

NBER WORKING PAPER SERIES

THE EFFECT OF TAX PREFERENCES ON HEALTH SPENDING

John F. Cogan
R. Glenn Hubbard
Daniel P. Kessler

Working Paper 13767
<http://www.nber.org/papers/w13767>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2008

We would like to thank Joe Antos, William Gentry, Michael Smart, George Zodrow, two anonymous referees, and seminar participants at Stanford and the Brookings Institution for helpful comments; Peter Banwarth for exceptional research assistance; Kate Bundorf for the idea of using the Community Tracking Study (CTS); and Peter Granda of the ICPSR for help in obtaining the CTS data. Kessler gratefully acknowledges support from the National Institute on Aging through the NBER. Contact authors at fkessler@stanford.edu. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2008 by John F. Cogan, R. Glenn Hubbard, and Daniel P. Kessler. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Effect of Tax Preferences on Health Spending
John F. Cogan, R. Glenn Hubbard, and Daniel P. Kessler
NBER Working Paper No. 13767
January 2008, Revised March 2011
JEL No. H2,I1

ABSTRACT

In this paper, we estimate the effect of the tax preference for health insurance on health care spending using data from the Medical Expenditure Panel Surveys from 1996-2005. We use the fact that Social Security taxes are only levied on earnings below a statutory threshold to identify the impact of the tax preference. Because employer-sponsored health insurance premiums are excluded from Social Security payroll taxes, workers who earn just below the Social Security tax threshold receive a larger tax preference for health insurance than workers who earn just above it. We find a significant effect of the tax preference, consistent with previous research.

John F. Cogan
Hoover Institution
Stanford University
Stanford, CA 94305
cogan@hoover.stanford.edu

R. Glenn Hubbard
Graduate School of Business
Columbia University, 101 Uris Hall
3022 Broadway
New York, NY 10027
and NBER
rgh1@columbia.edu, ws2187@columbia.edu

Daniel P. Kessler
Hoover Institution
Stanford University
434 Galvez Mall
Stanford, CA 94305
and NBER
fkessler@stanford.edu

I. INTRODUCTION

The tax preference for employer-sponsored health insurance, under which employer contributions to employee health insurance are deductible to the employer and non-taxable to the employee, is perhaps the most salient feature of postwar health policy in the United States. By making insured health spending appear less costly than it is, the tax preference gives employees the incentive to take compensation as health insurance rather than cash. This incentive has had two main effects. First, it has increased insurance coverage, particularly coverage through employers (see Jonathan Gruber, 2002, for a review of work identifying this effect). Second, it has distorted the quantity and type of health insurance that people choose in a manner that has increased health spending. At the most basic level, the tax preference increases demand for health services by reducing their price relative to other goods and services. The tax preference also induces people to channel their spending on health services through insurance rather than spend out of pocket. This tax bias favoring the purchase of health insurance leads people to choose health plans with lower deductibles and coinsurance rates, which in turn further increases spending because of the moral hazard inherent in health insurance.¹ Although economists differ about the consequences for social welfare of the effect of the tax preference on coverage, virtually all since Martin Feldstein's (1973) seminal article have agreed that the overspending on health care induced by the tax preference has contributed to the inefficiency of the U.S. health system.

¹ Even though individuals' health plan choices are mediated through employers, there are several reasons to expect a link between an individual's tax rate and his plan's generosity. First, most individuals with employer-sponsored health insurance are offered a choice of health plan. According to the Kaiser Family Foundation/Health Research and Educational Trust (2005), 67 percent of workers in 1996 and 63 percent in 2005 could choose from at least two plans. Second, even employers that offer only one plan choose its parameters based on their workers' preferences. In particular, M. Kate Bundorf (2002) shows that worker characteristics, including income (and therefore tax rate), are statistically significant determinants of plan generosity.

Assessments of the magnitude of the effect of the tax preference on health spending have therefore been at the center of health policy debates for more than 30 years. Yet, for two reasons, few studies have convincingly identified the size of this effect. First, because tax rates depend on income, and income and health spending are jointly determined outcomes of an incompletely observable process, disentangling the effect of the tax preference from other, unobserved determinants of income is difficult. Second, because researchers only observe imperfectly measured proxies for the true post-tax price for insurance, simple regression estimates of the effect of the tax preference are biased toward zero.

In this paper, we use the fact that Social Security taxes are only levied on earnings below a statutory threshold to identify the tax preference's impact. Because employer-sponsored health insurance premiums are excluded from Social Security payroll taxes, workers who earn just below the Social Security tax threshold receive a larger tax preference for health insurance than workers who earn just above it. By comparing health care spending of individuals in families with workers just above the payroll tax threshold with that of individuals in families with workers who are just below it, we can identify the effect of the tax preference for insurance.

Using data from the Medical Expenditure Panel Surveys Household Component (MEPS-HC) from 1996-2005, we estimate a significant impact of the tax preference for insurance on health spending, consistent with results in the previous literature. The paper proceeds in the next five sections. Section II reviews previous work concerned with estimating the effect of tax policy on health spending and describes our methodological approach. Section III discusses the data we use for our analysis, and presents tabular

results which show evidence of a large impact of the Social Security payroll tax on health care spending. People in families with a holder of employer-sponsored insurance who earns just above the tax threshold spend 24 percent less than people in families with an employer-sponsored-insurance policyholder who earns just below the tax threshold. Section IV imbeds the payroll tax in a more general framework for analyzing the impact of tax policy on health care spending and provides an estimate of the tax preference's effect. Section V uses data from the Community Tracking Study from 1996-1998 to investigate the mechanism through which this effect occurs. We find that holders of employer-sponsored health insurance who earn just above the tax threshold have lower premiums, and thus presumably less extensive health insurance coverage, than those who earn just below it. Section VI makes some concluding observations.

II. EFFECTS OF TAX PREFERENCES ON HEALTH SPENDING: PREVIOUS RESEARCH

As Mark Pauly's (1986) classic review highlights, conventional estimates of the effects of tax policy on health spending suffer from two sources of potential bias. First, because marginal tax rates rise with income, and individuals' incomes and health spending are jointly determined by an endogenous process, simple regression estimates of the effect of the tax rate on spending are in general biased and inconsistent. Second, because researchers only observe proxies for the after-tax price of insurance, and these proxies measure the true after-tax price with error, estimates of the effect of the tax rate are biased toward zero. As Pauly points out, these problems are of more than theoretical concern. The wide range of early estimates of the elasticity reported by Pauly, -0.14 to -

1.9, reflect the failure of traditional approaches to deal robustly with these sources of bias.

More recent research has taken two novel approaches to address this shortcoming. One arm of this research uses theoretical models to simulate the effects of tax policy based on existing empirical estimates of the price elasticity of demand for health services.² William Jack and Louise Sheiner (1997) use this parameter as an input to a theoretical model in which consumers choose insurance contracts and then face random health shocks. They calculate how much health spending would be reduced by repeal of the tax preference under various assumptions about individuals' coefficient of relative risk aversion. They also propose conditions under which extending the tax preference to out-of-pocket spending can actually reduce health spending.³ Jonathan Gruber (2002) uses a simulation model that highlights the role of firms' offering decisions to show how spending and employer coverage would change in response to complete and partial repeal of the tax preference. John Cogan, Glenn Hubbard, and Daniel Kessler (2011) show that the effect of the tax preference for insurance on health spending can be written as the product of two parameters: the price elasticity of demand for health services, and the

²The most comprehensive study of the elasticity of demand for health services remains the RAND Health Insurance Experiment, which found that a one percent increase in the out-of-pocket price led to a 0.2 percent decline in spending (Willard Manning, et al., 1987; Joseph Newhouse, 1993). In a more recent study using observational data, Matthew Eichner (1998) found that a one percent increase in the price led to approximately a 0.7 percent decline in spending.

³Expanding deductibility has two opposing effects on health care spending. First, it lowers the overall price of health care relative to other goods and services, thereby increasing health care spending. Second, it raises the price of purchasing health care through insurance relative to out-of-pocket. The second effect induces people to shift to health plans with higher deductibles and coinsurance rates which, in turn, lowers health care spending. Because the base of out-of-pocket spending is so much smaller than that of insured spending, the net change in the overall price of health care from expanding deductibility is much smaller than the change in the relative price of out-of-pocket versus insured care. This fact leads the second effect to dominate the first and, in turn, leads to a net decrease in spending.

effect of the tax preference on the coinsurance rate.⁴ They use the simulation results from Jack and Sheiner (1997) and Martin Feldstein and Bernard Friedman (1977), along with empirical estimates of the price elasticity of demand, to provide a range of estimates of the effect of repealing the tax preference.

A second arm of this research seeks to identify the effects of tax policy empirically using arguably exogenous variation in the tax rate.⁵ Jonathan Gruber and Michael Lettau (2004) estimate the effect of the tax preference using simulated tax rates as instruments for people's actual tax rates, where the simulated tax rates are functions only of interactions between state, year, and an individual's decile in the earnings distribution. They find an elasticity of firm-level insurance spending with respect to the after-tax price of -0.7. Michael Smart and Mark Stabile (2005) examine the consequences of the Canadian Medical Expense Tax Credit (CMETC), which allows taxpayers to claim a credit for 17 percent of medical expenses in excess of \$1,637 (2000 tax year) or 3 percent of net income, whichever is smaller. Qualifying expenses include both insurance premiums and out-of-pocket spending on hospital and physician services, prescription drugs, and long-term care. Using discontinuities in the CMETC formula, and variations in the effective formula across provinces and over time, they find moderate to large tax-price elasticities of demand for health services, between -0.7 and -1.0.

Each of these studies has important limitations. The consistency of Gruber and Lettau's (2004) estimates depend on the validity of their assumption that interactions

⁴ For simplicity, we use the term "coinsurance rate" to refer to all cost-sharing -- that is, payments made under an insurance policy's deductible and its copayment schedule.

⁵ A related set of studies examine the effect of tax policy on firms' decision to offer insurance, employees decision to take up insurance, and insurance plan choice (Jonathan Gruber and Ebonya Washington, 2005; William Jack, Arik Levinson, and Sjamsu Rahardja, 2006). These studies find small elasticities of offer and takeup, but larger elasticities of plan choice, with respect to the tax preference.

between state tax progressivity and individuals' incomes are uncorrelated with individuals' taste for health insurance or medical care. Smart and Stabile's approach requires a similar assumption, and in any event may not be applicable to the effect of the tax preference on health spending in the United States. On the one hand, the CMETC only applies to spending on products and services that are not already covered by publicly financed basic provincial health insurance. Because uncovered services may be more discretionary in nature, their elasticity of spending may be larger than the elasticity of spending on all health care. On the other hand, the CMETC applies to both insured and out-of-pocket spending, while the tax preference in the U.S. applies only to insured spending. Thus, the U.S. tax preference would tend to induce people to shift to health insurance plans with less coinsurance, whereas the CMETC would not.⁶ This effect, in turn, would increase moral hazard and increase health spending, which implies that the elasticity of spending with respect to the CMETC would be smaller than the elasticity of spending with respect to a tax preference for insurance.

As Gruber and Lettau (2004) point out, their estimated elasticities, as well as those of Smart and Stabile (2005), imply that repeal of the tax preference would lead to very large reductions in spending on employer-sponsored insurance. Accurate assessment of the effect of the tax preference is therefore a critical tax policy issue. In this paper, we offer an alternative strategy for identifying the effect of interest. We use the fact the Social Security program limits the annual amount of wage earnings that are subject to Social Security payroll taxation. Under the Social Security program, employers and employees each pay a 6.2 percent payroll tax on earnings below the

⁶ Smart and Stabile show that, in practice, the CMETC does not affect the demand for insurance on the intensive margin.

maximum taxable wage (the "wage base"). The wage base is set by law and is automatically adjusted each year by the average growth in Social-Security covered wages. (Table 1 reports the wage base for each of our sample years 1996-2005.) For workers who earn below the wage base, earnings are subject to Social Security payroll taxes while employer-sponsored health insurance premiums are not; above the wage base, neither earnings nor premiums are subject to tax..

This differential tax treatment of earnings creates a differential tax preference for insurance. Following Gruber and Lettau (2004), we specify the tax price of insurance, TP , as

$$TP = \frac{(1 - \tau - \tau_{ss} - \tau_{mc})}{(1 + \tau_{ss} + \tau_{mc})},^7$$

where τ is the marginal federal income tax rate, τ_{ss} is the Social Security tax rate, and τ_{mc} is the Medicare tax rate. Thus, for workers who earn below the wage base, the tax price of insurance is

$$TP = \frac{(1 - \tau - \tau_{ss} - \tau_{mc})}{(1 + \tau_{ss} + \tau_{mc})},$$

but for those above the wage base, the tax price is

$$TP = \frac{(1 - \tau - \tau_{mc})}{(1 + \tau_{mc})}.$$

In principle, by comparing health care spending of individuals in families with a employer-sponsored insurance policyholder who earns just below the maximum taxable wage to spending by individuals with a policyholder who earns just above this threshold, we can identify the impact of the tax preference on health care spending.

⁷ We compute the marginal federal tax rate by assuming each family takes the standard deduction. We cannot compute the state tax rates because the MEPS does not have information on the state of residence.

III. DATA

We use data from the MEPS-HC from 1996-2005. The MEPS-HC collects data from a sample of families and individuals that is drawn from a nationally representative subsample of households that participate in the National Health Interview Survey. The panel design of the MEPS seeks to collect information on each individual in the sample for two full calendar years. In particular, the MEPS-HC collects the amount of spending on health care services, wage earnings, and various demographic and economic information. The MEPS-HC (through the person-round-plan file) also links each individual with employer-sponsored health insurance to the worker who holds the policy that is the source of coverage. We limit our main analysis to individuals covered all year through the same full-year worker who holds a full-year employer-sponsored health insurance policy. Since in any given year only about 15 percent of all workers have earnings that exceed Social Security's maximum taxable wage, our sample pools together ten years of annual MEPS samples to obtain as many observations on high-wage workers as possible.

A. Results Using Sample Means

Because health care is a normal good, we expect, all else held constant, that health care spending would rise with wages. But if the tax preference provided by the Social Security payroll tax significantly affects health care spending, its impact should be manifested in lower health care spending among individuals in families in which the policyholder earns just above the payroll tax threshold relative to individuals in families in which the policyholder earns just below the threshold. Thus, if we array individuals

according to the earnings of their family's policyholder, we would expect to observe a monotonic relationship between earnings and health care spending across the earnings spectrum, except around the Social Security payroll tax threshold.

Table 2 groups individuals in the 1996-2005 pooled MEPS by the earnings of their family's policyholder, expressed as a percentage of the Social Security wage base in their sample year. The table reports the average health care spending level (in constant 2004 dollars, deflated with the CPI) of all individuals within each wage interval. As the table shows, health care spending generally rises with the earnings level of the family's policyholder. Individuals in families with a policyholder earning less than 70 percent of the Social Security wage base spend on average \$1,953 per year in 2004 dollars. As the wage and salary income of a family's policyholder increases, so does each member's health spending, up to 90-100 percent of the wage base (\$2,406 per year).

However, individuals in families with a policyholder earning 100-110 percent of the wage base spend 23.7 percent less on average -- \$1,836 per year -- than their counterparts in families with a policyholder earning 90-100 percent of the wage base. Given that the tax price of health insurance rises from 0.597 below to 0.692 above the wage base (see Table 3), this implies a raw elasticity of -1.823,⁸ which is highly suggestive of a large impact of the payroll tax preference on health care spending. The fact that as the wage and salary income of a family's policyholder increases beyond 100-110 percent of the wage base, health spending begins rising again, to \$2,110 for persons in families in which the policyholder earns more than 130 percent of the taxable threshold, lends further support to the hypothesis that the drop in health care spending just above the Social Security tax threshold is the result of the tax preference.

⁸ $-1.823 = [(1836 - 2406) / (0.5*(1836 + 2406))] / [(0.692 - 0.597) / (0.5*(0.692 + 0.597))]$.

Table 3 shows that differences in the characteristics of individuals and families are unlikely to account for the differences in spending. The table compares the characteristics from families in which the policyholder earns 90-100 percent of the Social Security taxable wage threshold to the characteristics from families in which the policyholder earns 100-110 percent of the threshold. As the table shows, the characteristics are quite similar, except for their tax price of insurance. Individuals in the higher-income group come from families with only slightly more educated policyholders (15.3 years, as opposed to 15.1 years). The occupation, industry, age, gender, and marital status distributions show no systematic differences that would explain the observed differences in health care spending between the two groups.

The total income of families whose policyholder earns just below the taxable earnings threshold is 8.0 percent (\$8,130) lower than family incomes whose policyholder earns just above the threshold. This difference is not surprising, but, importantly, rules out the likelihood that higher family incomes among families in the lower wage group generated by non-wage sources or other family member earnings explain their higher health care spending levels.

Figure 1 presents two breakdowns of the difference in spending between individuals in families with a policyholder just above and just below the Social Security tax threshold. The top panel of the figure shows that individuals in families in which a policyholder earns 90-100 percent of the wage base spend more on both inpatient and outpatient services (although not on outpatient dental and vision services, which are traditionally not covered by insurance) than individuals in families with a policyholder earning 100-110 percent of the wage base. The bottom panel shows average spending by

wage income quintile, again for individuals in families with a policyholder earning 90-100 versus 100-110 percent of the wage base. It shows that spending is higher in every quintile among those below versus above the wage base, statistically significantly higher for quintiles 1, 3, 4, and 5 ($p < 0.01$ for quintiles 1, 3, and 5, $p = 0.03$ for quintile 4). The difference in spending is most pronounced among individuals in the top quintile: those just below the wage base spent \$8,973 on average as compared to those just above the wage base, who spent \$6,442.

IV. ESTIMATING THE LINK BETWEEN TAX PREFERENCES AND HEALTH SPENDING

Although these results are suggestive, they do not formally control for differences in individuals and families that may affect both tax rates and health spending. To do so, we specify a model of the health spending of individuals $i = 1, \dots, N$ covered by full-year employer-sponsored health insurance in year t . Each individual is a member of a family $j = 1, \dots, J$ that consists of a policyholder who works outside the home for wages or salary, and that person's spouse and/or children, if s/he has them. An individual has characteristics X_{ijt} that include age, age squared, gender, and marital status. A family has characteristics Z_{jt} that include the number of covered family members and the educational attainment, occupation, and industry of the policyholder. Educational attainment is an indicator variable denoting whether the policyholder is a college graduate; occupation is an indicator variable for whether the policyholder is a professional or technical worker (all other occupations are the omitted group); industry is a series of indicator variables for whether the policyholder is employed in agriculture or mining, construction,

manufacturing, financial and business services, or public administration (all other industries are the omitted group). Finally, we characterize each family by its total income Y_{jt} and the wage and salary income of its policyholder W_{jt} .

A. Models

Our models specify health spending H_{ijt} as a function of the characteristics of individuals and families, X_{ijt} and Z_{jt} ; a function of family income and the wage and salary of the family's policyholder, $g(Y_{jt}, W_{jt})$; the tax price of insurance TP_{jt} ; a family-specific error term η_j ,⁹ and an individual-specific error term ε_{ijt} :

$$(1) \quad H_{ijt} = \theta_t + \lambda TP_{jt} + X_{ijt}\beta + Z_{jt}\gamma + g(Y_{jt}, W_{jt}) + \eta_j + \varepsilon_{ijt}$$

For the reasons discussed above, simple OLS estimates of (1) are likely to be inconsistent. For example, because TP_{jt} implicitly depends on Y_{jt} and W_{jt} , any misspecification of $g(Y_{jt}, W_{jt})$ would mean that TP_{jt} would be correlated with either η_{jt} and ε_{ijt} or both. We therefore estimate (1) by instrumental variables, using $ABOVE_{jt}$ as an instrument, where $ABOVE_{jt} = 1$ if the wage and salary income of the policyholder is greater than 100 percent of the wage base.

We estimate our model in both levels and logs. When we specify the model in levels, we assume that $g(\cdot)$ is a cubic function of household income and the wage and salary income of the policyholder:

$$g(Y_{jt}, W_{jt}) = \pi_1 Y_{jt} + \pi_2 Y_{jt}^2 + \pi_3 Y_{jt}^3 + \rho_1 \omega_{jt} + \rho_2 \omega_{jt}^2 + \rho_3 \omega_{jt}^3.$$

where ω_{jt} is the wage and salary income of the policyholder, expressed as a percentage of the wage base. This specification is the simplest one for a "fuzzy" regression

⁹ This specification of the error structure allows unobserved determinants of health spending to be arbitrarily correlated both within families and over time.

discontinuity model, as described by Angrist and Pischke (2009). In log form, (1)

becomes:

$$(2) \quad \ln(H_{ijt}) = \theta_t + \lambda \ln TP_{jt} + X_{ijt} \beta + Z_{jt} \gamma + g(Y_{jt}, W_{jt}) + \eta_t + \varepsilon_{ijt}$$

where $g(\cdot)$ is a log function of household income:

$$g(Y_{jt}, W_{jt}) = \pi_1 \ln(Y_{jt}) + \rho_1 \omega_{jt} + \rho_2 \omega_{jt}^2 + \rho_3 \omega_{jt}^3 .$$

B. Results

Table 4 presents OLS and instrumental variables estimates of the effect of the tax preference for insurance on health spending from (1) and (2). The OLS estimates from (1) and (2) (in columns (1) and (2)) suggest that the tax price of insurance has a small and statistically insignificant effect on spending on health services by individuals in families with full-year employer coverage. The IV estimates (in columns (3) and (4)), in contrast, show that decreases in the tax price (*i.e.*, increases in the tax preference) increase spending. At the average level of health spending and the average tax rate, the IV estimate in column (3) translates into an elasticity of -0.601. This estimate is slightly smaller (in absolute value) than the elasticity of -0.741 from the log specification in column (4). The elasticity from the linear model is statistically significantly different from zero at the 10 percent level; the elasticity from the log model is statistically different from zero at the 5 percent level. Hausman tests reject the consistency of the OLS estimates of both (1) and (2) at the 5 percent level (for the linear model, $p = 0.01$; for the log model, $p = 0.04$), suggesting that endogeneity of TP is a practical as well as theoretical concern.

Table 5 presents results from six alternative specifications of (2) to investigate the robustness of our main result. Column (1) of the table repeats the estimates from the log model in Table 4, column (4) for ease of comparison. Column (2) presents estimates from the model based only on individuals in households with policyholders earning between 50 and 150 percent of the wage base. Narrowing in on the discontinuity leads to an increase in the point estimate of the coefficient of interest (to -1.280, statistically significant at the 5 percent level), although the standard errors are large enough that the estimate on the restricted sample is not statistically different from the estimate on the full sample. Column (3) reports results from a model that omits everyone under age 18; the estimated effect changes hardly at all, although the standard error rises slightly. Columns (4) and (5) report the effect of the tax preference on inpatient and outpatient spending separately; the fact that inpatient spending is less price-responsive than outpatient spending is consistent with the results from the RAND experiment and other work that suggests that inpatient spending is less discretionary. Finally, columns (6) and (7) replicate the analysis on the earlier and the later five-year intervals covered by the sample, respectively. Although the effect of the tax preference is not statistically significantly different from zero in the later subsample, it is of roughly the same magnitude in both as in the full, pooled sample. A t -test of the difference between the two subsamples fails to reject that they are equal ($p = 0.42$).

We estimated five other versions of (2) to investigate further the sensitivity of our results. First, instead of narrowing the window around the discontinuity (as is conventional in regression discontinuity study designs), we estimated (2) using only observations that were *distant* from the discontinuity. Excluding the 4,948 individuals

from families with a policyholder who earned 90-110 percent of the Social Security wage base leads to an estimated elasticity of -0.347 (standard error 0.477).¹⁰ Second, to confirm that our result was not due to a few outliers, we excluded individuals with health spending over \$100,000. Excluding the 76 observations that met this condition, we estimate an elasticity of -0.728 (standard error 0.374). Third, to investigate whether our results might be due to temporary health shocks and reverse causality (running from health status to income), we excluded individuals from families with a policyholder who switched from above to below (or below to above) the wage base in their two years in the MEPS panel. Excluding these 5,494 observations, we estimate an elasticity of -0.750 (with a standard error of 0.443). Fourth, to investigate the possibility of reverse causality further, we reestimated the model using only family members other than the wage earner who holds the insurance policy. Including only these 48,507 observations, we estimate an elasticity of -1.404 (standard error 0.463). Fifth, we reestimated the model using placebo instruments that were 10 percent below and 10 percent above the true discontinuity. Using these instruments, we estimate elasticities of -0.321 (standard error 0.786) and -0.253 (standard error 0.576), respectively.

We also investigated whether the tax price affects *coverage*, to explore the extent to which this channel might accentuate the effect of the tax price conditional on having coverage. We estimated a model similar to (1), but on a sample of all employed people whether or not they had full-year employer-sponsored insurance (N = 111,139). We used as the dependent variable the percentage of the year that the individual had coverage. In

¹⁰ The fact that the estimated effect of the tax preferences is statistically insignificant based on observations that are distant from the Social Security wage base can be explained by two facts: that the difference in Social Security tax rates above versus below the wage base is partially undone by differences in income tax rates, and that there is a positive income elasticity of demand for health insurance that is not fully captured by our regression specification.

elasticity terms, the instrumental variable estimate of the tax-price effect on coverage is -0.096 (with a standard error of 0.041, significant at the 5 percent level). We discuss the economic importance of this effect in Section VI below.

V. UNDERSTANDING THE MECHANISM THROUGH WHICH THE TAX EFFECT OCCURS

It is possible that our results could be generated by an unobserved process that is correlated with a threshold effect of policyholders' wage and salary income but not caused by tax incentives. To investigate this concern, we test whether earning above versus below the Social Security tax threshold affects individuals' choice of insurance policy, as economic theory suggests that it should. This choice may occur at the employer level, by affecting the set of plans that are offered, or at the employee level, by affecting the plan that is taken up. In either event, if the estimated effect of the tax preference affects health spending by altering choice of insurance policy, then individuals with a greater tax subsidy should be covered by insurance that has higher premiums,¹¹ and this effect should only occur for individuals whose premiums are paid with pre-tax dollars.¹²

To investigate whether this is true, we use data from the 1996-1998 Community Tracking Study (CTS). The CTS consists of three surveys: an Employer Survey, a

¹¹ The premium for an insurance policy is identically related to the policy's generosity, as measured by its actuarial value. Given that a policy's actuarial value, AV , is defined as the share of health spending that the policy would cover for the average person, its premiums is $(1 + L) * AV * E(S)$, where $E(S)$ is the expected spending of the policyholder(s) and L is the administrative cost, or load, of the insurance.

¹² Purchases of employer-sponsored health insurance can be made with pre-tax dollars under two conditions: if the premiums are (nominally) paid by the employer, or if the premiums are paid by the employee through a Section 125, or "cafeteria," plan. However, premiums paid by employees (even through withheld salary) occur after-tax if the employer does not have a plan that qualifies under Section 125 of the Internal Revenue Code.

Household Survey, and a Followback Survey. The Employer Survey, conducted in 1997, collects information about the health insurance policies that the respondent employers provide (if any); the total premium per employee paid for single-person coverage, including both employee contributions and the amount paid by the employer on the employee's behalf; the share of the total premium that is paid by the employer on the employee's behalf; and employer characteristics (industry and firm size). The Household Survey, administered in 1996-97, is a national sample of households that can be matched to the Followback Survey. The Followback Survey, administered in 1997-98 to organizations that offered or administered private health insurance policies covering Household Survey respondents, collects detailed information on cost sharing, gatekeeping, and other characteristics of insurance policies (such as whether the policies are HMOs, PPOs, or point-of-service plans). For each household with a respondent holding employer-sponsored health insurance, the matched Household/Followback Survey sample thus contains detailed characteristics of each respondent's health insurance policy, similar (although not identical) to the information collected by the Employer Survey; of each respondent household's insurance policyholder, including his or her wage and salary income, age, gender, and educational attainment; of each respondent household's insurance policyholder's firm (industry and firm size); and of each respondent household (household income and household size). The variables from all three surveys are described in an Appendix (available from the authors).

Ideally, to test whether the tax preference affects choice of insurance policy, we would know for each policyholder of employer-sponsored insurance, the wage and salary income of the policyholder, the generosity of the insurance policy, and the share of the

policy that was financed with pre-tax dollars. However, although the matched Household/Followback Survey contains information on wage and salary income, it does not contain information on generosity, and none of the three surveys contains information on the tax status of payments for employer-sponsored health insurance.

We approximate this ideal in the following way. Using the Employer Survey, we run regressions of total premium per employee, P_k , and the share of premiums paid by the employer (for firms providing employer-sponsored health insurance), S_k , on the firm and insurance policy characteristics, U_k , where $k = 1 \dots K$ firms. These regressions are of the following form:

$$P_k = \varphi_1 U_k + \varepsilon_k$$

$$S_k = \varphi_2 U_k + \eta_k.$$

Then, we construct a vector of firm and insurance policy characteristics U_j for each individual j in the matched Household/Followback Survey who holds an employer-sponsored insurance policy, where U_j is defined identically to U_k . With this information, we calculate for each employer-sponsored insurance policyholder his or her predicted premium, $\hat{P}_j = \hat{\varphi}_1 U_j$, and predicted share of premium paid by employer, $\hat{S}_j = \hat{\varphi}_2 U_j$, using estimates of φ_1 and φ_2 from the Employer Survey. Finally, using the matched Household/Followback Survey, we run regressions of predicted premiums on the characteristics of the insurance policyholder and his or her household listed in the appendix, V_j ; polynomials in the policyholder's household income, Y_j , and his or her wage and salary income ω_j (expressed as a percentage of the 1997 Social Security wage base); and an indicator for whether the policyholder's wage and salary income was greater than 100 percent of the wage base, $ABOVE_j$:

$$\begin{aligned}\hat{P}_j &= \delta V_j + h(Y_j, W_j) + \iota_j \\ &= \delta V_j + \sigma_1 Y_j + \sigma_2 Y_j^2 + \sigma_3 Y_j^3 + \nu_1 \omega_j + \nu_2 \omega_j^2 + \nu_3 \omega_j^3 + \nu_4 ABOVE_j + \iota_j,\end{aligned}$$

where $h(\cdot)$, Y_j , and ω_j are defined as they are in the health spending models. We estimate this model on the full sample of all respondents with employer-sponsored insurance and on the half of the sample with above the median value for the predicted share of premium paid by employer (85 percent).

Estimates of ν_4 from this model, the effect of earning above the Social Security tax threshold on predicted premiums, are a lower bound on the effect on actual premiums. First, the estimate only captures the impact of the tax threshold on the portion of individuals' premiums that are correlated with their firm and insurance policy characteristics. Second, using predicted premiums instead of actual premiums leads to conservative estimates because predicted premiums do not depend on any of the characteristics of individuals in the firm. Therefore, by construction, estimates of the effect of the tax threshold on premiums obtained in this way cannot be due to any unobserved characteristics of individuals that may be correlated with wage and salary levels and health status.

Selected parameter estimates from this model (and its logarithmic analog) are in Table 6. The leftmost two columns of the table present results from the full sample. Holders of employer-sponsored insurance who earn above the Social Security threshold have lower predicted premiums, holding constant a polynomial in their wage and salary income, their household income, and their other characteristics, although this effect is not statistically significant. The rightmost two columns present results from the half of the sample who, based on the characteristics of their insurance policy and their firm, have at

least 85 percent of their premium paid by their employers. For individuals in this subsample, premiums are almost entirely paid out of pre-tax dollars. The effect on predicted premiums of earning above the Social Security threshold is more negative and significant at the 5 percent level.

VI. CONCLUSIONS AND IMPLICATIONS FOR FUTURE RESEARCH

As health policy analysts have long observed, the tax preference for health insurance is likely a key cause of inefficiency in the U.S. health care system. By making health spending in general, and insured health spending in particular, appear less costly than they are, the tax preference gives employees the incentive to take compensation as health insurance rather than cash -- even if the value of the spending to the employee is significantly less than its cost to society. In this paper, we use a new identification strategy to estimate the magnitude of the tax preference's effect on health spending. Using the fact that Social Security taxes are only levied on earnings below a statutory threshold to identify exogenous variation in tax rates, we find a statistically and economically significant effect of the tax preference for insurance on spending.

Our estimates of the tax-price elasticity of demand for health care spending of -0.601 and -0.741, depending on specification, are squarely within the range reported in earlier work. Gruber and Lettau (2004), for example, find an elasticity of firm-level insurance spending with respect to the tax price of -0.7; Smart and Stabile (2005) find tax-price elasticities of the demand for health services in the range of -0.7 to -1.0. The fact that our estimates, derived using a very different method than earlier work, are in the

same range should give health policy analysts greater confidence that the true effect lies in this range.

The elasticities in Table 4 can be used to simulate how much health spending by privately-insured individuals would fall if the tax preference were repealed. Given that the average federal marginal tax rate (including Social Security and Medicare taxes) in our sample is 35.9 percent (authors' unpublished calculation), repealing the tax preference would increase the after-tax price of insured expenditures from $0.641p$ to p , where p is the before-tax price of health care. Evaluated at the average after-tax price, this amounts to a 43.8 percent price increase ($=0.359p / (0.5*(0.641p + p))$). At an elasticity of -0.601 (table 4, column (3)), this increase translates into a 26.3 percent decrease in health spending; at an elasticity of -0.741 (Table 4, column (4)), this increase translates into a 32.5 percent decrease ($1.002*0.438$) in health spending.

These effects, while quite large, are consistent with the results from other simulation models, including our own work. Among individuals with nonzero health spending in our sample, the average coinsurance rate, as measured by the share of health spending that is out-of-pocket, is 32.3 percent. According to Jack and Sheiner (Table 2, 1997), at a tax preference for insurance of 32.8 percent, an effective coinsurance rate of this magnitude implies a coefficient of relative risk aversion of approximately 1.5. At a coefficient of relative risk aversion of 1.5, repealing the tax preference in their model leads to an approximate doubling of the coinsurance rate; this result is similar to the effect of the tax preference simulated by Feldstein and Friedman (1977). A doubling of the coinsurance rate from 32 to 64 percent, in turn, would lead to a decline in health spending of between 13.3 percent (at the RAND experiment's estimate of the elasticity of

demand for health services of -0.2) and 46.7 percent (at Eichner's (1998) elasticity of demand for health services of -0.7).¹³ Gruber (Table 5, 2002) reports that removing all tax subsidies for health insurance would result in a 35.4 percent decline in health spending among individuals who are offered insurance by their employer, expressed as a percentage point change from its initial value. Expressed as a percentage-point change at the average (in order to make his estimate comparable to the others), this difference amounts to a 43.0 percent decline in health spending.¹⁴

Simulations of the effect of repealing the tax preference based on our empirical elasticity estimates, however, may either overstate or understate the true effect. On the one hand, if spending becomes less responsive as the tax preference is phased out and the level of spending is reduced, then the effects of repeal will be smaller than observational data would predict. On the other hand, a nationwide change in tax policy may lead to a larger change in coinsurance rates and health spending than observational data based on small differences in marginal tax rates would predict. Amy Finkelstein (2007), for example, shows that the change in hospital spending associated with the introduction of Medicare was far greater than the elasticities from the RAND Experiment would have predicted. In addition, our estimates are based on the elasticity of behavior of individuals at the upper end of the income distribution (*i.e.*, around the Social Security threshold); the elasticity at the mean income may be either larger or smaller.

Our estimates of the effect of the tax preference on coverage are quite modest, consistent with the findings of other work (Gruber and Washington, 2005). As discussed above in Section IV, the elasticity of the number of days of insurance coverage with

¹³ That is, $0.133 = 0.2 * (0.32 / (0.5 * (0.32 + 0.64)))$ and $0.467 = 0.7 * (0.32 / (0.5 * (0.32 + 0.64)))$.

¹⁴ That is, $0.43 = (0.354 / (.5 * (0.354 + (1 - 0.354))))$.

respect to *TP* is -0.096 in this population, which means that even complete repeal of the tax preference would lead to only a 4.2 percent decrease in the number of covered days;¹⁵ at an average coverage rate of 92 percent, this translates into a decrease of 14 days (authors' unpublished calculation). This finding has two implications. First, subject to the caveats above, it implies that repealing the tax preference would likely have little effect on coverage for the full-year employed population that is the subject of our study. Second, it implies very small consequences for health spending. Using the estimates of the effect of coverage on health spending from Jack Hadley et al. (2008), decreasing coverage by 14 days would decrease expenditures by around \$85 per year per person, or approximately 2 percent.¹⁶

Future work might seek to investigate the effect of the tax preference on health outcomes. The RAND Experiment found no significant effect of the tax preference on insurance generosity for the average adult, which, combined with our estimates, would suggest that the tax preference for insurance leads to massive inefficiency in the allocation of resources. However, the RAND Experiment did find some effects of health spending for vulnerable subpopulations, and in any event is now based on medical technology that is considerably out-of-date. Without better information about how health spending currently affects health, translating the findings of this paper into implications for social welfare must necessarily be left to future research.

¹⁵ From above, repeal of the tax preference represents a 43.8 percent increase in *TP*, which implies a 4.2 percent decrease in the number of covered days because $-0.042 = 0.438 * -0.096$.

¹⁶ Hadley et al. (2008) calculate that spending per year for a full-year insured person is \$3,915, or \$10.73 per day; spending per year for a full-year uninsured person is \$1,686, or \$4.62 per day. Thus, decreasing a person's coverage by 14 days would lead to a decline in spending of $\$85 = 14 * (\$10.73 - \$4.62)$, or 2.4 percent on a base of \$3,593 ($= \$3,915 * (335/365)$).

REFERENCES

- Angrist, Joshua, and Jorn-Steffen Pischke, 2009. *Mostly Harmless Econometrics*. Princeton University Press, Princeton, NJ.
- Bundorf, M. Kate, 2002. "Employee Demand for Health Insurance and Employer Health Plan Choices." *Journal of Health Economics* 21, 65-88.
- Cogan, John F., R. Glenn Hubbard, and Daniel P. Kessler, 2011. *Healthy, Wealthy, and Wise* (2d ed.). Hoover Institution/AEI Press, Stanford, CA.
- Eichner, Matthew, 1998. "The Demand for Medical Care: What People Pay Does Matter." *American Economic Review* 88(2), 117-121.
- Feldstein, Martin, 1973. "The Welfare Loss from Excess Health Insurance." *Journal of Political Economy* 81(2), 251-80.
- Feldstein, Martin, and Bernard Friedman, 1977. "Tax Subsidies, the Rational Demand for Insurance, and the Health Care Crisis." *Journal of Public Economics* 7(2), 155-178.
- Finkelstein, Amy, 2007. "The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare." *Quarterly Journal of Economics* 122(1), 1-38.
- Gruber, Jonathan, 2002. "Taxes and Health Insurance." In James Poterba (ed.) *Tax Policy and the Economy, Volume 16*, MIT Press, Cambridge, MA.
- Gruber, Jonathan and Michael Lettau, 2004. "How Elastic is the Firm's Demand for Health Insurance?" *Journal of Public Economics* 88, 1273-93.
- Gruber, Jonathan and Ebonya Washington, 2005. "Subsidies to Employee Health Insurance Premiums and the Health Insurance Market." *Journal of Health Economics* 24, 253-76.
- Hadley, Jack, John Holahan, Teresa Coughlin, and Dawn Miller, 2008. "Covering the Uninsured in 2008: Current Costs, Sources of Payment, and Incremental Costs." *Health Affairs* 27, w399-w415.
- Jack, William, Arik Levinson, and Sjamsu Rahardja, 2006. "Employee Cost-sharing and the Welfare Effects of Flexible Spending Accounts." *Journal of Public Economics* 90, 2285-2301.
- Jack, William and Louise Sheiner, 1997. "Welfare-Improving Health Expenditure Subsidies." *American Economic Review* 87(1), 206-221.

Kaiser Family Foundation/Health Research and Educational Trust, 2005. "Survey of Employer Sponsored Health Benefits: 2005." Kaiser Family Foundation, Menlo Park, CA, <http://www.kff.org/insurance/7315.cfm>.

Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, Arleen Leibowitz, and M. Susan Marquis, 1987. "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment." *American Economic Review* 77(3), 251-277.

Newhouse, J.P., and the Insurance Experiment Group, 1993. *Free For All? Lessons from the RAND Health Insurance Experiment*. Harvard University Press, Cambridge, MA.

Pauly, Mark V., 1986. "Taxation, Health Insurance, and Market Failure in the Medical Economy." *Journal of Economic Literature* 24(2), 629-75.

Smart, Michael and Mark Stabile, 2005. "Tax Credits, Insurance, and the Use of Medical Care." *Canadian Journal of Economics* 38(2), 345-65.

Figure 1
 Distribution of spending by type and by quintile, 1996-2005
 (Individuals in families with insurance policyholder earning 90-100 percent and 100-110 percent of Social Security wage base)

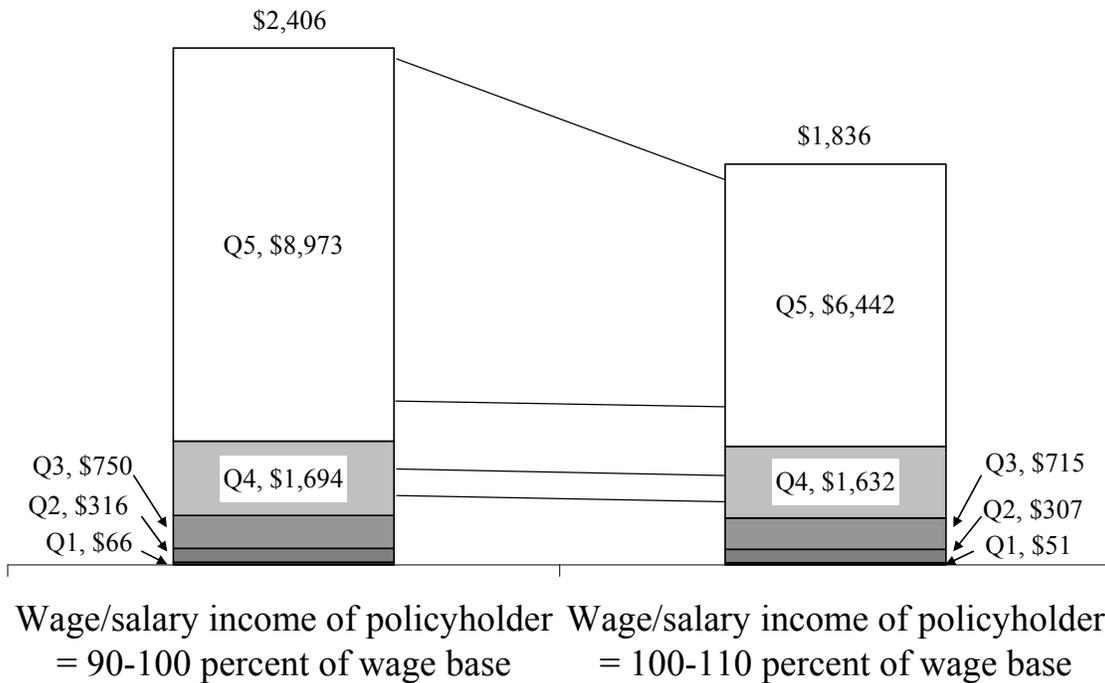
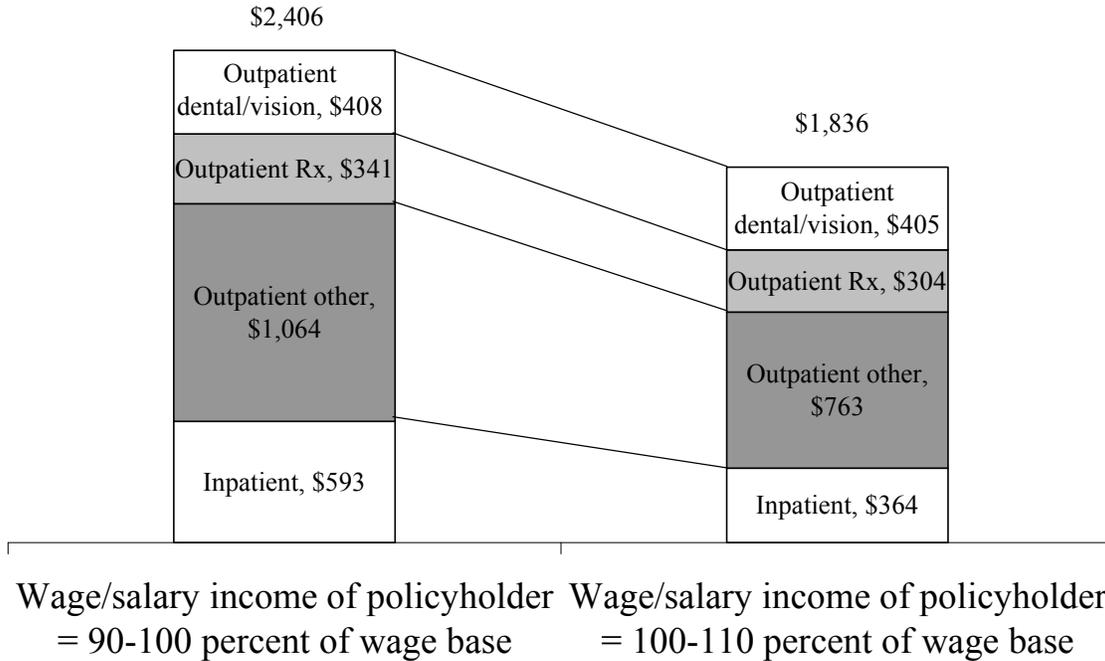


Table 1
Social Security wage base, 1996-2005

Year	Wage Base Amount (\$)
1996	62,700
1997	65,400
1998	68,400
1999	72,600
2000	76,200
2001	80,400
2002	84,900
2003	87,000
2004	87,900
2005	90,000

Table 2
Average health spending (2004 \$) of individuals in families with a full-year employed worker and full-year employer insurance coverage, 1996-2005
by wage and salary income of insurance policyholder

Wage/salary income of insurance policyholder	Average health spending (standard deviation)	Sample size [annual population size]	Average tax price
<70% of Social Security wage base	1,953 (5,046)	75,700 [81,440,000]	0.675
70-80% of wage base	1,937 (4,443)	5,943 [7,518,000]	0.618
80-90% of wage base	2,065 (4,726)	4,175 [5,428,000]	0.605
90-100% of wage base	2,406 (7,852)	2,977 [3,787,000]	0.597
100-110% of wage base	1,836 (3,848)	1,971 [2,572,000]	0.692
110-120% of wage base	2,110 (4,957)	1,811 [2,331,000]	0.689
120-130% of wage base	2,016 (5,929)	448 [626,000]	0.678
>130% of Social Security wage base	2,122 (4,526)	3,780 [4,953,000]	0.659
All wage/salary income levels	1,982 (5,073)	96,805 [108,700,000]	0.665

Note: Averages calculated using MEPS sample weights.

Table 3
 Characteristics of families and individuals,
 families with a full-year employed worker and full-year employer insurance coverage,
 1996-2005, by wage and salary income of insurance policyholder

	Wage/salary income of policyholder is...	
	90-100% of wage base	100-110% of wage base
Family characteristics		
Total income (2004 \$)	\$101,325 (33,811)	\$109,455 (31,869)
Tax price of insurance	0.597 (0.024)	0.692 (0.023)
Education	15.197 (1.997)	15.349 (2.008)
Occupation: managerial/technical	0.690	0.691
Industry: financial services	0.264	0.268
Industry: manufacturing	0.230	0.247
Industry: construction	0.043	0.046
Industry: agriculture/mining	0.007	0.007
Industry: public administration	0.100	0.093
Individual characteristics		
Age	31.934 (17.874)	31.812 (17.658)
Female gender	0.465	0.471
Married	0.504	0.501
Number of observations	2,977	1,971

Notes: Averages calculated using MEPS sample weights. Family education, occupation, and industry are defined to be those of its policyholder.

Table 4
Effect on health spending (2004 \$) of the after-tax price of insurance, families with a full-year employed worker and full-year employer-sponsored insurance

	Dependent Variable			
	Spending (1)	ln(spending) (2)	Spending (3)	ln(spending) (4)
Tax price of insurance				
TP	704 (491)		-1794 * (1064)	
ln(TP)		-0.044 (0.174)		-0.741 ** (0.374)
Effect of wage/salary income of policyholder				
Percentage of Social Security wage base	-551 ** (279)	0.209 (0.158)	-731 ** (306)	0.134 (0.166)
(Percentage of Social Security wage base)^2	437 * (233)	-0.168 (0.131)	714 ** (295)	-0.065 (0.149)
(Percentage of Social Security wage base)^3	-67 (62)	0.026 (0.034)	-125 (82)	0.006 (0.040)
Effect of household income				
Income x 10 ⁻⁴	153 *** (35)		72 (44)	
Income ² x 10 ⁻⁹	-104 *** (25)		-69 ** (27)	
Income ³ x 10 ⁻¹⁵	183 *** (50)		131 ** (53)	
Ln(income)		0.411 *** (0.036)		0.329 *** (0.053)
Estimator	OLS	OLS	IV	IV
Number of (person-year) observations	96,805	96,805	96,805	96,805
Number of families	29,037	29,037	29,037	29,037

Notes: *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively. Heteroscedasticity-consistent standard errors allowing for within-family correlation are in parentheses. One observation is one individual; observations are weighted using MEPS sample weights. Education, occupation, and industry are defined to be those of the policyholder. Estimated by instrumental variables, excluded instrument for tax rate is whether wage/salary income of primary worker is greater than 100 percent of the wage base.

Table 5
Effect on health spending (2004 \$) of the after-tax price of insurance, alternative models

	Dependent Variable						
	ln(spending) from table 4, column (4) (1)	ln(spending) (2)	ln(spending) (3)	ln(inpatient spending) (4)	ln(outpatient spending) (5)	ln(spending) (6)	ln(spending) (7)
After-tax price of insurance ln(TP)	-0.741 ** (0.374)	-1.280 ** (0.580)	-0.739 * (0.435)	-0.026 (0.251)	-0.774 ** (0.367)	-1.268 ** (0.515)	-0.689 (0.508)
Effect of family income ln(income)	0.329 *** (0.053)	0.040 (0.102)	0.268 *** (0.059)	-0.089 *** (0.034)	0.342 *** (0.052)	0.253 *** (0.071)	0.365 *** (0.074)
Exclusions from sample	None	Households with policyholders below 50% and above 150% of wage base	Age ≥ 18	None	None	1996-2000 only	2001-2005 only
Number of observations	96,805	37,262	69,931	96,805	96,805	43,254	53,551
Number of families	29,037	11,874	29,021	29,037	29,037	13,204	17,424

Notes: See notes to Table 4. All specifications are based on model (3) from Table 4 with the listed changes. In regression models but not reported in the table are coefficients on year, age, gender, marital status, family size, education, industry, and occupation.

Table 6

Effect on monthly employer-sponsored predicted insurance premiums (1997 \$) of earning above versus below the Social Security wage base, full sample and individuals with at least 85 percent of premiums paid by their employer

	Dependent Variable			
	Full Sample		Sample with >85% Premium Employer-Paid	
	Premiums	ln(Premiums)	Premiums	ln(Premiums)
Effect of wage/salary income of policyholder				
Indicator for above 100% of Social Security wage base	-1.90 (1.50)	-0.014 (0.010)	-4.09 ** (1.97)	-0.027 ** (0.014)
Percentage of Social Security wage base	8.90 (5.42)	0.112 *** (0.034)	-8.58 (7.25)	0.008 (0.048)
(Percentage of Social Security wage base) ²	-14.07 ** (6.13)	-0.138 *** (0.039)	1.81 (8.06)	-0.046 (0.054)
(Percentage of Social Security wage base) ³	4.14 ** (1.72)	0.036 *** (0.011)	-0.02 (2.23)	0.011 (0.015)
Effect of household income				
Income x 10 ⁻⁵	13.30 ** (6.64)		11.20 (9.04)	
Income ² x 10 ⁻¹⁰	-3.20 (10.00)		0.49 (13.40)	
Income ³ x 10 ⁻¹⁵	-2.50 (4.29)		-5.00 (5.70)	
Ln(income)		0.010 *** (0.003)		0.008 ** (0.004)
Number of observations	12,498	12,498	6,887	6,887

Notes: *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively. Heteroscedasticity-consistent standard errors are in parentheses. One observation is one employer-sponsored-insurance policyholder. Based on matched Household and Followback Survey data from the 1996-1998 Community Tracking Study. Predicted premiums and the percentage of predicted premiums paid by the employer imputed based on the characteristics of the policyholder's insurance policy and the policyholder's employer (industry and firm size). In regressions but not reported in table are age fixed effects, CTS site fixed effects (approximately equal to MSA fixed effects), and coefficients on gender, marital status, family size, and education.