Volume Title: Empirical Foundations of Household Taxation

Volume Author/Editor: Martin Feldstein and James Poterba, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-24097-5

Volume URL: http://www.nber.org/books/feld96-1

Conference Date: January 20-21, 1996

Publication Date: January 1996

Chapter Title: Tax Subsidies to Employer-Provided Health Insurance

Chapter Author: Jonathan Gruber, James M. Poterba

Chapter URL: http://www.nber.org/chapters/c6239

Chapter pages in book: (p. 135 - 168)

5 Tax Subsidies to Employer-Provided Health Insurance

Jonathan Gruber and James M. Poterba

The value of employer-provided health insurance is excluded from an individual's federal and state taxable income and from the social security tax base. These exclusions provide an incentive for individuals and firms to structure compensation arrangements so that employees receive employer-provided insurance, rather than cash compensation that they may ultimately use to finance their health care or health insurance purchases. This incentive has important economic implications: medical care financed by insurance will generally be overconsumed because of low copayment rates under traditional insurance policies. Tax incentives for employer provision of health insurance have therefore been cited, for example by Feldstein (1973), Pauly (1986), and Phelps (1992), as encouraging overinsurance and ultimately overconsumption of medical services.

Given the central role of this tax incentive in the medical economy, it is important both to measure it and to analyze how it would be affected by various policy reforms. This task is a complicated one because the tax system subsidizes medical care purchases in two ways. The first is the exclusion from income and payroll taxes of all employer insurance premium payments, as well as some fraction of employee payments for employer-sponsored insurance. The second is the deductibility of individual expenditures on medical care and medical insurance that exceed some minimum threshold, currently 7.5 percent

Jonathan Gruber is the Castle-Krob Career Development Associate Professor at the Massachusetts Institute of Technology and a faculty research fellow of the National Bureau of Economic Research. James M. Poterba is director of the Public Economics Research Program at the National Bureau of Economic Research and professor of economics at the Massachusetts Institute of Technology.

The authors are grateful to the National Institute on Aging and the National Science Foundation for research support, to Todd Sinai for exceptional research assistance, to David Bradford, Len Burman, David Cutler, Peter Diamond, Martin Feldstein, Louis Kaplow, and Roberton Williams for helpful comments and discussions, and to Dan Feenberg for assistance with the NBER TAXSIM model.

of adjusted gross income (AGI). The net tax incentive for insurance purchase depends on the subsidy for employer-provided insurance *relative* to that for health care purchase if an individual self-insures. Most previous analyses of the tax incentive for employer-provided health insurance have focused only on the first tax incentive.

This paper presents new evidence on the net tax subsidy to employer-provided health insurance, as well as new estimates of the likely effects of various tax policy reforms. We do so by combining information from the 1987 National Medical Expenditures Survey (NMES) with data from the U.S. Treasury Individual Tax Model file to estimate how the tax system affects the after-tax price of health insurance relative to the after-tax price of out-of-pocket health care spending. We aggregate respondents in the NMES into health insurance units then use the NBER TAXSIM model to estimate the tax saving to each from employer provision of health insurance. Our procedure preserves the rich cross-sectional variation in household spending on medical care better than approaches that impute insurance and medical care outlays to households in other data sets.

This paper is organized as follows. In section 5.1, we sketch the analytical framework that we use to measure the net tax subsidy to employer-provided health insurance. We define the tax subsidies to employer-provided insurance, employer-sponsored insurance that is paid for by the employee, and out-of-pocket spending on health insurance and health care. We then measure the net tax price of employer-provided insurance as a function of these subsidies. Section 5.2 describes the data sets we analyze and outlines our algorithm for measuring the tax subsidies.

Section 5.3 reports the basic results of our analysis. We begin by providing estimates of the marginal subsidy to additional insurance purchases since this is the margin of overinsurance of most concern to health policy analysts. We then describe the recent evolution of this subsidy. We compare the period before the Tax Reform Act of 1986 (TRA86), when the top marginal tax rate under the personal income tax was 50 percent, to the late 1980s, when the 1986 tax rate reductions were fully phased in, and to 1994, after several increases in marginal tax rates had brought top marginal tax rates to nearly 40 percent. We thereby illustrate how changes in the tax structure can affect the magnitude of the tax subsidy to employer-provided health insurance.

Section 5.4 describes the effect of various tax policy reforms on the net tax subsidy to health insurance purchase and reports illustrative calculations of how such reforms might affect the demand for health insurance. We consider capping the value of insurance benefits that are exempt from federal income taxation, as well as including the full value of employer-provided health insurance in both the FICA and federal income tax bases. We describe how these changes would affect the marginal subsidies to employer-provided health insurance, and under plausible assumptions about the price elasticity of demand for this insurance, we illustrate the effect of such reforms on insurance de-

mand. Although there is no definitive empirical evidence on the price elasticity of demand for health insurance, we present calculations using several values spanning results in the existing literature. Section 5.5 summarizes our findings and outlines several directions for future research.

5.1 Tax Subsidy to Employer-Provided Health Insurance

Employer-provided insurance is one of many ways of financing medical care services. It is therefore important to distinguish between subsidies to the purchase of health *insurance* and subsidies to the consumption of health *care* more generally. A change in the after-tax price of insurance can alter the financing of any given set of medical services, and since it changes the composite price of medical care, it may also affect the level of health care services consumed.[1] Our analysis is limited to the former effect, the impact of taxation on the financing of medical care. We assume that changes in the level of employer-provided health insurance would be offset by similar changes in household out-of-pocket spending, with little or no change in the level of health care consumed. We therefore understate the effect of tax reforms on the demand for health insurance since shifts in the aggregate demand for medical care would reinforce changes in the level of insurance demand following a tax reform.

We define the tax subsidy to insurance purchase in terms of the relative after-tax price of financing health care with insurance, and without insurance on an out-of-pocket basis. Our approach does not consider how the tax subsidy to employer-provided insurance affects the after-tax price of medical care, which prevents us from analyzing how insurance tax reforms would affect aggregate medical care spending.[2]

The current U.S. tax system subsidizes both employer-provided health insurance and out-of-pocket medical spending. Employees with employer-provided health insurance are not required to include the value of this insurance in their taxable income for federal and state income taxation, or in their wage tax base for the payroll tax.[3] The Joint Committee on Taxation estimates that in fiscal 1994, federal revenues from the personal income tax and payroll tax were nearly $90 billion lower as a result of these exclusions (U.S. Congressional Budget Office [CBO] 1994a). The tax system also subsidizes out-of-pocket

1. This distinction parallels a familiar analysis in the taxation of corporate capital income. Changing the tax treatment of debt would lead both to a shift in financing, i.e., differential use of debt and equity, as well as a shift in the ultimate level of real investment.

2. The after-tax price of medical care at the time of consumption depends on whether the patient is insured, the copayment rate and deductible level for the patient's health insurance (if insured), and whether the patient itemizes tax deductions and claims the medical expense deduction. For those who purchase insurance, there is also an ex ante price of medical care, distinct from the price at the time of consumption, that includes the price of purchasing insurance.

3. Employer-provided health insurance was encouraged by the 1942 Stabilization Act, which placed limits on wage increases but allowed employers to offer insurance plans to their employees. Scofea (1994) provides an introduction to the history of employer-provided health insurance in the United States.

spending on health insurance and medical services by allowing an itemized deduction for medical expenses. Itemizers can deduct expenditures on medical care and directly purchased health insurance in excess of 7.5 percent of AGI from their federal taxable income. The revenue cost of this provision, less than $4 billion in 1994, is much smaller than that for employer-provided insurance.

Our definition of the tax subsidy to employer-provided insurance considers both the after-tax cost of employer-provided insurance and the after-tax cost of out-of-pocket medical spending. We do not consider individually purchased health insurance, on the grounds that higher load factors and less-favorable tax treatment than for employer-provided insurance make this a dominated option for those who seek insurance.

5.1.1 After-Tax Cost of Employer-Provided Insurance

We consider an individual with a federal marginal income tax rate on earned income of τ, a net-of-federal-tax state income tax rate of τ_s, and employer and employee rates of payroll tax each equal to τ_{ss}. We assume that labor income taxes and payroll taxes are fully borne by labor, so that when an employer provides insurance that costs E dollars, the employee's wage is reduced by $E/(1 + \tau_{ss})$.[4] The employer is indifferent between purchasing $1 of insurance or paying wages of $1/(1 + \tau_{ss})$, since each dollar of wages requires a payroll tax payment as well. The change in the employee's after-tax wage income per dollar of employer-provided insurance, dw_{AT}/dE, is therefore

$$\frac{dw_{AT}}{dE} = \frac{1 - \tau - \tau_s - \tau_{ss}}{1 + \tau_{ss}}. \tag{1}$$

Many previous studies of taxation and employer-provided health insurance, including Feldstein and Allison (1974), Taylor and Wilensky (1983), Holmer (1984), and Burman and Williams (1994), have used dw_{AT}/dE or some variant of it to define the tax subsidy to employer-provided insurance. A parallel assumption is made in the literature on taxation and the demand for fringe benefits more generally.[5]

While the reduction in after-tax wages per dollar of employer-provided insurance is a key factor determining the after-tax price of such insurance, it is not the only one. We identify two other factors that affect the after-tax cost of employer-provided health insurance, and that consequently affect the relative price of this insurance vis-à-vis self insurance. First, because insurance firms include a load factor in their policy prices, the expected value of medical care outlays from $1 of spending on medical insurance is less than the expected

4. Several recent studies, notably Gruber and Krueger (1991) and Gruber (1994), support this assumption with respect to various types of employer mandates.

5. Woodbury and Hamermesh's (1992) study of how the TRA86 affected the demand for fringe benefits vs. wage income at universities is a recent example in this tradition. Earlier studies that adopt similar approaches but sometimes omit either the state tax or payroll tax include Sloan and Adamache (1986) and Long and Scott (1982).

value from $1 of out-of-pocket medical spending. The load factor, λ, reflects costs of administering an insurance plan, the profits of the insurer, and any other expenses incurred in minimizing the health risk of a given group to the insurer. This load factor affects the after-tax cost of employer-provided insurance relative to self-insurance of medical care costs.

Second, contrary to the assumption of complete employer provision of insurance above, employees pay a substantial and rising fraction, currently about 15 percent, of the premiums for employer-provided insurance. Blostin, Grant, and Wiatrowski (1992) report that in 1989, nearly half of the employees who received employer-provided health care benefits contributed to the cost of individual coverage, while for two-thirds of these workers, contributions were required for family coverage. Approximately three-quarters of these employee premiums are paid after tax, and paying them is a requirement of taking advantage of the favorable tax treatment of employer-provided insurance.[6] Employees who must make after-tax contributions to their employer-provided insurance receive favorable tax treatment on a smaller fraction of their health insurance than those employees whose insurance is fully provided by the employer. Recognizing employee contributions to the cost of employer-provided insurance therefore raises the after-tax price of this insurance. In defining the after-tax price of insurance, we use G to denote employee payments for employer-provided group insurance and E to denote employer payments. We assume that a fraction δ of employee premiums can be paid for on a pretax basis through cafeteria plans and other tax-favored arrangements.

One question that arises in considering employee payments for health insurance is why employers structure health plans with such payments, despite their tax inefficiency. There are at least two possible reasons. First, within any workplace, different workers will place different values on the benefit of health insurance coverage. Unless employers can selectively lower the wages of only those employees who value insurance coverage, employers who pay the full cost of insurance will disproportionately attract workers with a high value of insurance. Employers may view this outcome as unattractive, for example because the workers who value insurance the most may be less healthy and therefore less productive workers. Cost sharing can be an effective mechanism for reducing the selection effects associated with health insurance provision.

A second reason for employers to require some employee contributions relates to employee choice of health care plan. Many employers offer a choice across plans of differing generosity and cost. Employers may not be able to

6. The U.S. Bureau of Labor Statistics (1993, 1994) reports that approximately 33 percent of employees of firms with more than 100 employees, and 20 percent of employees of firms with fewer than 100 employees, can deduct their own premium payments from taxes. These are employees who can pay their premiums through cafeteria plans provided by their employers. We are not able to identify which employees can make pretax premium payments in the data below, so we randomly assign individuals to the pretax employee premium group with a probability of .25. This is a weighted average of the probabilities for small and large firms.

pay lower wages to employees who choose higher cost plans, and cost sharing can be used to induce choice of cost-effective insurance.

Recognizing both the load factor on employer-provided insurance and the existence of employee contributions to such insurance yields the following expression for the after-tax price of employer-provided insurance:

$$(2) \quad P_{HI} = \left[\left(\frac{1 - \tau - \tau_s - \tau_{ss}}{1 + \tau_{ss}}\right)*\left(\frac{E + \delta*G}{E + G}\right) + \frac{(1 - \delta)*G}{E + G}\right]*(1 + \lambda).$$

We define the tax subsidy to employer-provided insurance by comparing this after-tax price with the after-tax cost of self-insurance.

5.1.2 After-Tax Cost of Out-of-Pocket Medical Spending

It is widely recognized that the income tax code provides a form of insurance against large medical costs by permitting a deduction against taxable income for medical expenditures above a certain share of AGI. This provision of the tax code discourages insurance purchase since it lowers the after-tax cost of paying high medical expenses out of pocket. Bradford (1984) and more recently Kaplow (1991, 1992) discuss the implicit insurance in the tax system, but none of the previous studies of the tax subsidy to employer-provided insurance have considered this aspect of the income tax code.

The tax subsidy to out-of-pocket medical expenses depends on whether a taxpayer itemizes. For a nonitemizer, the after-tax cost of such spending is \$1. For itemizers, however, the after-tax cost of the marginal dollar of out-of-pocket medical spending is $1 - \alpha\tau$, where τ is the federal marginal tax rate and $\alpha = 1$ if the marginal dollar of spending exceeds the AGI floor and zero otherwise. We assume that medical expenses cannot be deducted in computing state taxable income.

For an individual considering the purchase of insurance, α is unknown. It is determined by the individual's taxable income and realized need for medical services during a tax year. If F denotes the AGI threshold above which medical expenses are deductible and T the individual's total medical spending, then the probability that the last dollar of health expenditures will be tax deductible ($\alpha = 1$) equals the probability that $T - F > 0$.[7] This is the probability that the marginal dollar of health costs covered by employer-provided insurance would have been deductible if it had been incurred on own account. We have no direct information on how individuals form expectations of α in contemplating insurance purchases. We therefore assume rational expectations about actual spending during the year, calculate actual values of α for all households in our data

7. Total medical spending is $E + G + O$, where E is the value of employer-provided insurance, G is personal spending on group insurance premiums, and O is individual out-of-pocket health care spending. We assume that out-of-pocket spending for those with employer-provided insurance would not be affected by a shift to self-insurance. The tax-deductible share of the *additional* medical spending that would result from reduced employer-provided insurance depends on the probability that $E + G + O - F > 0$, rather than the probability that $E + G - F > 0$.

sample, and use these values in place of expected values in calculating the after-tax price of insurance.[8]

Individuals with health insurance typically face lower marginal costs of health care services at the time of consumption than individuals without such insurance. This may affect their demand for medical services, and it suggests that total medical outlays, T above, may be a function of an individual's insurance regime. We address this by computing α under two different assumptions about the link between price at time of consumption and medical spending. We first assume that total medical spending is unaffected by the presence or absence of health insurance. Our second case assumes that the price elasticity of demand for medical care services is -0.33.[9] Our findings are relatively insensitive to our assumption about the link between insurance status and T because for most households α is zero.

The foregoing discussion focuses on the after-tax cost of a marginal dollar of health care spending, which we label *marginal* α. This should be distinguished from the fraction of insured spending that would be tax deductible if the individual were not insured, $(T - F)/T$, which we label *average* α. Both marginal and average α range between zero and one. Marginal α describes the after-tax cost of the medical expenses that an individual would incur if employer-provided health insurance coverage were reduced by \$1 and the resulting drop in insured medical care were replaced with out-of-pocket spending. Marginal α will only equal unity if the insured individual already has out-of-pocket spending in excess of the AGI floor. Relatively few insured individuals are in this situation. Average α corresponds to the after-tax cost of replacing all insured medical expenditures with out-of-pocket spending. Average α will be positive whenever total medical spending, including insurance, exceeds the AGI floor. When we tabulate the tax subsidy to employer provided insurance in tables 5.4 and 5.5 below, we use average α in our expressions for the after-tax price. For analyzing the effect of tax caps on insurance spending, however, we use marginal α in our calculations since individuals are adjusting insurance purchases on the margin. Using marginal α may cause us to overstate the subsidy to insurance and therefore the reduction in insurance that results from tax caps, if tax policy changes are not marginal.

5.1.3 Relative After-Tax Price of Insurance

We define the relative after-tax price of employer-provided insurance as the ratio of the after-tax price of this insurance to the after-tax cost of out-of-pocket medical spending:

8. Newhouse et al. (1989) report that the best predictor of current medical spending is past spending.

9. We measure the marginal copayment rate for those with insurance as the ratio of their out-of-pocket medical spending to total medical spending. This is likely to be closer to the average than to the marginal copayment rate, but we do not have any further information in our data set to improve this imputation.

$$(3) \qquad P_{rel} = \frac{\left[\left(\frac{1 - \tau - \tau_s - \tau_{ss}}{1+\tau_{ss}}\right)*\left(\frac{E+\delta*G}{E+G}\right) + \frac{(1 - \delta)*G}{E+G}\right]*(1+\lambda)}{1 - \alpha*\tau}.$$

If the tax code treated insurance premia and medical expenditures symmetrically—for example, if neither were deductible from taxable income or if both could be excluded from federal and state taxable income and from the payroll tax wage base—then the cost of insurance relative to the direct outlays on medical care would be $P_{rel} = 1 + \lambda$. We therefore consider the tax-induced distortion in the relative price of insurance to be $[P_{rel}/(1 + \lambda) - 1]$, where P_{rel} is given by equation (3).

Our measures of the after-tax price of health insurance and the relative price of insurance suffer from at least four limitations. First, we fail to distinguish between marginal purchases of incremental employer-provided insurance and the discrete decision to purchase such insurance. The load factors on marginal insurance purchases may be lower than average loads if these loads in part reflect administrative costs that do not rise when a policy becomes more extensive.

Second, we assume that when expenditures on employer-provided insurance fall, employer (E) and employee (G) spending decline in equal proportion. In fact, many employers contribute a flat amount to their group health insurance plans, and employees contribute the differential cost between the plan that they choose and the lowest cost option. In such cases, if G is not tax deductible, then there is no tax subsidy to insurance on the margin. Thus, our results below will overstate the average tax subsidy to workplace insurance.

Third, our formation ignores the possibility that individuals may be able to pay for their out-of-pocket medical costs with pretax dollars, as, for example, with medical spending accounts that are provided in some cafeteria plans. This will also lead us to overstate the tax subsidy to insurance by understating the tax benefit of self-insurance. Unfortunately, we have no data on the structure of employer contributions or the availability of such pretax out-of-pocket arrangements.

Finally, we exclude any possible link between changes in the tax treatment of employer-provided insurance and the aggregate level of health care spending. We emphasize the relative cost of employer-provided insurance versus out-of-pocket spending, but our estimates of the after-tax price of employer-provided health insurance, P_{HI} in equation (2), could also be used to assess the effects of insurance tax treatment on the demand for medical care.

5.2 Data Sources on Medical Care Spending and Tax Rates

This section describes our methodology for estimating the various parameters, such as marginal tax rates, load factors, and probabilities that medical expenditures are deductible from income taxes, that enter our expression for

the relative price of employer-provided health insurance. Because our analysis requires detailed information on the pattern of health care expenditures as well as the tax circumstances of individuals and households, we use the U.S. Treasury Individual Tax Model and the NBER TAXSIM program to impute tax rates to family units in the 1987 NMES. We then draw on the information on health insurance and health care spending in this database to analyze the effect of tax subsidies to employer-provided health insurance.

5.2.1 NMES Sample

We are not aware of any data set that includes detailed information on health insurance coverage, health care spending, and federal income tax status. The NMES is the best available household-level database on health care spending. This is a nationally representative household survey that followed roughly 20,000 families during 1987. It gathered information on the demographic and economic characteristics of both family units and individual family members, including information on labor force attachment and income by source. It also collected detailed data on insurance plans, and these data were cross-checked against information collected from insurance sources such as employers or insurance companies. The NMES includes information on expenditures on a variety of types of medical care. Most of this information was also cross-checked by interviews with medical providers.

To impute tax information such as marginal tax rates and itemization status to survey respondents in the NMES, we aggregate individual NMES respondents into "health insurance units" (HIUs). These units include the family head, his or her spouse, any children under age 19, and full-time students until they reach age 23. There may be multiple family heads within a household, for example, when elderly parents live with a younger nuclear family. We limit our sample to employed individuals and exclude the self-employed, families with someone who is aged 65 or over and therefore eligible for Medicare, families with anyone who is eligible for Medicaid, and families with missing information on insurance status.

Table 5.1 shows the quantitative importance of the various data restrictions that we have imposed in selecting our sample. The NMES universe contains 20,028 HIUs who represent a total of 168.5 million family heads and spouses. Since family heads and spouses are the only relevant decision makers for insurance purchases, we use only their sample weights in making our calculations. We define families as employed if either the head or spouse is employed and as self-employed if both the family head and spouse are self-employed.

We define families as insured if both the family head and that person's spouse report that they are insured in the fourth wave of the survey and if they report some spending, either by their employer or by themselves, on employer-provided group health insurance. Our analysis excludes some employed families who purchase individual insurance only since we are focusing on the tax incentives for employer-provided insurance. We define families as uninsured

Table 5.1 Sample Size and Sample Limitations from the 1987 NMES

Restriction	Families in NMES Sample	Weighted to Represent U.S. Population[a] (millions)
Total NMES sample	20,028	168.5
− Households with anyone over age 64 or oldest member under age 18	(5,688)	(32.6)
Subtotal	14,340	135.9
− Medicaid households	(617)	(3.5)
Subtotal	13,723	132.4
− Nonrespondents to insurance status questions on NMES Wave IV	(2,716)	(11.4)
Subtotal	11,007	121.0
− Families that are neither "insured" nor "uninsured" by our definitions	(3,856)	(41.9)
Subtotal	7,151	79.2
− Self-employed families	(136)	(1.6)
Subtotal	7,015	77.5
− Unemployed families	(770)	(6.5)
Subtotal	6,245	71.1
− Families with zero weight	(284)	(0.0)
Final sample	5,961	71.1

Source: Authors' tabulations using 1987 NMES.

[a]Weighted totals are weighted by sum of head and spouse weights.

if they report both the family head and spouse to be uninsured and have no employer-provided insurance. Since the NMES does not report insurance plan information for all persons, our definition excludes relatively more insured persons from the sample than uninsured persons.[10] Our final sample has 5,961 HIUs, representing a total of 71.1 million household heads and spouses.

Table 5.2 presents information on the insurance status of the individuals in the employed family units in our NMES sample. Just over 82 percent of our population-weighted sample, or 58.4 million household heads or their spouses, are part of an employed household and have employer-provided insurance. Our analysis suggests that 12.7 million employed individuals are uninsured. This translates into a higher fraction of employed individuals classified as uninsured than some other sources, but the disparity is explained by our stringent criteria for defining a household as insured.[11] Within the subsample that reports some employer-provided insurance, 42 percent have employer contributions only for insurance, and 52 percent have both employer and own contributions. Only 5

10. We also exclude NMES families with sample weights of zero. These are families that were added to the NMES during the survey year, for example, because they moved into an existing NMES household as a subfamily.

11. Our analysis yields estimates of the number of uninsured employed individuals that are similar to those in other studies.

Table 5.2 **Health Insurance Status of Employed Individuals in the 1987 NMES**

Status	Number (million)
Only employer premiums ($E>0$, $I+G=0$)	24.8 (1,936)
Only individual premiums ($E=0$, $G>0$)	3.0 (228)
Employer and individual premiums ($E>0$, $G>0$)	30.6 (2,360)
Uninsured (E=G=0) and self-reported uninsured	12.7 (1,437)
Total	71.1 (5,961)

Source: Estimates based on authors' tabulations from the 1987 NMES. HIUs, weighted to reflect the sampling probabilities of family heads and spouses, are the basis for tabulations. Calculations exclude individuals in households with anyone over age 65. Values in parentheses are number of NMES households in each group.

percent of the households in our sample report paying the full cost of employer-provided insurance themselves.

Table 5.3 reports the distribution of spending on tax-subsidized employer-provided insurance in our NMES sample. The subsample used to construct this table is the set of all individuals who are employed and who we classify as covered by employer-provided insurance. Column (1) shows the value of employer contributions for insurance (E), and column (2) shows the value of both employer contributions and pretax contributions by employees ($E + \delta G$). While the NMES figures are measured in 1987 dollars, all results in this paper have been inflated to 1994 dollars using the growth in personal health spending over the 1987–94 period.[12]

Table 5.3 shows that the mean value of employer-provided health insurance is $4,249. The mean value of employer and pretax employee spending is approximately 5 percent higher ($4,483). The distribution of this tax-subsidized spending is somewhat skewed, as is revealed by the lower median values. Nearly 10 percent of the sample reports employer contributions worth more than $8,000 per year, and more than one-quarter report values of less than $2,000.

5.2.2 Tax Rate Estimation

To estimate the marginal federal tax rate facing each NMES household, we must estimate each household's federal taxable income. The NMES reports

12. The CBO (1993) presents data on private health insurance expenditures in 1987 ($155 billion), along with forecasts for 1993 ($289 billion) and 1995 ($343 billion). We estimate 1994 expenditures by interpolating between the 1993 and 1995 forecasts; this yields $316 billion. We then use the ratio of 1994 to 1987 spending, $316/$155, or 2.039, to impute the distribution of 1987 spending to 1994. The estimates in table 5.3 may overstate actual employer contributions for insurance, given the rising role of insurance cost shifting to employees.

Table 5.3 **Distribution of Expenditures on Employer-Provided Health Insurance in the 1987 NMES (converted to 1994 dollars)**

Percentile	Employer Spending (1)	Employer Plus Pretax Employee Spending (2)
5th	0	602
10th	1,020	1,250
25th	1,877	2,044
50th	3,816	4,130
75th	5,872	6,021
90th	7,950	8,130
95th	9,920	10,159
Mean	4,249	4,483

Source: Authors' tabulations using 1987 NMES data. Estimates for 1994 are based on the ratio of total private health insurance spending in 1994 to that in 1987, as projected in CBO (1993).

information on a variety of family income flows, such as wage income, dividend income, and interest income, although it does not contain nearly as much detail on income sources as a tax return. In particular, it does not report capital gains income, which we set equal to zero.[13]

The NMES also asked respondents whether they itemized deductions for income tax purposes. This is a critical input to our calculation of both tax rates and the probability of deducting out-of-pocket medical expenditures. The NMES figures for itemization differ from those in the Treasury tax model, although we did not find any systematic pattern in the differences. In 1987, 58.1 percent of the joint filers in the Treasury tax model database, excluding those who claimed any household members aged 65 or over, itemized deductions; the figure is 48.5 percent in the NMES. This pattern is reversed for nonjoint filers, with the NMES itemizing share 18.3 percent and the Treasury tax model share 16.9 percent, again excluding those aged 65 and over. If the NMES understates the incidence of itemization, our results will tend to overstate the tax subsidy to employer-provided insurance.

We conduct our analysis using the reported 1987 data on medical care and insurance spending (inflated to 1994 levels, as described above), but we consider the tax subsidy to health insurance under several different income tax regimes. We ask what the tax subsidy to the level of health insurance purchased in 1987 would have been if households faced the tax rates that they faced in 1986, 1989, and 1994. To do this we "age" taxable income in the NMES database to 1986, to analyze the tax code in effect before TRA86, to 1989, when

13. Imputing capital gains income to individuals by income category made little difference to the estimates of after-tax insurance prices in Gruber and Poterba (1994). Besides capital gains, TAXSIM uses a number of other income items reported on a tax return, such as contributions to IRAs and Keogh accounts and self-employed business deductions, to compute taxable income and marginal tax rates. We set any tax return item for which we do not have information in the NMES equal to zero in estimating taxable income.

TRA86 was fully effective, and to 1994, the most recent tax year. Our aging procedure is a simplified version of that developed by Lindsey (1987). We assume that each family's AGI changes from year to year in the same way that average per capita AGI changes between years.

Our estimate of the marginal state tax rate facing each household is less precise than our estimate of the federal tax rate. The NMES does not report the respondent's state of residence, but rather reports four census regions. We therefore create 12 "stylized families," 6 joint filers and 6 single filers, at different income levels. We use TAXSIM to estimate state-specific marginal tax rates for each of these stylized families, and we then compute state-population-weighted averages within each census region for each of these family types. We then assign one of these averages to each NMES household by identifying them as similar to one of the stylized households.

Our calculation of each individual's marginal social security tax rate follows Feldstein and Samwick (1992a). They note that the statutory social security tax rate is not the true marginal rate since prospective benefits are linked to taxes paid. This linkage varies according to income and gender since the benefit formula is redistributive and since on average women live longer and therefore receive a higher present discounted value of benefits than men. The social security system also redistributes across households with different configurations of primary and secondary earners since secondary earners receive the higher of their benefits and one-half of the benefits of the primary earner. Feldstein and Samwick find that effective marginal tax rates vary from negative values (subsidies to labor supply) for some households to the statutory marginal tax rate for other households. As a result, our tax prices would be misspecified if we used the statutory social security tax rate.[14] We estimate the effective marginal social security tax rate facing each NMES respondent using data that Feldstein and Samwick (1992b) present on tax rates by age, sex, and family labor supply.[15]

Individuals also pay payroll taxes for disability insurance and for Medicare hospital insurance. We accounted for the former by grossing up the Feldstein-Samwick net tax rates by the ratio of the statutory tax rate for *both* OASI and disability insurance to the statutory OASI tax rate, under the assumption that disability insurance tax-benefit linkages follow the same pattern as those for OASI.[16] For individuals with labor income in excess of the taxable social secu-

14. Using statutory tax rate without adjustment ignores the fact that higher social security tax collections today will be offset by higher social security benefit payouts in the future.

15. We assign the Feldstein-Samwick (1992b) effective tax rate for single men to single men or to families where the wife earns more than one-half of the husband's earnings. We assign their rate for women to single women, and their rate for men with dependent spouses to those families where the wife's earnings are less than one-half the husband's earnings. One limitation of this approach is that we are using point-in-time labor supply to proxy for the relative earnings of husbands and wives over their lives, but that is all that is available in the NMES.

16. This is a crude assumption since disability insurance benefits accrue to a different population, with many more young men, than OASI benefits. In addition, there are other limitations to

rity maximum, we set τ_{ss} equal to zero. In the latter case, there is no tax-benefit linkage since all citizens receive Medicare at age 65 regardless of their work history. We therefore use the statutory Medicare tax rate in our calculation.

5.2.3 Comparison with Earlier Studies

Our methodology differs from that in the CBO's (1994b) study of tax subsidies to health insurance, the most prominent recent study of related issues, in using the NMES as the central database for analysis. The CBO study imputed information from *both* the NMES and individual tax records to a third data set, the Current Population Survey (CPS). This has the advantage of providing more detailed income information than the NMES, as well as a larger sample of respondents and the associated opportunity for more precise within-group analysis. But it has the disadvantage of compressing the substantial heterogeneity across individuals in their health care spending, and the correlation of that spending with health insurance circumstances, in the process of imputation. In addition, the CPS does not report itemization status, while the NMES does. Given our limited objective in analyzing tax subsidies to employer-provided health insurance, and our focus on national aggregates, we would not gain substantially from access to the additional information that is potentially available in the CPS.

Our analysis is closer in spirit to Taylor and Wilensky's (1983) study of tax incentives and employer-provided health insurance than to the recent CBO (1994b) study. Although Taylor and Wilensky (1983) did not consider a number of the factors we described above, such as the role of insurance loads, state taxes, and the share of employee-paid health insurance premiums, they did use an earlier data set similar to the NMES, the 1977 National Medical Care Expenditure Survey (NMCES), as the basis for their study. They imputed information on tax status to households in the NMCES and computed a variety of summary statistics on tax subsidies to health insurance.

5.3 New Estimates of the Tax Subsidy to Employer-Provided Health Insurance

We summarize our analysis of the tax subsidy to employer-provided insurance by reporting average values of the after-tax relative price of employer-provided insurance and out-of-pocket medical spending, P_{rel}, under two different assumptions about the effect of insurance status on total medical spending. Recall from the discussion above that these summary statistics use information

the Feldstein-Samwick net tax rate calculation, such as the fact that some social security benefits accrue to dependents of contributors, and the fact that social security may substitute for an imperfect private market for real annuities and therefore have an above-market value to recipients.

on average, rather than marginal, α. To permit comparison with earlier studies, we also report the sample average of the change in after-tax employee wage income for each dollar of employer-provided health insurance (dw_{AT}/dE).

Table 5.4 presents these summary statistics when each of the NMES families is assigned its federal marginal tax rate for 1986, 1989, and 1994. We report both the average value of the relative after-tax price of insurance and out-of-pocket care, as well as the cross-sectional standard deviation of this price. We distinguish two subgroups of the employed population, those with and without insurance, and tabulate results separately for each.

The results for insured employed families, using the 1994 federal tax code, illustrate our general findings. The average value of the reduction in after-tax wage income per dollar of employer spending on health insurance is 0.682. This implies that federal income and payroll taxes, and state income taxes, place a 31.8 percent tax wedge between the after-tax cost of all other goods ($1) and the after-tax cost of employer-provided health insurance. There is substantial disparity across households in the magnitude of this subsidy: the standard deviation of dw_{AT}/dE is .117. Comparing the results for 1986 with those for either 1989 or 1994 demonstrates that the tax rate reductions in TRA86 raised the after-tax wage cost of employer-provided benefits by an average of about four cents per dollar.[17] The changes in the structure of tax rates between 1989 and 1994, because they were concentrated on a small group of high-income households, did not substantially affect the average after-tax wage cost.

The two lower rows of table 5.4 present our estimates of P_{rel} from equation (3). This ratio is substantially higher than the after-tax wage cost of employer-provided health insurance benefits because it multiplies the after-tax wage cost by $1 + \lambda$ to reflect the insurance load factor, because it includes less favorably taxed employee contributions for health insurance in the numerator, and because the denominator $(1 - \alpha\tau)$ is less than 1.

The average value of P_{rel} for the 1994 tax code is 0.837 if we assume that medical spending is unaffected by whether or not an individual is insured; it falls to 0.811 if we assume that spending would be reduced, since the tax subsidy to self-insurance is then less valuable. There is somewhat more heterogeneity in these measures than in the simple after-tax wage cost measures in the first row; the standard deviation of P_{rel} is .151 when insurance status does not affect medical care needs and .147 when we allow this type of feedback. Moreover, the average value of P_{rel} rises less between 1986 and 1989 than the after-tax wage cost of health insurance. This is because TRA86 reduced marginal tax rates for some households but also raised the AGI threshold for deducting medical expenses from 5 to 7.5 percent. This reduced α, thereby raising the

17. This is consistent with Hausman and Poterba's (1987) finding that TRA86 actually raised marginal tax rates for more than one-third of taxpayers. For most of those who received rate reductions, these reductions were relatively small.

Table 5.4 **Tax Subsidies to Employer-Provided Health Insurance**

	Insured Employed			Uninsured Employed		
	1986	1989	1994	1986	1989	1994
dw_{AT}/dE	0.643	0.683	0.682	0.804	0.827	0.828
	(.123)	(.111)	(.117)	(.143)	(.157)	(.194)
P_{rel}						
$\eta=0$ to impute health spending if uninsured	0.814	0.839	0.837	0.926	0.945	0.945
	(.170)	(.149)	(.151)	(.136)	(.144)	(.173)
$\eta=-0.33$ to impute health spending if uninsured	0.779	0.813	0.811	0.932	0.949	0.950
	(.165)	(.145)	(.147)	(.137)	(.143)	(.171)

Source: Authors' tabulations based on imputation of tax rates to households in 1987 NMES. Each entry reports the average price weighted by family insurance spending. Column headings indicate which year's federal income tax and payroll tax schedule was used in constructing marginal tax rates.

Note: Numbers in parentheses are cross-sectional standard deviations.

after-tax cost of out-of-pocket medical care and partly offsetting the change in marginal rates.

We have also assessed the sensitivity of these results to variation in other tax parameters. We analyzed the effect of changing the AGI threshold for medical expenses, assuming 1994 tax rates, to illustrate how this aspect of the tax code affects incentives for insurance purchase. Lowering the AGI threshold from 7.5 to 2 percent, the pre-1983 level, raises the average value of P_{rel} for the insured employed from 0.837 to 0.902 when we assume a zero price elasticity of medical care demand, and from 0.811 to 0.868 when we assume an elasticity of -0.33. This policy change therefore has a much larger effect than the change in marginal tax rates under TRA86. It suggests that our recognition of the role of tax subsidies to out-of-pocket spending *can* be an important determinant of the after-tax price of insurance, even though at the current level of the medical expense deduction threshold, this effect is relatively small.

We also tried replacing the Feldstein-Samwick social security tax rate with the statutory tax rate. This induced a fine percentage point decline in the relative insurance price, highlighting the importance of accounting for tax-benefit linkages in the analysis of social security.

Table 5.4 makes it possible to compare the average relative prices for the insured and uninsured employed groups. The uninsured employed face higher average prices for employer-provided insurance than do their insured counterparts. This is because they are on average in lower income groups, and therefore face lower marginal tax rates, than the insured employed. Some of the observed relationship, of course, could reflect a demand curve for health insurance: those who face higher prices are less likely to buy insurance. Without more detailed analysis, however, it is impossible to disentangle the heterogeneity and demand curve effects.

For the uninsured employed, using statutory social security tax rates in place of the Feldstein-Samwick adjusted rates also causes a substantial increase in the measured tax subsidy to insurance. This is because a substantial fraction of the uninsured employed have earnings in the range over which the present discounted value of the social security benefit increment associated with an additional dollar of earnings offsets a substantial fraction of social security tax payments.

One reason for measuring the tax subsidy to employer-provided insurance is to estimate the efficiency cost associated with this tax expenditure. Table 5.5 presents the first step in any such calculation, our estimate of the price distortion induced by the tax system. Recall that if there were no tax distortions, $P_{rel} = 1 + \lambda$. We therefore use our estimate of P_{rel} with the actual tax system to compute $P_{rel}/(1 + \lambda) - 1$. The average value of this distortion declined from approximately 0.29 in 1986 to 0.26 by 1989; it has changed relatively little since then. It is notable that this estimate of the tax distortion is *smaller* than $1 - dw_{AT}/dE$, the distortion that is associated with the standard analysis of the tax incentive for fringe benefit provision (first row). In 1994, for example, the

Table 5.5 **Tax Distortions in the Relative After-Tax Price of Employer-Provided Health Insurance**

	1986	1989	1994
dw_{AT}/dE	0.357	0.317	0.318
P_{rel}			
$\eta=0$ to impute health spending if uninsured	0.286	0.261	0.263
$\eta=-0.33$ to impute health spending if uninsured	0.321	0.287	0.289

Source: Authors' tabulations based on imputation of tax rates from various years to households in 1987 NMES. Each entry reports the difference between the price of insurance with no tax subsidy (1 for the first row and 1.101 for the second two rows) and the price with the tax subsidy. Sample is employer-insured individuals.

estimated tax subsidy is 20 percent smaller than the traditional measure of the tax price. Since the deadweight loss from this subsidy rises with the square of the size of the subsidy, our calculations imply that the deadweight loss from tax subsidization is roughly 40 percent less than would be implied using the change in after-tax wages.

Our estimate of the distortion based on P_{rel} changes less between 1986 and 1989 than the average value of dw_{AT}/dE because we recognize the role of employee contributions to employer-provided health insurance, which dampen the tax subsidy, and because we include the 1986 increase in the AGI threshold for deducting out-of-pocket medical expenses in our analysis of the relative price of insurance. Our estimates therefore imply that the reduction in deadweight loss from the package of tax changes under TRA86 is much smaller than might be supposed based only on a comparison of dw_{AT}/dE at different points in time.

5.4 Capping the Amount of Tax-Exempt Employer-Provided Health Insurance

The revenue loss associated with the tax expenditure for employer-provided health insurance, and a perception that overinsurance has contributed to the rise in U.S. health care costs during the last two decades, has led to numerous proposals to alter the current tax treatment of employer-provided health insurance. One of the most common reform proposals, and one discussed at length in CBO (1994b), is capping the value of employer-provided insurance that could be excluded from taxable income. One special case of such caps would be complete inclusion of the value of employer-provided insurance in employee taxable income. In this section, we use our augmented NMES database to explore how various tax caps would affect the after-tax price of employer-provided health insurance, the demand for such insurance, and tax revenues.

There are many ways to tax employer-provided health insurance. Burman and Williams (1994) provide a detailed discussion of several options, including changes in the corporate tax deductibility of such insurance payments, as well as including some or all of the value of these benefits in the personal income tax base. Our analysis focuses on the case in which employers report the value

of employer-provided insurance benefits along with an employee's wage income and these benefits are then incorporated in the federal and state income tax base and the federal payroll tax base. In principle, there is no reason that employer-provided insurance needs to be taxed in all three forms, and one could disentangle the revenue effects associated with different types of incremental tax changes. We discuss the importance of the relative sources of revenue below.

By considering the case in which tax caps are implemented through the personal income tax, our analysis applies to tax caps that are specified for tax-filing units rather than individuals. Such family-level caps would be very difficult to implement through any system that relied on changes in the corporate rather than personal income tax. If caps were applied to individuals rather than families, two-earner couples in which both earners had an opportunity to receive employer-provided insurance could reduce the impact of the caps by choosing two individual policies rather than a single family policy. Caps on tax-unit health insurance benefits such as those we consider provide a strong incentive for households with two earners to eliminate duplicative insurance coverage since such insurance is likely to provide relatively little health benefit but could lead to a substantial increase in taxable income.

5.4.1 Analyzing Tax Caps

At the outset, we should recognize several basic points about tax caps. First, the cap should not affect the behavior of anyone who receives employer-provided health insurance benefits worth *less* than the cap.[18] Second, absent income effects, no one whose employer-provided health insurance exceeds the cap value prior to imposition of the cap should reduce his insurance outlays to less than the capped level. A system of tax caps would provide strong incentives for employers and employees to restructure benefits packages to reduce the fraction of health insurance value that exceeds the cap and to maximize the chance that the cost of coverage above this cap is paid by the employee, who may have an opportunity to deduct some insurance costs as itemized medical deductions.[19] One example of such a reaction would be scaling back the set of services covered by the employer-provided insurance plan, while introducing a cafeteria plan to allow workers to pay some of these costs with pretax dollars. We ignore any such responses to tax caps in computing the revenue and behavioral effects below, but they could be important in practice.

We estimate how tax caps of various dollar amounts would affect the average after-tax relative price of health insurance and out-of-pocket spending by

18. It is possible that caps on the excludable amount of employer-provided health insurance may reduce the demand for generous coverage from some employees who previously received benefits worth more than the cap and that this will work through the negotiation process that results in a benefits and wage package to reduce the level of health benefits.

19. We consider a tax cap that applies to federal and state income taxes, as well as payroll taxes. If the cap were only applied to federal income taxes, the incentive to reduce health insurance value above the capped amount would be smaller.

setting the relative tax price for any NMES family with employer-provided insurance above the cap to $(1 + \lambda)/(1 - \alpha\tau)$. This is just the expression for the after-tax cost of health insurance in equation (3), with $\tau = \tau_s = \tau_{ss} = 0$ in the numerator. Since the families who are most likely to be affected by any cap are those with high incomes, high marginal tax rates, and therefore high values of the tax subsidy before the cap, the change in the average after-tax relative price of insurance can be substantial even if the number of households affected by the cap is small. Throughout our tax cap analysis we use marginal α in evaluating equation (3) since this is the appropriate parameter for evaluating a marginal reduction in spending on employer-provided insurance as would be associated with a tax cap.

After describing the change in the after-tax relative price of employer-provided health insurance associated with the tax caps, we present illustrative calculations of how these caps would affect the demand for employer-provided insurance. We assume that caps would apply to all employer-provided insurance that was subsidized before the cap was enacted, regardless of whether this insurance was paid for by employers or employees making pretax contributions.[20]

For each NMES family, we compute the marginal after-tax relative price of employer-provided insurance under the status quo ($P_{rel,0}$) and under the assumption that employer-provided insurance above the cap is included in taxable income ($P_{rel,1}$). If these two prices are identical, we assume that the individual would not change his demand for employer-provided insurance (E_0). If the two prices are different, however, we estimate the individual's demand for health insurance at the new price as

$$E_1 = \max (C, E_0 *[1 + \eta*(P_{rel,1}/P_{rel,0} - 1)]), \tag{4}$$

where C denotes the level of the cap. The parameter η is the uncompensated price elasticity of demand for health insurance. We do not consider any income effects on the demand for insurance that might be associated with the introduction of tax caps. If E_1 is greater than C, then we take E_1 as the new level of employer-provided insurance. If E_1 is less than C, however, we assume that $E_1 = C$ and that the individual will locate at the kink point on the budget set. To find the aggregate change in the demand for employer-provided insurance as a result of a cap on the value of excludable benefits, we compute the sample-weighted sum of the changes in E_1 across all NMES households.

Our calculations make the strong assumption that each household affected by the tax cap can adjust the quantity of employer-provided health insurance that it receives in response to this tax policy change. This assumption is unrealistic since most workplaces offer only a few discrete choices with respect to health insurance coverage. Moreover, since individual employees cannot deter-

20. Excluding pretax employee contributions from the tax cap would result in a simple tax avoidance strategy. Firms would reduce their employer-provided health insurance but permit employees to purchase equivalent insurance on a pretax basis. This would circumvent the tax caps.

mine what benefits package their employer will offer, changes in the tax circumstances of an individual worker may not be reflected in a differential level of employer-provided insurance. Recognizing the important heterogeneity in tax preferences and insurance demand within workplaces, and incorporating this into the analysis, is therefore an important direction for future work.

A critical parameter in our calculation is η, the price elasticity of demand for health insurance. There are relatively few estimates of this parameter, and available estimates differ substantially (see Gruber and Poterba [1994] for a detailed review). There are also many different margins along which employers might alter their health insurance offerings, and it is not clear that elasticities of demand would be the same on all margins. For example, employers could reduce the value of insurance coverage provided to their workers by limiting the set of services covered, by raising copayment rates or deductibles, or by requiring a higher employee contribution for a given insurance policy. Previous studies, and our analysis below, treat adjustments on all of these margins as equivalent.

Previous cross-sectional studies of the price elasticity of demand for health insurance can be grouped into three types.[21] The first set of studies compare the quantity of health insurance demanded by high- and low-income households that face different marginal tax rates; these studies have produced a wide range of elasticity estimates.[22] The second set of studies consider evidence from hypothetical offers of supplemental insurance to participants in the RAND Health Insurance Experiment, reported in Marquis and Phelps (1987). This randomized experiment assigned individuals to plans with different copayment rates, with an out-of-pocket maximum of up to $1,000. At the end of the experiment, individuals were presented with hypothetical offers for supplemental insurance to lower their out-of-pocket exposure; the price of these offers varied across participants. The resulting elasticity of demand for the quantity of supplemental insurance was -0.6.

The third source of information on the price elasticity of demand is evidence from the take-up of price subsidies that were offered to small firms under experimental pilot projects. Thorpe et al. (1992) found an elasticity of demand of insurance coverage of between -0.07 and -0.33 for these firms. Gruber and Poterba (1994) suggested a price elasticity of demand for insurance coverage of -1.0 or greater in absolute value for self-employed individuals, focusing on tax changes to identify shifts in the after-tax price of insurance for this group. In light of this variation, we set $\eta = -0.5$ in our baseline case, and we also report analyses using values of -1.0 and -0.2.

The final aspect of the tax caps that we consider is their effect on total reve-

21. There are also a number of time-series studies, such as Turner (1987).

22. Examples of other studies that estimate the price elasticity of demand include Taylor and Wilensky (1983), who report an elasticity of -0.2; Woodbury (1983), who reports -1.7 to -3.5; Holmer (1984), who reports -0.16; Sloan and Adamache (1986), who report -0.6; and Woodbury and Hamermesh (1992), who report -2 to -3.

nue collections. We combine FICA, federal income tax, and state income tax revenue in our tabulations. Tax caps affect tax revenues in two ways. First, they collect taxes directly on employer-provided health insurance benefits that are valued at more than the tax cap. In addition, however, if some employees decide to reduce their demand for employer-provided health insurance as a result of the tax cap and its associated increase in the marginal cost of insurance, then their *taxable wages* will rise as their employer-provided health insurance benefits decline. We assume that any reduction in employer-provided insurance will be reflected dollar for dollar in pretax wage payments to workers.

The relative importance of the taxes collected on insurance benefits worth more than the cap and on increased taxable wages depends on the price elasticity of demand for health insurance. The *total* revenue collected as a result of the tax cap is independent of this elasticity, however, and just depends on the total value of employer-provided health benefits above the tax cap in the precap setting, that is, on the sum of $E_0 - C$ across households.

Finally, one important caveat to the results below is that we are using the total insurance expenditures of the HIU to identify the effect of tax caps. For some of the HIUs in our sample, insurance expenditures reflect employer-provided insurance coverage to both spouses. A cap that was imposed on each spouse separately would therefore have smaller effects than those estimated below.[23] For our base case described below, only 16 percent of HIUs (21.5 percent of couples) have multiple insurance policies. Among those couples who face binding tax caps, however, 32 percent have more than one policy (this is 24 percent of our total sample). This calculation overstates the effect of dual policies since, in some families with dual coverage, both spouses may have policies that exceed the cap. The problem of dual policies therefore does not appear to be an important limitation in applying the calculations reported below.

5.4.2 Results on Tax Caps

Table 5.6 presents our basic findings on tax caps for the case with a price elasticity of demand for health insurance equal to -0.5. We consider four tax caps, all denominated in 1994 dollars.[24] Our base case, shown in column (2), follows the caps suggested by CBO (1994a): $4,000 per year for joint filers, $1,600 per year for single filers, and $3,400 per year for heads of household. We then show the effects of (a) doubling these caps (col. [1]), (b) halving these

23. If the cap were imposed on insurance spending by tax-filing unit, as it would be if it were implemented through the individual income tax system, then the cap would apply to total family insurance spending. This would create a strong incentive for families to drop duplicative insurance policies and might lead to a larger response in the quantity of insurance demanded as a result of the tax cap.

24. By inflating 1987 expenditures in the NMES by the deflator for personal health care spending, we effectively index the tax caps to the medical cost deflator. If tax caps were indexed to the consumer price index rather than an index of medical care costs, the caps would become more stringent over time if health care inflation continues to outpace overall inflation.

Table 5.6 **Effect of Capping Employer-Provided Health Insurance Deduction Assumptions: 1994 Tax Code, $\eta = -0.5$**

	Level of Tax Caps (1994 $): Joint Filer/Single Filer/Household Head			
	8000/3200/6800 (1)	4000/1600/3400 (2)	2000/800/1700 (3)	0/0/0 (4)
1. P_{rel}	0.870 (.155)	0.954 (.146)	1.044 (.120)	1.111 (.060)
2. *Employed workers affected* (%)	0.119	0.509	0.704	0.821
3. *Employed insured affected* (%)	0.144	0.619	0.857	1.000
Changes in insurance demand[a]				
4. Average change in E	−306	−985	−1,267	−1,308
5. Average change in E if a change	−2,123	−1,590	−1,477	−1,361
Tax increase per insured employee[a]				
6. Tax on insurance benefits	30.0	190.5	532.2	1,067.0
7. Tax on higher wages	119.2	374.1	479.9	495.2
8. Total	148.2	564.6	1,012.1	1,562.2
Aggregate Revenue Raised[b]				
9. Tax on insurance benefits	1,690	11,123	31,075	62,303
10. Tax on higher wages	6,762	21,846	28,023	28,915
11. Total	8,652	32,969	59,097	91,218

Notes: Insurance market responses assume a price elasticity of demand −0.5 for employer-provided insurance. The base case value for the relative after-tax price in the first row, for the case with unlimited tax exclusion, is 0.837 (.151) as shown in table 5.4. Revenue effects on wage taxes assume that wages rise by the full amount of any reduction in employer-provided insurance. See text for further details.

[a]In dollars

[b]In million dollars.

caps (col. [3]), and (c) setting the caps to zero, so that all employer-provided health insurance benefits are included in taxable income (col. [4]).

The results show that even tax caps that affect relatively few households can have substantial effects on the average relative after-tax price of employer-provided health insurance. The 8000/3200/6800 cap, which would have been binding for 14.4 percent of employed insured workers in 1987, raises the average value of P_{rel} from 0.837 under the status quo to 0.870. The cap analyzed by the CBO has an even larger effect, with P_{rel} rising to 0.954.[25] This cap would affect over one-half of insured employees. Introducing caps first increases the variance of the relative after-tax price in the population, but as the share of

25. The relative price of insurance can in principle be greater than 1 because of the loading factor on insurance and the subsidization of self-insurance.

households affected by the cap rises and more and more households face $(1 + \lambda)/(1 - \alpha\tau)$ as their relative tax price, the variance declines. In column (4), which corresponds to eliminating the tax subsidy for employer-provided insurance, the variance of after-tax relative prices falls substantially.

Table 5.6 next reports our estimates of the change in the level of employer-provided insurance associated with each set of tax caps (rows 4 and 5). For the base case, for example, we calculate that the average reduction in insurance spending will be $565 per insured employee, or $33 billion. Because tax caps affect only a fraction of those employees with employer-provided health insurance, the decline in insurance levels for those who are affected by the caps is substantially larger than the average decline for all employees.

The entries in column (4) of rows 4 and 5 warrant particular note. Our estimates with a price elasticity of insurance demand of -0.5 suggest that eliminating the tax exemption for employer-provided health insurance would reduce the aggregate value of this insurance by 30.8 percent. This corresponds to an average per capita reduction of $1,308.[26]

Rows 6–8 present information on the revenue effects of changing the tax treatment of employer-provided insurance. We present the total revenue collected per insured employee (row 8) as well as the decomposition of this revenue between the tax on insurance premia above the cap and the tax on higher wages that result from reductions in employer-provided insurance. Rows 9–11 report the aggregate revenue consequences of each of these policies. We report the total increase in federal income tax and payroll tax revenue, as well as the small increase in state income tax revenue.

The entry in row 11 of column (4) shows that we estimate that elimination of the tax exemption for employer-provided insurance would have raised $91.2 billion (1994 dollars). More than two-thirds of this revenue is raised from taxes on the insurance that remains in force after the tax subsidy is removed. This estimate is about one-fifth higher than the estimate presented in CBO (1994b), even though the CBO includes both the revenue collected by taxing employer-provided insurance and the revenue collected from higher wage taxes. The CBO excludes some employer-provided insurance which may not be employment-related and so begins with a smaller annual flow of employer-provided insurance than we do.[27] Our estimates are quite similar to the Joint Committee on Taxation's estimates

26. Table 5.6 indicates that the average change in employer-provided insurance for those affected by the elimination of the tax exclusion is larger than the average change for all employed insured. There are 256 NMES respondents who report that they are insured by their employers and who have some out-of-pocket spending on insurance, but who report zero employer contributions for their insurance. The three-quarters of this group for whom out-of-pocket insurance spending is not tax preferred will be unaffected by the repeal of the tax exemption.

27. Another potential difference between our estimates and those of the CBO relates to our estimate of the prereform distribution of insurance spending. We use the actual reported distribution of employer-provided insurance premiums in the NMES, while the CBO made adjustments that lowered the estimated expenditures for high-income (and high tax rate) families. We are grateful to Roberton Williams for suggestions with regard to these disparities.

of the total revenue cost of the tax expenditure for employer-provided insurance.[28]

The revenue estimates for various tax cap proposals provide an indication of how much revenue could be raised by each alternative. Our base case 4000/1600/3400 cap raises roughly one-third as much revenue as the total elimination of the tax exclusion for employer-provided insurance. A much higher tax cap of 8000/3200/6800 only raises $8.6 billion.

It is also quite interesting to consider the implications of alternative plans for insurance expenditures and revenue raising. The base case plan, which affects only 62 percent of insured workers, reduces insurance expenditures by 75 percent as much as removing the tax exclusion altogether. This is because, due to the somewhat skewed distribution of insurance spending, the 38 percent of HIUs that are not affected by the cap do not spend much on insurance. On the other hand, fully removing the tax exclusion raises almost three times as much revenue. Thus, as the tax cap is tightened, there will be smaller marginal gains in terms of reducing "overinsurance," but larger gains in terms of revenues.

The estimates in table 5.6 assume that the price elasticity of demand for health insurance is -0.5. This is not a behavioral parameter that commands a strong empirical consensus, so we also present estimates of the change in insurance demand and the mix of increased revenues for two alternative elasticity estimates: -0.2 and -1.0. Table 5.7 presents these results. With a price elasticity of -1.0, we find that eliminating the tax exclusion for employer-provided insurance results in a decline of $2,609 in the quantity of insurance purchased, which is a 61 percent reduction. With a elasticity of -0.2, not surprisingly, the quantity adjustment is much smaller and corresponds to approximately a 10 percent decline in the value of employer-provided insurance. Whether this entire reduction in employer-provided insurance translates into a greater share of medical care being purchased on an out-of-pocket basis depends on whether individuals replace some employer-provided insurance with directly purchased insurance, an issue that we have not yet explored.

The source of increased revenue, whether taxation of insurance premia or taxation of higher wages, also is sensitive to our assumed elasticity. Using the elasticity of -1.0, more than half of the new revenue generated from eliminating the tax subsidy comes from taxing wages, while with an elasticity of -0.2, almost 90 percent of the revenue comes from the tax on insurance benefits. As we noted earlier, the total revenue collected is not sensitive to our elasticity assumptions, only the decomposition across revenue sources.

We have also explored the sources of increased revenues under our tax cap

28. There is one reason to suspect that our results may *underestimate* the change in taxes from tax caps. We calculate the revenue effects of taxing health insurance spending by multiplying changes in taxable income by the taxpayer's current marginal tax rate, ignoring any movements across tax brackets that might result from taxation of employer-provided insurance. Since the tax code is progressive, this should lead our calculations to underestimate the actual revenue gain.

Table 5.7 **Sensitivity of Results on Capping Employer-Provided Health Insurance Deduction to Assumptions about Price Elasticities of Insurance Demand**

	Level of Tax Caps (1994 $): Joint Filer/Single Filer/Household Head			
	8000/3200/6800	4000/1600/3400	2000/800/1700	0/0/0
	Elasticity = −1.0			
Changes in insurance demand				
Average change in E	−383	−1,437	−2,318	−2,609
Average change in E if a change	−2,654	−2,320	−2,704	−2,715
Tax increase per insured employee[a]				
Tax on insurance benefits	1.9	26.7	143.3	575.1
Tax on higher wages	146.2	538.0	868.8	987.1
	Elasticity = −0.2			
Changes in insurance demand				
Average change in E	−156	−437	−511	−523
Average change in E if a change	−1,083	−706	−596	−544
Tax increase per insured employee[a]				
Tax on insurance benefits	86.6	398.1	819.6	1,364.1
Tax on higher wages	61.5	166.6	193.6	198.1

Notes: Revenue effects on wage taxes assume that wages rise by the full amount of any reduction in employer-provided insurance. See text for further details.

[a]In dollars.

plans. For our base case plan (4000/1600/3400), approximately 56 percent of the revenues raised are federal income tax revenues. Another 33 percent are raised by the social security and Medicare taxes, with the remaining 11 percent being raised by state taxes. The distribution is very similar for alternative tax caps.

One final aspect of taxing employer-provided health insurance that our data can inform concerns the distribution of binding tax caps across income classes. Table 5.8 presents summary statistics on this issue. Each column corresponds to a different set of tax caps, from the previous tables, but now the entries show the fraction of NMES families in a given income category that would be constrained by each cap. For our base case 4000/1600/3400 cap in column (2), for example, the table shows that the cap would bind for 27.1 percent of the families with incomes between $10,000 and $20,000, compared with 67.8 percent of those with incomes between $75,000 and $100,000. The sample sizes for high-income groups in the NMES are relatively small, but the results at least illustrate the general pattern across income classes. All of the tax caps except the highest are binding for the majority of HIUs by approximately $35,000 in family income; removing the tax exclusion entirely binds for the majority of families with more than $10,000 of income.[29]

29. One potentially puzzling feature of table 5.8 is the failure to observe "100 percent" for each of the entries in col. (4) of panel B. This disparity arises because there are some NMES respon-

Table 5.8 Distribution of Binding Tax Caps by Family Income Class

AGI Class (thousand 1994 $)	Level of Tax Caps (1994 $): Joint Filers/Single Filers/Household Heads			
	8000/3200/6800 (1)	4000/1600/3400 (2)	2000/800/1700 (3)	0/0/0 (4)
	A. *All Persons*			
Under 10	3.4	17.9	26.5	30.3
10–20	5.9	27.1	43.7	50.9
20–30	10.4	47.6	68.9	75.8
30–40	10.3	53.4	75.9	86.1
40–50	13.0	59.9	82.8	92.9
50–75	14.2	63.0	84.9	94.2
75–100	19.9	67.8	86.1	94.1
Over 100	19.0	66.4	85.5	95.3
Total	11.8	51.1	70.6	78.8
	B. *Insured Persons Only*			
Under 10	9.9	53.2	78.6	89.9
10–20	10.9	50.0	80.5	93.9
20–30	13.0	59.0	85.4	94.0
30–40	11.5	59.8	84.9	96.4
40–50	13.6	63.1	87.1	97.6
50–75	14.6	64.6	87.1	96.7
75–100	20.6	70.3	89.2	97.5
Over 100	19.3	67.5	86.9	96.9
Total	14.4	62.2	85.9	96.0

Source: Authors' calculations using 1987 NMES data. Each entry shows the percentage of employed individuals who would be affected by tax caps of the magnitudes indicated. Results in col. (4) of panel B are not equal to 100.0 because there are 181 NMES respondents reporting employer-provided insurance but no spending ($E = G = 0$).

The rising incidence of binding caps at higher income levels reflects both the rising probability of having employer-provided insurance at higher income levels and the rising value of average premiums conditional on such insurance. In order to separate these factors, panel B of table 5.8 repeats these calculations for those with insurance. Here we can see that all except the most generous tax caps bind for the majority of insured persons at any income level, and the gradient with respect to income is much less steep.

5.5 Conclusion

Our analysis emphasizes two aspects of the current tax subsidy to employer-provided health insurance and presents new evidence on the economic effects

dents who have employer-provided health insurance, but who report zero employer expenditure (E). They do report positive out-of-pocket spending (G) on insurance, but given our assumption that only one-quarter of households with such expenditures make them on a pretax basis, even a zero tax cap does not bind for 75 percent of these households.

of various tax reforms. The conceptual points we emphasize suggest that the current federal tax code subsidizes employer-provided insurance less than many previous analyses would suggest. This is because a substantial and growing share of employees who receive employer-provided insurance must pay for part of this insurance with their own after-tax dollars, and because the tax code also provides a deduction for extreme medical expenses, thereby to some degree discouraging individuals from purchasing health insurance. Our empirical analysis of the effect of capping the value of employer-provided health insurance that could be excluded from taxation, or eliminating the exclusion entirely, suggests that these reforms could have substantial effects on the level of employer-provided insurance.

There are a number of important issues associated with both the determinants of the level of employer-provided insurance and the effect of tax reforms on this insurance that we have not addressed. One issue is the role of joint decision making in workplace benefits. We have not considered how to aggregate the heterogeneous changes in tax incentives for employer-provided insurance that would accompany many tax cap plans into decision rules for firms. We have also stopped short of asking whether changes in tax incentives would lead to different combinations of workers into firms or health insurance units. If tax reform led to greater heterogeneity in worker tastes for employer-provided health insurance, employers might respond by offering larger menus of insurance policies. This could be important for revenue estimation and could also have welfare implications.

A second important issue concerns general equilibrium effects in the health insurance markets. If a substantial number of currently insured workers decide not to purchase insurance under some of the tax reforms we consider, it is possible that the load factor facing those who remain in the insured pool may change. This could affect the demand for health insurance even by those who do not face tax caps.

A third issue is modeling the appropriate demand response to changes in the tax price of insurance. We have assumed a constant elasticity demand function and applied this elasticity equally to "looser" and "tighter" caps. In fact, individuals may be quite elastic with respect to insurance coverage on the margin, but less elastic when it comes to dropping their insurance entirely. Extending this analysis to consider a richer range of responses to different tax caps is an important step for future research.

Finally, this paper represents a strictly positive exercise. We have not considered any of the interesting normative issues surrounding the tax treatment of health insurance. One important argument for subsidizing workplace insurance is that workplace pooling, which is largely exogenous to underlying health, avoids the classic adverse selection problems in individual insurance markets. Rothschild and Stiglitz (1976) discuss the theoretical possibility that private insurance markets may fail, but there is little empirical evidence on the extent of such failures and their welfare consequences. Such an argument would

imply welfare losses if the removal of this tax subsidy led to the breakup of workplace pools. What is not clear, however, is the extent to which the tax subsidy, as opposed to other gains from pooling, is responsible for holding workplace pools together. If they can be measured, these pooling gains must be weighed against the distortions from excess consumption of medical care in deciding on the optimal level of tax subsidies.

More generally, the question of whether private insurance purchases should be subsidized depends on a host of unresolved issues, such as the degree to which uninsured individuals consume uncompensated care, how the costs of such care are shifted to paying health care consumers, the role of health insurance in affecting labor market behavior, and the other positive and negative externalities that a more-insured population may provide. All of these issues require further investigation.

References

Blostin, Allan P., Robert B. Grant, and William J. Wiatrowski. 1992. Employee payments for health care services. *Monthly Labor Review* 117 (November): 17–32.

Bradford, David F., and the U.S. Treasury Tax Policy Staff. 1984. *Blueprints for basic tax reform*, 2d ed. Arlington, Va.: Tax Analysts.

Burman, Leonard E., and Roberton Williams. 1994. Tax caps on employment-based health insurance. *National Tax Journal* 47 (September): 529–45.

Feldstein, Martin S. 1973. The welfare loss of excess health insurance. *Journal of Political Economy* 81 (March–April): 251–80.

Feldstein, Martin S., and Elisabeth Allison. 1974. Tax subsidies of private health insurance: Distribution, revenue loss, and effects. In *The economics of federal subsidy programs*, U.S. Congress, Joint Economic Committee, 977–94. Washington, D.C.: Government Printing Office.

Feldstein, Martin S., and Andrew Samwick. 1992a. Social security rules and marginal tax rates. *National Tax Journal* 45:1–22.

———. 1992b. Social security rules and marginal tax rates. NBER Working Paper no. 3962. Cambridge, Mass.: National Bureau of Economic Research.

Gruber, Jonathan. 1994. The incidence of mandated maternity benefits. *American Economic Review* 84 (June): 622–41.

Gruber, Jonathan, and Alan B. Krueger. 1991. The incidence of mandated employer-provided insurance: Lessons from workers' compensation insurance. In *Tax policy and the economy*, vol. 6, ed. David Bradford, 111–44. Cambridge: MIT Press.

Gruber, Jonathan, and James M. Poterba. 1994. Tax incentives and the demand for health insurance: Evidence from the self-employed. *Quarterly Journal of Economics* 109 (August): 701–33.

Hausman, Jerry A., and James M. Poterba. 1987. Household behavior and the Tax Reform Act of 1986. *Journal of Economic Perspectives* 1:101–19.

Holmer, Martin. 1984. Tax policy and the demand for health insurance. *Journal of Health Economics* 3:203–21.

Kaplow, Louis. 1991. The income tax as insurance: The casualty loss and medical expense loss deductions and the exclusion of medical insurance premiums. *California Law Review* 79 (December): 1485–1510.

———. 1992. Income tax deductions for losses as insurance. *American Economic Review* 82 (September): 1013–17.

Lindsey, Lawrence. 1987. Individual taxpayer responses to tax cuts, 1982–1984: With implications for the revenue-maximizing tax rate. *Journal of Public Economics* 33:173–206.

Long, James E., and Frank A. Scott. 1982. The income tax and nonwage compensation. *Review of Economics and Statistics* 64:211–19.

Marquis, M. Susan, and Charles E. Phelps. 1987. Price elasticity and adverse selection in the demand for supplementary health insurance. *Economic Inquiry* 25:299–313.

Newhouse, Joseph P., et al. 1989. Adjusting capitation rates using objective health measures and prior utilization. *Health Care Financing Review* 10 (Spring): 41–54.

Pauly, Mark. 1986. Taxation, health insurance, and market failure in the medical economy. *Journal of Economic Literature* 24:629–75.

Phelps, Charles E. 1992. *Health economics.* New York: Harper Collins.

Rothschild, Michael, and Joseph Stiglitz. 1976. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics* 90:629–49.

Scofea, Laura A. 1994. The development and growth of employer-provided health insurance. *Monthly Labor Review* 117 (March): 3–10.

Sloan, Frank, and Killard Adamache. 1986. Taxation and the growth of nonwage benefits. *Public Finance Quarterly* 14:115–39.

Taylor, Amy, and Gail Wilensky. 1983. The effect of tax policies on expenditures for private health insurance. In *Market reforms in health care: Current issues, new directions, strategic decisions,* ed. Jack Meyer. Washington, D.C.: American Enterprise Institute.

Thorpe, Kenneth L., et al. 1992. Reducing the number of uninsured by subsidizing employment-based health insurance: Results from a pilot study. *Journal of the American Medical Association* 267:945–48.

Turner, Robert. 1987. Are taxes responsible for the growth of fringe benefits? *National Tax Journal* 40:205–20.

U.S. Bureau of Labor Statistics. 1993. *Employee benefits in medium and large private establishments, 1991.* Bulletin no. 2422. Washington, D.C.: U.S. Department of Labor.

———. 1994. *Employee benefits in small private establishments, 1991.* Bulletin no. 2441. Washington, D.C.: U.S. Department of Labor.

U.S. Congressional Budget Office (CBO). 1993. *Projections of national health expenditures, 1993 update.* Washington, D.C.: U.S. Congress, Congressional Budget Office.

———. 1994a. *Reducing the deficit: Spending and revenue options.* Washington, D.C.: U.S. Congress, Congressional Budget Office.

———. 1994b. *The tax treatment of employment-based health insurance.* Washington, D.C.: U.S. Congress, Congressional Budget Office.

Woodbury, Stephen A. 1983. Substitution between wage and nonwage benefits. *American Economic Review* 73:166–82.

Woodbury, Stephen A., and Daniel S. Hamermesh. 1992. Taxes, fringe benefits, and faculty. *Review of Economics and Statistics* 73:287–96.

Comment

David F. Bradford

This paper presents measurements of the incentive effected by tax rules for employers to provide health insurance benefits to their employees. The problem poses conceptual and empirical challenges. The bulk of the work of the paper describes the methods by which the authors tease quantitative estimates of average subsidy rates out of, mainly, two disparate data sources (the National Medical Expenditure Survey and the U.S. Treasury Individual Tax Model file). But I propose to focus my comments on the conceptual side of the story.

I hope that my taking up conceptual issues is not taken to imply a lack of appreciation for the ingenuity and perseverance that the authors have devoted to the quantitative estimates. It is an extraordinary job, taking into account a great many fine points of the data sets and institutional setting.

The principal conceptual issues I would raise relate to the interpretation of the empirical work. The authors describe their analysis as limited to the question of the determination of how the tax system influences the division between employer-provided payments and out-of-pocket outlays by the employee in financing a *given* package of health care services. The distinction is made between this problem and that of determining the impact of the tax system on the level of health care services demanded. It is taken for granted that, from the perspective of the employee, the payments by the employer take the form of insurance, whereas the own out-of-pocket payments constitute the lack of insurance coverage that we describe as self-insurance.

The critical distinction drawn in this paper, then, is not between subsidized health care and unsubsidized health care, but between health care covered by insurance and that not covered by insurance. Although a richer set of variations is recognized in the analysis of policy experiments, the basic analysis seems to treat "covered by insurance" as a zero-one variable. Either one is covered by an employer-sponsored plan or one is self-insured.

Presumably, the reason for being interested in this distinction is the *moral hazard* associated with insurance: Put simply, a person covered by insurance has insufficient incentive to economize on health care services. Moral hazard is of particular importance in the case of health care coverage because the insured-against event has a large subjective element. One can usefully think of moral hazard as arising under health insurance at two points: First, if I am insured, my incentive is reduced to preserve my health by, for example, eating well and exercising regularly. Second, as typically structured (and it, perhaps, need not be so), health insurance alters my incentive to economize on the services deployed to deal with health problems that arise. Having concluded that I am sick, I have the option of a large range of treatments, ranging from letting

David F. Bradford is professor of economics and public affairs at the Woodrow Wilson School at Princeton University.

nature take its course to consulting a world-famous specialist. Insurance reduces, perhaps to zero, the marginal cost of choosing relatively expensive methods of dealing with any given health problem. I would guess that the second sort of moral hazard is the far more serious problem quantitatively.

If, in addition to subsidizing expenditures on health care, the tax system induces people to be covered by insurance, there is an "extra moral hazard" cost imposed on top of the usual deadweight loss. The authors discuss this overinsurance effect in connection with their simulation of the impact of various caps on employer-provided health care allowed to be excluded from the employee's taxable income.

I find myself in some doubt about the adequacy of the implicit model of insurance market equilibrium employed in the paper. (I use the term "some doubt" advisedly, given the complexity of the market and the needed analysis.)

First, as to the effect of a subsidy on the terms of insurance. Private insurance contracts may involve any of a variety of methods of internalizing the moral hazard problem. Classic approaches are to specify a deductible amount (so as to eliminate the moral hazard in the choice of service level in situations in which the loss to the insured is relatively small) and co-insurance rate (so that the extra cost to the insured of choosing expensive treatment is raised above zero). Other methods include monitoring and regulation by the insurance company, to reward my choice of a healthy life-style, to limit my freedom to identify myself as sick, and to restrict my choice of treatments according to some criterion of medical necessity.

Economic theory predicts that such devices to moderate moral hazard will be equilibrium phenomena. That is, they will be chosen by mutual agreement between insured and insurer. The insured will accept the various limits on coverage in return for the saving on the cost of the insurance. The saving will be greater than the expected cost of "necessary" treatment that will be borne by the insured, because the insured will have an incentive (or be obliged by the monitoring regime) to limit treatment to what really is necessary (as distinct from what would be chosen if there were no limit or cost to the insured).

A second classic problem of insurance is *adverse selection.* The possibility of adverse selection is due to asymmetric information, whereby I know I am at high risk of incurring health care expenses, but the insurance company cannot observe this fact. If insurance is priced to break even with average risks, I will tend to choose relatively extensive insurance coverage. The insurance company will then lose money. Adverse selection problems can also arise as a result of regulatory requirements that limit the ability of insurance companies to vary the premium charged according to the health characteristics of the insured. Adverse selection seems a likely explanation for the often-noted high loading and limited availability of individual health insurance policies.

The devices to deal with adverse selection include physical exams, limits on preexisting coverage, and similar techniques to break into the information advantage of potential insureds. These measures can be costly. An important

technique that economizes on such underwriting procedures is to bundle the offering of insurance with some other choice that is likely to dominate the insurance in importance. The choice of employment is a prime example. (The point is recognized by the authors.) The fact that a particular employer offers a particular health insurance plan is not very likely to play a large role in the determination of who will become an employee. To the extent this is true, employers can offer health insurance on terms that may be better than those obtainable by employees in a private, individual insurance market. Other features of the terms of employment may also present opportunities to mitigate adverse selection. For example, good health benefits may attract high health risks, but putting a substantial part of the compensation in the form of retirement annuities will deter those who do not expect to live long.[1]

Thus, economic theory gives us grounds to expect that employer-based health insurance would be observed in unregulated equilibrium, even without any tax incentives. If I approach this analysis with the expectation that (1) most people are going to prefer some insurance coverage to self-insurance and (2) employer-based plans are very likely to be common even in the absence of a subsidy, it seems to me that I am led to a different interpretation of the results.

To isolate the point (and I cannot hope to settle it), imagine that employer-based insurance (or something economically equivalent) is exactly what we would predict in the absence of a tax subsidy. Then the subsidy rates developed by the authors are, as a first approximation, those that apply to the purchase of health care, and only secondarily to the characteristics of the insurance contract involved.

To make this point clearer, I have played devil's advocate with my tax lawyer colleagues, asking whether I could write a contract with my employer with the following terms: I will take my pay in the form of health care services to be purchased by my employer as needed, as determined by me, subject to the IRS rules about what constitutes health care. If (as I hope and expect) the employer's outlays fall short of a specific level (my current salary) the employer will make up the difference in cash. The income tax law makes explicit provision for the exclusion from the employee's taxable income of health care services supplied by the employer. *If* the plan I described to my lawyer friends were feasible from the standpoint of the tax law, it would leave me in the position of obtaining the full advantage of the exclusion of health care benefits from tax, but without any insurance beyond the implicit insurance provided by the tax rate. (Even that might be eliminated if we worked out that the cash settlement varied by *more* than the employer's outlays.) But you would not be able to observe this fact from data on employee health care outlays.

There are two problems. First, my lawyers have pretty well convinced me that my plan would not work. Second, who would want such a plan? There

1. Sherry Glied, who pointed this out to me, suggests that the empirical magnitudes make this a plausible balance.

may, however, be schemes as yet undiscovered that would move a good distance in the direction I described (such as a commitment to experience rating of my coverage) and that would pass legal muster. This question deserves investigation, with due cognizance of the second point: The choice is not between no insurance and full insurance. The question is how much, if any, excess moral hazard is induced by the tax law. My suspicion is, a lot, but I think the exploration of this question requires a closer modeling of the specifics of the insurance plans than has yet been addressed by the authors.

There is a second point at which I think the analysis would benefit from greater attention to the characteristics of equilibrium insurance contracts. It is assumed in the paper (and by most analysis of this issue) that one could administer a cap on employer-provided health insurance, with insurance in excess of the cap to be taxed as wages to the employee. Insurance market considerations suggest two possibilities that may merit closer analysis: First, to the extent that the employer-provided insurance is similar in anti-moral-hazard characteristics to what would be predicted in the absence of tax subsidy, we would expect the *insurance* equilibrium to be, not the locus of adjustment to the change in the rules, but rather the amount of health care services purchased. Second, since the implicit insurance policy provided to an employee depends greatly on the characteristics of the employee (coverage of an older person, e.g., is predictably much more costly than coverage of a younger person), and on the pool of employees in a plan (those in a firm with many old people, or many sick people, will face a higher average premium), it seems to me seriously doubtful that simple caps, taking no account of such differences in characteristics (let alone preferences), would constitute viable tools in practice.