Microsimulation Estimates of the Effects of Tax Subsidies for Health Insurance

Abstract *- The continued rise in the number of non–elderly Americans without health insurance has led to considerable interest in tax–based policies to raise the level of insurance coverage. This paper describes a detailed microsimulation model that has been developed to evaluate such tax–based policies, and its findings for the impact of policies on government costs and insurance coverage. I find that while tax subsidies could significantly increase insurance coverage, even very generous tax policies could not cover more than a sizable minority of the uninsured population. But there are several design features that can clearly make tax policy more effective: using tax credits rather than deductions; making credits refundable; and addressing the timing mismatch between when insurance purchases are made and tax refunds are received. I also document a clear tradeoff between the scope of tax subsidies and their efficiency.*

INTRODUCTION

The dramatic rise and high rates of uninsured people in the U.S., despite an economic boom that has had only one interruption in 15 years, is striking. In 1987, 14.8 percent of non–elderly Americans were without health insurance. Over the next decade, the percentage of the non–elderly population without insurance coverage grew by nearly 25 percent to 18.3 percent, so that in 1997 there were over 43 million uninsured Americans (Employee Benefit Research Institute, 1999). Particularly troubling is the significant increase in the number of uninsured children in the U.S.; despite dramatic increases in the expansion of public health insurance through the Medicaid program since the mid–1980s, the share of children without health insurance has grown by over 10 percent since 1987.

These trends have motivated considerable policy discussion at both the Federal and state levels. At the Federal level, they were one motivation for the ultimately unsuccessful attempt of the Clinton Administration to promote comprehensive reform of our health care system. The failure of this attempt has returned the policy focus to incremental reforms. Following the passage of the Child Health Insurance Program in 1997, most Congressional discussions have centered on the idea of using the tax system to subsidize the purchase of insurance by individuals. Tax–based approaches to expanding insurance coverage have a certain intuitive appeal. They

Jonathan Gruber
Massachusetts Institute of Technology, Cambridge, MA 02139 and NBER, Cambridge, MA 02138

National Tax Journal
Vol. LIII, No. 3, Part 1

would provide tax equity by providing financial benefits to those purchasing coverage individually; these benefits are now enjoyed only by the self-employed or those with employer-sponsored coverage. They would also rely on the private insurance system rather than a government sponsored program that might carry stigma for some people. And finally, they can be seen as providing a tax cut rather than creating a more politically controversial new spending program.

Yet while the tax equity argument is compelling—especially given the estimated $100 billion that is now spent each year providing federal tax subsidies for the purchase of employer-sponsored health insurance[1]—the ability of tax subsidies to meaningfully reduce the number of people uninsured remains uncertain and unproven. Moreover, the spectrum of tax-based approaches that have been proposed either formally or informally is quite large, ranging from deductibility of insurance costs for individuals to refundable tax credits that might cover most or all of the cost of typical health insurance policy.

In this paper, I assess the potential implications of a range of tax-based approaches using a new micro-simulation model developed specifically for this purpose. I examine how different characteristics of these proposals are likely to affect such outcomes as: the overall cost to the federal government, the number of the uninsured who would gain coverage, which income groups would benefit from the subsidies, and how those who now have employer-sponsored coverage would be affected.

TAX POLICY TOWARDS HEALTH INSURANCE—PROPOSED APPROACHES

The past few years have seen a variety of proposed reforms to the tax treatment of health insurance expenditures. There are a large number of potential approaches to expanding tax subsidies to health insurance. Any detailed proposals must, at least, address the following list of questions.

Deduction or Credit?

Currently, the employer exclusion is a deduction, through which individuals' taxable income is effectively reduced by the amount of their health insurance expenditure. The actual reduction in taxes (i.e., the subsidy) depends on the individual's tax rate. The alternative is a credit, through which an individual's taxes are directly reduced by a fixed dollar amount, regardless of the individual's tax rate. One key difference between these approaches is that a credit provides the same amount of subsidy to all income-eligible taxpayers, while a deduction provides a subsidy that rises with the tax rate (so that it is higher for higher income taxpayers). Another difference is that the credit can provide a 100 percent subsidy to the cost of insurance for some individuals, while a deduction only provides a subsidy rate equal to the individual's tax rate (e.g., someone in the 15 percent tax bracket would receive a subsidy equal to 15 percent of the insurance premium). This is an important difference for health insurance subsidies because, of those uninsured who have positive tax liabilities, 90 percent are in the 15 percent tax bracket.[2]

Refundability?

A key limitation of tax policy in increasing insurance coverage is that 45 percent of the uninsured do not pay any taxes against which any subsidy can be applied.[3] If a tax credit was made refund-

[1] Sheils and Hogan (1999).
[2] Author's tabulations from the March 1997 CPS.
[3] Author's tabulations from the March 1997 CPS.

able, as with our current Earned Income Tax Credit (EITC), it would reach even those potential participants with no tax liability. This would mean that the individual would receive a refund from the IRS equal to the amount of the credit.

Cap on the Subsidy?

The current employer exclusion is unlimited, applying to all expenditures by an employer on health insurance. An alternative, for either a deduction or a credit, would be to cap the amount of insurance expenditure that is eligible for tax subsidization. This can be done by only qualifying a certain percentage of spending for the subsidy (as is currently done for the self–employed), or by capping the amount of spending that is eligible at some dollar level, or some combination of the two. This would lower the costs of the tax policy, but would also lower its value to potential participants and as a result limit use of the subsidy.

Income Limitations?

Insurance status is fairly well correlated with income; 85 percent of the uninsured have incomes below the median household income level for their family structure.[4] As a result, the ability of tax policy to target a given amount of public dollars to the uninsured is enhanced if the availability of tax subsidies is income limited to some extent.

Which Populations?

As noted earlier, there are three potential groups that can benefit from new tax subsidies: those who are not employed in an incorporated firm; those who work for a firm that does not offer health insurance; and those who work in a firm that does offer health insurance, but for whom the employee portion of health insurance contributions are made on an after–tax basis (e.g., no Section 125 plan is available). Tax subsidies can, in theory, be offered to only the first population, by restricting the new subsidy only to those not offered insurance; to the first and second, by qualifying only non–group (or non–employer provided) health insurance expenditures for subsidies; and to all three, by qualifying any out of pocket (non–employer) spending on health insurance for a subsidy.

Other Policies?

Another important question is whether tax policies targeted towards non–group coverage should be accompanied by insurance market reforms that would make coverage more accessible for individuals buying coverage on their own. These could include requiring insurers to offer insurance to anyone regardless of health status or to restrict variations in premiums based on an individual's health.

OVERVIEW OF THE SIMULATION MODEL CREATED FOR THE ANALYSIS

In this section, I provide a very brief overview of the simulation model employed for this analysis. A more detailed description is provided in Gruber (2000).

This model uses as its base micro–data on a nationally representative sample of individuals from the Current Population Surveys (CPS) for February and March, 1997. The former has information on employer insurance offering, while the latter has information on insurance coverage from all sources, income, firm characteristics, self–assessed health status and demographics. As well, in recent years, the March CPS has also provided a detailed calculation of taxable income and tax rates for each family in the sample. We

[4] Author's tabulations from the March 1997 CPS.

supplement the CPS with data from KPMG–Peat Marwick, which provide information by region and firm size on employer premiums, employee premium shares, and whether employee premiums are made on a pre–tax basis, data from the Community Tracking Survey, and quotes from non–group insurers on the costs of non–group insurance policies.

We then use the data to simulate the impact of alternative tax policies on insurance coverage. This involves assessing how the policies affect individuals and employers in different circumstances—based on such factors as their insurance status, income, and tax rate—and how those individuals and employers respond. For each policy, the simulation model can estimate effects such as: the overall cost to the federal government, how many and what types of people become insured, and how many employers currently offering coverage drop it. In doing this type of analysis, a number of assumptions must be made about how individuals and employers will respond to tax subsidies, including:

- The extent to which the currently uninsured will purchase the newly subsidized insurance coverage
- The extent to which those with non–group coverage will take up subsidies to their insurance spending
- The extent to which those with group coverage will switch to non–group coverage if it is subsidized
- The extent to which firms will react to the availability of subsidized non–group coverage by dropping their offering of insurance to their employees, or by cutting back on employer premium contributions to insurance
- The extent to which firms will react to the availability of subsidized employee premium contributions by reducing their premium contributions.
- The extent to which those employees dropped from group coverage will then take up subsidized non–group coverage
- The extent to which insured employees facing higher premium contributions will drop group coverage, and to which uninsured employees facing lower premium contributions through tax subsidization will take up group coverage.
- The extent to which those on Medicaid will switch back to non–group coverage if it is made available on a subsidized basis

These assumptions—which are detailed in the technical appendix to Gruber (2000)—are based on other published studies, where available, as well as consultations with economic, actuarial, and policy experts.

We consider below a variety of tax policy options. It is very useful, however, to start with a common "base case," from which the implications of changing policy parameters can be considered. Our base case for the analysis below is a tax credit for health insurance spending that is capped at $1,000 per year for single filers and $2,000 per year for joint/head of household filers. As we detail in the Appendix, this amounts to roughly 43 percent of the costs of a typical non–group policy for an uninsured individual, and about 31 percent of the costs of a family policy for the typical uninsured family. The credit is refundable, and the full amount of the credit is available only to joint filers with taxable incomes of $75,000 or less, phasing out to zero credit at taxable incomes of $100,000; the limits are $45,000 and $60,000, respectively, for single filers. It is available only for non–employer provided insurance, so that it cannot be used towards the purchase of employer health insurance premiums; but it is available to all persons, even those offered employer–provided health insurance.

TAX POLICY: INSURANCE COVERAGE AND COST IMPLICATIONS

Base Policy

The impacts of this base case policy on insurance coverage and costs is presented in Tables 1A and 1B. The first table shows the total cost of the policy; the take–up of the subsidy by various groups, categorized by their pre–subsidy insurance status; and the net change in the size of these groups from before to after the subsidy. We present all population estimates both in absolute millions of persons and as a percentage of the size of the group before the policy impact; all group sizes include only the non–elderly. We explore in particular, for the employer–insured, the avenues that lead to the net change in this group. The second table shows a distributional analysis of the impacts of the policy. We consider the division of the population into those below the federal poverty line ($17,274 for a family of four), those between 100 percent of the federal poverty line and 200 percent of that amount, those between 200 and 300 percent, those between 300 and 400 percent, and those over 400 percent of the federal poverty line; this last cutoff is about 33 percent more than the median family income of $50,000 for families of this size. For each group, we show: the net cost and the percent of costs attributable to the group; subsidy take–up in absolute and percentage (relative to group size before the policy impact) terms; the change in the uninsured in absolute and percentage terms; and the cost per newly insured person (total dollars spent on that group relative to the reduction in the uninsured).

Our key findings are:

- The total cost of this policy is $13.3 billion dollars per year (in 1999 dollars).
- Almost 18.4 million persons take–up the subsidy, which is 8.2 percent of the total non–elderly population.
- Of those taking up, 4.7 million were previously uninsured (11 percent of the uninsured), 8.6 million were previously covered by non–group insurance (57 percent of those covered by non–group insurance), 4.7 million were previously covered by employer–provided insurance (3.2 percent of those covered by employer–

TABLE 1
REFUNDABLE $1,000/$2,000 CREDIT FOR NON–GROUP INSURANCE, ALL ELIGIBLE

	Number of Persons (Millions)	Percent of Insurance Category	Net Cost ($1999 Millions)
Total Cost in 1999 dollars	—	—	13,285
Total Take–up of Subsidy	18.37	8.2	—
Previously non–group	8.60	57.2	7,006
Previously uninsured	4.72	11.1	4,655
Previously employer–insured	4.68	3.2	1,824
Previously Medicaid	0.36	1.8	–200
			—
Total Change in Population Size			
Non–group	9.77	65.0	—
Uninsured	–4.03	–9.5	—
Employer–Insured	–5.37	–3.7	—
Firm dropped to non–group	–1.05	–0.7	—
Firm dropped to uninsured	–0.12	–0.1	—
Switch to non–group	–3.64	–2.5	—
Uninsured due to decreased contributions	–0.57	–0.4	—
Medicaid	–0.36	–1.8	—
Cost per Newly Insured ($1999)	—	—	$3,296

provided insurance), and 0.4 million were previously covered by Medicaid (1.8 percent of those covered by Medicaid).

- On net, the number of uninsured falls by slightly more than 4 million, which is 9.5 percent of the uninsured population.
- On net, the number of persons with non–group insurance rises by 9.8 million, which amounts to a rise of two–thirds in the size of this group.
- On net, the number of persons with employer–provided insurance falls by 5.4 million, which is 3.6 percent of the size of this group. This change is comprised of:
    - 1.1 million persons whose firms stop offering group insurance, so that they move to the non–group market;
    - 0.1 million persons whose firms stop offering and they become uninsured;
    - 3.6 million persons who switch from group to non–group insurance;
    - and 0.6 million persons who become uninsured because their firms are raising the employee share of insurance premiums and they decide to drop coverage.

While this policy lowers the number of uninsured, it also induces a substantial shift from group to non–group coverage. Moreover, almost one–half of those taking up the subsidy are persons who are currently already purchasing non–group insurance. As a result, the *net cost of the policy per newly insured person is almost $3,300*, which is substantial. By comparison, on average in our sample, employer–provided insurance costs $1,860 per person covered, and non–group insurance costs $2,100. That is, due to imperfect targeting, the government is paying 50 percent more than the cost of the typical non–group policy per person newly insured.

It is interesting to note that most of the government cost of imperfect targeting of this subsidy arises primarily through take–up by the existing non–group insured, 57 percent of whom take advantage of this new subsidy, not through dropping and switching among the existing employer–insured. This is because, while those on employer insurance who drop or switch cost the government money through their take–up of the subsidy, they also save the government revenues by dropping their currently tax subsidized employer coverage. For example, for those workers whose firms drop their health insurance coverage, we assume that their wages will rise to reflect the fact that their employer is no long paying for health insurance, and can therefore afford higher wages. These higher wages will then be taxed, raising new revenues, and offsetting the cost of their take–up of the new insurance subsidy. For those who switch from group to non–group insurance, we assume that the cost savings to the employer is passed back to workers on average in the form of higher wages (although not specifically to the switching employees), once again raising revenues. And revenues also rise since employers react to this policy, to some extent, by lowering their pre–tax contributions for health insurance, and once again raise wages to compensate for this.

Distributional Analysis

Given the strong correlation between insurance status and income, it is important to consider not just the aggregate impacts of this subsidy, but its distributional implications as well. This is done in Table 2, using the income groupings relative to the poverty line described above. There are several findings of note from this distributional analysis:

- The lowest income group, which contains 45 percent of the uninsured, receives about 26 percent of the net

TABLE 2
DISTRIBUTIONAL ANALYSIS

Group	Net Cost ($1999 Millions)	Percent of Costs	Subsidy Take-up (Millions)	Percent of Group	Change in Uninsured (Millions)	Percent of Unins.	Cost per Newly Insured ($1999)
<100% of FPL	$3.489	26.2	4.39	8.6	–1.27	–6.6	$2,739
100–200% of FPL	$4,012	30.2	5.31	11.6	–1.64	–13.1	$2,447
200–300% of FPL	$2,478	18.7	3.50	9.2	–0.71	–13.1	$3,506
300–400% of FPL	$1,466	11.0	2.20	7.7	–0.24	–11.3	$6,040
>400% of FPL	$1,840	13.9	2.97	4.8	–0.17	–5.3	$10,956

spending on this policy. Only about 1.3 million of the uninsured in this group gain coverage (6.6 percent of the uninsured below the poverty line); this is about one-third of the total number of uninsured who gain coverage across all income groups. Overall, this policy is more efficient for this subgroup than for the full population, with a cost of $2,740 per newly insured. This is primarily because there are few non-group insured taking up the policy in this income range, relative to the number of uninsured taking it up.

- Those between 100 and 200 percent of poverty, a group that contains another 30 percent of the uninsured, receive about 30 percent of the net spending from this policy, and there is a decline in the uninsured of about 1.6 million. Spending is even more efficient in this group than for the lowest income group, with a cost of $2,500 per person newly insured, since there is an even higher ratio of uninsured to non-group insured in this income range.

- Those between 200 and 300 percent of poverty receive almost 20 percent of the net spending from the policy, but there is a decline in the uninsured of only 0.7 million. As a result, spending is less efficient for this group, with a cost per newly insured of over $3,500.
- Those above 300 percent of poverty receive 24 percent of the net spending of this policy, but there is only a very small change in the number of uninsured, in large part because there are so few uninsured in this income group. As a result, spending is much less efficient at these higher income levels. For those between 300 and 400 percent of poverty, there is a cost of over $6,000 per newly insured. For those above 400 percent, there is a cost of almost $11,000 per newly insured.

Thus, a majority of spending under this policy (56 percent) is targeted to those below 200 percent of the poverty line, and three-quarters is targeted to those below 300 percent of the poverty line. But the spending that is done on those above 300 percent of the poverty line is very inefficient, with a total of $3.3 billion spent on this group for a reduction in the number of uninsured of only 400,000. Overall, while this policy has a high cost per person newly covered, it is providing a large tax break that is mostly targeted to those

below 300 percent of the poverty line; these distributional gains should be weighed against any inefficiencies of this policy relative to alternatives.

Alternative Policies

While the base policy mimics a number of proposed tax subsidies, there are at the same time a host of alternative structures that have been proposed. While we cannot do justice in this limited space to the full variety of alternatives available to policymakers, we consider several alternative approaches to provide a flavor of how the effects of tax policy change as the structure of the program is altered. We present the key findings for each of these alternatives in Table 3, including: the take-up of the subsidy; the cost; the change in the uninsured, non-group insured, and employer-insured; the cost per newly insured; and the percentage of benefits that flow to those with incomes below 200 percent of the poverty line.

Making the Credit Non-Refundable

One option that will lower costs substantially, and simplify administration, is to make the subsidy non-refundable. On the other hand, this will severely limit the benefits of this subsidy for the uninsured, more than 60 percent of whom have tax liabilities less than $1,000 (meaning they can only partially benefit from a non-refundable credit).

The impacts of a non-refundable $1,000/$2,000 credit are presented in the second row of Table 3. This does indeed lower the costs of the subsidy, which fall to almost half the cost of the non-refundable credit ($7 billion). But the impact on the size of the uninsured population falls even more, with fewer than two million uninsured gaining coverage (only 4.3 percent of the uninsured). As a result, the cost

TABLE 3
ALTERNATIVE POLICIES

	Total Take-up (M)	Total Cost ($ M)	Change in Uninsured (M)	Change in Non-Group Insured (M)	Change in Employer Insured (M)	Cost per Newly Insured Person	% of Benefits for <200% FPL
Base Policy	15.45	10,390	-4.00	9.80	5.41	$2600	52.1
Non-Refundable Credit	9.58	5,957	-1.89	5.65	-3.70	$3157	21.5
Deduction	5.12	450	0.03	1.33	-1.34	—	26.5
Limit to Non-Offered	9.17	5,502	-2.09	6.40	-4.12	$2635	66.5
$500/$1,000 Credit	8.85	2,263	-1.68	4.06	-2.15	$1349	58.9
$2,000/$4,000 Credit	29.00	33,391	-7.67	22.25	-13.99	$4353	45.8
50% of Costs Subsidized	11.79	5,227	-2.72	6.37	-3.46	1,924	57.1
Phase-out from $30,000 to $50,000	12.42	8,295	-3.66	7.76	-3.73	2,269	65.2
No liquidity constraints	17.03	11,622	-5.41	11.38	-5.41	2,147	56.0
$2,000/$4,000 cap and no liquidity constraints	33.87	39,211	-12.04	27.13	-13.98	3,257	52.1

per newly insured person is even higher than with the refundable credit ($3,800), largely a function of the fact that such a high share of the dollars are going to the previously non–group or employer–insured. Moreover, the distributional consequences of this approach are much less attractive. Only 23 percent of the spending through this policy goes to those below the twice poverty line, and only 2 percent goes to those below the poverty line.

There are a number of political and administrative arguments against refundability, most significantly the question of whether net tax refunds to low income families are hidden forms of "welfare" payments. But the results here speak clearly: refundability is critical for appropriate targeting of tax incentives to the low income uninsured.

Using a Deduction

Another alternative that can limit costs further is to use a deduction rather than a credit, but this approach has problems similar to non–refundability in reaching the uninsured. Moreover, of the half of the uninsured that do pay taxes, 90 percent are in the 15 percent tax bracket, so that a subsidy in the form of a deduction is worth relatively little to them.

The results of an unlimited deduction of non–group health insurance costs are presented in the third row of Table 3; we assume that this is an "above the line" deduction that is available to all taxpayers and not just those that itemize. The costs of this policy are dramatically lower than for the alternatives (only $870 million per year). But the impact on insurance coverage is also much more modest, with only 250,000 uninsured gaining coverage. This is because there is only modest overall take–up of this subsidy by the uninsured to begin with (600,000 persons), and much of this is then offset by firm dropping and reduced coverage due to firm contribution reductions. Estimating with precision the change in the number of uninsured in the range around zero is difficult, but it is clear that effects of deductibility will likely be minimal on both costs and coverage. At the same time, this policy has much worse distributional characteristics, with less than 30 percent of the benefits flowing to those below 200 percent of the poverty line.

Note that the cost that does arise from this policy is not due to take–up by the previously employer–insured; the government actually makes money on this population, with the government revenue from higher wages due to firm dropping and contribution reductions outweighing the government cost of subsidy take–up. Rather, the inefficiency arises primarily from the fact that *three–quarters of those taking up this subsidy were already non–group insured.*

Limiting the Credit to Those Not Offered Employer Coverage

One alternative to try to better target the subsidy is to limit the credit to those not eligible for employer insurance coverage. There are of course difficult administrative issues associated with implementing and enforcing such a policy, as discussed in Meyer, Silow–Carroll, and Wicks (1999). But the advantage is that being offered insurance by one's employer is tightly related to being covered by insurance, so this policy provides a device for better targeting subsidy dollars to the currently uninsured.

We consider the impact of a refundable $1,000/$2,000 credit that is limited to those not offered employer insurance in the next row of Table 3. The total cost of this option is much lower than the base policy, at only $5.5 billion per year, although the number of persons newly insured falls as well (to 2.1 million). The efficiency of this alternative is almost identical to the base case, at $2,535 per newly insured person. This reflects the cancellation of two effects, relative to the base case. On the one hand, there are savings from much lower take–up of this policy by the existing non–group insured, since many

of them are offered employer–provided insurance. On the other hand, there is a much larger increase in the uninsured pool from firm dropping (there is no switching here, since offered individuals can't take the subsidy); we estimate that 3.2 million persons are dropped by their firms, and 700,000 of them remain uninsured.

This policy is somewhat more distributionally attractive than the base policy, with over two–thirds of the benefits flowing to those below 200 percent of the poverty line. These modest distributional gains, however, must be balanced against the costs and difficulty of enforcing this administratively awkward restriction (which we have not accounted for in the estimates).

Changing the Scale of the Subsidy

While we have chosen a credit of $1,000 for singles and $2,000 for marrieds as our base case policy, one could consider less or more generous alternatives as well. In the next two rows of Table 3, we consider first halving, then doubling, the generosity of this policy. We find that smaller credits cover fewer people, but do so in a more targeted way. At a credit of $500 for singles and $1,000 for marrieds, we estimate costs that are only one–fifth of the base case, but the reduction in the uninsured is almost one–half as large. As a result, the spending per newly insured person is only $1,350, which is substantially below even average group costs per person. On the other hand, at a credit of $2,000 for singles and $4,000 for families, which would approximate the full cost of insurance for these populations, we estimate that costs rise three–fold, but the number of newly insured less than doubles, so that the spending per newly insured rises to $4,353 per person. At the same time, the small credit covers fewer than 2 million new persons, while the larger credit covers over 7.7 million.

The smaller subsidy also targets its spending more directly at the bottom of the income distribution, with over 60 percent of the dollars flowing to those below 200 percent of poverty. On the other hand, the $2,000/$4,000 credit spends less than half its dollars on those below 200 percent of poverty. This worsening of distributional impacts as generosity rises reflects the dramatic increase in take–up by both the (relatively high income) non–group insured and employer–insured. On the other hand, while only 3 percent of the uninsured below poverty and 5 percent between 100 percent and 200 percent of poverty gain coverage with the smaller credit, the larger credits results in 11 percent of the uninsured below poverty and 25 percent of those between 100 percent and 200 percent of poverty gaining coverage.

Thus, there is a clear tradeoff as the generosity of the tax credit is changed. Modest credits cannot deliver a very large change in the uninsured population, but the newly insured that are covered tend to be the lowest income and are low cost. Very large credits can induce substantial changes in the uninsured population, but only at a very steep cost per newly insured.

Easing Liquidity Constraints

A key issue in implementing tax credits is the mismatch between the flow of tax subsidies and the flow of insurance premium payments. Low income households who would like to take advantage of tax credits during a given year, but who only receive their credit the next spring, may face liquidity problems. If the government can find a solution to this timing mismatch, it can increase the propensity of the uninsured to take–up tax subsidies. A variety of analysts have proposed solutions to this problem, such as paying tax credits directly to insurers (Ethridge, 1999). But our track record with the Earned Income Tax Credit (EITC) suggests caution in assuming that this problem is

easily overcome: while individuals can claim their EITC throughout the year, and presumably for many individuals it would be of some value to do so, over 99 percent of claimants receive the credit as a lump sum the next spring (Leibman, 1998).

While we have assumed that liquidity constraints reduce take–up in our base case calculations, it is important to assess the impact of easing them by assuming that the government solves the liquidity problem. As shown in Table 3, easing liquidity constraints increases by $1.3 billion the cost of the base policy (absent any additional interest or other costs to the government of easing these constraints), and results in an additional 1.4 million newly insured persons, for a total of 5.4 million newly insured. This implies a substantial increase in the efficiency of the policy, with a cost of only $2,150 per newly insured. Moreover, for larger tax credits, the impacts of easing liquidity constraints are also heightened (as shown in the following row of Table 3). For a $2,000/$4,000 credit, the costs increase by $6 billion per year, but the number of newly insured rises to over 12 million.

CONCLUSIONS

Federal policymakers continue to look to tax policy as a politically attractive vehicle for addressing the problems of the uninsured in the U.S. As a result, it is critical to carefully assess the cost, insurance coverage, and equity implications of alternative approaches to tax subsidization. While point estimates of the effects of any major change in health financing cannot be estimated with perfect precision, simulation analyses using common assumptions are particularly useful for comparing and contrasting the effects of alternative proposals.

We have compared alternative tax policy designs using a consistent set of measures, including: the overall cost to the federal government, the number of uninsured who gain coverage, the federal cost per person newly insured (which is a measure of how efficiently federal dollars are being used), and the proportion of benefits that flow to those below 200 percent of the poverty level (which is a measure of the degree to which the policy targets those with low incomes).

There are several clear conclusions from this analysis. First, it is difficult to design a tax policy which delivers a modest cost per newly uninsured person, while insuring a large number of new persons. The base policy considered here—a refundable credit of $1,000 for singles and $2,000 for families—is more generous than many of the proposals being considered by federal policymakers, and yet still subsidizes less than half of the estimated cost of a non–group insurance for a typical person. While it would decrease the ranks of the uninsured by an estimated 4 million persons (less than 10 percent of the uninsured population), the average cost per person newly insured is $3,300. Raising the value of the credit or allowing tax subsidies to be used towards the purchase of employer sponsored coverage would insure more people, but also raise the cost per person newly insured significantly. Lowering the value of the credit would be more efficient—meaning that the cost per person insured would be lower—but the result would be an even smaller dent in the number of Americans uninsured.

Second, there are clearly more and less efficient ways to cover a given number of uninsured. We find in particular that non–refundable credits are much more expensive per uninsured person covered, while covering fewer of the uninsured. We find as well that income limits on eligibility that more tightly target the policy towards the lower part of the income distribution in which the uninsured are concentrated can significantly lower costs with essentially no impact on the number of persons newly insured. Finally, we find that poli-

cies that can more tightly match the timing of tax subsidies with the timing of insurance payments can improve both the scope and efficiency of tax policy, especially for low-income people.

Third, different approaches to tax subsidies vary not only in the efficiency with which they reach the uninsured, but also in how effective they are at targeting resources to those with low incomes. For example, a policy that targeted refundable credits of $1,000 for singles and $2,000 for families towards people with incomes of less than $50,000 would provide 69 percent of its benefits to those below 200 percent of the poverty level. In contrast, a policy that allowed people to deduct non-group insurance premiums would provide just 30 percent of its benefits to people below twice the poverty line, and a credit that was not refundable would target an even smaller portion of aid to the poor or near poor.

Finally, tax-based subsidies—particularly those whose subsidies are most expansive—would likely lead to reductions in the number of people with employer-based coverage. For example, we estimate that the base case—a refundable credit of $1,000 for singles and $2,000 for families—would reduce the number of people with employer coverage by 5.4 million. Most of these people (3.7 million) would switch from employer coverage to non-group insurance because they would find the new tax subsidies more attractive than their current situations. However, the remainder would either be dropped by their firms—and then either purchase non-group insurance or go without coverage—or become uninsured because their employers increased the amount they must pay for insurance. Policies that mitigate firm dropping of coverage or switching to non-group insurance by employees (e.g., by allowing the credit to be used towards the purchase of employer coverage) tend to cost more in total and also per person newly covered.

This paper does not discuss in detail a number of important additional issues to be considered with tax subsidies. Four in particular stand out. The first, mentioned briefly earlier and discussed in more detail in Meyer, Silow-Carroll, and Wicks (1999), is administrative complexity. This is particularly relevant given the importance of surmounting liquidity constraints in improving the efficiency of tax policy. Another difficult administrative issue is geographic adjustment of credit levels. There is substantial regional variation in the cost of insurance, and even very large variation within states; for example, in the data we use on employer premium averages across even large regions of the U.S., there is a 50 percent difference between the lowest and highest premium regions. We have assumed that this is not reflected in credits that are provided, but policymakers may choose to target the credit to local cost levels. This is an advantage of an uncapped deduction, which more naturally respects local cost variation (although only in a limited way due to the small resulting subsidy rates).

The second is the erosion of the base of employer-provided health insurance. The past decade has seen a steep decline in employer-provided insurance coverage, and tax subsidies to non-group coverage would only exacerbate this decline. If there are pooling advantages to having individuals obtain their insurance through the workplace, then this is a potential concern with policies targeted only to non-group coverage. Moreover, those leaving the employer pool will be the healthiest employees, leading to a rise in costs per covered person among those remaining in the pool. On the other hand, however, doubling the size of the non-group market (as we estimate would occur in the base policy) could substantially improve the functioning of this market, both in terms of administrative efficiency and reduced adverse selection. And de-linking insurance from the workplace

could improve the functioning of the labor market by reducing insurance–induced immobility across jobs, or "job lock" (Gruber, 1999).

Third, this paper has focused almost exclusively on gains in insurance coverage, but it has not differentiated the *kinds* of insurance that individuals are holding. If tax policy leads, either through switching or employer dropping, to fewer individuals covered by very generous employer policies and more covered by substantially less generous non–group policies, then there are implications for quality of health care that become potentially relevant. Whether the differences in quality of plan are actually relevant for health is in fact disputable. But the impact on quality of coverage remains an important concern for tax policy.

Fourth, we have not considered in any detail either pricing reactions in the non–group market nor state and/or federal regulatory reactions to this type of policy. It is possible, as noted earlier, that the substantial increase in enrollment in the non–group market could lead to reductions in prices. Moreover, non–group insurance plans might design policies targeted specifically to the available level of the credit, further increasing take–up from what is modeled here (although this takeup might be in plans with significantly less generous benefits than are typical today).

At the same time, this analysis assumes that policies in the individual market are universally available (at health risk adjusted prices). While such "guaranteed issue" in the individual market is required in some states, most states allow insurers to exclude people who are in poor health, which could reduce take–up. It is possible that state and/or federal regulators could accompany tax subsidies with individual market regulations to limit such practices, but these regulations are controversial. And, the net impact of insurance market reforms in the context of tax subsidies is uncertain, as it would raise costs for the most healthy and lower them for the least healthy. It is also possible that regulators could accompany subsidies with requirements on product quality in the non–group market, making it harder for individuals to buy the lower quality product (relative to employer–provided insurance) that is available in the non–group market, mitigating the impact of tax policy on insurance take–up.

In summary, tax policy does hold some promise as a means of providing health insurance to some of the uninsured. But providing coverage to substantial numbers will require very large expenditures, both overall and per person newly covered. Even the most effective policy considered here, a $2,000/$4,000 credit that is accompanied by a solution to liquidity problems, costs almost $40 billion per year and covers only 30 percent of the uninsured. Thus, tax policy can likely be most useful as one part of an overall strategy to address uninsurance in the U.S., as opposed to a solution in and of itself.

Acknowledgments

I am grateful to Barrett Kirwan, Robin McKnight, and Melissa Schettini for outstanding research assistance, and to Linda Bilheimer, Gary Claxton, Judy Feder, Sherry Glied, Julie James, Jim Mays, Mark Merlis, Mark Pauly, Diane Rowland and particularly Larry Levitt for helpful comments, and to the Kaiser Family Foundation for research support.

REFERENCES

Employee Benefits Research Institute.
"Sources of Health Insurance and Characteristics of the Uninsured: Analysis of the March 1998 Current Population Survey." Washington, D.C.: EBRI, 1999.

Etheredge, Lynn.
"Tax Credits for Uninsured Workers." Health Insurance Reform Project of the George Washington University. Mimeo, 1999.

Gruber, Jonathan, and James Poterba.
"Fundamental Tax Reform and Employer–Provided Health Insurance." In *Economic Effects of Fundamental Tax Reform*, edited by Henry J. Aaron and William G. Gale, 125–70. Washington, D.C.: Brookings Institution, 1996.

The Kaiser Project on Incremental Health Reform.
"The Difference Different Approaches Make: Comparing Proposals to Expand Health Insurance." Available at www.kff.org, prepared by Judith Feder, Cori Uccello, Ellen O'Brien, 1999.

Leibman, Jeffrey B.
"The Impact of the Earned Income Tax Credit on Incentives and Income Distribution." In *Tax Policy and the Economy, 12*, edited by James Poterba, 83–120. 1998.

Manning, Willard, et al.
"Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment." *American Economic Review* 77 (1987): 251–77.

Meyer, Jack, Sharon Silow–Caroll, and Elliot Wicks.
"Tax Reform to Expand Health Coverage: Administrative Issues and Challenges." Report prepared for the Kaiser Family Foundation, 1999.

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

Sheils, John, and Paul Hogan.
"Cost of Tax–Exempt Health Benefits in 1998." *Health Affairs* 18 (1999): 176–81.