Volume Title: Tax Policy and the Economy, Volume 19

Volume Author/Editor: James M. Poterba, editor

Volume Publisher: MIT Press

Volume ISBN: 0-262-16236-9

Volume URL: http://www.nber.org/books/pote05-1

Conference Date: October 7, 2004

Publication Date: September 2005

Title: Tax Policy for Health Insurance

Author: Jonathan Gruber

URL: http://www.nber.org/chapters/c0164

Tax Policy for Health Insurance

Jonathan Gruber, *MIT and NBER*

Executive Summary

Despite a $140 billion existing tax break for employer-provided health insurance, tax policy remains the tool of choice for many policymakers in addressing the problem of the uninsured. In this paper, I use a microsimulation model to estimate the impact of various tax interventions to cover the uninsured, relative to an expansion of public insurance designed to accomplish the same goals. I contrast the *efficiency* of these policies along several dimensions, most notably the dollars of public spending per dollar of insurance value provided. I find that every tax policy is much less efficient than public insurance expansions: while public insurance costs the government only between $1.17 and $1.33 per dollar of insurance value provided, tax policies cost the government between $2.36 and $12.98 per dollar of insurance value provided. I also find that targeting is crucial for efficient tax policy; policies tightly targeted to the lowest income earners have a much higher efficiency than do those available higher in the income distribution. Within tax policies, tax credits aimed at employers are the most efficient—and tax credits aimed at employees are the least efficient—because the single greatest determinant of insurance coverage is being offered insurance by an employer, and because most employees who are offered such an option already take up that insurance. Tax credits targeted at non-group coverage are fairly similar to employer tax credits at low levels, but they are much less efficient at higher levels.

Federal, state, and local governments in the United States intervene in health insurance markets in a number of ways. Most prominent are the major public insurance programs for the elderly and disabled (Medicare) and for low income groups (Medicaid). In 2004, the

Medicare and Medicaid program each spent about $300 billion. Close behind is a much less known federal program that spends over $140 billion per year subsidizing the private purchase of health insurance. This program is larger than Unemployment Insurance, workers' compensation, the Temporary Assistance for Needy Families (TANF) cash welfare program, and Earned Income Tax Credit (EITC) wage subsidies combined. Yet it is little known and even less understood by the general public and many politicians.

This program is the tax exclusion of employer-provided health insurance expenditures. When employers pay their employees in cash, that compensation is taxed by federal, state, and sometimes local income taxes, as well as by federal and state payroll taxes. Yet when employers pay those same employees in health insurance, that compensation is completely untaxed. For a worker in the District of Columbia who faces an income tax rate of 25 percent, a Social Security and Medicare payroll tax rate (combined employer and employee shares) of 15.3 percent, and a DC income tax rate of 9.5 percent, this amounts to an almost 50 percent subsidy to employer-provided health insurance relative to cash compensation. In total, estimates suggest that these subsidies total over $140 billion in the United States in 2004.

Despite the large amount that governments in the United States spend on health care, major access problems remain. Forty-five million Americans lack health insurance, which results in limited access to many basic health services and reduced health. For many politicians, the answer to the access problem is a simple one: further expansion in the tax subsidization of health insurance.

In this paper, I analyze broadly the possibilities for tax policy as a means of addressing our health care problems. I begin by discussing the role of the existing tax exclusion. I then discuss a host of additional tax policies that might be used to increase health insurance coverage in the United States, ranging from tax subsidies to the purchase of health insurance plans by individuals, to targeted tax subsidies for employers.

To formalize these discussions, I rely on an extensive microsimulation model that has been developed to analyze the implications of a wide variety of health insurance reform options. This model incorporates the best available evidence from the health economics literature to illustrate how individuals, families, and firms respond to changes in the insurance environment. By incorporating these responses, I am able to compute dynamic estimates of the impact of health insurance

reforms on the distribution of health insurance coverage, government costs, and private health care burdens.

The paper proceeds as follows. I begin, in Part 1, with a detailed description of existing and proposed tax policies toward health insurance, and a brief review of the relevant literature on their impacts. In Part 2, I briefly describe the microsimulation model that forms the basis for my analysis. For comparison to later tax policy analyses, in this section I discuss the analysis of a prototypical expansion in public health insurance. Part 3 evaluates a host of alternatives for increasing insurance coverage. Part 4 concludes.

1. Background

1.1 The Employer Exclusion

As noted in the introduction, the third largest health care "program" in the United States is the exclusion of employer-provided health insurance expenditures from taxation. The subsidy to employer-provided health insurance is generally not well understood. This is not a subsidy to employers, but rather a subsidy to employees for insurance purchased in the employment setting. From the employer's perspective, whether she pays a worker in wages or health insurance is irrelevant; either way, a dollar of employer spending has the same effect on the firm's bottom line (and thus on corporate tax payments). From the worker's perspective, however, there is a large difference: by being paid in health insurance, rather than wages, the worker is saving taxes. So if the government wanted to end the tax subsidy, it would not involve corporate taxation; rather, the subsidy would be ended by including employer spending on health insurance as part of taxable compensation to the individual employee.

This subsidy has been extended not only to employer spending but to employee contributions to health insurance plans as well. In firms that have established a Section 125 plan, employee contributions for health insurance can be made on a tax-free basis. The cost of this subsidy is included in the $144 billion estimate noted above.

For many years, the exclusion of employer-provided insurance from taxation was seen as discriminatory toward the self-employed. Beginning in 1986, the health insurance premiums of the self-employed were made partially tax-deductible, and these premiums are now fully tax deductible. This additional exclusion costs the government $3.3 billion per year in foregone tax revenues.

1.2 Proposed Tax Policies Toward Health Insurance

Despite the existence of these large subsidies to employer-provided insurance, there have been many proposals in recent years to use the tax code to provide additional subsidies to private insurance to reduce the ranks of the uninsured. Most prominent among these are proposals to provide individuals with tax credits to purchase non-group insurance on their own (not through employers). A typical example of such a *non-group credit* is that proposed by President Bush in his 2004 budget. The key features of his plan are:

- It is fully available to all singles up to $15,000 of (modified) adjusted gross income (AGI), and phases out for singles by $30,000 of AGI.
- It is fully available to all families up to $30,000 of AGI, and phases out for families by $60,000 of AGI.
- The credit amount is $1,000 for each adult and $500 for each child, up to a maximum of $3,000 per family.
- Families can use this credit against up to 90 percent of their non-group insurance costs.
- The credit can be used only for non-group (specifically non-employer-provided) insurance.

An alternative, or supplement, for non-group credits are credits for employees to take up the insurance that they are offered, or an *employee tax credit*. Such tax credits are motivated by the fact that a large share of the uninsured, roughly one-quarter, are offered health insurance by their employers but do not take up that offer. The goal of employee tax credits is to subsidize those who are offered insurance—and thus entice them to take up that offer.

A final alternative for tax policy is to expand the existing subsidy to employers for their spending on health insurance through *employer tax credits*. For example, in his campaign for the 2004 democratic presidential nomination, Congressman Richard Gephardt proposed a 60 percent credit for the cost of health insurance for all firms, in addition to their existing tax exclusion. There have also been a large number of congressional proposals for targeted employer credits to small and/or low-wage firms.

1.3 Efficiency Implications of Tax Policies

Given the budgetary limitations on any public approach to expand health insurance coverage, a key concept that drives reform is the

efficiency of the policy. There are several different means of defining efficiency, which I will review below, but the basic concept is the extent to which new public spending is directed to those who would otherwise be uninsured, as opposed to "buying out the base" of existing insured individuals. The issue that is central to all of these definitions is *targeting*. If individuals were indelibly labeled as "insured" or "uninsured," then the government could easily target new tax subsidies to those labeled "uninsured," with no spending on those labeled "insured." In fact, this is not the case: insurance status is a choice of the individual and can respond to government policy in a way that causes the policy to have lower efficiency.

It is useful to think about the uninsured as tuna and those who already have insurance as dolphins. The goal of the environmentally conscious fishing industry is to catch as many tuna as possible in its nets, while minimizing the number of dolphins who are caught by those nets (which happens because tuna and dolphins swim together in the ocean). If the uninsured tunas are swimming in a separate ocean than are the insured dolphins, the problem is minimized. And if the uninsured tunas greatly outnumber the insured dolphins, then there is also a minimal dolphin catch. But, in reality, the 45 million uninsured tunas mostly swim in a part of the ocean where there are 180 million insured dolphins, making it difficult if not impossible for policymakers to design insurance nets to capture the tuna without pulling in the much more numerous dolphins.

There are three sources of inefficiency with tax credits. The first is spending on those who already have coverage through the subsidized form of insurance. The group that would benefit most clearly from subsidies for non-group insurance, for example, is those already holding non-group insurance. Yet the use of subsidies by this group does nothing to reduce the number of the uninsured.

The second is the crowdout of other forms of insurance through subsidizing a particular form of insurance. For example, when the government subsidizes non-group insurance, it can lead those with group insurance to move to the non-group market, either by their decision (switching out of employer-provided insurance) or by their employer's decision (dropping the offering of insurance at the firm). This crowdout may or may not lead to inefficiency in government spending, however. For example, when individuals lead employer-provided insurance for non-group insurance, they increase spending on new non-group subsidies but decrease spending on the existing

exclusion of employer-provided insurance purchases. On net, the rise in government spending is unclear.

The third source of inefficiency is the possible reduction in coverage for those who are insured before the policy is put into place. For example, suppose that a firm has a workforce that is predominantly, but not universally, eligible for a non-group credit. This firm might decide to stop offering health insurance because the majority of its employees can use the credit instead. The minority of employees that cannot use the credit is then out of luck, however, because they have lost their employer insurance with no subsidized alternative, and these individuals may become uninsured. This rise in the number of those who are uninsured offsets the reductions in the ranks of the existing uninsured, thus reducing the efficiency of the program by raising spending per newly insured person.

Based on this discussion, there are several different means of measuring efficiency. The traditional measure is the "buck for the bang": dollars of public spending per person newly insured. Another measure of interest is the extent of crowdout: the reduction in employer-provided insurance when other forms of insurance are subsidized (or when the existing subsidies to employer-provided insurance are reduced). A third measure of interest is the (gross) rise in the uninsured due, for example, to firm dropping. A fourth measure is the share of beneficiaries of any intervention who were previously uninsured, as opposed to receiving subsidies to remain insured.

A final measure incorporates the *type of uninsured* who are affected by reform. Simply counting the dollars per newly insured is not satisfactory when different reforms may appeal to very different populations. For example, a reform that significantly increases insurance coverage among children will have much lower costs than one that has the bulk of its effects in the much higher cost adult population. But this is an unfair comparison because the latter reform is essentially extending more valuable insurance coverage than is the former. So a better measure of efficiency is the spending per dollar of insurance value provided, which incorporates both the numerical increases in coverage and the cost of the individuals who are provided coverage.

2. Microsimulation Modeling

The analysis in this paper relies on microsimulation modeling using a model I have developed over the past five years (as first described in

Gruber and Levitt, 2000). This microsimulation model has several components, which I describe in this section.

2.1 Data

The data base for this analysis is the February and March 2001 Current Population Survey (CPS). The March survey contains data on family demographic characteristics, income, and health insurance coverage, while the February survey adds information on employer insurance offering. Importantly, the March survey also contains data on taxable income and marginal tax rates.

These data are matched to information on health insurance premiums and health costs. Data on the premiums for employer insurance, and the distribution of premiums between employers and employees, comes from the annual Kaiser/HRET national survey of employers. These data are matched by state (or state group for small states) and are assumed to fall with firm size. For non-group insurance, a premium for a healthy 40-year-old male is assigned based on analyses from the Community Tracking Survey and the Medical Expenditure Panel Survey (MEPS), and data on premiums collected by the Commonwealth Fund, the Health Insurance Association of America, and ehealthinsurance.com. This premium is then adjusted by age, sex, and health status using factors provided by an actuarial consulting firm.

Finally, data on underlying medical expenditures comes from the MEPS. Total medical expenditures of those with employer-provided health insurance are estimated as a function of age, sex, and health status. These estimates are then reduced by 15 percent to account for administrative costs of private health insurance. The resulting costs are used in two ways: as a measure of the cost of the individual if they were on public insurance; and as a measure of the underlying value of insurance provision. All cost data in the model has been updated to 2004 dollars.

2.2 Modeling Individual Behavior

These data are used to develop a microsimulation model that computes the effects of health insurance policies on the distribution of health care spending and private and public sector health care costs. This model takes as inputs both the data sources described above and the detailed parameterization of reform options. The model first turns these policy rules into a set of insurance price changes; for example, if the policy intervention is a tax credit for non-group insurance, then the

model computes the implied percentage change in the price of non-group insurance for each individual in the model. These prices changes are then run through a detailed set of behavioral assumptions about how changes in the absolute and relative price of various types of insurance affect individuals, families, and businesses.

The key concept behind this modeling is that the impact of tax reforms on the price of insurance continuously determines behaviors such as insurance take-up by the uninsured and insurance offering by employers. The model assiduously avoids so-called knife-edge behavior, where some critical level is necessary before individuals respond and beyond which responses are very large. Instead, behavior is modeled as a continuous function of how policy changes (net of tax) insurance prices.

In doing this type of analysis, a number of assumptions must be made about how individuals will respond to tax subsidies, through their effect on the price of insurance. These assumptions have been developed based on the available empirical evidence, as reviewed in detail in Gruber (2002). Some of the key assumptions are described in the following subsections.[1]

2.2.1 Take-up of Subsidized Non-group Insurance Among the Uninsured I calculate take-up of such subsidies by the uninsured by applying both a price elasticity and a correction for the burden of premiums relative to income. For the base price elasticity, I use −0.625. I then augment this with a correction factor of the form: $[1 - (X/\text{income})]^2$, where X is the post-subsidy non-group premium for one-half of the population, and X is the pre-subsidy non-group premium for one-half of the population. This term accounts for two factors that are likely to lead to take-up but then fall with income. The first is the fact that as income falls, individuals are less likely to take up subsidies that are less than 100 percent because disposable income is needed for other expenditures that may be perceived as more urgent (such as food and housing). The second is liquidity constraints: insurance expenditures are made throughout the year, but any credits or deductions are received the next April. This is a much larger problem for lower income individuals who have both little savings and potentially poor access to credit markets. I assume that, due to administrative efforts to address this "advanceability" problem, it arises for only one-half of the sample.[2] The quadratic form of the expression captures the fact that both of these effects are likely to operate very strongly toward the

bottom of the income distribution. On average, the take-up elasticity for the uninsured is −0.45 to −0.5.[3]

2.2.2 Switching from Group to Non-group Policies I assume that individuals compare their out-of-pocket costs of group insurance with the subsidized costs of non-group insurance in making their switching decisions. In particular, switching from group to non-group is a function of the post-subsidy non-group premium minus the post-subsidy employee cost of health insurance, divided by the full cost of group insurance (the value of the insurance), with an elasticity of −0.33.

2.2.3 Price Sensitivity of Employee Take-up of Employer-Provided Insurance One of the clearest lessons from health economics over the past decade is that the decision of employees to take up insurance provided by their employers is not very sensitive to price.[4] As a result, for those with insurance whose employers raise the cost of contributions, I compute the ratio of changes in employee contributions to insurance relative to the full price of employer-provided insurance and assume that there is only a −0.1 elasticity of take-up of employer-provided insurance with respect to this ratio. For those without insurance whose employers lower the cost of employee contributions, I compute the percentage change in employee contributions and assume an elasticity of −0.067 for changes of less than 75 percent, and an elasticity that rises to −0.75 for changes between 75 and 100 percent.

2.3 Modeling Firm Behavior

A key aspect of modeling health insurance policy is appropriately reflecting the decisions of firms because 90 percent of private health insurance is provided by employers. Economists tend to model firm decision-making as reflecting the aggregation of worker preferences within the firm. The exact aggregation function is unclear, as I reviewed earlier (Gruber, 2002); in my model, I assume that the mean incentives for the firm (e.g., the average subsidy rate for non-group insurance) is what matters for firm decision-making.

The fundamental problem faced by individual-based microsimulation models is that data on individuals does not reflect the nature of their co-workers; thus, it is impossible to compute exactly concepts such as the average non-group subsidy in a worker's firm. I address this problem by building synthetic firms in the Current Population Survey (CPS) and assigning each CPS worker a set of co-workers

selected to represent the likely true set of co-workers in that firm. The core of this computation is data from the Bureau of Labor Statistics that show, for workers of any given earnings level, the earnings distribution of their co-workers separately by firm size, region of the country, and health insurance offering status. Using these data, I randomly select 99 individuals in the same firm-size/region/health-insurance-offering cell as a given CPS worker to replicate statistically the earnings distribution for that worker's earnings level. These 99 workers then become the co-workers in a worker's synthetic firm.

These synthetic firms then face three decisions about insurance: offering (whether to offer if now not offering, or whether to drop if now offering); the division of costs between employer and employees; and the level of insurance spending. I model each of these decisions as subject to pressures from government interventions, in particular:

- Subsidies covering outside insurance options (non-group insurance or public insurance) exert pressures on firms offering insurance to drop that insurance and to raise the cost of employee contributions.
- Subsidies covering employer spending on insurance cause firms that don't now offer insurance to be more likely to offer, cause firms to pick up a larger share of the cost of insurance, and cause a rise in employer spending on insurance.
- Subsidies covering employee spending on insurance also raise the odds that firms offer insurance and raise employer spending on insurance, but they *lower* employer contributions to insurance.

Modeling the firm reactions to these pressures involves once again making a number of assumptions about the behavior of these synthetic firms. Some of the key assumptions are described in the following subsections.

2.3.1 Firm Offering/Dropping I key firm offering/dropping responses to the price elasticities of insurance demand for firms estimated in Gruber and Lettau (2004) as follows: −0.69 for firms with fewer than 100 employees, −0.2 for firms with 100–999 employees, and −0.1 for firms with more than 1,000 employees. For a firm offering in response to employer subsidies, I compute the ratio of subsidies to existing employer spending and apply these elasticities: for a firm offering in response to employee subsidies, I reduce this by 0.7 to account for the fact that only about 70 percent of employees take up insurance; for a firm dropping in response to a non-group subsidy, I

compare the extent of the non-group subsidy to the existing tax subsidy to employer insurance: when the non-group subsidy is below the existing group tax subsidy, I apply only a fraction of the Gruber-Lettau elasticities, rising from 50 percent to 100 percent to the point where non-group subsidies and existing tax subsidies are equal; and from that point on, I simply apply the Gruber-Lettau elasticity.

2.3.2 Employee Contributions Government subsidies covering spending on employer-provided insurance affects the distribution of spending across employer and employee. If the subsidy goes to the employer, I assume that 30 cents of each dollar of subsidy is spent in buying back employee contributions.[5] Likewise, if the subsidy goes to the employee, I assume that employers raise employee contributions to offset 70 cents of each dollar of (average across the firm) subsidies to employees. When a subsidy covers non-group insurance, I assume that the firm raises employee contributions by 15 percent of the subsidy rate to encourage non-group insurance take-up.

2.3.3 Employer Spending If the government offers an open-ended percentage credit, I assume an elasticity of spending with respect to the credit amount of 50 percent, but if there is a flat dollar credit, I assume that only 20 cents of each dollar goes to higher spending. For employee credits, spending reacts in the same way, but it is scaled down by 0.7.

Finally, a key assumption for this type of modeling is the assumption on the wage incidence of changes in employer insurance spending. In earlier work (Gruber, 2001), I reviewed the literature on incidence and concluded that there is strong evidence for full shifting to wages of firm-wide changes in insurance costs, with some evidence of shifting to sub-groups within the workplace as well. I make a mixed incidence assumption for this model. Any firm-wide reaction, such as dropping insurance or lowering employee contributions, is directly reflected in wages. Yet any individual's decision, such as switching from group to non-group insurance, is not reflected in that individual's wages; rather, the savings to the firm (or the cost to the firm) is passed along, on average, to all workers in the firm.

2.4 An Example: Expanding Public Health Insurance

It is difficult to interpret the results for tax policy that come from this model without some baseline, so in this section I illustrate the results

from this model for two examples of expansions in our existing safety net of public insurance programs. The first example is the introduction of free public health insurance for all persons in the United States with income below 100 percent of the poverty line; in this range, children and many parents are already eligible for public insurance, but single adults and most parents are not. The second is the introduction of public health insurance for all persons in the United States with income below 225 percent of the poverty line; once again, most children in the United States are eligible for much of this range, but most adults are not.

The effects of these three policies on health insurance coverage, costs, and the income distribution are summarized in Table 1, which presents some key outputs from these model runs. The first row shows the number of uninsured persons who take up the Medicaid expansions, which is 3.13 million persons for the 100 percent expansion, and 8.75 million persons for the 225 percent expansion. The next row shows a

Table 1
Public Insurance Expansions as a Benchmark

Policy	Expand to 100% of poverty	Expand to 225% of poverty
	Changes in population (millions of persons)	
Uninsured takeup	3.13	8.75
Uninsured share of takeup (%)	85.1%	69.6%
Uninsured increase	0.07	0.82
Net decrease in uninsured	3.06	7.93
Net change in employer insured	−0.31	−3.88
Net change in non-group insured	−0.34	−1.19
Net change in publicly insured	3.68	12.57
	Costs ($2004 billions)	
Total cost per year	$11,350	$28,670
Cost per newly insured ($)	$3,707	$3,615
	Targeting	
Average age of newly insured	37.5	36.3
Percentage of newly insured in fair/poor health	21.3	15.5
Average cost of newly insured	$3,154	$2,673
Spending per dollar of insurance provided	$1.17	$1.33

first measure of policy targeting: the percentage of individuals taking up public insurance who were formerly uninsured. For the 100 percent expansion, 85 percent of those taking up public insurance were formerly uninsured, and only 15 percent were leaving other forms of insurance to join public insurance. For the 225 percent expansion, 70 percent of those taking up public insurance were formerly uninsured. As we will see below, these are incredibly well-targeted policies.

The next row shows the rise in the number of uninsured that occurs through one of two channels: individuals who lose insurance when their firms stop offering, or individuals who drop employer-provided insurance when the cost of contributions rise in response to the public insurance expansion (the majority of this total is accounted for by the former group). This number is very modest for the 100 percent expansion, 0.07 million, but it is more sizeable for the 225 percent expansion, 0.82 million. The next row shows the net change in uninsured, approximately 3 million and 8 million, respectively.

The remaining rows in the first panel show the change in the size of other insurance groups from this policy. There is a small reduction in the number of employer insured from the expansion to 100 percent as those with employer insurance switch to public insurance (since it is free) and some employers stop offering employer insurance. There is also a small migration from non-group insurance to public insurance. When eligibility is expanded to 225 percent of poverty, there is a much larger reduction in the number of employer insured (almost 4 million) and non-group insured (over 1 million).

The total cost of these policies is shown in the next panel. The expansion to 100 percent of the poverty level costs almost $11.4 billion per year, while the expansion to 225 percent costs almost $29 billion. Both policies have a comparable cost per person who is newly insured of roughly $3,700.

The next panel of the table focuses on the targeting of the policy, in terms of which types of formerly uninsured individuals are helped by this intervention. For these expansions, the average age of the uninsured person who gains insurance is roughly 37 years; 15–20 percent are in fair or poor health, and 18–20 percent are in excellent health. For comparison, among all uninsured persons, the average age is 30 years; 9 percent are in fair or poor health, and 29 percent are in excellent health. Thus, the set of individuals now covered by these expansions is in much worse health, and it is therefore more costly to insure, than is the set of uninsured individuals not affected by the policy.

This point is summarized by the next row in the table, which shows the average cost associated with insuring the uninsured who gain coverage through these initiatives. For this calculation, I have imputed to each person in the data the cost of insuring them through public insurance (which is roughly 15 percent below the cost of private insurance) as a measure of the true insurance cost. For all uninsured persons, this average cost is $2,100; for those gaining insurance through these expansions, the cost is 25–50 percent higher, at $2,700–$3,200.

The last row of Table 1 shows government spending per dollar of insurance value provided, which is the ratio of government spending to the sum of the insurance value provided to the uninsured. This figure is $1.17 for the expansion to 100 percent of the poverty level, indicating that for each dollar of insurance that the government is providing, it is spending $1.17. Thus, roughly speaking, the deadweight loss of this approach to providing insurance, relative to an ideal that gave insurance only to those who were otherwise uninsured, is 17 cents. This small deadweight loss arises from the small amount of substitution from other forms of insurance into public insurance that accompanies the reduction in the uninsured. For the larger expansion, the inefficiency rises somewhat, to 33 cents.

3. Alternative Expansions of Tax Subsidies

In this section, I consider the implications of alternative types of expansions of the tax code to increase insurance coverage. To make our results comparable to the runs shown earlier for public insurance, I have chosen tax policy parameters designed to meet two goals: reduce the number of uninsured by 3 million persons, and reduce the number of uninsured by 8 million persons. The cost per person covered from these policies typically falls as the number of persons covered rises, so it is important to compare these policies on a comparable basis.

One overall note of importance: for all of the analysis considered here, I assume that any tax policies are fully refundable. Roughly half of the uninsured do not pay taxes, so any non-refundable tax policy will have very limited impact.

3.1 Non-group Credits

I begin with non-group tax credits of a simple form. Individuals and families are eligible for a credit of one size for single coverage and a credit of another size for family coverage. This eligibility is restricted

on income in one of two ways. First, I consider a tightly targeted non-group credit that is fully available to single persons with income below $15,000, or families with income below $30,000, and that phases out as income rises, with eligibility ending at $30,000 for singles and $60,000 for families. Second, I consider a loosely targeted non-group credit that is fully available to single persons with income below $25,000, or families with income below $50,000, and that phases out as income rises, with eligibility ending at $50,000 for singles and $100,000 for families. Given these income restrictions, the value of the tax credit is then set to hit the targets of 3 million and 8 million reduction in the uninsured.

The results of this analysis are shown in Table 2. Several differences from Table 1 are immediately apparent. First, there is a much larger

Table 2
Non-group Credits

Targeting	Tight (15–30 K/30–60 K)		Loose (25–50 K/50–100 K)	
Target	3 million	8 million	3 million	8 million
Credit amount	$1,460/ $3,650	$4,750/ $11,875	$1,685/ $4,210	$4,350/ $10,875
	Changes in population (millions of persons)			
Uninsured takeup	5.08	12.38	7.88	15.08
Uninsured share of takeup (%)	37.4%	45.0%	34.0%	36.2%
Uninsured increase	2.08	4.37	4.88	7.07
Net decrease in uninsured	3.00	8.01	3.00	8.01
Net change in employer insured	−5.29	−13.50	−13.16	−25.97
Net change in non-group insured	8.69	22.57	16.42	34.85
Net change in publicly insured	−0.40	−1.06	−0.26	−0.87
	Costs ($2004 billions)			
Total cost	$11,222	$56,458	$21,468	$85,065
Cost per newly insured ($)	$3,741	$7,048	$7,157	$10,619
	Targeting			
Average age of formerly uninsured	25.6	27.4	26.2	27.9
Percentage newly insured in fair/poor health	2.3	4.0	2.4	3.8
Average cost of newly insured	$1,481	$1,675	$1,488	$1,668
Spending per $ of insurance	$3.24	$4.63	$9.26	$7.50

gross increase in the uninsured, which offsets the gross reduction in meeting our targets of 3 and 8 million. For example, to hit the target of an 8 million person reduction with tight targeting requires take-up by 12.4 million persons because there is a 4.4 million person rise in the uninsured after firms drop insurance. Second, there are very large reductions in the number of employer insured due both to employee switching and firms dropping insurance; roughly 20 percent of this reduction comes from switching, and the remaining 80 percent is due to firm dropping. Morever, a much smaller share of the recipients, between 34 and 45 percent, were previously uninsured; the majority are using this subsidy while retaining insurance coverage.

This may seem like a lot of firms dropping insurance in response to (in particular in the first column) fairly small non-group credits, but it is important to remember that this reduction in employer insured is from a very large base of over 160 million employer-insured. The 5.29 million reduction in employer-insured in the first column, for example, represents just over 3 percent of the employer-insured in the United States. Even with the enormous non-group credit shown in the second column, $4,750 for singles and $11,875 for families, only 8.2 percent of those with employer-provided insurance drop that insurance. This is partly because the credits are targeted to only a subset of employees so that, on average, the pressure on employers to stop offering insurance (due to the erosion of the employer tax advantage) is small. When the credit is more loosely targeted, in the final two columns, the reduction in the employer-insured is much larger, rising to almost 16 percent of the employer-insured in the final column.

Third, this approach is by and large more expensive than public expansions. For the 3 million target, with the tightly targeted credit, the cost is very similar to the public expansion. The other approaches, however, are much more expensive, both overall and (by definition) per person newly insured. Indeed, a loosely targeted credit designed to cover 8 million persons costs over $85 billion per year, or more than $10,000 per person newly insured.

Another striking difference between public expansions and non-group credits is the targeting of the spending. In contrast to public expansions, the set of uninsured who gain coverage through non-group credits is much healthier than the average uninsured person, with an average age of 25–28 years and only 2–4 percent in fair or poor health. The average cost of insuring the newly insured is only $1,500–$1,800 per year.

The reason for these low costs is that these types of partial subsidies to non-group insurance are much more attractive to the healthy individuals, for whom the lower cost of non-group insurance makes these subsidies a larger percentage. As a result, the value of insurance provided by these policies is much less than it is for a public expansion. Indeed, as the last row shows, it takes more than $3 to almost $10 of government spending to provide just $1 of insurance coverage through these policies. Thus, by this measure, non-group tax credits are much less efficient than are public insurance expansions.

A clear lesson from this analysis is that the efficiency of non-group credits is much higher if they are tightly rather than loosely targeted. Given the low incomes of the uninsured, focusing credits on the lowest income groups in society leads a higher share of the expenditures to be directed towards those who would be otherwise uninsured. This highlights the value of tightly targeting health care interventions, but it is not clear how politically realistic such targeting will be given the broad reach of recent tax policy changes.

3.2 Employee Credits

Another tax policy alternative that has received substantial attention is the use of tax credits to offset the costs to employees of purchasing their employer-provided health insurance. The motivation for these credits is the "low hanging fruit" of the large number of uninsured who are already offered employer-provided health insurance. Because these individuals are in an arena where it is easy to obtain health insurance, the reasoning goes, and because employers already pay the majority of insurance costs, it should be relatively cheap to subsidize the cost of these uninsured taking up insurance.

There are three problems with this argument, however. First, if employee contributions become tax-subsidized by the government, then employers have an incentive to shift the costs of insurance to employees—or at least they no longer have a disincentive to do so. As noted earlier, Gruber and McKnight (2003) find a substantial negative response of employee contributions to the tax subsidy covering employer-provided insurance; presumably, tax subsidies covering employee contributions would have the opposite effect. Second, this is a very poorly targeted policy: the vast majority of those offered employer-provided health insurance take up that insurance. Indeed, of those offered employer-provided insurance, only about 7 percent are uninsured; even among the population below the poverty line, 75

Table 3
Employee Credits

Targeting	Tight (15–30 K/30–60 K)		Loose (25–50 K/50–100 K)	
Target	3 million	8 million	3 million	8 million
Credit amount	$800/ $2,000	$3,600/ $9,000	$610/ $1,525	$2,000/ $5,000
	Changes in population (millions of persons)			
Uninsured takeup	3.47	9.21	3.55	9.29
Uninsured share of takeup (%)	5.8%	13.0%	3.5%	8.1%
Uninsured increase	0.47	1.22	0.54	1.27
Net decrease in uninsured	3.00	7.99	3.01	8.02
Net change in employer insured	5.45	13.64	5.53	13.95
Net change in non-group insured	−0.74	−2.07	−0.86	−2.43
Net change in publicly insured	−1.7	−3.57	−1.66	−3.50
	Costs ($2004 billions)			
Total cost	$35,153	$177,460	$58,671	$209,077
Cost per newly insured ($)	$11,707	$22,198	$19,501	$26,073
	Targeting			
Average age of newly insured	27.6	28.5	27.6	28.7
Percentage of newly insured in fair/poor health	8.5	8.5	7.5	8.0
Average cost of newly insured	$1,986	$2,005	$1,944	$1,981
Spending per $ of insurance	$5.82	$10.92	$9.91	$12.98

percent of those offered insurance are insured. Finally, as noted above, a sizeable literature now documents the fact that the decision to take up employer-provided insurance, if offered, is not price sensitive.

Table 3 shows the results of running several employee tax credit policies through my microsimulation model. Once again, this table considers tightly and loosely targeted credits, with the same income cutoffs as for non-group insurance: targeted to hit 3 million and 8 million person reductions in the uninsured. Along some dimensions, employee tax credits look more attractive than non-group credits. There is only a small rise in the number of uninsured, which offsets the takeup by the previously insured; this small increase comes from individuals leaving their employer-provided insurance because the firm has raised the cost of contributions, but they are not one of the subsidized employees. There is also a rise rather than a fall in employer-provided

insurance. Finally, as the last panel shows, these policies are somewhat better targeted to the uninsured in poor health than are non-group credits, although the targeting is still much less than with public insurance expansions.

Employee credits perform much worse, however, on one key dimension: cost. To cover 3 million persons with a tightly targeted employee tax credit would cost over $35 billion per year, a cost of almost $12,000 per newly insured; if the targeting is looser, the cost rises to almost $20,000 per newly insured. These higher costs arise because there is enormous expenditure on the vast majority of those who already have insurance, so this is simply a subsidy to existing behavior: in most cases, fewer than 10 percent of those using this subsidy were previously uninsured. As the final row of Table 3 shows, the government would spend between $6 and $13 per dollar of insurance provided if it pursued these types of policies.

Thus, while employee tax credits create much less disruption in insurance markets than do non-group credits, employer-insured individuals can take them without changing their existing insurance arrangements. These tax credits cost much more because so much of the spending is an inframarginal subsidy to those who already have employer-provided insurance. Indeed, only about 10 percent of the individuals taking advantage of this credit were previously uninsured.

3.3 Employer Credits

The final type of tax credit that I consider is credits paid to employers offering insurance to their employees. Once again, the parameters of these tax policies are chosen to hit targets of 3 and 8 million person reductions in the uninsured, and once again there is more tightly and more loosely targeted versions of these policies. The more loosely targeted version is credits that are provided to all employers with fewer than 50 employees. While typical of many proposals to subsidize employer-provided coverage, this type of subsidy structure has two disadvantages. First, the majority of firms with fewer than 50 employees still offer health insurance: non-offering is concentrated in the very smallest firms, and in those firms with the lowest-wage jobs. Second, a cutoff, or a so-called cliff, at 50 employees can provide disincentives for firms to grow beyond the critical 50 employee level.

To address these concerns, a more tightly targeted version of these credits would make three changes: focus the subsidy dollars on the smallest firms, target subsidies to firms with the lowest wage

Table 4
Employer Credits

Targeting	Tight (25–50 employees; 20–30 K average earnings)		Loose (<50 employees)	
Target	3 million	8 million	3 million	8 million
Credit amount	$1,050/ $2,625	$3,100/ $7,750	$700/ $1,750	$2,370/ $5,925
	Changes in population (millions of persons)			
Uninsured takeup	3.00	7.95	2.97	8.02
Uninsured share of takeup (%)	20.2%	35.1%	8.6%	18.2%
Uninsured increase	0	0	0	0
Net decrease in uninsured	3.00	7.95	2.97	8.02
Net change in employer insured	4.65	12.14	4.79	12.74
Net change in non-group insured	−1.03	−2.68	−1.16	−3.33
Net change in publicly insured	−0.62	−1.51	−0.66	−1.39
	Costs ($2004 billions)			
Total cost	$14,078	$58,865	$20,695	$86,829
Cost per newly insured ($)	$4,685	$7,407	$6,997	$10,829
	Targeting			
Average age of newly insured	35.0	30.7	30.3	30.8
Percentage of newly insured in fair/poor health	8.6	8.5	8.1	8.1
Average cost of newly insured	$1,986	$2,004	$1,961	$1,978
Spending per $ of insurance	$2.36	$3.70	$3.57	$5.47

employees, and phase out the subsidy as both firm size and wages rise to avoid cliffs. In my more tightly targeted version of these employer credits, phase out starts at a firm size of 25 (and ends at a firm size of 50), and credits are provided in full only to firms with average earnings below $20,000 and are phased out by firm average earnings of $30,000 (the average for small firms).

The results for these more tightly and loosely targeted credit versions are presented in Table 4. The first noticeable implication of this approach is that the gross and net change in the number of uninsured are identical: there is no crowdout with subsidies to employers. Nevertheless, once again, a sizeable share of the dollars are delivered to those who already have health insurance; as the final row of the second panel shows, only between 9 and 35 percent of the subsidy recipients were formerly uninsured.

In terms of total costs, and therefore costs per newly insured, the employer credit is comparable to the non-group credit, and both remain much less costly than the employee credit. One striking difference between the non-group credit and the employer credit, however, is that the employer credit covers a substantially higher cost population. The average cost of the individuals gaining insurance is around $2,000 for the employer tax credits, while the average cost is around $1,500 for the non-group tax credits. As a result, government spending per dollar of insurance delivered is much lower for the employer credit than for the non-group credit.

Table 4 once again highlights the important role of targeting. The tight employer credit is fairly well targeted, at least relative to other tax policies, with 20–35 percent of the benefits going to the uninsured. The loose employer credit, however, is very poorly targeted because the majority of those in firms with less than 50 employees already have health insurance, so that only 9–18 percent of the benefits go to the uninsured.

3.4 Comparison

Table 5 provides a comparison of these various policy options, along the various measures of efficiency noted earlier:

- Induced increase in the number of uninsured.
- Change in employer-insured.
- Dollars of spending per newly insured.
- Percentage of beneficiaries formerly uninsured.
- Average cost of those gaining insurance.
- Dollars of spending per dollar of insurance provided.

The best measure of efficiency of government spending is the last item in the list. It is immediately clear that expanding public insurance vastly outperforms tax policy along this dimension. The most efficient tax policies along this dimension spend three times as much per dollar of insurance provided as do public expansions.

Within tax policies, several lessons are apparent. First, employer tax credits are the most efficient outcome, followed fairly closely by non-group credits, with employee tax credits clearly the worst. Second, tightly targeted tax credits are much more efficient than are loosely targeted tax credits, particularly for non-group insurance. Third, efficiency almost universally declines as the size of the credit grows because (1) the uninsured are receiving more per person and (2) the

Table 5
Policy Comparison

	Uninsured increase	Share of recipients formerly uninsured	Change in employer insured	Cost per newly insured	Average cost of newly insured	Spending per dollar of insurance
			Covering 3 million persons			
Public	0.07	85.1	−0.31	$3,707	$3,154	$1.17
Tight non-group	2.08	37.4	−5.29	$3,741	$1,481	$3.24
Loose non-group	4.88	34.0	−13.16	$7,157	$1,488	$9.26
Tight employee	0.47	5.8	5.45	$11,707	$1,986	$5.82
Loose employee	0.54	3.5	5.53	$19,501	$1,944	$9.91
Tight employer	0	20.2	4.65	$4,685	$1,986	$2.36
Loose employer	0	8.6	4.79	$6,997	$1,961	$3.57
			Covering 8 million persons			
Public	0.82	69.6	−3.88	$3,615	$2,673	$1.33
Tight non-group	4.37	45.0	−13.50	$7,048	$1,675	$4.63
Loose non-group	7.07	36.2	−25.97	$10,619	$1,668	$7.50
Tight employee	1.22	13.0	13.64	$22,198	$2,005	$10.92
Loose employee	1.27	8.1	13.95	$26,073	$1,981	$12.98
Tight employer	0	35.1	12.14	$7,407	$2,004	$3.70
Loose employer	0	18.2	12.74	$10,829	$1,978	$5.47

credits become more attractive to the insured. Thus, one cannot compare directly two policies that cover very different numbers of persons because the policy covering more persons will be less efficient. Fourth, the efficiency of any policy is determined by several interactive factors: the size of the credit required to achieve the targeted reduction in the number of uninsured, the share of benefits going to the uninsured, and the targeting of the benefits in terms of the health of the uninsured.

For example, compare the tight non-group and tight employer credits that cover 3 million persons. In this case, the size of the credit is about 50 percent bigger for non-group credits, the share of recipients previously uninsured is almost twice as large for non-group credits, but the recipients are much healthier for non-group credits. As a result, on net, employer credits are more efficient. As another example, compare employee to employer credits. These credits are similar in terms of the size of the credit and the average health of the formerly uninsured recipients, but employer credits deliver a much larger share of benefits to the formerly uninsured.

From Table 5, we can outline the weaknesses of each tax policy relative to each other, and relative to the benchmark of a public insurance expansion. Non-group credits have the highest share of recipients who are formerly uninsured; for loosely targeted credits, this share is much higher than it is for either employee or employer credits. Yet they are much less efficient than employer credits because the uninsured who take the credits are much healthier and because a much larger credit is required to achieve the net reduction in the number of uninsured. This larger credit is required, in turn, because the non-group credits cause the largest increase in the number of uninsured. Thus, there is a vicious cycle with non-group credits; to cover many uninsured people requires a larger credit, but the larger the credit, the more the erosion of the employer market and the larger gross rise in the number of uninsured that must be offset by uninsured take-up of the non-group credit.

Employer credits feature a small required credit amount, and the uninsured who take up coverage are of average health. But these credits deliver a relatively small share of their benefits to the formerly uninsured, particularly if the credit is loosely targeted. The least attractive option, from an efficiency perspective, is credits to employees because such a very small share of benefits accrue to the formerly uninsured, particularly for loosely targeted employee credits.

4. Conclusion

It is clear from the analysis in this paper that if the goal is to cover 3–8 million uninsured persons, expanding public insurance is a more efficient option than any tax policy considered so far. Despite this fact, tax policy will continue to be the avenue of choice for expanding health insurance coverage in the United States for politicians of many stripes. Thus, it is critical to understand the strengths and weaknesses of alternative tax policy approaches.

Several lessons for tax policy are clear from this analysis. First, and probably most important, regardless of which tax policy option is considered, targeting is key: tightly targeted tax policies dramatically outperform loosely targeted policies in terms of efficiency. This conclusion is important to emphasize because targeting comes with political costs; it is much more politically expedient to allow a larger group of individuals to benefit from a policy than to restrict those benefits to a smaller (low-income group). Yet widening the income range of tax policies comes at great cost in terms of their effectiveness.

Second, one cannot compare directly two policies that cover very different numbers of uninsured because the efficiency of any tax policy falls as its scope increases. Finally, what matters for the efficiency of tax policy is not only the targeting of benefits in terms of the share of individuals who are uninsured, but also which individuals are covered. Providing coverage to very young and healthy individuals provides less insurance value per dollar of spending than does providing coverage to higher cost groups.

Notes

I am grateful to Charles Stoecker and William Lincoln for excellent research assistance, and to Jim Poterba and conference participants for helpful comments.

1. There is obviously variation in the possible assumptions that could be made here. See Remler et al. (2002) for a broad review of the assumptions made in microsimulation models such as these.

2. This assumption may be generous given that the government's only existing experience with advanceability, advance claiming of the earned income tax credit, has only a 1 percent take-up rate.

3. At the average income correction factor in our sample of the uninsured, this produces an elasticity of −0.5. This estimate is lower than that of Gruber and Poterba (1994), who suggest elasticities of −1 or greater (in absolute value). The upper bound elasticity is sim-

ilar to recent estimates by Royalty (2000). The average elasticity is somewhat higher than the range of −0.33 to −0.4 presented in Marquis and Long (1995).

4. See Gruber and Washington (forthcoming) for a review of the literature on this point and additional evidence.

5. Gruber and McKnight (2003) estimate an elasticity of employee contributions with respect to the tax price of employer-provided insurance of 0.3.

References

Gruber, Jonathan (2001). "Health Insurance and the Labor Market." In *The Handbook of Health Economics*, Joseph Newhouse and Anthony Culyer (eds.). Amsterdam: North Holland, 645–706.

Gruber, Jonathan (2002). "Taxes and Health Insurance." In *Tax Policy and the Economy 16*, James Poterba (ed.). Cambridge: MIT Press, 37–66.

Gruber, Jonathan, and Michael Lettau (2004). "How Elastic Is the Firm's Demand for Health Insurance?" *Journal of Public Economics* 88:1273–1294.

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

Gruber, Jonathan, and Robin McKnight (2003). "Why Did Employee Health Insurance Contributions Rise?" *Journal of Health Economics* 22:1085–1104.

Gruber, Jonathan, and James Poterba (1994). "The Elasticity of Demand for Health Insurance: Evidence from the Self-Employed." *Quarterly Journal of Economics* 109:701–734.

Gruber, Jonathan, and Ebonya Washington (forthcoming). "Subsidies to Employee Health Insurance Premiums and the Health Insurance Market." Forthcoming in *Journal of Health Economics*.

Marquis, Susan, and Stephen H. Long (1995). "Worker Demand for Health Insurance in the Non-Group Market." *Journal of Health Economics* 14:47–63.

Remler, Dalia, Joshua Graff Zivin, and Sherry Glied (2002). "Modeling Health Insurance Expansions: Effects of Alternative Approaches." NBER Working Paper no. 9130. September.

Royalty, Anne (2000). "Tax Preferences for Fringe Benefits and Workers' Eligibility for Employer Health Insurance." *Journal of Public Economics* 75:209–227.