Volume Title: Tax Policy and the Economy, Volume 16

Volume Author/Editor: James M. Poterba, editor

Volume Publisher: MIT Press

Volume ISBN: 0-262-16210-5

Volume URL: http://www.nber.org/books/pote02-1

Conference Date: October 30, 2001

Publication Date: January 2002

Title: Taxes and Health Insurance

Author: Jonathan Gruber

URL: http://www.nber.org/chapters/c10862

TAXES AND HEALTH INSURANCE

Jonathan Gruber
MIT and NBER

EXECUTIVE SUMMARY

A common prescription for reducing the number of uninsured is to increase the tax subsidization of health insurance in the U.S. Yet, we already provide over $100 billion per year in tax subsidies to health insurance. This paper provides an assessment of the past and potential impacts of taxation on health insurance coverage and costs. I begin by reviewing the central facts on health insurance and taxation. I then provide a framework for assessing the impacts of tax policies on health insurance coverage and costs, and I review the existing empirical evidence on the key behavioral parameters required to model these impacts. I conclude with the policy implications of these findings for tax policies to expand insurance coverage.

1. INTRODUCTION

Uninsurance is one of the worst social problems in the U.S., and it has continued to worsen as the economy has improved throughout the 1980s and the 1990s. In 1987, 14.8 percent of non-elderly Americans were without health insurance. Over the next decade, the non-elderly population without insurance coverage grew by nearly 25 percent, to 18 percent, before falling for the first time in two decades last year. Still,

despite the recent good news, over 42 million Americans lack health insurance.

The problem of the uninsured has been a major focus of policy debate throughout the 1990s. The most prominent example was the proposed national health insurance plan of the Clinton Administration, which was resoundingly defeated in 1994 (Cutler and Gruber, 2001). Tax incentives to expand health insurance coverage have also been considered continually through the last decade, and received particular attention during the presidency of George H. Bush and again in the presidential election campaign of 2000. In recent years, proposals have been made to offer tax credits to individuals to buy insurance in the non-group market (by President Bush); to expand those credits to cover the cost to individuals of their group insurance policies (by the Progressive Policy Institute); and to provide credits to firms to induce them to offer insurance (by a number of members of Congress).

But the existence of these proposals should not be taken to imply that the U.S. doesn't already dramatically subsidize the provision of insurance in the workplace. In fact, the exclusion of employer (and some employee) health insurance expenditures from the income tax base costs the government over $100 billion in lost revenues annually (Shiels and Hogan, 1999). Indeed, some tax-based proposals would end or limit this exclusion of health insurance from the income tax base, and use the resulting funds to offer tax subsidies to individuals for insurance purchase.

Disentangling the costs and benefits of these alternative approaches to tax subsidization of health insurance is difficult, and revolves centrally around a series of behavioral parameters that determine how individuals and firms will respond to changes in the tax treatment of health insurance. An incomplete list of such parameters includes: the price sensitivity of the decision of firms to offer health insurance; the price sensitivity of the take-up decision, conditional on offering, of employees; the price sensitivity of insurance demand among those not offered insurance; the influence of subsidies on the structure of employer-provided insurance plans, such as employee premium sharing. Despite the importance of these issues, however, we have remarkably little evidence on the key behavioral elasticities, and the evidence that does exist is often contradictory.

The purpose of this paper is to lay out a framework for researchers and policymakers to think about how tax policies might affect the level and distribution of health insurance coverage in the U.S. I begin by reviewing the key facts that are relevant to thinking about health-insurance policy. I then turn to a discussion of the central parameters that we need to know to fully model both the effect of the existing tax subsidy and the effects of tax-based approaches to increasing insurance coverage in the

U.S. I then discuss what we know about these parameters. Finally, I discuss the implications of the facts and our existing knowledge for the design of tax policy toward health insurance in the U.S.

2. THE FACTS

There are a number of key facts that must be considered when modeling the effect of taxes on health insurance in the U.S.

2.1 90% of Insurance Coverage Is Employer-Based

The predominant source of insurance coverage in the U.S. is employer-based insurance. This is shown in Table 1, from EBRI (2000), which shows the sources of insurance coverage in the U.S. over time. A fairly constant feature of insurance coverage has been that more than 9 in 10 of those who are privately insured derive their insurance from an employer, generally their own, their spouse's, or their parents'.

Why is the employment setting the predominant source of insurance coverage? There are three potential reasons. First, there may be substantial economies of scale in administering insurance which increase the value of pooling mechanisms. Second, the major problem facing providers of insurance is adverse selection, so that insurers are constantly searching for means of pooling large groups of individuals along dimensions exogenous to health in order to ensure a predictable distribution of medical costs. Workplaces provide just such a pooling mechanism.

Finally, the U.S. tax code subsidizes health insurance purchase through the firm relative to the non-group market by excluding the value of that insurance from an individual's income, for both income and payroll tax purposes. This leads to a very large effective subsidy to the cost of health insurance for workers. The result of this subsidy is that there is a lower "tax price" of insurance:

$$TP = \frac{1 - \tau_f - \tau_s - \tau_{ss} - \tau_{mc}}{1 + \tau_{ss} + \tau_{mc}},$$

where τ_f is the federal income tax marginal rate, τ_s is the state income tax marginal rate, τ_{ss} is the marginal payroll tax rate for the OASDI program (the 6.2-percent tax rate that is levied equally on employees and employers); and τ_{mc} is the marginal payroll tax rate for the Medicare HI program.[1] I differentiate the last two programs because, beginning in the

---

[1] The reason that the payroll tax rate is additive in the denominator is that the employer is indifferent between purchasing one dollar of benefits and paying wages of $1/(1 + \tau_{ss} + \tau_{mc})$, since each dollar of wages requires a payroll tax payment as well.

TABLE 1

Insurance Coverage of Non-elderly Population over Time

| | Proportion of population (%) | | | | | | | | |
| | Employment-based | | | | Public | | | | |
Year	Total	Own name	Dependent	Indiv. purchased	Total	Medicare	Medicaid	Tricare/ CHAMPVA	None
1987	69.2	33.8	35.4	6.7	13.3	1.4	8.6	4.0	14.8
1988	69.0	33.9	35.0	6.3	13.3	1.5	8.7	3.8	15.5
1989	68.6	33.9	34.7	6.6	13.2	1.5	8.8	3.6	15.7
1990	67.0	33.1	33.8	6.5	14.5	1.6	10.2	3.6	16.1
1991	66.3	32.8	33.5	6.1	15.5	1.6	11.1	3.5	16.3
1992	64.7	31.8	32.9	6.5	16.0	1.7	11.8	3.3	17.0
1993	63.5	32.9	30.7	7.3	16.7	1.6	12.7	3.3	17.3
1994	63.6	32.7	30.9	7.1	16.9	1.6	12.5	3.8	17.1
1995	63.8	32.7	31.1	6.9	16.6	1.8	12.5	3.2	17.4
1996	64.0	32.9	31.2	6.8	16.0	2.0	12.1	2.9	17.7
1997	64.2	32.8	31.5	6.7	14.8	2.0	11.0	2.8	18.3
1998	64.9	33.1	31.7	6.5	14.3	2.0	10.4	2.9	18.4
1999	65.8	33.4	32.4	6.6	14.2	2.0	10.4	2.7	17.5

early 1990s, the taxable maximum for the HI program was increased above that for the OASDI program (and was eventually removed altogether); the marginal rate is zero above the taxable maximum for payroll taxation. For a typical worker in the 15-percent tax bracket, facing a 5-percent state tax rate and a 15.3-percent combined payroll tax rate, this tax price is roughly 0.65; a dollar of health insurance costs 35 cents less than a dollar of other goods purchased with after-tax wages.

In addition, there are also tax subsidies available to employees for their spending on employer-provided health insurance, under section 125 of the Internal Revenue Code. Section 125 generally provides that an employee in a cafeteria plan will not have an amount included in gross income solely because the employee may choose among two or more benefits consisting of cash and qualified benefits. A qualified benefit generally is any benefit that is excludable from gross income under the Code, including health insurance, group-term life insurance, 401(k) plans, child care, and adoption assistance.

While employee contributions can therefore be excluded from taxation with a section 125 plan, protection of employee contributions is far from complete. The data on the prevalence of such arrangements are sketchy. The most recent available data, from Kaiser Family Foundation (2000), suggest that half of all workers are in firms that offered such flexible benefit plans. The reason for less than full coverage of this generous tax benefit is unclear, but some of it may have to do with extensive IRS regulation of these arrangements to ensure that they are not abused. For example, the regulations state that no more than 25% of the benefits of a plan can be attributed to any "highly compensated" employee, essentially ruling out the availability of section 125 plans for very small firms. Moreover, there are strict and complicated rules that limit the flexibility of employees to switch sources of insurance coverage during the year if they are paying their health insurance contributions on a pre-tax basis.

Perhaps as a result of these inherent advantages of the group insurance market, the non-group insurance market has not provided a very hospitable environment for insurance purchase. Load factors in this market are high, and the generosity of the typical policy is much lower than in the group market (Gruber and Madrian, 1996). A recent study by the Kaiser Family Foundation found that, for those individuals in less than perfect health, it was often not possible to get coverage that was fully comprehensive, the particular illness of the individual often being underwritten out of the policy. Prices were also very variable in this market, making it difficult to effectively anticipate the cost of insuring oneself.

There is an existing tax subsidy to the non-group market itself, for a particular group: the self-employed. Beginning in 1986, the self-employed

were allowed to deduct 25% of their insurance premiums from their taxable income. This share has grown over time and is slated to reach 100% by 2003.

2.2 Employer-Based Coverage Has Been Eroding, with Public Insurance Picking Up Some of the Reduction

The notable time trend in Table 1 has been the steady erosion of employer-provided insurance coverage in the post-1987 period. The share of the population with private coverage fell from 69.2 to 63.5 percent in 1993, before rising again slowly to 65.8 percent. This trend was partially offset in 1987–1993 by a sharp rise in the share of the population with Medicaid coverage, due to extensive expansions in that program, particularly for children. But the slow rise in employer coverage after 1993 has also been offset by a sharp decline in public coverage, particularly from 1996 to 1998, which may attribute to an unintended consequence of welfare reform (see Gruber, 2001, for a review of Medicaid program issues).

A striking feature of the erosion in employer-provided insurance coverage is that it was *not* occurring through a decline in employer offering of health insurance, but rather through a decline in employee take-up of that insurance, conditional on it being offered (Cooper and Schone, 1997; Farber and Levy, 1999). A central, and unresolved, mystery is why we saw this decline in take-up. This period was marked by a rise in the share of health insurance costs borne by employees (Gruber and McKnight, 2001). But all available evidence, as I will review below, suggests that take-up of insurance by employees is not very sensitive to its price. Indeed, takeup of employer-based insurance in general remains quite high, as I point out below.

2.3 Most of Those Offered Insurance Take It

An key fact for designing tax policy towards the uninsured is that most of those who are offered insurance by their employers take it up. This is illustrated in Table 2, which represents tabulations of insurance take-up rates from a 1997 Robert Wood Johnson survey of employers. This table cross-tabulates firm size against average earnings in the firm, and in each cell lists the take-up rate of insurance, which is computed as the number of covered employees divided by the number of employees eligible for coverage.

What is striking about this table (in particular in reference to the next table we will see) is the high and uniform rate of take-up across cells.

TABLE 2
Takeup Rates by Firm Size and Average Earnings Categories

Average earnings ($)	Takeup rate[a]				
	1–9 employees	10–24 employees	26–49 employees	50–99 employees	>100 employees
<15,000	0.81	0.70	0.70	0.71	0.76
15,000–30,000	0.84	0.83	0.79	0.78	0.82
30,000–50,000	0.88	0.84	0.83	0.90	0.88
>50,000	0.89	0.84	0.89	0.75	0.88

[a] From author's tabulations of 1997 Robert Wood Johnson Foundation survey of employers.

While take-up does rise somewhat with firm size and with earnings, it is quite high in every cell in this table. If employees are offered insurance, they appear to take it up uniformly at quite high rates.

2.4 Insurance Offering Is Highly Correlated with Firm Size and Average Wage

The offering of insurance, on the other hand, does very quite significantly across firms, along both of these dimensions. Table 3 cross-tabulates (using somewhat different categories for emphasis) firm size and average earnings again, this time summarizing in each cell the average rate of offering health insurance. Offering of insurance is in fact quite low for the smallest firms, even at high wage levels, and for larger firms at the very lowest wage levels.

The correlation with firm size corresponds to the non-tax arguments made above for why firms would offer insurance: economies of scale and predictability of insurance expenditures both rise with firm size. Indeed, the Congressional Research Service (1988) reports that the loading factors on insurance are roughly 35% higher for the smallest than for the largest firms. The correlation with average wages may reflect preferences across firms for insurance offering; but it is notable that when insurance is offered in these low-age firms, employees then take it up at a high rate. This suggest some disconnect between the offering and take-up decisions, which I will return to in the evidence section below.

Given these last two pieces of evidence, it should not be surprising the the primary correlate of being uninsured is not being offered insurance. Data from the Current Population Survey data show that over three-quarters of the uninsured are not offered health insurance on their jobs.

TABLE 3
Insurance Offer Rates by Firm Size and Average Earnings Categories

Average earnings ($)	Offer Rate[a]				
	1–9 employees	10–24 employees	25–49 employees	50–99 employees	>100 employees
<10,000	0.24	0.45	0.63	0.81	0.95
10,000–15,000	0.32	0.55	0.76	0.88	0.93
15,000–20,000	0.43	0.70	0.83	0.93	0.98
20,000–25,000	0.50	0.77	0.86	0.95	0.97
25,000–30,000	0.55	0.83	0.92	0.91	0.98
>30,000+	0.61	0.88	0.94	0.95	0.98

[a] From author's tabulations of 1997 Robert Wood Johnson Foundation survey of employers.

2.5 The Uninsured Are Quite Mixed with the Insured

Finally, any solution to address the problem of the uninsured, tax-based or not, must recognize a fundamental conundrum: the uninsured are not an isolated and easily identified sub-population. This is illustrated in Table 4, also from EBRI (2000), which shows the income distribution of the uninsured. The second column shows the number of uninsured in each income category listed in the first column; the third column shows the percentage of the uninsured in each income category; and the final column shows the percentage of that category that is uninsured. The last column is the least surprising: the share of any income group that are uninsured declines with increasing income. What is more striking is the second column: there are many uninsured who are not poor or even near poor. Indeed, almost a quarter of the uninsured live in families with incomes over $50,000 per year.

This table highlights the difficulty of targeting programs to cover the majority of the uninsured. To cover the majority, you must go fairly high up in the income distribution. But, as you do so, you enter ranges where only a small fraction of the group is uninsured.

3. HOW DOES TAX POLICY AFFECT INSURANCE COVERAGE?

To fully understand how alternative tax policies might affect insurance coverage requires knowledge of a wide variety of key behavioral parameters. In this section, I review the parameters that must be measured to fully assess the range of tax policy effects.

TABLE 4
### *Income Distribution of the Uninsured*[a]

Income category ($)	Number of uninsured (in millions)	Percentage of uninsured in income category	Percentage of income category that is uninsured
<5000	4.8	11.4	44.3
5000–9999	3.3	7.8	31.0
10,000–14,999	4.5	10.6	34.6
15,000–19,999	4.5	10.7	32.0
20,000–29,999	6.8	16.1	24.6
30,000–39,999	5.3	12.5	19.4
40,000–49,999	3.3	7.9	13.9
>50,000	9.7	22.9	8.5
Total	42.1	100	17.5

[a] From EBRI (2000).

Figure 1 summarizes the channels through which tax policies might affect coverage, and the resulting behavioral responses that are key to assessing tax-policy effects. The leftmost part of this chart shows the three possible types of tax subsidies to insurance: subsidies to employers to offer coverage (either reforming the existing subsidy, or offering new subsidies); subsidies to employees to take up coverage (once again, either reforming the existing section 125 option, or offering new subsidies); and subsidies to individuals for non-group insurance coverage. Each of these types of tax subsidies has effects on both firms and on individuals, as shown in the middle of the figure. Moreover, the decisions of both firms and individuals feed back to each other. Finally, for sizable tax interventions, there may be effects on the insurance market itself which will affect the decisions made by firms and workers.

In terms of firm decisionmaking, the key element of response is the decision to offer insurance in response to an increased subsidy to employer coverage, or to drop insurance in response to a reduction in the net subsidy to employer coverage. For increases or reductions in the existing tax subsidy to employer-provided insurance, these responses are likely to be symmetric. But a key unresolved question is the extent to which changes in other subsidies would impact employer decisions. For example, would a 10% subsidy to employees for their expenditures on employer-provided insurance have the same impact on insurance offering as an additional 10% subsidy to employers for their insurance spending? In terms of simple economic theory, the answer is clearly yes. But there may be differences in practice that make these responses asymmetric.

FIGURE 1. *Channels for Effects of Tax Policies on Coverage*
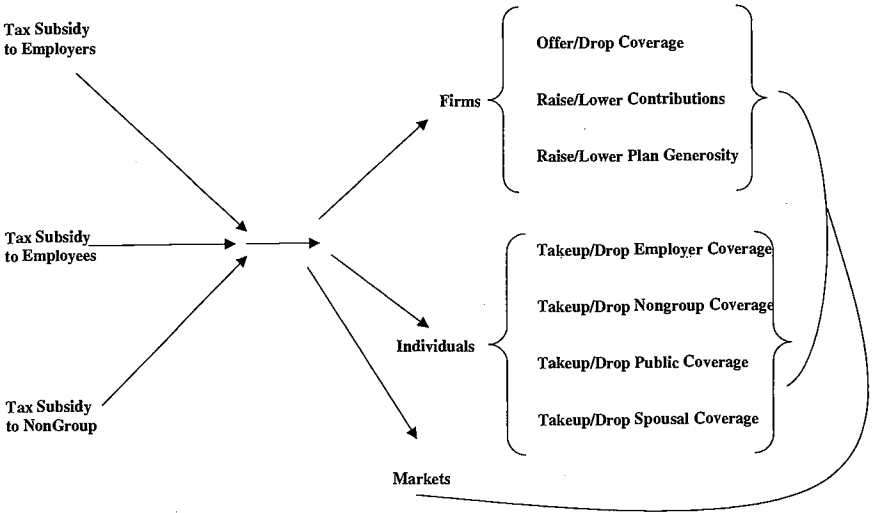
For example, subsidies to employers may be targeted by firm size or some other firm characteristic, but are available to the firm as a whole. But subsidies to employees that are targeted will likely leave some employees in a firm eligible and others not. While the subsidy amounts may be the same on average (e.g., a 50-percent subsidy to the firm or a 100-percent subsidy to 50 percent of the workers), in practice the impacts may be quite different, depending on how employee preferences are aggregated in the determination of benefits.

Similarly, in terms of economic theory, subsidies to non-group and to group insurance should have similar and opposite effects. But, once again, in practice the reactions of firms may not be symmetric, since workers may not view the (currently) inhospitable non-group market as an effective substitute for the group market.

Firms also have other margins of response besides the decision to offer (or not to offer) insurance. One important margin is the decision on how much to contribute to insurance costs. Currently, firms pay about 75 percent of the costs of insurance, but this has fallen dramatically over the past 15 years (Gruber and McKnight, 2001). If higher employee premiums lead employees not to take up their employer coverage, then this shift in costs from employee to employer could be significant. Similarly, employers may react to tax subsidies by reducing the generosity of their insurance plans along a variety of dimensions, such as shifting from fee-

for-service to managed care plans, raising patient copayments, or restricting benefit coverage. Indeed, it is this subsidy on the margin to insurance costs that led Martin Feldstein to criticize the tax subsidy to employer-provided insurance in the early 1970s. He claimed that this subsidy would lead to overly generous insurance coverage, which through moral hazard would then raise spending on medical care, increasing further the demand for insurance coverage, and thereby leading to a spiral of rising medical costs.

Individuals can also respond directly to tax subsidies either to employee purchases of insurance or to non-group subsidies. There are four dimensions along which individuals can respond. The first is to change their take-up of employer-provided insurance, conditional on its being offered. The second is to move into or out of non-group insurance coverage. The third is to move into or out of public insurance coverage. This channel may seem more controversial, but the majority of those made eligible for public insurance over the past 15 years have been eligible as well for employer-provided insurance, so this is a margin of potentially active substitution. Finally, married couples can shift insurance coverage from one spouse's job to the other's as the relative subsidy to one spouse or the other changes.

In addition, the decisions of employers and workers can have feedback effects on each other. As employers change their offering of insurance, this will affect employee take-up. Likewise, changes in employer-provided insurance generosity can affect take-up and decisions to move across spouses or into public or non-group coverage. Moving the other direction, when economists model firms' benefit decisions, they do not think of a firm as a distinct entity, but rather as an aggregation of its workers. So individuals' responses to tax interventions can also feed back to firms' decisionmaking. An important research question, alluded to above, is how worker preferences are aggregated; once again, this aggregation may differ across types of tax subsidies and along the margin of employer response (e.g. offering vs. employee contributions). I discuss the scant evidence on this question below.

Finally, tax policy can affect the insurance market directly. So long as the market supply for insurance is upward sloping, any major intervention that increases demand will lead to a partially offsetting pre-tax rise in insurance prices. On the other hand, many have argued that the high and unstable prices in the non-group market reflect the "thinness" of this market, and that a major subsidy to non-group policies which led to more purchase could reduce inefficiency and lower prices in that market. These changes in pretax prices will obviously mitigate or exacerbate any direct response by firms and individuals to tax incentives.

Thus, there are an enormous number of margins that must potentially be modeled to assess the effect of a full range of tax subsidy options on insurance coverage. In principle, there are many direct channels from tax policy to individual and firm behavior, as well as a large number of potential feedback effects from firms to individuals, from individuals to firms, and from the market to both. Given the impossibility of estimating all of these responses, a key question is the extent to which response symmetry can be called upon to apply behavior from one type of subsidy to another.

4. WHAT DO WE KNOW ABOUT KEY BEHAVIORAL PARAMETERS?

There is a large literature devoted to estimating some of the behavioral parameters that were discussed in section 3. In this section, we review that literature. We begin by reviewing what is known about firm behavior, then turn to individual behavior.

4.1 *Firm Offering Decisions*
The behavioral response on which there has been the most work is the elasticity of insurance offering by firms. There have been several approaches to estimating this elasticity. The first approach discussed here is to use variation in the premiums faced by firms to identify the price sensitivity of their offering decision. Two examples of this work are Feldman, Dowd, Leitz, and Blewett (1997), who use information from 1993 for a sample of small firms in Minnesota to estimate price elasticities of −3.9 (single coverage) to −5.8 (family coverage); and Marquis and Long (1999), who use data from 1993 for 10 states to estimate a much smaller price elasticity of only −0.14. A key problem with this approach, however, is that one only observes premiums for the firms that do offer insurance, and they must be imputed to firms that do not. Thus, instruments must be found that are correlated with the price of insurance but not firm demand, and previous articles have not used firm characteristics that are likely to meet this criterion (e.g. whether the firm is unionized).

The third approach is to use variation in taxation to identify the price elasticity of offering, in essence asking whether those firms with higher tax-related subsidies to insurance purchase are more likely to offer insurance. Leibowitz and Chernew (1992) use variation in tax rates across states to examine the effect of after-tax prices on insurance offering by small firms, as well as using variation in premium quotes across locations obtained from small-group insurers. They separately estimate the

response to premiums and subsidies, and obtain an elasticity between −0.8 (premiums) and −2.9 (subsidies). Royalty (1999) also uses cross-state variation in marginal tax rates to estimate an elasticity of firms' insurance offering across all employers at −0.63. Gentry and Peress (1994) study cross-city differences in the average share of workers offered health insurance benefits, as a function of cross-state differences in after-tax prices of insurance. They find that for each percentage-point increase in the price of health insurance, the fraction of blue-collar workers covered by employer-provided insurance declines by 1.8 percentage point, implying an elasticity of −1.4; however, there is no statistically significant effect for white collar workers.

These types of studies have the advantage that differences across cities and states in tax rates should be independent of insurance-offering decisions. But they may not be entirely independent: cities and states with substantial taste for insurance may be the ones that offer the largest tax breaks, which would lead to a strong relationship between price and offering. This criticism is addressed in recent work by Finkelstein (1999), who studies the removal of the large (25%) tax subsidy to supplemental private health insurance in Quebec in 1993, and finds an elasticity of −0.42 to −0.54 for employer offering. But it is somewhat unclear how to apply the elasticity of offering of supplemental insurance for a national health insurance scheme to the decision of U.S. firms to offer full private health insurance plans.

A third approach comes from running small-scale subsidy pilot programs for small businesses and evaluating the response of firms to subsidized prices. These pilot programs have the advantage of essentially providing a randomized intervention. Two such pilot programs are evaluated in Helms, Gauthier, and Campion (1992) and in Thorpe et al. (1992). The former study finds a wide variety of price responsiveness across sites, with sites such as Utah offering 40% discounts and seeing only 4% enrollment among uninsured firms (an elasticity of only −0.1) and other sites such as Arizona offering 10% discounts and seeing 4–11% enrollment (an elasticity of −0.4 to −1.1). The latter finds very weak response to a program that provided a 50% subsidy to the price of insurance for small firms in New York, with an elasticity of only −0.07 to −0.33. But it is unclear whether the small elasticities estimated here are because of the temporary experimental nature of these subsidies; firms may be reluctant to set up insurance plans based on subsidies that will only last for a short time. There could be much larger responses to more permanent changes in the after-tax price of insurance.

A final approach is to use responses of firms to hypothetical questions about changes in the price of insurance. Morrisey, Jensen, and Merlock

(1994) use the response to such hypothetical questions to estimate a price elasticity of insurance offering among small firms of −0.92. But it is unclear whether firms respond in the same way to hypothetical questions as they do when faced with an actual insurance purchase decisions.

A final issue with this literature is that, with firm-level data, one does not observe the characteristics of the employees to which the firm is responding in making its benefit decisions: who is the *marginal worker* whose preferences determine the firm's decisions? One article which attempts to address both the identification and marginal-worker issues is Gruber and Poterba (1994). They study how the self-employed responded to the Tax Reform Act of 1986, which introduced a subsidy to the insurance purchases of the self-employed. This natural experiment provides exogenous variation in the after-tax price of insurance. Moreover, for the self-employed, there is no issue of deciding who is the marginal worker. They find significant increases in the insurance coverage of the self-employed relative to the employed over this period, with an implied price elasticity as large as −1.8. Unfortunately, however, it is unclear how generalizable these results are to firms, which must aggregate the preferences of all their workers in making benefit decisions.

A more recent approach to surmounting these problems is Gruber and Lettau (2000). In that paper, the authors use data from the Employee Compensation Index (ECI) dataset, which collects information both on firm insurance provisions and on the characteristics of a sample of workers in the firm. The latter feature allows them to directly measure the distribution of tax rates within the firm, and to assess which tax rate seems to be most central in driving benefit provision decisions. They also introduce a new identification strategy, extending the previous tax-based work to rely on state tax progressivity and changes in state taxes over time, while controlling for mean differences across states that are likely to be correlated with tastes for insurance. They estimate an elasticity of insurance offering of −0.3 to −0.4, towards the lower end of the previous literature. Gruber (2001) recently applied the identification strategy of Gruber and Lettau's paper to data from the Current Population Survey, which gathers data on a random sample of workers but not on the distribution of workers in a firm. He finds a higher elasticity of offering of −0.7 in these data. But the ECI estimates seem more reliable in view of the higher quality of the data.

4.2. Employer Insurance Spending

There is also a sizable literature on the effect of after-tax prices on employer insurance spending. Estimates of this elasticity come from three types of studies. The first is time-series evidence on how total spending

on employer-provided health insurance responds to changes in federal tax rates, presented in Long and Scott (1982), Vroman and Anderson (1984), and Turner (1987). These studies typically yield estimates of the price elasticity of demand between 0 and −0.5. But the results are hard to interpret, as there are many things changing in the time-series data; for example, the fact that health insurance coverage fell in the 1980s may be the result of declining marginal tax rates, but it may also be the result of a shift in the job base towards service jobs that are less likely to provide insurance.

A second set of studies, including Taylor and Wilensky (1983), Holmer (1984), and Sloan and Adamache (1986), analyze cross-sectional data on individuals or firms and ask whether those with higher tax-related subsidies to insurance purchase spend more on insurance coverage. But a potential problem with these studies is that differences across individuals in their tax rates arise in part from differences in the underlying behavior of individuals or firms, such as differences in labor supply, family structure, or the nature of the work force. It is impossible to tell whether differences in observed insurance coverage are due to taxes or these behavioral differences. A wide range of estimates emerge from these studies; Pauly (1986) summarizes the consensus range as −0.2 to more than −1.0.

The final approach that attempts to overcome the problems inherent in the previous cross-sectional literature examines how demand for insurance responds to plausibly independent legislated tax differences. Woodbury and Hammermesh (1992) analyze all fringe-benefit expenditures around the Tax Reform Act of 1986 in a panel data set of colleges and universities. They conclude that tax reform substantially reduced the demand for fringes, with an estimated elasticity in excess of −2. But this is not focused on health insurance spending per se, so it is difficult to disentangle the impact on health insurance.

Gruber and Lettau (2000), in the study described above, also examine the impact of tax subsidies on employer-provided insurance spending, using similar variation across states in their tax systems to Woodbury and Hamermesh. They also find a quite large elasticity of insurance spending: −0.94.

4.3 *Employer Contributions to Health Insurance*

A particularly important margin of response to tax subsidies is how employers change their contributions to health insurance for employees. This response would be subsumed in the spending elasticities cited above, but it is important to break it out distinctly due to the potential impact of changing contributions on employee take-up.

There is only one study of which I am aware on this response. Gruber and McKnight (2001) use the Current Population Survey to study the impact of tax changes on the decisions of employers to pay all of the cost of their health insurance plans. The share of employers paying all those costs fell from 44 percent in 1982 to 28 percent by 1998. They model this variable as a function of a variety of actors, including the tax subsidy to employer-sponsored insurance. They find that the tax subsidy did have a very important impact on employer contribution policy. Their central findings suggest that for each 10-percentage-point reduction in the tax subsidy, the share of employers paying all of the costs of health insurance falls by 1.7 to 3.8 percent.

4.4 Employee Take-up of Employer-Provided Insurance

A central parameter for evaluating the effects of tax subsidy policy is the elasticity of employee take-up of employer-provided insurance if it is offered. As highlighted earlier, take-up among those offered is fairly high for all firm sizes and average firm wage levels. But does take-up respond to the prices charged to employees for insurance?

The answer, to date, appears to be no. Two studies have examined how employee take-up responds to the price charged for insurance (Chernew, Frick, and McLaughlin, 1997; Blumberg, Nichols and Banthin 2001). Both papers find that firms that charge more for insurance have no lower take-up of their insurance policies. These papers do run into the possible problem that firm premium decisions are endogenous to employee tastes. The direction of bias here is unclear, and depends on whether firms set low employee premiums when there are tastes for insurance, or when there are not (because of paternalism or to satisfy insurance-company conditions for high employee take-up). Gruber (2001) also investigates this question by assessing whether the existing employer tax subsidy affects coverage, conditional on offering, as it should for those with a section 125 plan (roughly half of employees by the most recent estimates). But I find no such effect.

The lack of elasticity of employee take-up is very striking in view of the time-series trends discussed earlier. Over the mid-1980s through the late 1990s, the trend was towards rising employer contributions and falling employee take-up. This is shown in Figure 2, from Gruber and McKnight (2001), which shows the share of workers who have employer-sponsored insurance over time, and the share of workers whose employers pay all of the costs of insurance over time. As those authors note, there is a striking correspondence between these series, and a price elasticity of employee takeup of 0.4 would explain the time-series trend in take-up. But this appears to be well above the best estimates to date of the take-up elastic-

FIGURE 2. *Group Health Insurance vs. Employer Pays All*
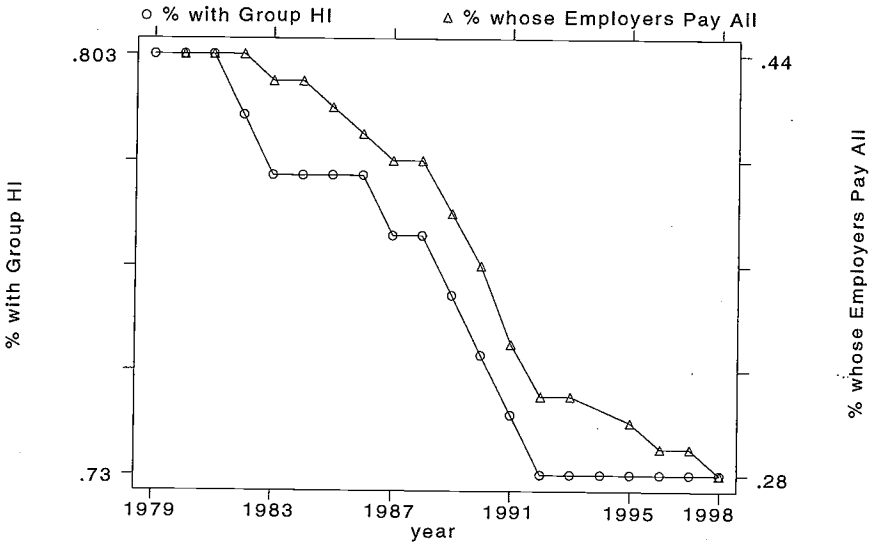
ity. Thus, the cause of this trend towards declining take-up remains a mystery.

On the other hand, these findings are consistent with a growing body of evidence which suggests that it is the decisions that employers make for their workers that are most important in determining worker benefit provision, not active decisions taken by those employees. The most striking example is Madrian and Shea (2001), who find that when a firm moved the default investment option and contribution level for its 401(k) plan, the vast majority of workers did not move from that default, despite it being a worse choice for many of them.

4.5 Substitution between Forms of Insurance Coverage

As the lower part of Figure 1 illustrates, a key issue for modeling the effect of tax policy on individual insurance coverage is the substitutability across different forms of insurance coverage. The margin of substitutability for which there is the best evidence is substitutions between private and public coverage. There is a large literature on *crowdout* over the past 5 years which examines the question of whether those made eligible for public insurance will drop their private insurance to take it up. The earliest estimates of crowdout suggested it was quite large, with one person losing private insurance for every two gaining public insurance in the late 1980s and early 1990s (Cutler and Gruber, 1996). But

subsequent estimates suggest that the effect may be more modest, on the order of 10–20 percent; see Dubay (1999) for a review.

There is less evidence on the other key margins of substitutability, between own employer coverage, non-group coverage, and spousal coverage. Gruber (2001) did recently extend the analysis of Gruber and Lettau (2000) to consider the impact of changing the employer-provided insurance tax subsidy on all insurance margins, not just on the margin of insurance offering. We find that a larger subsidy to employer-provided insurance causes a reduction in public insurance, consistent with the crowdout literature. But we find little substitutability with non-group coverage.

To date, there is little evidence on the substitutability of own and spousal insurance coverage. On net, I find that the effect on offering is almost directly translated to a bottom-line effect on being uninsured, suggesting that employer offering is the key margin of response to existing subsidies.

4.6 Substitution across Spouses

There is enormous scope for substitution of health insurance coverage across spouses. For example, of the 28 million male married workers aged 21–64 who are insured on the job, 15 million have wives that are offered health insurance. More than 8 million report that their wives actually take up the coverage on their jobs. Of the 16 million married female workers aged 21–64 who are insured on the job, 12 million have husbands who are offered insurance, and almost 9 million of those husbands take up that offer.

With such a large number of spouses who are jointly offered insurance, it would seem that substitution across spouses as policy changes would be a real possibility. There is only one article that investigates this issue (Monheit, Schone, and Tayor, 1999), using 1987 data, and their results suggest that the decision to take double coverage (health insurance coverage through both the husband and wife), conditional on both spouses being offered insurance, is very sensitive to incentives. For example, they find that the odds of having double coverage rise by more than a third if both plans are free to the husband and wife. They also find that double coverage take-up is higher among those families in ill health, and that it is more common when it results in a more comprehensive set of benefit coverage than does take-up by one spouse alone.

This set of findings suggest that substitution across spouses who are offered insurance may be quite fluid as tax policy changes. How much substitution there would be depends on how universal the policy changes are, and how similar the jobs are that wives and husbands hold.

For an income-targeted subsidy, for example, it is possible that a husband who works with many low-income workers will lose his offer of health insurance, while his wife who works with many high-income workers will not.

4.7 The Big Mystery: Take-up among Those Not Offered

At least three-quarters of the uninsured are not offered health insurance. Thus, in focusing on the effects of tax subsidies on non-group insurance, the key question is how price-sensitive this group will be in their take-up decisions. Unfortunately, we have essentially no evidence on this critical question. The one relevant paper is by Marquis and Long (1994), who estimated the demand for non-group insurance coverage among workers not offered employer-sponsored coverage as a function of the area-specific price of non-group coverage. They estimate an elasticity of non-group coverage of -0.3. The problem with this approach is that the price of insurance reflects not only true price differences in insurance (differences in the load factor, or the premium cost relative to expected claim expenses), but also differences in medical costs and differences in underlying tastes for insurance. Both of these latter two factors will bias downward any estimate of the effect of area insurance prices on demand. Thus, we are left with only a lower bound on the elasticity of insurance take-up among those not eligible for employer coverage.

4.8 How Are Employee Preferences Aggregated?

As noted above, appropriate modeling of the implications of subsidies to employer-provided insurance requires an understanding of the mechanism for aggregating employee preferences in the firm's benefits decisions. The best discussion of this issue in the context of benefit provision is in Goldstein and Pauly (1976). They conclude that the equilibrium benefit determination could arise in one of two ways. One is through the collective choice of the existing set of workers, through an insider-outsider or union mechanism. In this case, through standard voting arguments, the benefits chosen will reflect the tastes of the median worker. The second is through the choices of employers, whose goal is to minimize their total labor costs, and will therefore design their benefits packages to reflect the average preferences of their workers.

If there is a perfect Tiebout equilibrium across firms, with workers sorted completely by their tastes for insurance, then the average and median tax prices will be everywhere the same and the distinction between these models will not be important. However, as Gruber and Lettau (2000) discuss, there remains considerable dispersion between these measures within firms; almost 10 percent of firms have a median

and a mean tax price which differ by 5 percentage points or more. If there is imperfect sorting, then these models can have very different implications, depending on the difference between the median and the mean. Further complications arise when considering the fact that both mobility and influence within the firm differ across workers. Firms may not consider the average of all workers' preferences in making benefits, but may weight more highly the preferences of either mobile workers or the "most influential" workers in the workplace.

Assessing the appropriate model of within-workplace benefit determination is not purely an academic concern. Modeling the implications of tax reform may depend critically on appropriately capturing the structure of how tax prices throughout the firm affect benefit decisions. If, for example, the median tax price is the only one that matters, then tax reforms such as that in 1993 which raised tax rates only at the top of the distribution will have essentially no impact on insurance decisions. But if other movements of the distribution matter, then such reforms may have larger impacts.

Gruber and Lettau investigate this issue by drawing on the strength of their ECI data to include several moments of the distribution of tax prices in their insurance demand model. They find that the median tax price explains benefit determination significantly better than the average. But they also find that there is an important additional role for the highest-paid worker in the firm (the worker with the lowest tax price). So the appropriate voting model in their context appears to be one with a decisive median voter but some extra influence for the highly paid worker.

Unfortunately, however, it is not clear if these results, estimated in the context of an unlimited tax deduction, extend to other tax-subsidy structures, such as targeted tax credits. Consider a very generous non-group credit that pays the full costs of insurance for 49 percent of a firm's workers, but is zero for those at the median and above. It seems unlikely that the firm would not respond at all to such an outside option. In these types of cases, it may be the mean incentive across all workers which better captures how preferences are aggregated.

4.9 Market Responses

Another mystery area is how markets will respond to tax interventions. As noted above, tax subsidies could lead to rises or falls in group or non-group market insurance prices. A particularly critical question is whether, for non-group subsidies of a given value (such as the Bush plan), non-group plans will emerge that are targeted to that dollar value. Advocates

of non-group subsidies point to the availability of very cheap insurance plans over the Internet.

But, even if such plans are offered, there is the question of whether they will be demanded. After all, even in the face of existing low-cost policies, there remain over 42 million uninsured persons. An outstanding mystery in insurance markets is why there is not more demand for low-cost, catastrophic policies, given that they insure the substantial risks we should be most worried about but can reduce costly moral hazard. But given this lack of demand, even the offering of such low-cost policies tied to subsidy amounts may not significantly increase take-up.[2]

5. POLICY IMPLICATIONS

The discussion of section 3 and the evidence presented in section 4 have several clear implications for tax policy towards insurance. Before presenting these implications, it is important to consider the goals of subsidizing insurance through the tax code. Presumably, the most important goal is to raise insurance coverage in the U.S., and in the most cost-efficient manner. Thus, the discussion below will focus on the *efficiency* of various alternative policies, measured as their cost per person newly insured. But a secondary goal of tax policy may be horizontal equity, or redistribution to those already paying for their insurance without the tax subsidies available to others. This is not a goal that will receive attention in the discussion below, but it is important to recognize that if it is the goal of policy, it may be well served by some of the alternatives I dismiss as "inefficient" below.

5.1 Reforming the Existing Tax Subsidy

I first consider the implications of reform of the existing tax subsidy to employer-provided health insurance. One alternative here would be to end this subsidy, perhaps redistributing the dollars to other forms of insurance subsidization. Gruber and Lettau (2000) perform some simulations using their estimates to assess the implications for employer offering and insurance spending. The results are summarized in Table 5, which shows the impact of a several alternative reforms on the rate of insurance offering, the level of insurance spending conditional on offering, and the overall level of spending. The ranges of estimates reflect the

---

[2] Moreover, the low prices of catastrophic policies in today's market may reflect positive selection: such policies will only be demanded by the healthy, so that prices can stay low. When subsidies are available that make such policies cheaper, then there may be more demand for them by the sick, which would lead to price rises.

TABLE 5
Implications of Reform of Existing Tax Subsidy

Reform	Change in offering[a] (%)	Reduction in employer-insured[b] (million $)	Change in spending among offered[a] (%)	Total change in spending[a]
Remove all tax subsidies	−14.1	22.3	−35.4	−50.3% [−$1231]
Remove income tax subsidy	−8.6	13.6	−23.7	−32.8% [−$815]
10% tax rate cut	−0.9	1.4	−2.8	−3.7% [−$95]

[a] From Gruber and Lettau (2000).

[b] Result of applying reductions in preceding column to national total for employer-insured.

alternative models estimated in this paper; to be conservative, in this discussion I focus on the lower-bound estimates.[3]

This paper finds that there would be very significant implications of reducing the existing tax subsidy. Completely removing the existing subsidy, with respect to all (federal and state) income taxes as well as with respect to the payroll tax, would lower the rate of employer insurance offering by 14%. Assuming (following the results in Gruber, 2001) that all of those dropped by their employer would become uninsured, this implies that 22 million Americans would become uninsured. This would represent a roughly 50% rise in the number of uninsured. Moreover, this paper also finds a very large reduction in spending on those who have insurance, with the net result being a 50% reduction in the dollars spend on employer-provided insurance.[4]

These are enormous impacts, and the potential for 22 million more uninsured Americans should lead anyone to pause before removing the existing employer subsidy. But it is also important to remember that the existing subsidy costs over $100 billion per year in forgone revenue. So this says that we are spending about $5000 per person to insure these 22 million persons, a quite high level. So the question becomes whether

[3] It is of course important to recognize that these projections are only as precise as the underlying estimates; the central estimates on which they are based are a probit coefficient on insurance offering of 5.424 (2.017), and an OLS coefficient on log spending of −1.465 (0.404).

[4] The calculations in the final column are conjectural in that they assume that the firms that stop offering insurance are spending the average amount on insurance before dropping. It seems likely that those firms were spending less than average, so that there is a smaller reduction in overall spending than is implied in Table 5.

other alternatives are available at a more reasonable cost to insure large numbers of Americans.

Table 5 also shows the impacts of less dramatic reforms to our existing system. The first is to remove the existing subsidy in the income tax, but in for the payroll tax. This would still lead to very large reductions in insurance offering and spending, with the potential for 13.6 million more uninsured. The final column illustrates that even very modest reforms to the tax code can have large effects on employer-provided insurance. This column shows the impact of cutting all tax rates by 10%, roughly akin to the original proposal of the Bush Administration. Even this very small change could lead to 1.4 million more uninsured, and a reduction in total spending on employer-provided insurance of almost 4%.

One potential reform which would may less dramatic implications for coverage would be not to reduce the existing employer subsidy but to cap it, for example at the mean or median cost of a group insurance plan. As discussed in more length in Gruber and Poterba (1996), such a reform might temper the inflationary aspects of the tax subsidy highlighted by Feldstein, and significantly raise government revenues, but without causing a major displacement of the employer-insured. In principle, administering such a cap would be straightforward: employers would simply be asked to report, for tax purposes (either income tax alone or payroll tax as well), any spending they make on an employee's behalf for insurance beyond some cap level.

But, in practice, caps run into important administrative and political difficulties. First, there are very large regional disparities in the cost of health insurance, which would ideally be reflected in the cap level. But there has never been a regionally adjusted tax credit, and efforts in other arenas (e.g. poverty measurement) to have regional adjustment have run into daunting political difficulties. Second, a cap would penalize workers for having high-cost co-workers, since it would reflect average and not individual-specific insurance spending. This would lead to general redistributions from older, higher-cost industries to newer, lower-cost ones, raising further daunting difficulties.

5.2 New Subsidies to Employers

A more likely direction for reform is to offer new subsidies to employers. But doing so immediately runs into the type of efficiency considerations highlighted above. The majority of employers already offer health insurance. For this group, new subsidies are just redistribution, with no impact on insurance coverage, except through feedback effects on employees through reduced employee contributions or more generous insurance levels.

Thus, to be cost-effective, new subsidies to employers must be well targeted. Table 2 illustrates how such targeting would be most effective: if subsidies were largest for the smallest and lowest-wage firms. But there is a trade-off with targeting: the more targeted are subsidies, the more distortionary "cliffs" there may be to firms' decisions on pay and hiring. This was in fact an important criticism of the proposed subsidies to small firms embedded in the Clinton HSA plan.

But such "cliffs" need not in fact arise, so long as there are smooth phaseouts with respect to wages and firm size. Consider for example the following subsidy structure:

- There is a maximum subsidy rate of 0.4 for firms of 10 employees or fewer with average annual earnings of $10,000 or less.
- The subsidy is reduced by 0.01 for each extra employee above 10, so that it reaches zero at 50 employees.
- The subsidy is also reduced by the amount that earnings rise above $10,000 per year. This reduction factor is proportional to firm size— that is, the bigger the firm, the faster the subsidy is reduced as wages rise. The formula is:

  - Firms of 10 employees or fewer: subsidy reduced by 0.0222 for each $1000 rise in earnings above $10,000 per year.

  - Firms of more than 10 employees: subsidy reduced by (firm size)/ 450 for each $1000 rise in earnings in hourly wages above $10,000 per year. For example, for a firm of 25, the subsidy is reduced by 0.0555 for each $1000 rise in earnings above $10,000 per year.

I have simulated the effect of such a policy in the RWJF data used to compute Tables 2 and 3, drawing on the evidence presented above. The results suggest that for an annual cost of $3.8 billion, 2.2 million persons could be insured per year, for a cost of $1720 per newly insured. And there are very low implicit tax rates embedded in this gradual phaseout structure. On average, for every additional $100 paid out in wages, firms only lose $1.2 in subsidy; even at the maximum, the subsidy loss is only $15 per $100 paid out in wages, a modest distortion by the standards of our tax system. Similarly, for each worker hired, the subsidy loss is on average only $82; the maximum possible subsidy loss per hire is $720. So distortions need not be enormous to offer very targeted subsidies.

But one interesting feature is that the efficiency of such firm subsidies diminishes as they get larger. For example, an expanded version of the above subsidy plan that delivers a maximum subsidy of 50 percent

would cost 175 percent as much ($6.7 billion/year), but only cover 50 percent more persons (3.3 million), so that the cost per newly covered rises to $2070. This is a common finding in all simulation work on tax subsidies: their efficiency is inversely related to their scope. This is because as the subsidies get larger, they necessarily become less targeted and more attractive to those who are already providing insurance. Nevertheless, the efficiency of these firm subsidies seems quite high relative to the other policies considered below.

5.3 New Subsidies to Employees

The second major alternative discussed earlier was new subsidies to employees to take up employer coverage. The first point to note about such subsidies is that, as a device for targeting take-up per se, they are likely to be very inefficient. This is because fewer than 10% of those offered insurance are actually uninsured. Moreover, the work reviewed above suggests that the take-up decision is not very elastic with respect to price. These two facts are a recipe for an inframarginal subsidy that will serve only as redistribution and not to increase insurance coverage.

But, in a general economic model, such subsidies should also increase employer insurance offering. Indeed, there is no economic rationale for not treating them symmetrically with a subsidy to employers: both are subsidies to offering insurance through the workplace. In practice, however, their effects might differ somewhat. Employer subsidies are targeted to the characteristics of a firm, while employee subsidies are targeted to employees. This has the advantage that it may be possible to do better income targeting with employee subsidies, since even low-wage firms have high-wage workers. But it has the disadvantage that it may be harder to target low-wage firms by simply giving subsidies to low-wage employees. Many low-wage employees work in high-wage firms that already offer insurance. So even a very tightly targeted subsidy to employees is likely to result in little new offering, as most of the dollars flow to those already in firms offering insurance.

5.4 Non-group Subsidies

The final alternative is to subsidize non-group purchase of insurance. The prototypical proposal here, which is quite similar to that proposed by President Bush, would be:

- $1000 credit for individuals; $2000 for families
- Usable for non-group insurance purchase only
- Refundable and advanceable
- Phased out for upper-middle- and upper-income families

There are several difficulties faced by non-group credits. The first is a simple fact: roughly half of the uninsured do not pay taxes (Gruber and Levitt, 2000). This makes the tax code a problematic mechanism for delivering subsidies to the uninsured. In principle, this can be addressed by making credits refundable. In practice, this raises two difficulties. First, refundability is a very difficult political goal in recent years, as many conservatives view refundable credits as akin to cash welfare. There were enormous battles over the refundability of the child credit in 1997 and again in 2001, with the result being only very partial refundability of this credit. Similar battles are likely to occur with these tax credits.

Second, and more fundamentally, the usefulness of refundability is quite limited without effective advanceability. Insurance premiums are due from the beginning of year t, but tax refunds do not arrive until spring of year $t+1$. Thus, the uninsured must have sufficient resources to advance-fund their insurance purchases if credits are to be effective; but most uninsured, indeed, most Americans, do not have sufficient liquid assets to do such advance funding. In principle, this could be surrounded by legislating an advanceable credit. But our one experience with such a feature is not encouraging. Individuals can claim their Earned Income Tax Credit (EITC) throughout the year rather than the next spring, and it would be sensible for most claimants to do so; but fewer than 1% of claimants take advantage of this option (Liebman, 1998). The reasons for this low take-up are unclear, but the main conclusion is that low-income taxpayers appear to be reluctant to take any risk that they will end up facing a tax liability, rather than receiving a refund, on April 15. This will limit the effectiveness of advanceability for a non-group credit as well.

Finally, the major difficulty faced by such a credit is that non-group insurance is very expensive. Today, the typical non-group policy for a family costs $6,000 to $7,000 per year. Thus, even a sizable $2,000 credit leaves the family with $4,000 to $5,000 in costs to pay, which is enormous relative to the incomes of the working poor uninsured. This point ties to the earlier issue about market responses. If the insurance market can respond to the availability of this credit by delivering low-cost insurance products that are demanded by the public, then these affordability barriers may be overcome.

As a result of these limitations, my previous work suggests only modest impacts of nongroup credits on net insurance coverage (Gruber and Levitt, 2000; Gruber, 2000). This work is based on a major micro-simulation model that takes as its base the 1997 Current Population

Survey, and imposes on that base a complicated set of behavioral equations that make assumptions on all of the relevant margins in Figure 1 (as detailed in the appendix to Gruber, 2000). The base case is a $1000 or $2000 refundable credit which is not effectively advanceable, due to the limitations noted above. I assume that this credit is available to singles with incomes up to $75,000 and families with incomes up to $100,000. It is worth noting that subsidies that are more tightly income-targeted are more efficient, as for incomes above the median there are relatively few uninsured. But, once again, there is an important political question about whether it is feasible to have a truly targeted new credit in today's environment; all of the middle-class entitlements of the past 5 years have extended to income ranges similar to those used here (or higher).

I find that for this base-case policy there is a cost of $13.3 billion per year, and net reduction in the uninsured of 4 million persons, for cost of $3300 per newly insured, well above the costs cited above for employer subsidies. The effectiveness of the credit rises significantly if it is advanceable, but the cost per newly insured remains above $2500. Thus, non-group credits do not appear nearly as effective as group credits, due to the problems noted above.

6. CONCLUSIONS

Tax policy towards health insurance is likely to remain a topic of vigorous debate in the years to come. This paper has laid out some key facts, economic evidence, and policy simulations to help guide this debate. The key conclusion that I draw from existing facts and evidence is that policies targeted to firms are more likely to be effective than are policies targeted to individuals, as firms appear quite price-responsive in their insurance-offering decisions, and the actions of firms appear to be directly translated to individual coverage. But the prevalence of offering means that to be cost-effective such subsidies must be tightly targeted to the firms least likely to offer insurance: small and low-wage firms.

The more important conclusion to be drawn from this paper is that we still know remarkably little about a number of key parameters that determine the effectiveness of tax policy towards insurance coverage. Most notable among these is the responsiveness, to new subsidies to buying insurance, of the existing uninsured who are not offered insurance. But there are a variety of other unanswered questions as well that must be addressed by future research if we are to draw fully informed conclusions as to the efficacy of tax policy.

REFERENCES

Blumberg, Linda, Len Nichols, and Jessica Banthin (2001). "Worker Decisions to Purchase Health Insurance," mimeo, Urban Institute.

Chernew, Michael, Kevin Frick, and Catherine G. McLaughlin (1997). "The Demand for Health Insurance Coverage by Low-Income Workers: Can Reduced Premiums Achieve Full Coverage?" *Health Services Research* 32:453–470.

Congressional Research Service. *Costs and Effects of Extending Health Insurance Coverage.* Washington, DC: U.S. Government Printing Office, 1988.

Cooper, Phillip, and Barbara Schone (1997). "More offers, fewer takers for employment-based health insurance: 1987 and 1996," *Health Affairs*, 16, 142–149.

Cutler, David, and Jonathan Gruber (2001). "Health Policy in the Clinton Era: Once Bitten, Twice Shy," forthcoming in Jeffrey Frankel and Peter Drszag, eds., *Economic Policy During the 1990s.*

Cutler, David, and Jonathan Gruber (1996). "Does Public Insurance Crowd Out Private Insurance?," *Quarterly Journal of Economics*, 111(2), May 1996, 391–430.

Dubay, Lisa (1999). "Expanding Public Insurance Coverage and Crowd-Out: A Review of the Evidence," in *Options for Expanding Health Insurance Coverage: What Difference Do Different Approaches Make?*, co-edited by Judith Feder and Sheila Burke (Washington DC, Henry J. Kaiser Family Foundation).

Employee Benefits Research Institute (2000). "Sources of Health Insurance and Characteristics of the Uninsured: Analysis of the March 2000 Current Population Survey." Washington, D.C.: EBRI.

Farber, Henry, and Helen Levy (1999). "Recent Trends in Employer-Sponsored Health Insurance Coverage: Are Bad Jobs Getting Worse?" *Journal of Health Economics.* 19, 93–119.

Feldman, Roger, Bryan Dowd, Scott Leitz, and Lynn A. Blewett (1997). "The Effect of Premiums on the Small Firm's Decision to Offer Health Insurance," *Journal of Human Resources* 32:635–658.

Finkelstein, Amy (1999). "The Effect of Tax Subsidies to Employer-Provided Health Insurance on Workplace Pooling: New Evidence From Canada." Forthcoming in *Journal of Public Economics.*

Gentry, William and Eric Peress (1994). "Taxes and Fringe Benefits Offered by Employers." NBER Working Paper no. 4764.

Goldstein, G. S.; Pauly, M. V. (1976). "Group Health Insurance as a Local Public Good" in *The Role of Health Insurance in the Health Services Sector*, Richard N. Rosett, ed. New York: Neale Watson Academic Publications, p. 73–114.

Gruber, Jonathan (2000). "Medicaid," forthcoming in *Means Tested Transfer Programs in the U.S.*, Robert Moffitt, (ed.) University of Chicago Press, (still forthcoming).

Gruber, Jonathan (2001). "The Impact of the Tax System on Health Insurance Decisions." MIT. Mimeo.

Gruber, Jonathan, and Michael Lettau (2000). "How Elastic Is the Firm's Demand for Health Insurance?" NBER Working Paper no. 8021. November.

Gruber, Jonathan, and Larry Levitt (2000). "Tax Subsidies for Health Insurance: Costs and Benefits." *Health Affairs* 19:72–85.

Gruber, Jonathan, and Brigitte Madrian (1996). "Health Insurance and Early Retirement: Evidence from the Availability of Continuation Coverage," in *Ad-*

vances in the Economics of Aging David Wise, ed. University of Chicago Press, pp. 115–143.

Gruber, Jonathan, and Robin McKnight (2001). "Why Are Employee Health Insurance Contributions Rising?" MIT. Mimeo.

Gruber, Jonathan, and James M. Poterba (1994). "Tax Incentives and the Demand for Health Insurance: Evidence from the Self-employed," *Quarterly Journal of Economics* 109:701–733.

Gruber, Jonathan, and James Poterba (1996). "Tax Subsidies to Employer-Provided Health Insurance." In *Empirical Foundations of Household Taxation*, Martin Feldstein and James Poterba (eds.). Chicago: University of Chicago Press.

Helms, W. David, Anne K. Gauthier, and Daniel M. Campion (1992). "Mending the Flaws in the Small-Group Market." *Health Affairs* 11:8–27.

Holmer, Martin (1984). "Tax Policy and the Demand for Health Insurance." *Journal of Health Economics* 3:203–221.

Kaiser Family Foundation (2000). *Employer Health Benefits Survey.* Menlo Park, CA: Kaiser Family Foundation.

Liebman, Jeffrey B. (1998). "The Impact of the Earned Income Tax Credit on Incentives and Income Distribution," in James Poterba, ed. Tax Policy and The Economy 12, 83–120.

Liebowitz, Arleen, and Michael Chernew (1992). "The Firm's Demand for Health Insurance." In *Health Benefits and the Workforce*, U.S. Department of Labor. Washington: U.S. Government Printing Office.

Long, James E., and Frank A. Scott (1982). "The Income Tax and Nonwage Compensation." *Review of Economics and Statistics* 64:211–219.

Madrian, Brigitte C., and Dennis F. Shea (2001). "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics.* Vol. 116, No. 4, pp. 1149–1187.

Marquis, M. Susan, and Stephen H. Long (1994). "Worker Demand for Health Insurance in the Non-group Market." RAND Corp. Mimeo.

Marquis, M. Susan, and Stephen H. Long (1999). "To Offer or Not to Offer: The Role of Price in Employer Demand for Insurance." RAND Corp. Mimeo.

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.

Monheit, Alan, Barbara Steinberg Schone, and Amy Taylor (1999). "Health Insurance Choices in Two-Worker Households: Determinants of Double Coverage." *Inquiry* 26:12–29.

Morissey, Michael A., Gail A. Jensen, and Robert J. Morlock (1994). "Small Employers and the Health Insurance Market." *Health Affairs*, 13, 149–161.

Pauly, Mark (1986). "Taxation, Health Insurance and Market Failure in the Medical Economy," *Journal of Economic Literature*, 24, 629–675.

Royalty, Anne Beason (1999). "Tax Preferences for Fringe Benefits and Workers' Eligibility for Employer Health Insurance." *Journal of Public Economics.* Vol. 75, issue 2, February 2000, pp. 209–227.

Sheils, John, and Paul Hogan (1999). "Cost of Tax-Exempt Health Benefits in 1998." *Health Affairs* 18:176–181.

Sloan, Frank, and Killard Adamache (1986). "Taxation and the Growth of Nonwage Benefits." *Public Finance Quarterly* 14:115–139.

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, Jack Meyer (ed.). Washington: 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* 267:945–948.

Turner, Robert (1987). "Are Taxes Responsible for the Growth of Fringe Benefits?" *National Tax Journal* 40:205–220.

Vroman, Susan, and G. Anderson (1984). "The Effect of Income Taxation on the Demand for Employer-Provided Health Insurance." *Applied Economics* 16:33–43.

Woodbury, Stephen A., and Daniel S. Hammermesh (1992). "Taxes, Fringe Benefits, and Faculty." *Review of Economics and Statistics* 73:287–296.