

Modeling Employer Decisions to Offer Health Insurance

Sherry Glied and Joshua Graff Zivin
Department of Health Policy and Management
Mailman School of Public Health
Columbia University

Revised August, 2004

Abstract:

Very little is known about how firms that employ heterogeneous workers make decisions about offering health insurance coverage to their employees or about how they allocate the costs of that coverage across these heterogeneous workers. In this background paper, we examine how modelers address these questions. We then pull together empirical and theoretical evidence on determinants of the offer decision and on the incidence of health insurance costs and identify some stylized facts that characterize employer behavior. Finally, we assess modeling practices in light of this evidence.

We thank the Robert Wood Johnson Foundation for research support and Linda Bilheimer for encouraging us to pursue this project. We thank participants in a meeting in Washington, D.C. on July 14, 2004 for many helpful suggestions. Sangeeta Goyal provided excellent research assistance.

Under almost every proposal to expand health insurance coverage, some employed people (or their dependents) who currently hold private insurance through their employers would become eligible for subsidies for non-group coverage or for public coverage. If enough workers in a firm can obtain coverage at lower cost outside the firm, it seems likely that the firm will stop offering coverage. There is considerable controversy, however, about which workers and how many workers would need to prefer coverage outside the firm for this to happen. Thus, modeling the consequences of alternative proposals is closely linked to modeling employer behavior. Two aspects of employer behavior are critical—why employers offer coverage and how workers pay for it.

Why do employers offer health insurance to their workers? While there may be some workplace benefits of coverage through lower turnover or absenteeism, the most important reason that employers offer coverage is that workers want it. Workers, in turn, prefer to obtain coverage through their jobs for several reasons. Job-based coverage is tax-favored; it takes advantage of economies of scale in the administration of insurance; it provides a way to avoid adverse selection; it offers the possibility of multiple-year insurance coverage (see Glied, forthcoming). Overall, some 66 percent of all U.S. firms offer coverage to their workers (Kaiser Family Foundation 2003).

How do employees pay for this coverage? Economic theory suggests that it is paid for (primarily) through foregone wages and explicit employee contributions. But the process through which this happens is clearly complex. It is particularly complicated in firms where workers have diverse preferences about health insurance, as is often the case. For example, in firms where coverage is offered, some workers—about 17 percent—do not take it up (although many in this group obtain coverage elsewhere; Kaiser Family Foundation 2003). In these firms (and perhaps in others too), workers obviously have varied preferences about job-based coverage. Economic theory says that the costs of health insurance coverage are borne by workers, but all empirical studies—and everyday experience—suggest that workers such as these who do not take up offered coverage fail to receive higher wages in exchange.

Economic theory and empirical research provide little insight about how firms that employ heterogeneous workers make decisions about offering coverage or about how they

allocate the costs of that coverage across these heterogeneous workers. This is a surprising gap, since health insurance costs now accounts for 5.6 percent of total compensation costs (EBRI 2002). The question has renewed policy significance because of the growing interest in incremental expansions of health insurance coverage. These expansions would alter the relative costs and benefits of group coverage for a segment of the workforce. In effect, they would increase the degree of heterogeneity in health insurance preferences across the workforce.

In this background paper, we examine how modelers address these questions. Our focus is on modeling expansion strategies, such as tax credits or deductions, which have differential impacts on workers within a firm. We pull together empirical and theoretical evidence on determinants of the offer decision and on the incidence of health insurance costs and identify some stylized facts that characterize employer behavior. Finally, we assess modeling practices in light of this evidence.

I. Modeling Strategies

Three basic decisions distinguish employer behavior modeling strategies: whether to model based on individual characteristics or firm characteristics; how to aggregate preferences within the firm into a decision rule; and how to allocate the costs of health insurance across workers.

There are no available data sets that include detailed information about individuals, about the characteristics of their co-workers, and about health insurance coverage. Modelers must either treat each individual in a micro data set as representative of her firm or create synthetic firms.

Treating each worker as representative of her firm means that the modeler excludes information about the characteristics of other workers at the firm. This is an important exclusion under two conditions. First, this assumption cannot capture characteristics of workers other than the preferences of the average employee¹. Thus, it cannot reflect the situation where there are synergies in the provision of health insurance across workers. Second, this assumption does not mimic the behavior of firms if firm decision rules are a discontinuous function of average worker preferences. This would be true, for example, if the decision rule depended on whether the

average worker's benefit from dropping was above a threshold. The average worker within a firm might be above the threshold after a new program is put in place, yielding the result that the firm drops coverage. Some representative workers, however, might not be above the threshold, so that only a fraction of these workers would be modeled as losing offers. Under this modeling strategy, for example, high wage workers might never be predicted to lose coverage.

Modelers who use representative worker rules generally make firm dropping decisions a function of the individual's gains or losses from dropping coverage as well as of the size of the firm. In effect, such decision rules place little weight on individual worker preferences in large firms, and more on individual worker preferences in small firms. In the extreme case of a self-employed individual, presumably the firm's decision and the individual's decision are identical.

The creation of synthetic firms requires modelers to use outside information about the distribution of earnings, ages, and other characteristics within firms of different sizes to generate firms that resemble actual firms. There is relatively little information to go on in doing this. Some data sets contain information on employee composition (such as the RWJF employer survey and the BLS Employment Cost Index surveys). Even these surveys, however, lack information that would be pertinent to actual firms—such as the outside insurance coverage and coverage options available to their employees.

Modelers who construct synthetic firms can make use of the entire distribution of wages and price changes within the firm to model employer responses. They can then apply a decision rule to various aspects of this distribution. One option is to compute the average change in the situation of all workers in the firm after a policy change. This might mean, for example, computing the average change in the price of non-group coverage relative to group coverage across all workers. An alternative is to use a voting rule. Under a standard voting rule, the coverage preferences of the median worker within the firm (perhaps relative to a threshold) determine the firm's decision. Again, some modelers apply different decision rules for large and small firms. Finally, modelers can examine the total cost to the firm of offering coverage and compare these to the total costs of not offering coverage.

The distribution of wages is bounded below (either by zero or by the minimum wage). Thus, the median wage is virtually always below the average wage within a firm. The median

characteristics of employees will generally be “lower” than the average characteristic of employees. This implies that, in general, a model that uses a median worker or voting rule will weight the preferences of lower income employees more heavily than will a model using an average employee rule or a model that sums costs of compensation across all workersⁱⁱ.

The final critical element of modeling has to do with the allocation of health insurance savings (or costs) among workers within the firm. Models that incorporate individual worker decisions, such as median or average worker models, require the median or average worker to contrast his or her after-tax income net of health insurance costs if the firm drops coverage to net income if obtaining coverage through the employment group. In computing after-tax income in the firm-drop scenario, modelers must make assumptions about how firms will alter each worker’s wages when the firm foregoes payment of insurance premiums. Models that sum costs of compensation across workers allow the offer decision to be modeled independent of the allocation of costs across workers (for incremental policy changes)ⁱⁱⁱ.

In practice, all modelers assume that the decision to accept or reject fringe benefits offered by one’s employer has no effect on an individual’s own wage. Workers who turn down offered coverage save paying the employee share of premiums, but they get no other cash benefit. Some modelers do project some effects on overall wages from individual workers dropping. All modelers assume that the cost of health insurance will be added to workers’ wage compensation only if the firm ceases offering coverage. Currently, all modelers assume that the firm’s average cost of health insurance is subtracted from workers’ wages. To our knowledge, no modelers use differential group-specific incidence or alternative incidence assumptions.

II. Evidence on Employer Behavior

Employer behavior is a surprisingly understudied area in economics. The absence of datasets that link employer and employee characteristics helps to explain this general vacuum. The Census bureau has been working on the development of such datasets, but to date none of the matched data sets include information on health insurance (see Hellerstein and Neumark 2004 for a description of some of these datasets). Insights gleaned from these linked data may inform future modeling work. For now, however, modelers must depend on more indirect evidence.

Empirical Evidence on Firm Behavior

Several studies have reported price and non-price correlates of the firm's offer decision. Unionized firms, firms in the manufacturing sector, and public sector employers are more likely to offer health insurance than are other firms (Buchmueller et al 1999). Small firms are less likely to offer health insurance than are large firms (KFF/HRET, 2002; Gruber and Lettau, 2000). Small firms typically face much higher loading costs than do larger firms (Pauly 1997), so this offer difference suggests that prices of health insurance matter to the offer decision.

A reasonably substantial literature drawn mainly from health services research examines the price elasticity of firm offer decisions. The range of results is rather broad. Empirical estimates of this elasticity vary from a high of -5.8 (Feldman et al. 1989) to a low of -0.07 (Thorpe et al. 1992). The studies differ in how they measure the price of health insurance and in how they impute a price of coverage to firms that do not offer coverage.

It is difficult to compare these elasticities because they are derived from populations that vary in the share of the population uninsured. A similar percentage change from a larger or smaller base translates into a very different effective take-up rate in the uninsured population. Table 1 summarizes the results of all studies of firm offer and translates all price elasticities to take-up responsiveness in the uninsured population^{iv}. There is substantial variation in responsiveness even after conducting this translation.

There is also some evidence that the extent of responsiveness to price varies with other firm characteristics. Many studies find that small firms are much more responsive to changes in price (measured by state or provincial marginal tax rates) than are larger firms (Gruber and Lettau 2000; Stabile 1999; Finkelstein 1999; Leibowitz and Chernew 1992).

Estimates of the price elasticity of offer often do not incorporate information on the alternative sources and costs of coverage available to workers. Presumably, the firm's decision, which reflects some aggregation of workers' decisions, ought to reflect these alternatives. We should expect the firm's elasticity of offer to vary with the availability of spousal coverage or of public coverage (see Shore-Sheppard et al. 2000 for some evidence of this). In models that use state tax rates to identify price variation, the higher elasticities found in small firms likely reflect

the much smaller price advantage over non-group coverage in small firms relative to larger firms.

Unfortunately, the firm offer models provide little information directly useful to modeling insurance expansions that are based on individual tax credits or deductions. Insurance expansions alter the demand for health insurance within the firm—not the price faced by the firm in aggregate. Estimates of firm-level price elasticities are useful in estimating the response of the firm to a firm-level subsidy. Firm-level price elasticities may also be useful to the extent that firms offer coverage beyond the level demanded by their workers (for example, to reduce absenteeism).

Some literature examines the characteristics of workers within firms that do and do not offer coverage. It is clear that employee demand for health insurance is a central consideration in firm offer decisions. Firms that employ higher wage workers (who are likely to have a greater demand for health insurance) are more likely to offer coverage (Vanness and Wolfe 2002; Bundorf 2002; Cooper and Schone 1997; Gruber and Lettau 2000). Among small firms, those with higher average wages are more than twice as likely to provide health insurance for their workers as are low-wage small firms (82 percent. 39 percent) (Gabel et al. 1999). Firms whose workers are likely to have access to health insurance through their spouses are less likely to offer coverage (Dranove, Spier, and Baker 2000; Cooper and Schone 1997; KFF/HRET 2002).

One study, Gruber and Lettau (2000), examines the effect of labor force heterogeneity within the firm on the decision to offer coverage. Gruber and Lettau find that the offer decision at the firm level depends both on the prices faced by the median worker (rather than the average worker) and on the prices faced by highly compensated workers^v. Gruber and Lettau construct simulation models to describe how changes in the tax treatment of health insurance (which mainly affect more highly compensated workers) would affect offer decisions. The results of these simulations suggest that the impact of highly compensated workers can be rather substantial. The change in offer rates in models that use median changes only is only about 60 percent as large as the change in models that also incorporate the preferences of the most highly compensated workers.

All of these studies of firm behavior examine cross-sectional variation in offers. The time series evidence on firm offering presents a somewhat different picture. Over the period 1990–2000, while health care costs per capita rose 69 percent (adjusted for CPI) (Center for Medicare and Medicaid Services 2002) and between 1987–2000, the percentage of workers covered by job-based insurance fell from 71 percent to 68 percent, the rate of firm offering hardly budged. Offer rates have remained constant with all of the changes in coverage occurring through changes in take-up. In essence, the time series elasticity of firm offer appears to be very close to zero. This pattern is consistent with a model in which firm offer decisions depend on the relative price of group versus non-group coverage. Employer-based coverage has adapted to rising costs in other ways, particularly by raising the employee share of health insurance premiums (Farber and Levy 1998; Cooper and Schone 1997).

Incidence of Health Insurance Costs

A second set of studies examines the incidence of health insurance costs. These two strands of literature have not been linked. Theoretically, however, they are closely connected. If the incidence of health insurance costs is perfect, so that each worker pays exactly the cost of his or her own coverage, firm decisions to offer coverage to a group of workers can be treated as a series of transactions between the firm and each individual worker.

In the case of the firm with only one worker (or of a firm made up of identical workers), each worker would bear the value of the benefits selected in the form of lower wages. Wages will decline to compensate for the costs of benefits, as long as the worker values benefits at their cost (or at more than their cost)^{vi}. The theoretical literature on incidence does not provide clear predictions about the incidence of fringe benefits in a heterogeneous workforce.

At the national level, the evidence that health insurance costs are shifted to workers is quite strong. Fluctuations over time in employer payments for health insurance (or benefits more generally) are uncorrelated with the labor's share of the GDP, implying a one-for-one tradeoff between these two forms of compensation.

The empirical literature on compensating differentials at the individual level, however, is surprisingly weak (Morrisey 2002). It has proven very difficult to identify tradeoffs between

workplace conditions, whether risks or benefits, and wages. In simple analyses, the typical finding is that workers whose jobs are “better” in terms of their non-wage attributes also receive higher wages. The usual explanation for this finding is that those employed in good jobs are likely to differ from those in worse jobs on both observable and unobservable characteristics. The difference in the tax treatment of wage and non-wage compensation also means that highly compensated workers will also be those who most value non-wage compensation.

One strategy for identifying compensating differentials has been to use longitudinal data and track the same individuals as they move from job to job. This method is flawed, however, because job changers are likely to differ from those who remain at a job and because their reasons for changing jobs may be correlated with their productivity.

Another strategy has been to exploit natural experiments that alter fringe benefit provisions. The most prominent study in this vein is the examination of the incidence of mandated maternity benefits (Gruber 1994). This study finds that the incidence of mandated maternity benefits falls on the class of workers who benefit from the new mandate. The singular strength of this study is that the mandate was externally imposed on all workers. A weakness, however, in the context of modeling insurance expansions, is that the incidence of a mandated benefit may be more akin to the incidence of a tax than to the incidence of voluntarily-provided benefits, except to the extent that the mandate affects offer probabilities.

Other studies use naturally-occurring variation that is not a consequence of policy change. In a study comparing wages across cities that differ in their health care costs, Sheiner (1999) found that older, less healthy workers pay more for job-based coverage in high cost cities than do younger, healthier workers. Pauly and Herring (1999) find that returns to experience are lower among workers with health insurance than among those without coverage. Olson (2002) found that among married women working full time, those who had health insurance receive 20 percent lower wages than their counterparts without health insurance.

In weighing these studies, it is important to recognize the “file drawer phenomenon.” Empirical studies that do not identify statistically significant incidence (or compensating differentials) are much less likely to be published (or even completed) than are those that do. In that light, it is worth noting the large number of recent studies that fail to find evidence of

incidence. These include Levy and Feldman (2001), Simon's (2001) studies of job losers, and several natural experiment studies, including studies of the effect of increases in the minimum wage (Simon and Kaestner 2003 find no effect; Royalty 2000 finds small effects) and of the Medicare secondary payer legislation (Garret and Glied 2000). These studies suggest that incidence may operate only weakly at the level of the individual. Further evidence for weak individual incidence comes from the literature on job lock. By construction, the extent of job lock must be inversely proportional to the extent of individual incidence. If incidence were perfect, job lock would be non-existent (Gruber and Madrian 2002). The weak results of minimum wage studies are particularly striking. If workers bear the full cost of their own health insurance, workers who hold health insurance and have earnings near the minimum wage should be strongly affected by increases in minimum wage rates. The failure to systematically establish this result suggests that incidence is more complex.

A number of studies have explored other possible patterns of cost incidence. Ketsche and Custer (2004) suggest that incidence may operate differently for high- and low-wage workers. Garrett and Nichols (2004) suggest that incidence may operate differently in small firms, depending on the extent to which they compete for workers with large firms. In general, we would expect owners of firms that were inefficient purchasers of health insurance (including small firms that incur relatively high loading costs) to bear these incremental costs of coverage directly, rather than passing them along to workers within the firm (Danzon 1989).

Research on employee participation in health insurance plans provides additional evidence that individual incidence is weak. Several studies examine worker's decisions to participate in health insurance that is available to them. If incidence were perfect, workers' take-up decisions would respond to the total cost of coverage. The literature on employee participation decisions, however, strongly suggests that the allocation of health insurance premiums between workers and employers affects health insurance participation decisions.

The tax treatment of health insurance suggests that employers should never charge an employee premium. Doing so subjects payments to taxation (or forces the employer to develop and administer a tax-sheltered premium conversion plan). Today, however, fewer than 30 percent of workers obtain single coverage without paying a premium, and 10 percent of workers

can obtain family coverage without paying a premium. Even among workers who did not have a premium conversion plan available, and were required to pay taxes on the employee share of premiums, nearly 39 percent were required to pay a premium for single coverage in 1997 (Levy 1998).

Employers recognize the importance of employee premium share in take-up decisions. Studies find that employers increase employee premium shares when more workers have alternate sources of coverage (Dranove, Spier, and Baker 2000). This behavior serves to discourage married employees from obtaining coverage from both employers (so-called double coverage) and to shift coverage across employers. If incidence were perfect, there would be no point to this increase in shares. Employees covering their spouses would pay for that coverage themselves. Those with double coverage would pay double for it. While Dranove, Spier, and Baker (2000) and Levy (1998) show that employers alter premium shares to address heterogeneity of preferences, their specification of worker characteristics within the firm make their results difficult to convert to elasticities of contribution rates. Research that is currently ongoing examines how firms change employee premium shares in response to policy changes.

Most studies find that workers are much more responsive to the employee premium 'share' of coverage than to the total cost of coverage. Several studies find that employees are at least somewhat responsive to premium shares (Cutler 2002; Blumberg et al. 2001). A recent study by Gruber and Washington (2003), however, finds that the increase in out-of-pocket premiums induced by pre-tax premiums for all federal employees had minimal impact on participation. Studies virtually always find that the absolute amount an employer pays in premiums has no impact on the decision to take-up coverage (Chernew, Frick and McLaughlin 1997; Blumberg, Nichols and Banthin 2001). Again, these studies use elasticity measures, which are sensitive to the initial fraction of the population uninsured. We converted these elasticities into effective take-up rates for the uninsured population and found that responses to employee contributions were quite similar across these studies (see Table 2). Note that some of these studies lack information on alternative sources of coverage for the affected workers. We would expect responsiveness to employee contribution rates to vary with such alternative coverage sources.

Stylized Facts Based on the Empirical Literature

The empirical literatures on firm offers and on incidence suggest several stylized facts about employer behavior. The decision to offer coverage varies with the price of coverage, particularly for small firms; characteristics of the labor force, including average wage levels and the distribution of wages within the firm, affect the offer decisions; and that over time, firms have not responded to rising costs by ceasing to offer coverage. The costs of coverage fall on workers in aggregate, and perhaps at the group level, but higher cost individuals do not appear to receive lower wages in exchange for health insurance (with the possible exception of workers in very small firms). Worker take-up decisions are sensitive to the employee share of premiums and employers charge employees a share of premiums, even though this sometimes exposes them to greater taxation.

In sum, the extent and nature of heterogeneity of preferences over health insurance within the workforce has strong implications for whether firms offer coverage and whether individuals take-up coverage that is offered to them. Low wage workers who work with higher wage workers appear to receive greater compensation (measured as the sum of wages and employer health insurance payments) than do those who work only with other lower wage workers. Employers do appear to manipulate the employee share of premiums in order to reduce participation rates among workers who do not value benefits (at least up to a point). To date, there is little evidence on how employers vary these worker shares with policy changes.

III. Theory of Benefits Provision

The economic theory that describes the decision to offer health insurance (or other fringe benefits) draws on the broad compensating differentials literature. The theory of compensating differentials says that firms seek to minimize their total labor costs. Labor costs comprise both wage and non-wage benefits. Non-wage benefits include payment in kind (such as health insurance benefits), deferred compensation (such as pensions), and working conditions, including job safety. Employees trade off these job attributes for wages such that the utility derived from the compensation package at one job is equal to that derived from the compensation package at the best alternative job.

The compensating differential literature encompasses the case where employees' utility of benefits may exceed the cost to employers of providing these benefits – in that case, the tradeoff of wages for benefits may be more than \$1 of wages for \$1 in benefits^{vii} (Famulari and Manser 1989). The substantial economies of scale associated with group insurance coverage, combined with the favorable tax treatment of this coverage, suggests that this situation is likely to be the case for many, though not all, workers in the context of health insurance. Consider a worker who is willing to pay the non-group, after-tax price for health insurance. The dollar value of this worker's benefit from employer-sponsored coverage may exceed the cost of that coverage by as much as 100 percent. Other workers, however, may place a very low value on employer-sponsored health insurance. This group is likely to include people, such as some tax credit recipients, who have the option of obtaining free or very low cost coverage in the non-group market.

This basic economics literature relates individual utility curves—and resulting tradeoffs—to firm level decisions. If firms consisted of only one worker, this model would be sufficient to capture all salient employer behavior. In practice, however, firms make benefit decisions that encompass many workers. The theoretical literature is not very clear about how a firm, making a single firm-level tradeoff between fringe benefits and wages, weights the utility functions of its various workers, including workers who would place a very high and a very low value on group insurance coverage^{viii}.

Health Insurance as a Local Public Good

The main theoretical research examining the problem of firm decision making about health insurance in the presence of worker preference heterogeneity is Goldstein and Pauly (1973). Danzon (1989) offers a similar treatment of this problem. Goldstein and Pauly characterize the problem of offering coverage in a firm as one involving a “local public good”. Suppose a firm that offers health insurance must make coverage available to all workers of a given type upon similar terms. Like other public goods, once the offer of coverage is available to one worker, it is available to others as well. Moreover, as is the case for other public goods, firms cannot perfectly assess individual preferences for health insurance. They do not vary wages at the individual level to account for heterogeneous preferences. Thus, when individual workers

within a firm choose not to take-up coverage that is offered to them, they do not receive cash wages equal to the cost of the coverage they declined.

Several features of employer-sponsored health insurance generate links between workers. The apparent existence of economies of scale in insurance coverage suggest that the cost of covering additional workers in the same firm declines as the size of the group, and the participation rate, rises. Moreover, if the participation rate is not high enough, the firm may not qualify for group rates, imposing substantial costs on all workers who would choose to participate. In such situations, the cost of covering one additional worker will be less than the average cost of coverage.

Some—but not all—of the properties of health insurance coverage that lend it the character of a local public good are a consequence of regulations. The requirement to offer coverage on similar terms to different workers is a consequence of tax law. Non-discrimination rules place certain limits on the types of benefits that can be offered to highly compensated employees (although there are many ways around these rules). Likewise, the inability to adjust wages according to the cost of health insurance is, in part, a consequence of tax law. Economies of scale in size and participation, however, are inherent to group insurance.

The economic theory of local public goods follows Tiebout (1956). Tiebout argued that municipalities would compete to attract residents by offering public goods. In equilibrium, all residents of a given municipality would have identical preferences about provision of the public good. People with different preferences would move to other municipalities where their preferences would be met.

The Tiebout model, in the firm context, suggests that, in equilibrium, all workers in a firm would share the same preferences for health insurance. In that situation, firms themselves would differ in the combination of fringes and wages that they offered, and workers would sort themselves across firms according to their preferences. All workers in a given firm would desire the same combination of fringe benefits and wages.

There is, indeed, some evidence of sorting in the context of health insurance. Several studies show that workers with low preferences for health insurance are disproportionately

employed in firms that do not offer coverage (Marquis and Long 1995; Monheit and Vistnes 2001; Levy 1998). Similarly, some evidence suggests that increases in health care costs lead firms that offer coverage to increase their use of contract arrangements whereby they can avoid hiring workers with lower willingness to pay for such coverage (Levy 1998).

The extent of this sorting, however, is naturally limited. Firms appear to produce output most efficiently when they use a combination of workers whose preferences over fringe benefits are likely to differ. For example, most firms include both highly-paid managers and less-highly-paid line workers. These two groups are likely to differ systematically in their fringe benefit preferences. Lumpiness may also preclude perfect sorting. In many labor and product markets, there are not enough people to sort into homogenous firms.

If sorting were perfect, the problem of modeling a new public program or tax credit would be straightforward. The decision process of the firm could be conflated with the decision process of a single, representative worker. For example, the introduction of a tax credit available only to low-wage workers would lead to re-sorting of workers into firms. Only coverage of tax credit recipients would be affected by the introduction of the credits. Those employees who had selected firms offering coverage before the tax credit was introduced, and who were not offered credits, would again select firms offering coverage in the post-credit era. Some firms might change their offering status, but workers would then migrate among firms until they returned to their most preferred wage-benefit combination. In a world with perfect mobility and entirely flexible firms, few people would lose coverage through the introduction of a tax credit^{ix}.

It is much more difficult to characterize the decisions to offer local public goods of a firm with a heterogeneous workforce^x. Goldstein and Pauly consider this problem and argue that the firm should choose its fringe/wage tradeoff to respond to workers' preferences at the margin. The point at which the average of the marginal rates of substitution between wages and benefits across all workers in the firm is equal to the marginal cost of benefits yields the optimal level of benefits.

In the literature on union decision making, the decision of how much coverage to offer in the non-union context depends, similarly, on the preferences of the marginal worker. The marginal worker (or class of workers) is that worker (or class of workers) who is just indifferent

between this firm and an alternative firm. Unfortunately, this concept has limited practical application for modelers. It is not empirically straightforward to define the characteristics of the “marginal worker.” Policy changes, such as those contemplated here, would also likely change the identity of that worker. Finally, because the decision to offer or not offer coverage is dichotomous, rather than continuous, the margin may be quite thick, including many heterogeneous workers^{xi}.

There is no reason, in principle, to expect the average of workers’ marginal rates of substitution or the preferences of the marginal worker to represent the preferences of most or even many other workers, or of the median worker. The idea of the average worker comes closer to the notion of the marginal worker—but in this model, the “average” represents average marginal tradeoffs between wages and fringe benefits, not average health insurance costs, subsidies, or wages.

Goldstein and Pauly (1973) describes the firm’s decision about how much health insurance to offer. Another way to restate the local public goods model is to consider the employer’s decision as to whether or not to provide health insurance to the firm’s employees (Graff Zivin and Glied, in progress). If the employer offers coverage, it will then collect “taxes”, in the form of reduced wages, from the employees sufficient to pay the full cost of health insurance. The employer will offer coverage whenever the sum of the valuations of job-based coverage across all participating workers is greater than the cost of that coverage for all workers participating. Worker valuations of job-based coverage are likely to vary with the nature of subsidies to coverage (for example, the tax treatment of health insurance makes job-based coverage more valuable to higher wage employees), with the size of the firm (since larger firms have greater administrative advantages relative to smaller firms), and with worker-specific health care preferences^{xii}. This implies that a firm which contains only a small proportion of workers with a very high valuation of job-based coverage might well choose to offer coverage, even though the majority of workers, and the median worker, might not place a great value on such coverage (see also, Danzon 1989).

Union Models

Goldstein and Pauly (1973) contrast their average marginal rate of substitution model with a model of a unionized enterprise where decision-making is described in terms of voting. In a unionized firm, workers collectively select a wage-fringe tradeoff by voting for collective bargaining agreements or by electing union leadership. In voting models, the median voter—here, the median worker—drives decisions. Goldstein and Pauly, consistent with the theory of union decision-making, describe this median worker as a more senior and established worker than the marginal worker contemplated in traditional theory, who is described as a recently hired, young worker. As more senior and established workers are assumed to value fringe benefits relatively more highly than would younger workers, the model can explain the empirical finding that unionized workers receive a larger share of their compensation in the form of fringe benefits than do non-union workers. Goldstein and Pauly's main empirical finding is that unionized workers have higher fringe benefits as a share of total compensation than do other workers^{xiii}.

In the unionized firm context, workers earn super-normal compensation because of the presence of the union. Even unionized workers who would prefer a different mix of wages and non-wage benefits may remain at the firm if they value the total compensation package highly enough. If a tax credit for non-group coverage were made available in a unionized market, the decision about whether to take it or not would depend on the eligibility of the median worker (voter). As long as the median worker is not eligible for the tax credit, the firm's offer decision will be unaffected by the credit. Some tax credit eligible workers might then leave the firm in favor of jobs where compensation is weighted more heavily toward wages. If the median worker is eligible, the firm will stop offering coverage. Some ineligible workers might leave the firm for other firms where coverage continues to be offered.

This union worker scenario can be effectively modeled through a median worker model with synthetic firms. This model may also be applicable in industries with a large union presence, where the threat of unionization affects compensation decisions even at non-union plants. At present, however, only 12.9 percent (BLS 2004) percent of American workers are enrolled in unions, and unionized workers make up an even smaller fraction of low-income workers, the group most likely to be eligible for substantial tax credit subsidies. In practice, it would likely be prohibitively complicated to model unionized firm (or industry) decision-making differently from that in non-union sectors.

Who Bears the Cost of Coverage?

In a perfectly competitive, freely mobile labor market, firms would compete to offer health insurance to each worker desiring this benefit. In the resulting equilibrium, there would be perfect sorting and the compensating differential due to health insurance would be equal to the cost of health insurance for each individual worker.

In Goldstein and Pauly's model of local public goods, mobility is imperfect (production functions may require the use of a heterogeneous workforce). If all workers at the firm must be enrolled in the same health insurance option, but some workers value coverage more than others do, payment for coverage (measured compensating differentials) will vary among classes of workers according to each class of workers willingness to trade-off wages for the specified level of health insurance.

Another way to model this local public goods problem is the view the employer as allocating "taxes" across workers to fund coverage. The cost-minimizing tax scheme is the equivalent of Ramsey taxation, under which the employer would charge each employee a tax equal to that employee's valuation of job-based coverage. Ramsey taxation will be the only sustainable tax scheme if each worker has the option of defecting to a firm that does not offer health insurance (recall that there is no equilibrium except perfect sorting if workers may defect to a firm that does offer health insurance). In the context of job-based coverage, this arrangement implies that workers who place a high value on such coverage would make side-payments that effectively subsidize the cost of coverage for low valuation employees. This implies that low valuation workers who work with high valuation workers are subsidized in their purchases of group health insurance. Conversely, if these workers did not obtain coverage through their jobs, their wages would increase by less than the cost of health insurance coverage.

Under the union model, the median worker's willingness to pay for coverage determines the quantity of health insurance offered and the price paid for that coverage by the median worker. The model provides little information about the allocation of costs among all other workers. An equal allocation of health insurance costs across all workers is consistent with this model.

Employee Premium Shares

Goldstein and Pauly (1973), Dranove, Spier, and Baker (2000), and Levy (1998) consider the question of why firms require employee contributions to premiums (particularly when such contributions are not tax favored). In all three models, the authors conclude that if preferences over health insurance are quite heterogeneous, the total cost of compensation can be reduced if low valuation workers do not participate in the health insurance plan. At the same time, however, employers must retain a substantial fraction of employees in the plan to take advantage of group rates, so contribution rates (even when contributions are tax-sheltered) may be limited^{xiv}.

In this scenario, the effect of the introduction of a new non-group (or public) option on contribution rates is not obvious. It will depend on existing contribution rates, the degree of workers' preferences for the outside option, and the share of workers affected. In a firm that is not requiring contributions, the introduction of a new outside option available to some small share of workers would lead to the introduction of contributions, so that these low-valuation workers would drop coverage. In a firm already requiring contributions, the introduction of an enhanced outside option will lead more workers to drop coverage, holding contribution rates constant. If too many workers would drop coverage to maintain the firms' plan, and if the value of that plan to high valuation workers is very high, required contribution rates could actually fall, to stem defections from the plan. In theory, it is even possible that the introduction of an enhanced outside option available to low-valuation workers only, combined with an employee premium contribution requirement, could lead to more offers—offers that would then be taken up only by high valuation workers.

Implications of the Theoretical Literature

The theoretical literature provides a convincing description of the behavior of unionized firms. In such firms, the decision to offer coverage will be based on the preferences of the median worker. Moreover, the cost of coverage may be evenly split among all workers. Regrettably, the applicability of this model to most US workers is likely to be limited.

The theoretical literature also provides strong hypotheses about behavior in the case where there is complete individual incidence. In that case, firms offer coverage and individual

workers take it or leave it according to their own preferences. There is virtually no empirical support for this scenario.

The theoretical literature is also clear in the case where there is perfect worker mobility. In this case, workers sort themselves into firms that are homogeneous with respect to health insurance preferences and each worker then pays the full cost of her own coverage.

The local public goods model offers a description of the employer-based health insurance market if mobility is very low. This model suggests that employers choose the cost-minimizing combination of wages and benefits, recognizing the effect of benefit choices on the wages paid to each worker. These models yield the result that decisions are driven by the average of marginal rates of substitution (or by workers' willingness to pay for job-based coverage). If, for example, higher-wage workers have higher marginal rates of substitution of wages for health insurance than do low wage workers (they are more willing to trade-off wages for health insurance, because they face higher costs of obtaining coverage in the non-group market), their preferences will count more. Moreover, they will bear a higher share of the cost of coverage than will lower-wage workers.

The question of which model to use is ultimately an empirical question about the extent of benefit preference heterogeneity (or matching) in the labor market, and the extent to which such heterogeneity varies with policies. To date, we have, at best, estimates of what fraction of workers within a firm turn down coverage that is offered to them. This provides only one indication of extreme mismatches (and even here, the size of the mismatch depends on assumptions about how the wages of workers who turn down coverage are adjusted). It does not take account of those workers who do obtain job-based coverage but would prefer higher wages instead^{xv}. We have no information about the sensitivity of allocations to policy changes (for example, no studies to date have shown whether workers whose children became eligible for SCHIP switched to firms that do not offer family health insurance coverage.)

IV. Implications of Theory and Empirical Evidence for Modeling

Our review of the theoretical and empirical literature suggests that there are enormous gaps in our theoretical and empirical understanding of employer decision-making. The inter-linked

literatures on employer behavior and on compensating differentials both badly need additional research work. Thus, our conclusions should be viewed as highly tentative and exploratory.

First, given the uncertainty around how preferences are weighted within the firm, modelers should consider exploiting the full power of synthetic firms by summing benefits and costs within the firm under the scenario that health insurance is or is not offered. In effect, summing benefits and costs is similar to a well-specified average worker model or to a weighted median worker model. It does not impose specific assumptions about incidence across workers within the firm. The preponderance of evidence—both theoretical, descriptive, and empirical—suggests that higher income workers exert more influence within the firm than do lower income workers.

Second, although it is computationally convenient, there is little theoretical or empirical evidence to support the allocation of health insurance costs equally among workers. The evidence on employer behavior, by contrast, is consistent with the hypothesis that compensating differentials vary within the firm. Firms may vary incidence according to several categories of broad classes, such as wages, experience, or family size. If higher-wage workers pay a larger share of insurance costs, they would have a disproportionate influence on firm's decision-making in a Gruber-Lettau (2000) model. Likewise, older workers should have a disproportionate impact on firm decisions if, as seems likely, they pay a larger share of the firm's health insurance costs (corresponding to their own higher costs). Such heterogeneous incidence behavior would help explain the anomalous minimum wage results. A median worker model with heterogeneous allocation of the costs of health insurance within the firm may yield results similar to those of an average wage model.

As always, modelers should be as precise as possible in describing the methods used in their models. This is particularly important with respect to employer modeling both because of the uncertainties in this area and because it is computationally difficult to conduct sensitivity analyses around employer behavior. Modelers should specify how they are allocating the costs of premiums (or of premiums foregone). They should describe the decision rule used within firms (average worker, median worker, weighted average). Finally, they should describe how the parameters which are being "averaged" are constructed.

Several types of new research could be very useful in informing models. First, research is needed about how much misallocation of workers exists. Analyses that examine changes in matching after policy changes would be particularly valuable. A corollary of these empirical analyses would be an assessment of how to relate estimates of the extent of matching to decisions about whether to model firm decisions in terms of local public goods or as individual trade-offs within the firm.

Second, analyses of changes in synthetic firm behavior in response to price changes using alternative voting rules should be compared to firm-level estimates of price elasticities. This exercise of calibrating micro estimates to macro estimates may shed light on how best to characterize the process by which firms aggregate the preferences of their workers.

Third, it would be useful to better understand how firms pass along the costs of health insurance to workers. There are clearly limits on this process in the short run, but no research to date has characterized what these limits are^{xvi}. Furthermore, it is likely that firms that are inefficient purchasers of health insurance bear some costs directly, but this has not been shown empirically.

Finally, there is very little literature on the determinants of employee premium shares. Decisions about premium shares may depend on the availability of arrangements that permit employees to shelter their payments from taxes (such as flexible spending accounts). Since the literature documents considerable sensitivity of workers to these shares, it would be useful to better understand how they are determined and how they vary with policy.

References

Blumberg L., L. M. Nichols and J. S. Banthin (2001), "Workers Decision to Purchase Health Insurance," *International Journal of Health Care Finance and Economics*, 1(3/4): 305-325.

Bundorf, MK. Employee Demand for Health Insurance and Employer Health Plan Choices. *Journal of Health Economics* 21 (2002): 65-88.

Center for Medicare and Medicaid Services (<http://cms.hhs.gov>)

Chernew M., K. Frick and C. G. McLaughlin (1997), "The Demand for Health Insurance Coverage by Low-Income Workers: Can Reduced Premiums Achieve Full Coverage?," *Health Services Research*, 32(4): 453-470.

Cooper P. F. and B. Schone (1997), "More Offers, Fewer Takers for Employment-Based Health Insurance: 1987 and 1996," *Health Affairs*, 16(6): 142-149.

Cutler D. (2002), "Employee Costs and the Decline in Health Care Coverage," Paper prepared for NBER conference on Frontiers in Health Policy Research, May 2002. D. C.

Danzon, Patricia (1989). "Mandated Employer-Based Insurance: Incidence and Efficiency Effects." Working Paper.

Dranove D, Kathryn E. Spier and Laurence Baker, (2000), "'Competition' Among Employers Offering Health Insurance," *Journal of Health Economics*, 19, 121-140.

Employee Benefits in Private Industry 2003, Bureau of Labor Statistics, Washington, Employer Health Benefits Survey, 2003, Kaiser Family Foundation (www.kff.org)

Employer Spending on Benefits 2001, Economic Benefits Research Institute (EBRI), 2000.

Famulari, M. and ME Manser. Employer-provided Benefits: Employer Cost versus Employee Value. *Monthly Labor Review*. December 1989. 24-32.

Farber H. S. and Helen Levy (1998), "Recent Trends in Employer-Sponsored Health Insurance: Are Bad Jobs Getting Worse?" NBER # 6709.

Feldman, R., B. Dowd, S. Leitz, and L. Blewett. (1997). "The Effect of Premiums on the Small Firm's Decision to Offer Health Insurance." *The Journal of Human Resources* 32(4): 635-58.

Garrett, A.B. and S. Glied. "The Impact of the Zebley Decision on Children's Participation in SSI and AFDC." *Journal of Policy Analysis and Management* 19(2) 2000: 275-296.

Gentry, W. M. and E. Peress (1994), "Taxes and Fringe Benefits Offered by Employers." NBER # 4764

- Graff Zivin, J. and S. Glied. In progress. Employer Decisions to Offer Health Insurance. Columbia University.
- Goldstein and Pauly (1973), "Group Health Insurance as a Local Public Good," *The Role of Health Insurance in the Health Services Sector*, R. N. Rosett (ed), NBER.
- Gruber J. (1994), "The Incidence of Mandated Maternity Benefits," *American Economic Review*, 84(3): 622-641.
- Gruber J. (2000), "Tax Subsidies for Health Insurance: Evaluating the Costs and Benefits," NBER # 7553.
- Gruber J. and E. Washington (2003), "Subsidies to Employee Health Insurance Premium and the Health Insurance Market," ERIU Working Paper 15, University of Michigan, MI.
- Gruber J. and M. Lettau (2000), "How Elastic is the Firm's Demand for Health Insurance," NBER #8021.
- Gruber, J. and B. C. Madrian (2002). "Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature." NBER Working Paper No. w8817.
- Hadley, J. and J. D. Reschovsky (2002). "Small Firms' Demand for Health Insurance: The Decision to Offer Insurance." Inquiry 39(Summer): 118-137.
- Hellerstein, J. K. and D. Neumark. Production Function and Wage Equation Estimation with Heterogeneous Labor: Evidence from a New Matched Employer-Employee Data Set. NBER Working Paper 10325. February 2004.
- Helms, D., A. Gouthier and D. Campion. (1992). "Mending the Flaws in the Small-Group Market." Health Affairs: 7-27.
- Ketsche, P. and William Custer. The Effect of Employer-Based Health Insurance on Wages and Returns to Tenure: Evidence that High- and Low-Wage Workers Differ. Paper Presented at AcademyHealth Annual Research Meeting, San Diego, CA, June 6, 2004.
- Ketsche, PG An Analysis of the Effect of Tax Policy on Health Insurance Purchases by Risk Group. Journal of Risk and Insurance 71(1): 91-113 2004.
- Leibowitz, A. and M. Chernew (1992). "The Firm's Demand for Health Insurance," In U.S. Department of Labor, Health Benefits and the Workforce (Washington: U.S. Government Printing Office).
- Levy, H. and R. Feldman. 2001. Does the Incidence of Group Health Insurance Fall on Individual Workers? International Journal of Health Care Finance and Economics 1 (3-4): 227-247.

- Levy, Helen. 1998. Who Pays for Health Insurance? Employee Contributions to Health Insurance Premiums. NCPR Working Paper 222 – 04-10-2001.
- Marquis, M. S. and S. H. Long (2001). "To offer or not to offer: the role of price in employers' health insurance decisions." Health Services Research **36**(5): 935-58.
- Monheit A. C. and J. P. Vistnes (1999), "Health Insurance Availability at the Workplace: How Important are Workers Preferences," *Journal of Human Resources*, 34(4): 770-785.
- Morrissey, M. A.. 2001 Why Do Employers Do What They Do? Compensating Differentials. International Journal of Health Care Finance and Economics 1, 195-201.
- Morrisey, M. A., G. A. Jensen R. J. Morlock. (1994). "Small employers and the health insurance market." Health Affairs **13**(5): 149-209.
- Olson C. A. (2002), "Do Workers Accept Lower Wages in Exchange for Health Benefits," *Journal of Labor Economics*, 20(2), S91-S114.
- Pauly, M. and B. Herring (1999). Pooling Health Insurance Risks. Washington, DC, American Enterprise Institute Press.
- Royalty A. B. (2000), "Do Minimum Wage Increases Lower the Probability that Low-Skilled Workers Will Receive Fringe Benefits?" Manuscript, Indiana University-Purdue University Indianapolis and Stanford University.
- Royalty, A. (2000). "Tax preferences for fringe benefits and workers' eligibility for employer health insurance." Journal of Public Economics **75**(2): 209-227.
- Sheiner L. (1999), "Health Care Costs, Wages and Aging," Federal Reserve Board of Governors #99-19, Washington D. C.
- Shore-Sheppard, L. T. Buchmueller, G. Jensen. (2000) "Medicaid and Crowding Out of Private Insurance: A Re-examination Using Firm-Level Data," Journal of Health Economics, vol. 19, no. 1, January 2000, pp. 61-91.
- Simon K. I. and Robert Kaestner (2003), "Do Minimum Wages Affect Non-Wage Job Attributes? Evidence on Fringe Benefits and Working Conditions," NBER # 9688.
- Simon, K. Involuntary Job Change and Employer Provided Health Insurance; Evidence of a Wage/Fringe Benefit Tradeoff? International Journal of Health Care Finance and Economics 1 (3-4): 249-271.
- Tiebout, Charles. "A Pure Theory of Local Expenditures," *The Journal of Political Economy*, Vol. 64, No. 5. (Oct., 1956), pp. 416-424.

Thorpe, K., A. Hendricks, D. W. Garnick, K. Donelan and J. P. Newhouse. (1992). "Reducing the Number of Uninsured by Subsidizing Employment-Based Health Insurance." JAMA **267**(7): 945-48.

Vanness, DJ and BL Wolfe. 2002. Government Mandates and Employer-Sponsored Health Insurance: Who is Still Not Covered? International Journal of Health Care Finance and Economics 2: 99-135.

ⁱ It may also be an important exclusion if the offer rule is discontinuous in the preferences of the average employee.

ⁱⁱ In US data, the average wage by detailed industry is very closely predicted by about 60 percentweight of the median wage plus 30 percent of the 90th percentile wage.

ⁱⁱⁱ The independence of these decisions depends on there being no limits on the degree to which health insurance costs can be passed along in the form of reduced wages. In practice, modelers rarely incorporate such limits into their models.

^{iv} We thank Michael Chernew for this suggestion.

^v Bundorf (2002) also finds evidence that heterogeneity within the firm affects firm offering decisions.

^{vi} In a competitive labor market, firms will bid up total compensation until it is equal to the value of the marginal product of labor and the incidence of fringe benefits will equal their cost.

^{vii} Alternatively, if the provision of benefits differentially attracts certain types of workers to the firm or affects quit and layoff decisions, a firm might even be willing to accept a tradeoff of less than \$1 in wages for \$1 in added benefits.

^{viii} This literature also provides little guidance in the situation where firm preferences also figure into the offer decision. This could occur if, for example, the provision of health insurance reduced absenteeism, or if workers attracted by job-based coverage were more stable and had longer expected tenure than those who did not value this benefit. Similarly, firm-level factors, such as the firm's profitability, may affect its ability to bear unanticipated shocks in health insurance costs. Although these factors may be important in determining the firm's decision to offer coverage, they are not likely to be affected by the types of policy changes we are considering here.

^{ix} Some coverage losses might occur. If enough people became eligible for the tax credit, the size of firms made up entirely of workers who desired job-based coverage might fall, raising the cost of coverage in these firms, and leading to a decline in offers. Moreover, if incidence is different in groups than in the non-group market, adverse selection gains might disappear as people left groups.

^x The problem of modeling the provision of local public goods when mobility is limited has also been a concern in the literature on local government.

^{xi} Goldstein and Pauly focus on variations in the quantity of health insurance, not the dichotomous offer question.

^{xii} Note also, that the value of job-based coverage may also vary according to the rules associated with paying for that coverage. In particular, coverage that provides long-run protection against changes in health status may be valued more highly than coverage that adjusts taxes with health status.

^{xiii} Note that several other theories are consistent with this finding as well. For example, union leaders may prefer non-wage compensation because it ties workers more effectively to the firm and to the union.

^{xiv} Again, this implies no individual incidence, since if there were individual incidence and tax-sheltered premium contributions, the form of payment would not affect participation.

^{xv} It does not measure the degree to which workers do not receive coverage of optimal generosity.

^{xvi} Analyses of how health insurance offer decisions vary with profitability and with profit taxes may contribute to our understanding of how these costs are allocated between owners and workers in the short run.

Table 1: Elasticities of Firm Offer and Corresponding Offer Responsiveness Among Non-Offering Firms (or Workers in Such Firms)				
Source	Reported Elasticity	Estimated Percent Not Offered in Pop.	Take-up Rate Among Non-Offering Firms	Universe
Hadley & Reschovsky (2002)	-0.54	51.0%	1.06	Small-Firms in 1996
Marquis & Long (2001)	-0.14	52.7%	0.27	Small-Firms in CO, FL, MN, NM, NY, ND, OK, OR, VT, WA (1993)
Gruber & Lettau (2000)	range -0.31 to -0.41	22.3%	1.39 to 1.84	Full-Time Workers (1983-95)
Royalty (1999)	-0.63	32.4%	1.94	Full-Time, Non-elderly adult workers (1988 & 1993)
Feldman et al. (1997)	single cov: -3.91 family cov: -5.82	54.5%	single cov: 7.17 family cov: 10.68	2,000 Small-Firms in MN (1993)
Gentry & Peress (1994)	-1.8	30.3%	5.94	Blue-Collar U.S. workers (88-92)
Morrisey, Jensen, and Merlock (1994)	-0.92	51.0%	1.8	Small-Firms in 1993
Helms, Gauthier, and Campion (1992)	range by state: -0.1 to -1.1	50.5%	0.2 to 2.18	Small-Firms in AL, AZ, CO, FL, ME, MI, UT, TN (1991)
Leibowitz & Chernew (1992)	to premiums: -0.8 to subsidies: -2.9	28.8%	to premiums: 2.78 to subsidies: 10.07	U.S. Workers (1989)
Thorpe et al. (1992)	range -0.07 to -0.33	63.4%	0.11 to 0.52	Small-Firms in Albany, Poughkeepsie, and Brooklyn (1988)

We estimated the percentage not-offered based on the characteristics of the population studied in each paper using a matched February–March 2001 Current Population Survey.

<p>Table 2: Elasticities of Response to Employee Premiums and Implied Take-Up Responsiveness for Uninsured Population</p>
--

Source	Reported Elasticity	Estimated Uninsured Percent in Pop.	Implied Take-Up Responsiveness for Uninsured Population
Gruber & Washington (2003)	-0.02	6.41%	-0.31
Chernew, Frick & McLaughlin (1997)	-0.07	26.1%	-0.25
Blumberg, Nichols & Banthin (2002)	-0.04	26.1%	-0.15

We estimated the percentage uninsured based on the characteristics of the population studied in each paper using the March 2003 Current Population Survey.