls for labor earnings, allowing for the effect

to differ above and below the FICA discontinuity, to implement the regression discontinuity design. I also let the effect differ based on medical expenditures relative to m . It is also important to interact the labor earning variables with the medical expenditure variable. Thus, I estimate the following specification

$$\begin{aligned}
Q_y(0.5|\tau, L_i) = & \gamma_t + \beta_m \ln(1 - \tau_i) + \alpha_m \ln(1 - \tau_i) \times (m_i - m) \\
& + \delta_m^1 (L_i - 100) \times \mathbf{1}(L_i < 100) \\
& + \delta_m^2 (L_i - 100) \times \mathbf{1}(L_i \geq 100) \\
& + \theta_m^1 (L_i - 100) \times \mathbf{1}(L_i < 100) \times (m_i - m) \\
& + \theta_m^2 (L_i - 100) \times \mathbf{1}(L_i \geq 100) \times (m_i - m) \quad (19)
\end{aligned}$$

While this equation may look confusing, when one evaluates at $m_i = m$ and $L_i = 100$, we are left with

$$Q_y(0.5|\tau, L_i) = \gamma_t + \beta_m \ln(1 - \tau_i). \quad (20)$$

The other variables control for differences related to variation in labor earnings and medical costs. I use an IV-GMM framework with moment conditions

$$g_i(\beta_m) = w_i Z_i [\mathbf{1}(y_i \leq Q_y(0.5|\tau, L_i)) - 0.5] \quad (21)$$

where Z_i includes the labor earnings variables, the medical expenditure variables, their interactions, and two instruments. I consider the tax variable and the tax variable interacted with the medical care variable as endogenous. The two instruments account for this endogeneity. First, I instrument with a dummy variable for “below the FICA maximum.” This variable should shock the marginal net-of-tax rate, conditional on labor earnings. Second, I interact the “below the FICA maximum” dummy variable with $(m_i - m)$. This instrument is necessary to help identify the α_m term. The other variables are treated as exogenous. The weights are used in the moment conditions to weight observations closer to $m_i = m$ and $L_i = 100$ more heavily. Optimization uses simulated annealing. Standard errors are derived through bootstrapping.

Equation (16) offers the clearest interpretation of the results. We are interested in

the elasticity

$$\frac{\partial \ln \left(1 - \frac{\text{OOP}}{m} \right)}{\partial \ln(1 - \tau)}. \quad (22)$$

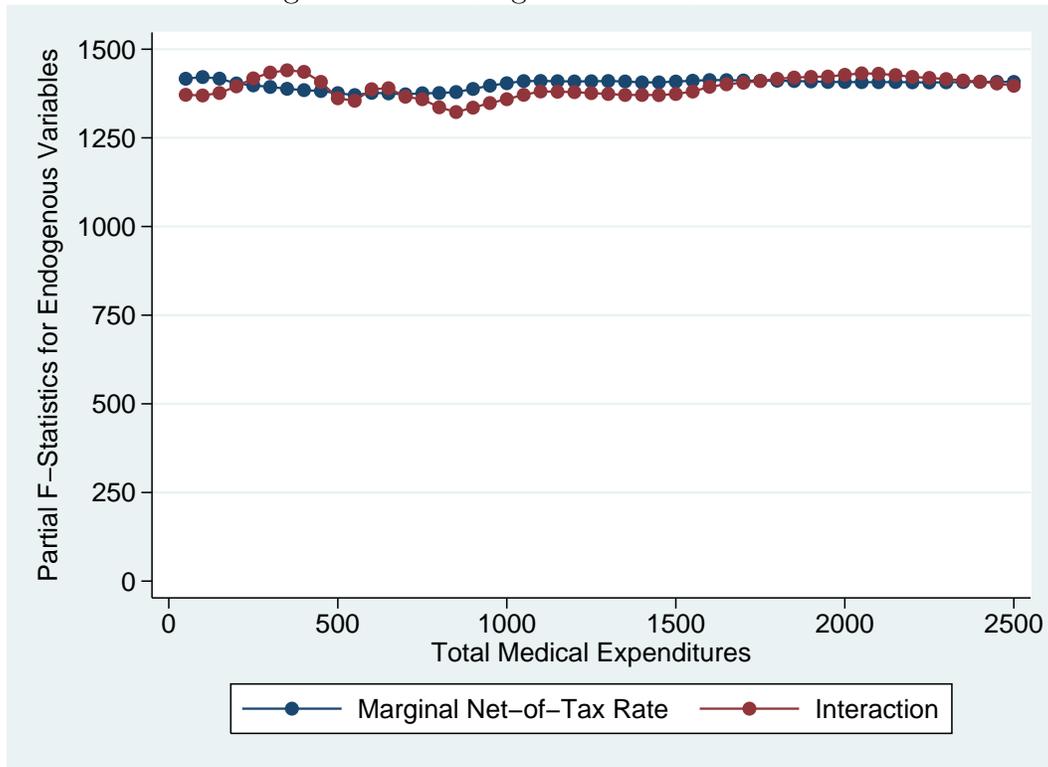
The coefficients estimated using equation (19) provide these estimates.

5 Results

5.1 First Stage

Estimation of equation (19) requires a strong relationship between the instruments and the endogenous variables. The partial F-statistics are very strong for both endogenous variables - the marginal net-of-tax rate variable ($\ln(1 - \tau_i)$) and the interaction term ($\ln(1 - \tau_i) \times (m_i - m)$). Since I estimate (19) for every \$50 of medical expenditures, I report partial F-statistics for both variables at each expenditure level in Figure 1.

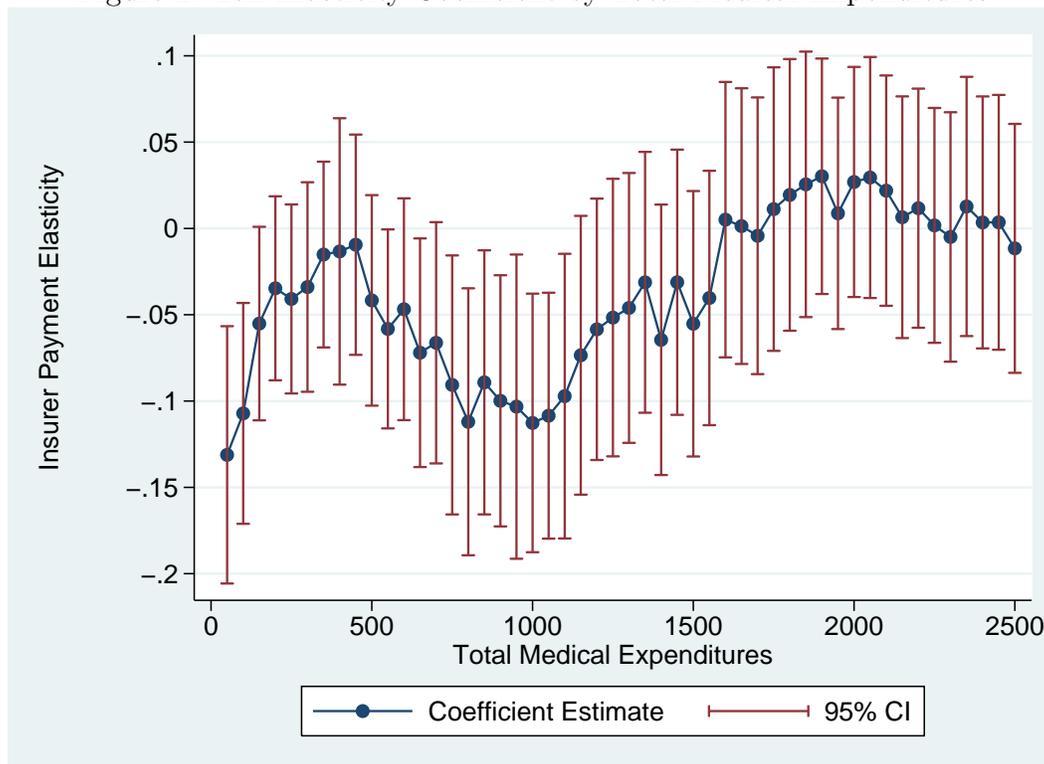
Figure 1: First Stage Partial F-Statistics



5.2 Main Results

The results can be presented in a simple figure. On the x-axis is total medical expenditures (m). On the y-axis is the coefficient estimated in equation (21). This coefficient is the elasticity of the amount paid by the insurer with respect to the marginal net-of-tax rate. As the marginal net-of-tax rate increases, we would expect the insurer to pay less (see equation (16)). Figure 2 presents the results.

Figure 2: Tax Elasticity Coefficient by Total Medical Expenditures



The results present some interesting evidence about the impact of the tax subsidy on cost-sharing in private health insurance plans. First, there is a clear and large effect at the bottom of the medical expenditures distribution, which can be interpreted in the model as discussed in section 3.3.1. At \$50 of medical expenditures, an increase in the marginal net-of-tax rate by 10% leads to a decrease in insurer payments 1.3%. The results are similar at \$100 of medical care, but the elasticity is not significant from 0 at levels directly above that, suggesting that these first \$100 of medical care are “pre-paid” medical expenditures for a fraction of the population. After the elasticity returns close to 0, the estimates provide the U-shape suggested by the model. The elasticity steadily decreases until hovering around

-0.1 between \$750 and \$1100 of medical payments. The estimates then steadily increase, take a large jump upward at \$1600 of medical care, and then bounce around 0 from that point on.

5.3 Discussion

The results above provide complementary evidence to the model. While these estimates may seem small, note that this paper is not interested in how changes in tax rates affect health insurance cost-sharing. Instead, variation in tax rates is used as a proxy for variation in the tax subsidy. This paper is interested in how the tax subsidy affects cost-sharing. Multiplying the estimates by 5 allows for interpretation of these estimates in this regard since elimination of the tax subsidy is equivalent to a 50% reduction in the mean value of the tax subsidy in the sample ($\ln(1 - .4) - \ln(1 - 0) = -.51$).

I should note some limitations of this empirical strategy. Most important, the model discusses the cost-sharing response to the tax subsidy as a function of underlying health risk. This health risk is unobservable in practice so I compared people with the same total medical expenditures. Variation in coinsurance rates due to the tax subsidy combined with moral hazard would suggest that people with the same total medical expenditures and different tax rates are actually different health types. A healthy household in a generous plan (partially due to the tax subsidy) may consume as much care as an unhealthy household in a less generous plan.

While this is a data limitation, the estimates may still be consistent as long as the ratio between out-of-pocket spending and total expenditures is not systematically different between these two households at a fixed level of medical care. Imagine a “typical” plan with a deductible, a coinsurance rate, and a stoploss point. Moral hazard pushes a household further along this nonlinear budget constraint, but their out-of-pocket expenditures would have been the same at medical expenditures m as a household which actually consumes m . In other words, if we could eliminate moral hazard and force the household to consume m , their out-of-pocket spending would be the same as the household actually consuming m . Plans, of course, are much more complicated in practice, but it is not clear that the type of care consumed by this household should be systematically different in terms of the trajectory of cost-sharing. While it would be ideal to observe the θ presented in the model, the results above should still provide parallel empirical evidence related to the model.

Next, tax rates could impact health insurance preferences in other ways than those modeled due to incentives within the tax code. In 2003, Health Savings Accounts (HSAs) were established as an incentive to enroll in High-Deductible Health Plans. High tax rate individuals are most likely to benefit from these accounts. In my data, I do not observe HSAs and, more importantly, I do not observe if out-of-pocket payments are made from these accounts. Consequently, the true out-of-pocket after-tax price may be lower for some medical payments. Furthermore, these incentives could potentially cause high tax rate individuals to enroll in less generous plans than the model would predict, biasing the strategy against finding any effect. Both of these concerns are theoretically important. During the 1996-2009 time period, however, High Deductible Health Plans with Saving Options (HDHP/SO) were relatively rare. Of all covered workers, only 8% had an HDHP/SO in 2009 (Claxton et al. [2011]), the highest percentage during my time period.

The tax code also allows itemizers to deduct out-of-pocket medical expenses that exceed 7.5% of their adjusted gross income for the year. This deduction can be thought of as a type of partial insurance for large financial losses due to medical costs. This insurance is larger for high tax rate households, reducing the benefit of insurance at high levels of annual medical expenditures. The implication for the results above is that the elasticities at the top of the expenditure distribution may be biased upward. High tax rate households could purchase slightly less generous insurance for high medical costs. I only observe out-of-pocket payments in the data, not whether those payments are later deducted.

6 Conclusion

This paper introduces a new model of optimal health insurance design. This framework allows me to study how the tax subsidy distorts cost-sharing, encouraging households to purchase more generous health insurance at relatively modest levels of annual medical care consumption. The model provides evidence that the impact of the tax subsidy should be very large for low levels of medical expenditures and then possibly U-shaped higher in the distribution. An empirical test finds supportive evidence of this theoretical result. The empirical results suggest that the tax subsidy decreases cost-sharing. The model, consequently, shows how this decreased cost-sharing should encourage additional medical care consumption and higher premiums.

A Appendix

Lemma 3.1 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: $\frac{\partial^2 V}{\partial \theta'^2} + \frac{\partial^2 V}{\partial \theta \partial \theta'} = 0$. Consequently, 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:

$$\frac{\partial c}{\partial \theta'} = -\frac{U_h}{U_c} \frac{\partial m}{\partial \theta'} \tag{23}$$

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 (23) then implies that $\frac{\partial c}{\partial \theta'} \leq 0$.

Proposition 3.1 The optimal coinsurance rate at θ is defined by

$$1 - \frac{\partial s}{\partial m} = \frac{1}{1 - \tau} \frac{U_c}{f(\theta)} \frac{1}{\nu} \{A(\theta) - [1 - F(\theta)] A(0)\} \tag{24}$$

Proof:

Equations (12) and (13) lead to

$$\lambda \left[\frac{U_h}{U_c} - (1 - \tau) - \tau \left(\frac{\partial s}{\partial m} \right) \right] f(\theta) = \mu(\theta) U_{hh} \quad (25)$$

$$f(\theta) - \lambda \left[\frac{1}{U_c} \right] f(\theta) = -\mu'(\theta) \quad (26)$$

Deriving $\mu(\theta)$ from (26) and using $\mu(0) = 0$,

$$\begin{aligned} \mu(\theta) &= 1 - F(\theta) - \lambda \int_{\theta}^{\infty} \frac{1}{U_c(t)} f(t) dt \\ \mu(0) &= 1 - \lambda \int_0^{\infty} \frac{1}{U_c(t)} f(t) dt = 0 \end{aligned}$$

We are left with

$$\begin{aligned} \mu(\theta) &= 1 - F(\theta) - \frac{A(\theta)}{A(0)} \\ \lambda &= \frac{1}{A(0)} \end{aligned}$$

I can replace $\mu(\theta)$ and λ in equation (25). I can 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

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

Replacing U_{hh} with $\frac{U_c}{\nu}$ completes the proof.

References

- David Auerbach and Stuart Hagen. CBO's health insurance simulation model: A technical description. Technical report, Congressional Budget Office, 2007. URL <http://www.cbo.gov/sites/default/files/cbofiles/ftpdocs/87xx/doc8712/10-31-healthinsur>
- Gary Claxton, Matthew Rae, Nirmita Panchal, Janet Lundy, Anthony Damico, Awo Osei-Anto, Kevin Kenward, Heidi Whitmore, and Jeremy Pickreign. Employer health benefits: 2011 annual survey. Technical report, The Kaiser Family Foundation and Health Research & Educational Trust, 2011.
- John F. Cogan, R. Glenn Hubbard, and Daniel P. Kessler. The effect of tax preferences on health spending. *The National Tax Journal*, 64(3):795, 2011.
- Peter A Diamond. Optimal income taxation: An example with a u-shaped pattern of optimal marginal tax rates. *American Economic Review*, 88(1):83–95, March 1998.
- Daniel Feenberg and Elisabeth Coutts. An introduction to the TAXSIM model. *Journal of Policy Analysis and Management*, 12(1):189–194, 1993.
- Amy Finkelstein. The effect of tax subsidies to employer-provided supplementary health insurance: evidence from canada. *Journal of Public Economics*, 84(3):305–339, June 2002.
- Jon Gruber and Emmanuel Saez. The elasticity of taxable income: evidence and implications. *Journal of Public Economics*, 84(1):1–32, April 2002.
- Jonathan Gruber and Michael Lettau. How elastic is the firm's demand for health insurance. *Journal of Public Economics*, 88:1273–1293, 2004.
- Jonathan Gruber and Larry Levitt. Tax subsidies for health insurance: Evaluating the costs and benefits. *Health Affairs*, 19(1):72–85, January/February 2000.
- Jonathan Gruber and James Poterba. Tax incentives and the decision to purchase health insurance: Evidence from the self-employed. *The Quarterly Journal of Economics*, 109(3):701–33, August 1994.
- James A Mirrlees. An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38(114):175–208, April 1971.
- Emmanuel Saez. Using elasticities to derive optimal income tax rates. *Review of Economic Studies*, 68(1):205–29, January 2001.