Expanding Coverage Via Tax Credits: Trade-Offs And Outcomes

It may not be possible to devise one tax credit to satisfy all. But why not begin with a limited group, and expand as we learn?

By Mark Pauly and Bradley Herring

ABSTRACT: In this paper we discuss various options for using refundable tax credits to reduce the number of uninsured persons. The effect of tax credits on the number of uninsured depends on the form of the credit scheme adopted. Moreover, since large subsidies for private insurance directed to low-income persons have never been implemented, there is considerable uncertainty about the effect of various tax credit proposals. We find that small credits will do little to reduce the number of uninsured but that credits covering about half of the premium for a benchmark policy might have a significant effect, especially if they take a fixed-dollar form and can be used for policies with few restrictions. Finally, we discuss the normative issues surrounding the "costs" of these credits schemes, and the policy issues raised by the uncertainty of the effects.

THE 16 PERCENT OF THE U.S. population that lacks health insurance defies simple characterization. It covers the spectrum of family structures, ages, work involvement, ethnicities, and tastes. However, there are two causes of the lack of insurance that all analysts accept: Many people's incomes are too low to allow them to afford insurance, and the premium they would have to pay is too high to make insurance purchasing attractive.

The "refundable tax credits" (hereafter, just "credits") approach to increasing the extent of insurance coverage targets the uninsured in these groups. (Using refundable credits in some fashion to help people afford insurance is a recurring proposal on the part of presidential candidates and members of Congress.) Credits do something exceedingly simple: They reduce the net premium a person would have to pay for at least some insurance policy. Refundable credits are used to offset other taxes or are paid to the household if there is no tax liability; they differ from tax deductions because they are not

Mark Pauly is the Bendheim Professor and chair of the Department of Health Care Systems at the Wharton School, University of Pennsylvania. Bradley Herring is a Robert Wood Johnson Foundation Health Policy Scholar at Yale University's Institution for Social and Policy Studies.

used to change taxable income and do not depend on the household's income tax rate. Thus, credits simultaneously increase the affordability and the value of coverage (compared to its absence). This key observation points to both the strengths and the weaknesses of the credits approach.

One of the strengths is simplicity itself. While their actual administration can be complex, tax credits are a simple, straightforward tool. They do not require changing the health care regulatory structure, negotiating with providers, reorganizing the delivery system, or altering the philosophy of medical treatment. However, for those who think that the problems of the uninsured go far beyond their need for insurance, or who think that the problems of the U.S. health care system go far beyond the problem of the uninsured, this sharp but narrow focus will be viewed as a weakness. Credits will not necessarily reduce waste, lower administrative costs, improve doctor/patient relationships, or cure dread diseases. Credits are intended to increase access to the current system, not to redesign the system. Still, the influx of millions of formerly uninsured persons now armed with health insurance probably would change the system—but that would be an indirect (although probably desired) effect, not a goal.

For some, further restructuring of the insurance markets in which credits can be used is also an objective. In both the individual and small-group markets, it would be desirable to lower administrative costs and protect insured persons against risk-related jumps in premiums, if ways can be found to achieve these goals that do more good than harm.[1] In this paper, however, we do not focus directly on such restructuring, but we highlight links between credits programs and reforms and discuss some aspects of reform.

What Do We Want; What Do We Want To Know?

Credit plans can take on many forms for a given population and can be targeted at a wide variety of subpopulations. Within the overall objective of reducing the number of uninsured persons, a range of impacts are possible. Especially if resources are limited, there are trade-offs and valuations that must be considered. For example, are "we" indifferent between strategies that reduce the head count of the uninsured, regardless of which type of uninsured population gains coverage? Are some uninsured persons more of concern than others: the poor versus the nonpoor, high risks versus low risks, people who reject opportunities to obtain coverage cheaply versus those with few good chances, rational consumers who plan ahead versus those who assume that they are immortal?

Two equally important effects of credit programs should be of

interest to policymakers: (1) the change in the number and mix of the uninsured (compared with no program); and (2) the effect on "disposable income" of those who use the credit (compared with no credit). The first measure describes potential impacts on the health care system; the second, potential impacts on the fiscal system and (eventually) on overall macroeconomic activity. The first impact occurs because the credit provides a targeted subsidy; the second, because the credit provides a targeted tax cut.

Because credit programs can have so many design variants, attempts to estimate the two effects necessarily cover a wide range. The proper answer to the question of the impact of credits is that it all depends on the size and form of the credit. In addition, because many plans direct subsidies to populations that have never before received the kinds or magnitudes of subsidies for purchasing the qualified private insurance plan of their choice, attempts to estimate what credits would do must necessarily cover a wide range. However, it is the range of likely values, not a point estimate of a particular researcher's best guess, that ought to be relevant for policy.

■ **Policy design in the face of uncertainty.** Virtually any program to reduce the number of uninsured persons or relieve the financial burden of paying for medical insurance will have uncertain impacts as long as people participate in them voluntarily and resources are not so lavish that any alternative to participating in the plan is economically irrational. Our view of proper policy evaluation is that this uncertainty should be front and center, not ignored or disguised to avoid "confusing" policymakers. The reason is that what seems the best policy if only we knew everything with certainty will be an inferior policy if our guesses turn out wrong; another policy may not be the unique best in any given setting but may offer a good outcome over a wide range of eventualities.

This kind of flexibility and adaptability are inherent benefits of tax credit schemes. They can be designed to fit a wide variety of equity definitions. Moreover, they do not require enormous investment in fixed administrative structures, which are difficult to change later in the face of experience, and their financial design parameters can be modified easily and quickly.

Analyzing The Design And Impacts Of Credits: A New Approach

The great bulk of analyses of tax credit approaches to health insurance have taken one of two forms: either a specific credit plan is proposed, costed out, and defended (or perhaps attacked); or general observations on the desirability of credits are provided with little detail.[2] A few quantitative analyses have made comparisons

"Current counts of the number of uninsured in working families reflect people's responses to the incentives to buy insurance."

within an eclectic set of proposals, usually chosen to span some spectrum of whatever is politically germane at the moment.[3] In this paper we adopt a different method, first specifying a set of "universal" design parameters for credit plans and then illustrating the trade-offs among those design parameters by estimating the effects of a variety of plans with specified parameters. Our primary goal is to illustrate these trade-offs, and we comment briefly on the confidence in estimates of impacts, including our own.

■ **Design parameters.** Besides eligibility, any credit scheme has three important design parameters: (1) the dollar or proportional subsidy the credit provides; (2) whether the credit is a set dollar amount or a set proportion of the premium; and (3) the premium for the lowest-cost insurance plan eligible for the credit. For instance, a credit might pay $1,000 toward a set of plans, the least costly of which carries a premium of $1,200 but whose premium could be as high as $2,500; the potentially insured person would then face a net premium ranging from $200 to $1,500. Or the credit might pay 50 percent of the premium; the person would then have to pay $600 for the least costly plan but $1,250 for the most costly.

To illustrate the differences among plan design features, we have estimated the reduction in the numbers of uninsured in working families and their distribution across people of different types for plans taking on a range of values for each of these three parameters.[4] Much of our discussion, though, highlights results that are subject to approximately the same "cost" of $1,000 in terms of tax reductions—either a flat-dollar credit equal to that amount or a proportional credit covering about 50 percent of average premiums.[5]

Using data from the 1996 Medical Expenditure Panel Survey (MEPS), we employ two different methods to estimate these effects. We do so because current knowledge is sufficiently imprecise that it is desirable to look for overlap among different approaches to estimation rather than basing conclusions on a single calculation or even a single method applied to a variety of data sets.

Problems of imprecise measurement. The reason current estimates are imprecise is instructive and relevant to our own choice of methods. Current counts of the number of uninsured in working families presumably reflect people's responses to the current set of incentives to buy insurance. (According to the MEPS data, more than 80 percent of the uninsured have an employed family member.) The

primary incentive is the current tax treatment of employer coverage. This tax law provision allows those who receive part of their total compensation in the form of a group insurance premium to exclude the amount the employer paid from taxable income for all federal and state taxes. The implicit subsidy is therefore equal to the worker's marginal tax rate, and could be as high as 50 percent for high-wage workers facing high state income taxes.

We do know the proportion of high-wage workers who obtain insurance when they are subsidized at such a high rate—virtually all of them. We also know that many—but by no means all—low-wage workers obtain insurance when the subsidy rate is low; the current subsidy rate would range from 8 percent to 23 percent for the low-wage workers with either zero or 15 percent marginal federal income tax rates. However, most tax credit proposals involve high subsidies—up to 50 percent or more—directed at low-income persons. So while we know what happens with high-wage workers and high subsidies, and with low-wage workers and low subsidies, we have no observations on actual behavior of low-wage workers with high subsidies, with the subsidy applying primarily to individual insurance.

We take specific account of the existence of current subsidies in both how we estimate behavior and how we calculate the effect of behavior on the proportion uninsured. Specifically, since many uninsured workers could have obtained tax-subsidized job-related insurance (especially in today's tight labor market) but chose not to do so, any new subsidy will have to be at least as large as the rejected subsidy. Small subsidies, on the order of 25 percent of premiums, will therefore generally only affect those who are self-employed—a minority of today's uninsured workers. To generate more relevant estimates, however, we can use the observations on the insurance purchase rates of low-wage employees who obtain both a tax subsidy and low administrative loading provided by large-group insurance. This comes closest to what might be achieved with moderately large subsidies for the purchase of individual insurance.

We therefore estimate an econometric model to predict the likelihood that workers and their dependents obtain insurance, where the "price" of insurance incorporates the administrative loading appropriate to their industry and the tax subsidy (applied to the average employer-paid portion of the premium) appropriate to their income level. This is intended to be a model of long-run behavior, so we assume, as empirical evidence and theory strongly suggest, that the incidence of both credits and employer premium payments is on workers' wages.[6] The loading and the subsidy are initially assumed to be proportional to the premiums in the group and individual

markets.

As might be expected, the administrative loading for individual insurance as well as the size of the proportional credit will determine the proportion of the population induced by the credit to obtain insurance. In the results presented below, we assume that such loading equals 30 percent of premiums. This appears to be what more efficient nongroup insurers can achieve.[7]

■ **Impact of credits for a benchmark plan.** Exhibit 1 shows the change in the proportion of uninsured for different proportional tax credits applied to a comprehensive benchmark plan. While the estimates we present assume that persons at all income levels are eligible for credits, results also are shown for persons either below or above 300 percent of the poverty level ($48,108 for a family of four in 1996). A small (25 percent credit) is claimed largely by the self-employed and causes only a small reduction in the number of uninsured persons. A moderate (50 percent) credit, in contrast, cuts the proportion uninsured in half, and a large (75 percent) credit converts more than four-fifths of the uninsured to insured. The empirical fact driving these results is that, even among low-wage persons,

EXHIBIT 1
Impact Of Proportional Tax Credits: Extrapolating Group Insurance Purchases

Income level	Percent insured	Percent reduction in uninsured
All income levels[a]		
Currently	80.0%	-
25% credit	82.6	13.0%
33% credit	84.0	20.2
50% credit	90.3	51.7
66% credit	95.1	75.5
75% credit	96.6	83.2
Below 300% of poverty		
Currently	69.3	-
25% credit	72.7	11.1
33% credit	75.5	20.0
50% credit	85.2	51.8
66% credit	92.1	74.2
75% credit	94.4	81.9
Above 300% of poverty		
Currently	89.7	-
25% credit	91.5	18.0
33% credit	91.8	20.8
50% credit	95.0	51.5
66% credit	97.8	79.1
75% credit	98.6	86.9

SOURCE: Authors' calculations using 1996 Medical Expenditure Panel Survey data.
NOTES: The federal poverty level was $16,034 for a family of four in 1996. A proportional tax credit of *x* percent reimburses an amount equal to *x* percent of the insurance premium to the insurance purchaser through decreasing one's tax liability—perhaps via a refund. The reduction in the uninsured results from a proportional tax credit equal to the specified amount applicable to individual insurance with 30 percent administrative loading. Individual-market premiums are assumed to be risk-rated by whatever mechanism exists in the employment-based market.
[a] Sample includes full-time workers and their dependents only.

only a tiny proportion of persons working in industries or occupations characterized by very large firms remains uninsured. A key issue is whether the behavior of low-wage employees in that setting can be extended to other workers with similar net premiums but in different types of insurance markets.

"Reservation price." Because of concern that extrapolating from the group insurance setting, although reasonable, may be overly optimistic, we implemented a second technique. This other simulation approach is based on expected-utility maximization for a population facing a given distribution of medical expenses. The idea is that we can observe both the mean and the variance of the total out-of-pocket costs faced by a certain proportion of the population; here, too, we use the large-sample 1996 MEPS data. Given an assumption about the level of risk aversion, we can estimate the maximum premium or "reservation price" such people would be willing to pay for a given insurance policy. If the reservation price exceeds the premium, net of the credit they will face, we then assume that they will choose to become insured. Presumably even a person who is not at all risk-averse will prefer insurance if the net premium is less than the expected value of out-of-pocket expenses.

The strength of risk aversion turns out to be a relatively minor determinant of what proportion of people buy insurance. What is much more important is the expected amount of out-of-pocket expense; this value is reduced considerably by access to free care by the uninsured. Indeed, the MEPS data indicate that the uninsured typically pay only about 30 percent of the actual cost of their care.[8]

Specifically, we used the MEPS data to estimate the expected cost of out-of-pocket expenses for adults, adding to that an estimate of risk aversion and the value of the additional care that insurance will induce them to consume. This then is the price they would be prepared to pay for insurance. The expected cost of out-of-pocket expense is assumed to include two components: the actual out-of-pocket payments the person makes, and a "disutility" cost associated with any free or bad-debt care (because of lower quality, less personal attention, or longer waits).

A lower bound for this reservation price results from assuming that currently uninsured persons will continue to use relatively low amounts of medical care if insured (even after adjusting for moral hazard); an upper bound results from assuming that the currently uninsured will use amounts equal to those used by currently insured persons with similar observable characteristics (such as age, sex, and health status). For our "central tendency" estimates we assume that the quantity is the average of these two levels, but we also show the range about this average. People are assumed to be willing to buy

insurance if their reservation price exceeded the premium for individual insurance, with 30 percent loading and adjustments for age and sex, as reduced by a tax credit.

In Exhibit 2 we assume that the great majority of persons who receive care whose cost is covered by charity or bad-debt forgiveness cannot pay the full cost for the care they need, but that they do attach disutility to this outcome. An important determinant of the effectiveness of credits is how large this disutility is. If we assume that one dollar of the bill for bad-debt or charity care carries disutility equivalent to thirty cents of out-of-pocket payment, the estimated reduction in the percentage uninsured for a 50 percent proportional credit is estimated to be 58 percent (range, 30–85 percent), as shown on line 3 of Exhibit 2. If instead we assume that this disutility equals fifteen cents on the dollar, the central estimate

EXHIBIT 2
Impact Of Tax Credits: Estimating Utility-Derived Reservation Prices

	Percent reduction in uninsured	
	Estimate	**Range[a]**
Proportional credits		
Higher "cost" of free care		
25% credit	38.9%	(12.0–65.9)
33% credit	44.7	(17.9–72.5)
50% credit	57.5	(30.3–84.7)
66% credit	73.1	(51.9–94.3)
75% credit	82.6	(68.4–96.9)
Lower "cost" of free care		
25% credit	13.3	(01.8–24.8)
33% credit	17.1	(02.6–31.6)
50% credit	27.4	(06.0–48.8)
66% credit	46.2	(19.5–72.9)
75% credit	58.6	(31.2–86.0)
Fixed-dollar credits		
Higher "cost" of free care		
25% credit	46.9	(27.6–66.3)
33% credit	53.5	(36.2–70.6)
50% credit	65.5	(53.0–78.0)
66% credit	71.9	(59.0–84.9)
75% credit	74.3	(60.9–87.8)
Lower "cost" of free care		
25% credit	26.0	(11.6–40.5)
33% credit	35.8	(22.9–48.8)
50% credit	48.1	(38.1–58.1)
66% credit	58.9	(54.7–63.1)
75% credit	61.2	(55.5–66.9)

SOURCE: Authors' calculations using 1996 Medical Expenditure Panel Survey data.
NOTE: Reservation prices were determined for a comprehensive insurance plan with a $200 deductible, 20 percent coinsurance, and a $1,500 out-of-pocket maximum. Premiums vary proportionately with respect to age- and sex-related expected expenses, and administrative loading is assumed to equal 30 percent of premiums.
[a] The lower bound of each estimate assumes that the currently uninsured continue to consume relatively low levels of medical care given their age, sex, and health status when insured—even after controlling for increased consumption due to moral hazard. The upper bound assumes that the currently uninsured consume the same amounts of medical care that the currently insured consume, given the same age, sex, and health status.

is 27 percent (range, 6–49 percent). Only if we were to assume that the uninsured attach no disutility at all to bad debt and charity care do credits have negligible effects.

A key determinant of the effectiveness of a proportional credit program, therefore, is how much access to free care the uninsured have and how they feel about it. It is obvious but painfully true that if uninsured persons have easy access to free care, that policy (even if unattractive) may for some be preferable to having to pay even half of the premium under a credit program.

■ **Fixed dollar or proportion?** The reduction in the number of uninsured persons also depends importantly on the form of the credit. Suppose that the same tax credit amount is used to provide a fixed-dollar rather than a proportional credit toward a benchmark plan. This means that persons who would have paid a low total premium will now have more of that premium covered by the credit, while those who would have paid a high premium will receive less. As shown in the bottom panel of Exhibit 2, the effect of substituting a fixed-dollar credit for a proportional credit of equal total "cost" is generally to increase the "effectiveness" of the credit in reducing the number of uninsured. The reason is simple: A fixed-dollar credit equal to about half of the total premium for a typical individual is equal or very close to the total premium for younger persons, especially males. In effect, the fixed-dollar credit can cover more persons than the proportional credit can because it covers younger, lower-risk persons (who also skimp on care to a large extent when they are uninsured). While one might well prefer that high-risk people (who will benefit more in terms of access to care—and possibly final health levels—and financial protection) become insured, the choice of credit type is not so obvious when the trade-off is between covering a few more high risks and many low risks.

■ **Altering the level of benefits.** The third policy design feature that can alter the effect of a credit is to change the cost of the plan for which the credit must be used. If the cost of the required plan is cut from $2,000 for an individual to $1,000, fixed-dollar credits bring in many more people. However, $1,000 obviously buys a less comprehensive policy than $2,000 does. In the limit, if the person eligible for a $1,000 credit could use it on any policy that costs at least $1,000, there should be virtually universal take-up of the credit, as long as insurance is worth anything and people know about the credit. Coverage could become near-universal in the sense that almost everyone would have some coverage, but it would not be comprehensive.

Evaluation obviously depends on estimating what the coverage a smaller credit would buy would look like. If it was catastrophic

"If we are not prepared to spend the money to cover everyone, a painful trade-off is unavoidable."

coverage, the deductible would have to be high; there would not be much additional financial protection. Much of the benefit then would go to just replace free care, which is disproportionately provided to persons who incur high expenses. If the premium were used to buy a first-dollar policy, the upper limit would have be set at less than $2,000, exposing the person to financial risk even in the case of moderate illness and covering many small claims with high administrative costs. However, most people will prefer neither pure catastrophic nor pure fixed-dollar coverage. Lower-income persons will want coverage with a modest deductible and a decent upper limit.

Effects of partial coverage. We use our second simulation technique to estimate what $1,000 policy people would choose to maximize their expected utility. If free care is available for large expenses, but the person prefers to avoid free care, we estimate that the policy the person would most prefer with a $1,000 premium for an average-risk individual will have a $525 deductible but an upper limit of $4,072. Compared with no insurance, the person will be protected against those moderate expenses he or she might have expected to have to pay. (Of course, if providers begin to charge for some of the catastrophic care formerly provided free, the financial protection—but not the utility gain from less charity care—is eroded.) An upper limit on benefits is chosen (against traditional economic intuition), since we have observed that the proportion of free care to total expenses increases dramatically as total expenses rise.[9]

The key policy issue obviously is how beneficial the improved access associated with this partial-coverage policy would be. The trade-off is then between having a population in which everyone has some but incomplete coverage and having one in which some have much better coverage than before but others still remain totally uninsured. If we are not prepared to spend the money to cover everyone, this type of painful trade-off is unavoidable.

Real Costs Versus Budgetary Accounting

Considerable effort has been made to estimate the effect of various tax credit proposals on the number of uninsured persons. The next step usually is to calculate the budgetary impact of the credit and then divide the reduction in the number of the uninsured by the net budgetary impact to calculate a so-called cost per newly insured person. However, this amount does not represent a cost in terms of

real resources used, even if one assumes that insurance premiums are exactly equal on average to the cost of resources consumed by persons with insurance. Usually the measured "cost" exceeds the premium because some of the credit is claimed by persons who would have been insured in the absence of the credit, in either the group or individual market. In the case of persons who would have received subsidized group insurance, the actual net "cost" is relatively small, since it is only the difference between the credit and the current tax subsidy. In contrast, most of those who buy individual insurance are not subsidized at all, and thus their entire credit is incorrectly counted as a "cost" to the tax system.

Our key point is that credits claimed by the formerly insured in excess of their current subsidies do not represent real resource costs. If the insurance is equivalent, the real resources used for their medical benefits are the same with or without a credit. The payment from the Treasury to such persons is just a transfer or, more accurately, is a tax reduction relative to their tax obligation before a credit program was introduced. What taxpayers as a group "lose" when the credit is offered they gain in the form of lower taxes. This is a hard message for policymakers to accept, not just because some feel that tax collections "belong" in some sense to the government, but also because discussions of the federal budget tend to treat credits as identical to government outlays for real resources (such as buying a tank), even though no resources are consumed by transfers.

The transfer represents substitution of a targeted tax cut for responsible behavior for some alternative tax cut. How can we make the distinction between real costs and transfers sharp enough to offset this accounting distortion? One way is to consider an example in which some people are eligible for credits to help them buy health insurance while others who are identical with regard to income, health status, and everything else but the value they place on insurance are declared ineligible. Specifically, the rule is that if your value of health insurance was higher than its cost, so that you had already obtained insurance (which you paid for by sacrificing money wages), you are ineligible. Would this be fair?

We think that the answer is clearly negative: No principle of equity of which we are aware which would make this distinction. In fact, no one has proposed to deny credits to those who have already bought individual insurance, although many propose to deny credits to persons who obtained the insurance they value highly by taking a less attractive job just because it comes with insurance. (Perhaps this is because such persons are under the mistaken impression that the employer is "giving away" insurance to them rather than taking the cost out of their wages.)

The other way to make the distinction is to think of credits explicitly as tax cuts. Suppose we proposed to reduce every worker's income and payroll taxes by $1,000 across the board but refused the tax reduction to those who failed to obtain health insurance. Especially if the credit or reduction is large enough that most people at every income level would claim it, this arrangement could hardly be labeled unfair.

How To Navigate Through The Uncertainty?

A number of studies have estimated the impact of tax credit schemes on the number of uninsured persons; results range widely. Why do analysts disagree? Our own estimates have already shown how different and equally plausible assumptions about parameters still unknown (for example, such as how people feel about charity and bad-debt expense) can lead to a wide range of estimates of outcomes. Since no one has observed relatively large subsidies for the lower-income populations who are the bulk of the uninsured, analysts are forced to extrapolate from behavior in slightly similar but by no means identical settings.

One such setting is the behavior of the self-employed in response to changes over time in the tax deductibility of their insurance premiums.[10] The problem here is generalizability: The self-employed for whom deductibility matters tend to have different incomes from all uninsured and to be less attracted by the fringe benefits that come from wage employment.

The other approach, used in the larger-scale simulations, is to estimate the rate of purchase of health insurance by assuming that a program of tax credits for all is equivalent to the reduction in premiums some employees face when they select jobs in which the employer pays part of the premium. Both cross-section and time-series variations have been studied.[11] The question therefore, is: If *x* percent of the workers at ABC Corporation are uninsured when the firm asks them to pay some proportion of the premium, would *x* percent of a random sample of workers also choose to be uninsured if a credit covered the same proportion of their premium? The answer would be affirmative only if workers were randomly distributed across jobs, independent of the values they attach to health insurance, which is surely not the case. But since employers provide insurance to attract those workers who value health insurance, it seems quite unrealistic to assume that even after controlling for workers' observable characteristics, each firm's workforce represents a random slice of the working population.

To be more specific, there are at least two differences between the choice setting under a credit and what is observed in the employ-

ment-based setting with nonzero employee premiums. The first difference is that employees usually have only a few if any choices, whereas persons with credits could choose the plan they like best from a wide range of plans. If I am asked if I want to pay half of the premium for the oppressive managed care plan my employer has selected, the fact that I turn it down hardly proves that there is no insurance plan I would select if I received a credit toward the plan I choose. Second, persons who value health insurance at less than they would have to pay for it under group provision presumably would most prefer to work for firms that pay higher cash wages and offer no employer-paid health insurance at all or who pay only a small fraction of its premium if they do pay something. Thus, the observed percentage of uninsured in firms with high employee premium shares is likely to overstate what the percentage would be among a random or average sample of all workers.

Other researchers have tried to estimate uninsurance rates based on experience in Medicaid programs, on the impact of variation in state income tax rates, and on the behavior of small employers who received subsidies.[12] It should be no surprise that the range of estimates is large, which appropriately reflects great uncertainty.

Other Design Features

As suggested by this discussion, there is more to the purchase of insurance than just the size of a credit. People will be more likely to use a credit of a given size if they can use it easily for a very wide range of insurance policies. The difficult experience of Medicaid and the State Children's Health Insurance Program (SCHIP) in some states in getting people to enroll in free coverage for which they are eligible is often and appropriately cited as evidence that a credit program may not reach all of the uninsured. Today's individual insurance market (as we have discussed elsewhere) is sorely in need of improvement, especially in terms of cost, ease of enrollment, and continuation of coverage.[13] While guaranteed renewability, which protects people against the onset of high-risk conditions, is now required for individual insurance, that type of insurance remains costly to administer. The economies associated with group choice of plan and group billing available in the group setting are not easily extended to a market in which persons choose individually whether to be insured, what plan to take, whether to pay their premiums, and whether they should complain.

■ **Simplified eligibility procedures.** Some features that are intrinsic to tax credit plans may help with some of these problems. How easy it is to establish eligibility depends on how complex the eligibility rules are. Suppose that every worker at a given (low)

family income level is declared to be eligible for a credit. Then there is no need to "apply." Once the government knows the person's income (which it must know in order to establish that there is little or no income tax liability), it could simply mail out or distribute credit certificates to all such persons, which they could then redeem against the premium for a previously qualified insurance plan or against their taxes. Millions of potential customers armed with large discount coupons might be enough to induce individual insurers to offer attractive and easy-to-buy policies. Such a program could greatly reduce selling costs and stimulate the individual health insurance market to transform itself (possibly with the help of electronic means) into a more efficient mass market for coverage.

■ **Links to employment-based coverage.** One of the most common sources of confusion about the use of credits is the belief that they will destroy the employment-based groups in which the great bulk of Americans now buy insurance. The goal for a credits scheme is clear: People should buy insurance in the most efficient setting for them. For those who work for large firms where the advantages of group insurance are strong, the availability of a tax credit need not change their choices, as long as the credit can be voluntarily substituted for the tax exclusion and can be used to pay for group as well as individual insurance. If the program denies the use of credits to persons who want to arrange their insurance purchases at the workplace, there will be an inappropriate negative effect.

There may be an appropriate negative effect at small firms with heterogeneous workforces, which were settling for the single policy that pleased few of them only because the tax exclusion made it a good deal. If the administrative cost savings over individual insurance are small, and the value of the wider range of choices potentially available in that market is large, it is desirable to terminate small-group purchasing in favor of individual insurance. Since group insurance does not pool risk much more effectively than does individual insurance with guaranteed renewability, not even older workers or those who become higher risks will stand to lose from such a transition.[14] If the credit is neutral with regard to the way in which coverage is obtained, people will not use it for alternatives to group insurance if those alternatives are more costly.

The workplace is the place at which all workers are required by law to specify the amount of federal tax they expect to have to pay and to begin to pay it in installments through the use of withholding. The withholding mechanism provides a tailor-made vehicle for financing health insurance: The person who expects to be eligible for a large credit can declare this fact and receive a reduction in withholding, which can then be used immediately to pay for part of

the monthly insurance premium. Some employers may be willing to act as collection agents for insurers (through the use of defined-contribution plans) to reduce transactions costs, which could be incorporated into Web-based billing and administration systems.

Concluding Comments

Since we are virtually certain that part of the reason why some people are uninsured is because they feel they cannot afford insurance, we can be virtually certain that any well-designed tax credit program that cuts net premiums will reduce the numbers of uninsured persons somewhat. However, because the great bulk of the uninsured could have taken advantage of a least a moderate tax subsidy for employer coverage but did not, we can also be virtually certain that a small (say, 25 percent or less) tax credit will not have much effect. It will reduce the taxes of the self-employed who were already purchasing or close to purchasing insurance, but they are a small proportion of all workers. A small credit makes more sense as tax policy (to achieve equity and efficiency on the choice between employment and self-employment) than as health policy.

Beyond these conclusions, predicting the effect of credits that are large enough to matter is fraught with uncertainty, both because there are many possible designs and because the behavioral responses are properly subject to a wide range of conjecture. The form of the credit matters: If it takes the form of a proportional credit for a comprehensive policy at less than 50 percent of the premium, the effect on both coverage and equity could be limited to a few higher risks. In contrast, a fixed-dollar credit for any policy that costs as least as much as the credit would be widely used. In between these two extremes, a fixed-dollar credit targeted toward a more comprehensive plan could cut the proportion of uninsured by a third to two-thirds—although those who switch would be healthier.

■ **Targeting the poor.** Another design issue concerns eligibility: Since affordability is a greater problem for those with low incomes (but not just for those with incomes close to the poverty line), it would seem sensible to target a limited credit to such households. A policy that provided a generous credit to the poor (enough to cover the full premium of a decent insurance policy) and then reduced the size of the credit as income rose would direct subsidies to where they are needed but would have some negative effects on work effort. The rate of reduction in the credit can be made smaller only by extending the credit to higher income levels. The question of how to design and administer a credit that varies inversely with income is one we have discussed at length elsewhere.[15]

■ **Limiting the credit.** Still another design issue is whether eve-

"An incremental strategy would not lead to ideal insurance for all, but it should get some insurance to almost everyone."

ryone at a given income level is eligible for the credit (as a substitute for the tax exclusion if they are already covered by group insurance), or whether it is limited to those who have not taken or are not eligible for such group insurance. In the latter case, there would still be an increase in the number of insured persons, but limiting the credit would exchange one kind of inefficiency for another. We now inefficiently subsidize people to get group insurance of their employer's choosing; if a generous credit could only be used in the individual insurance market, we would be subsidizing people to use an insurance with a wide range of choices but with higher administrative cost. (We do not think that there would be an appreciable impact on risk pooling, for reasons we discuss in our previous work.)[16] Generalizing the credit to all low-income persons would raise the level of additional "cost" per newly insured person, but that cost would only be a transfer to people who deserve it and whose employment and health insurance choices become less distorted. Transfer or not, however, it would still show up in the public sector's budget and take on a life of its own.

■ **Value of an incremental approach.** Given the uncertainties of estimation and the range of trade-offs that have barely been considered and that themselves lack crucial data (such as whether coverage improves health more per dollar for high risks or low risks), suggesting policy is a dangerous mission. If we were to accept this assignment, we would propose to begin in an incremental fashion and use observation of effects as real-time feedback. One sensible place to begin might be with a relatively moderate (but not modest) fixed-dollar credit to families with incomes no higher than the median (and possibly lower). The credit should be enough to cover most of the cost of a basic preferred provider–type policy for the poor and would decline in value as taxable income rises. The policies eligible for the subsidy would be subject to only mild limitations, and people would be permitted to use the credit for the "moderate-deductible/moderate-limit" form we think many would prefer. Such a strategy would not initially lead to ideal insurance for all, but it should get some insurance to almost everyone who is now uninsured. If the credits are about $800–$1,000 per individual, the annual gross value of this tax cut would be on the order of $79 billion, to be offset in part by higher income and payroll taxes on higher taxable incomes.[17]

THIS FIRST STEP would not achieve perfection in the view of many (even though it would be costly), since the coverage is less than comprehensive and makes painfully apparent the ambiguous role played by charity and bad-debt care. (Some states have been paying for stop-loss coverage for insurance for low-income persons.) But it could hardly be worse than staying where we are and probably seeing more growth in the ranks of the uninsured when the economy slows. We should perhaps not make the perfect we as taxpayers will not finance the enemy of the good we would cover. By seeing where credits stimulated coverage and where they were relatively ineffective, policymakers would be in a much better position to modify and fine-tune, to achieve adequate coverage for all Americans at a fair price.

This research was supported in part by a grant from the Leon Lowenstein Foundation.

NOTES

1. For more discussion, see M. Pauly et al., *Responsible National Health Insurance* (Washington: AEI Press, 1992); and M. Pauly and B. Herring, *Pooling Health Insurance Risks* (Washington: AEI Press, 1999).
2. For examples of the former, see D. Cox and C. Topoleski, "Individual Choice Initiatives: Analysis of a Hypothetical Model Act," EBRI-ERF Policy Forum (5 May 1999); G. Wozniak and D. Emmons, "Tax Credit Simulation Project Technical Report," American Medical Association Discussion Paper no. 00-1 (June 2000); and K. Thorpe, "New Estimates of the Federal Costs and Numbers of Newly Insured in Senator Bradley's Health Insurance Proposal" (Unpublished paper, Emory University, 8 November 1999). For examples of the latter, see S. Butler, "A Tax Reform Strategy to Deal with the Uninsured," *Journal of the American Medical Association* (15 May 1991): 2541–2544; M. Pauly and J. Goodman, "Tax Credits for Health Insurance and Medical Savings Accounts," *Health Affairs* (Spring 1995): 125–139; and L. Blumberg, "Expanding Health Insurance Coverage: Are Tax Credits the Right Tack to Take?" (Unpublished paper, Urban Institute, 12 August 1999).
3. These include J. Sheils, P. Hogan, and R. Haught, "Health Insurance and Taxes: The Impact of Proposed Changes in Current Federal Policy" (Paper prepared for the National Coalition on Health Care, 18 October 1999); and J. Gruber and L. Levitt, "Tax Subsidies for Health Insurance: Costs and Benefits," *Health Affairs* (Jan/Feb 2000): 72–85.
4. The explicit methodology used to obtain these estimates can be found in M. Pauly and B. Herring, "Cutting Taxes for Insuring: Options and Effects of Tax Credits for Health Insurance" (Paper to be published in the AEI Seminar Series in Tax Policy, presented 2 June 2000).
5. Data from the 1996–1997 Community Tracking Study's Household Survey indicate that the average single premium for individual insurance policies is approximately $1,500, with premiums actually paid varying around this value, primarily because of age and sex differences. However, individual insurance is generally thought to be less comprehensive than group insurance, so a credited "benchmark" generous plan obtained in the individual market may perhaps be more costly than these currently held policies. Jonathan Gruber's simulation model assumes that considerable adverse selection exists in this market with

individual insurance premiums set to average $2,214, so that a $1,000 fixed-dollar credit covers 43 percent of a single premium for a typical uninsured person. For details, see J. Gruber, "Tax Subsidies for Health Insurance: Evaluation of the Costs and Benefits" (Prepared for the Kaiser Family Foundation, January 2000). Empirical evidence that such systematic adverse selection is not observed in this market can be found in Pauly and Herring, *Pooling Health Insurance Risks.* In short, a $1,000 credit toward the purchase of individual coverage will cover, on average, 40–67 percent of the premium.

6. For more discussion, see M. Pauly, *Health Benefits at Work: An Economic and Political Analysis of Employment-Based Health Insurance* (Ann Arbor: University of Michigan Press, 1997).
7. M. Pauly, A. Percy, and B. Herring, "Individual versus Job-Based Health Insurance: Weighing the Pros and Cons," *Health Affairs* (Nov/Dec 1999): 28–44.
8. More detail regarding the amount, type, and sources of free care the uninsured receive can be found in B. Herring, "Access to Free Care for the Uninsured and Its Effect on Private Health Insurance Coverage" (Doctoral Dissertation, University of Pennsylvania, 2000).
9. Ibid., for the relationship between free care and the magnitude of total charges.
10. J. Gruber and J. Poterba, "Tax Incentives and the Decision to Purchase Health Insurance," *Quarterly Journal of Economics* (August 1994): 701–733.
11. For analysis extrapolating upon variation in employee-paid premiums across different firms, see M. Chernew, K. Frick, and C. McLaughlin, "The Demand for Health Insurance Coverage by Low-Income Workers: Can Reduced Premiums Achieve Full Coverage?" *Health Services Research* (October 1997): 453–470. For analysis extrapolating upon the growth in employee-paid premiums over time, see Sheils et al., "Health Insurance and Taxes."

12. Gruber assumes low responsiveness of the uninsured with very low incomes to tax credits based in part upon the less-than-universal take-up rates of free Medicaid. For more details, see J. Gruber, "Tax Subsides for Health Insurance"; A. Royalty, "Tax Preferences for Fringe Benefits and Workers' Eligibility for Employer Health Insurance" (Unpublished paper, Stanford University, May 1999); and K. Thorpe et al., "Reducing the Number of Uninsured by Subsidizing Employment-Based Health Insurance: Results from a Pilot Study," *Journal of the American Medical Association* (19 February 1992): 945–948.
13. Pauly et al., "Individual versus Job-Based Health Insurance."
14. Pauly and Herring, *Pooling Health Insurance Risks.*
15. See Pauly et al., *Responsible National Health Insurance.*
16. See Pauly and Herring, *Pooling Health Insurance Risks.*
17. Consider a $900 credit. There are four cost components we calculate. The first is a cost of $28.5 billion in credits going to 31.6 million currently uninsured persons with incomes below 300 percent of poverty; we assume a 95 percent take-up rate of at least a partial coverage policy equal to the credit. The second is a cost of $3.6 billion in credits to the 4.0 million low-income persons who are currently insured in the individual market. The third is a cost of $47.2 billion in credits to the 52.4 million low-income persons who are currently insured in the group market. These three costs comprise a "gross" value of $79.3 billion. The fourth "cost" is a gain of $7.3 billion in taxes collected from these group-insured persons who no longer receive the current subsidy equal to their marginal tax rate; thus, the "net" value of this tax credit would be approximately $72.0 billion. More detail of the above calculation is available from the authors upon request; e-mail Bradley Herring, <bradley.herring@yale.edu>.