

NBER WORKING PAPER SERIES

HOW ELASTIC IS THE FIRM'S
DEMAND FOR HEALTH INSURANCE?

Jonathan Gruber
Michael Lettau

Working Paper 8021
<http://www.nber.org/papers/w8021>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2000

We are grateful to Kokkeong Pua, Detelina Vasileva, and Suzanne Webster for research assistance, to Dan Feenberg for assistance with the NBER TAXSIM model, and to Doug Holtz-Eakin, Alan Krueger, Bruce Meyer, Jim Poterba, and seminar participants at MIT, the NBER, Harvard, and the Society of Government Economists meetings for helpful comments. Gruber acknowledges funding from the National Institute on Aging and the National Science Foundation. The views of the paper are those of the authors and do not necessarily reflect the view of the National Bureau of Economic Research, the U.S. Department of Labor or any of its staff members.

© 2000 by Jonathan Gruber and Michael Lettau. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How Elastic is the Firm's Demand for Health Insurance?

Jonathan Gruber and Michael Lettau

NBER Working Paper No. 8021

November 2000

JEL No. H51, I18

ABSTRACT

We investigate the impact of tax subsidies on the firms decision to offer insurance, and on conditional firm spending on insurance. We do so using the micro-data underlying the Employee Compensation Index, which has a major advantage for this exercise: the matching of very high quality compensation data with information on a sample of workers in the firm. We find that, overall, there is a modest elasticity of insurance offering with respect to after-tax prices (elasticity of -0.31 to -0.41), but a larger elasticity of insurance spending (elasticity of -0.66 to -0.99). We also find that the elasticity of offering is driven solely by small firms, for whom the elasticity is much larger, but that spending is more elastic in large firms. We provide some evidence on how the aggregation of worker preferences determines benefits provision decisions. In particular, we find evidence to support a median voter model of benefits determination, along with some additional influence for the most highly compensated workers in the firm. Our simulation results suggest that major tax reform could lead to an enormous reduction in employer-provided health insurance spending.

Jonathan Gruber

Department of Economics

MIT

E52-355

50 Memorial Drive

Cambridge, MA 02142

and NBER

gruberj@mit.edu

The dominant feature of the health insurance market in the U.S. is the provision of private health insurance through the workplace. In 1998, 91% of the privately insured non-elderly population, representing 65% of the total non-elderly population, received their insurance through the workplace (EBRI, 2000). There are a number of reasons why health insurance may naturally be provided through the workplace in the U.S. There may be substantial economies of scale in administering insurance which increase the value of pooling mechanisms. Workplaces provide a natural pooling mechanism along dimensions largely exogenous to health status. And, finally, the U.S. tax code subsidizes health insurance purchase through the firm relative to the non-group market by excluding the value of that insurance from an individual's income, a tax exclusion estimated to cost more than \$100 billion in foregone state and local tax revenues in 1999 (Shiels and Hogan, 1999).

This tax subsidization of employer-provided insurance has been criticized along a number of dimensions. First, the subsidy is regressive, providing the largest tax break for the most well off employed persons. Second, the subsidy subsidizes insurance purchases over purchases of other goods, leading potentially to "overinsurance" and, perhaps, ultimately to excess cost inflation in the health sector. Third, the subsidy may prop up an inefficient structure, in that health insurance provided through the workplace has been shown to induce substantial reductions in job-to-job mobility, or "job lock". As a result, a number of analysts and policy-makers have proposed replacing the current system of tax deductible employer payments for health insurance with individually-based insurance tax subsidies.

While the first of these criticisms is unassailable, assessing the validity of the second and third criticisms requires understanding the impact of the tax subsidy to employer-provided

insurance on the amount and nature of employer-provided insurance. As a result, over the past 25 years a large number of articles have assessed the responsiveness of employer-provided insurance to its after-tax price. Unfortunately, this literature has led to a wide range of estimates, from roughly zero to almost -6.

This wide range likely reflects two important limitations of the previous literature. One is that many articles in this literature have been unable to control for confounding factors correlated with both the after-tax price of insurance and the demand for insurance; this limitation has been addressed to some extent by more recent work. The second is that all of the work in this literature has either used data at the level of the firm, or at the level of randomly selected workers in firms (through individually-based micro data sets). The disadvantage of the first of these is that we do not observe the characteristics of the workers to whom the firm is responding. The disadvantage of the second is that, with any workforce heterogeneity, the firm may not be responding solely to the demand of the randomly-observed worker. While this weakness has been recognized, it has not been effectively addressed, due to the well-known (and well-lamented) absence of data with information on both firm benefits provision and details on a sample of workers in that firm.

In this paper we address these difficulties by drawing on an excellent administrative data source that provides high quality information on insurance offering, insurance spending, and, most importantly, information on the distribution of characteristics of workers each firm. These are the micro data of the National Compensation Survey, which underlie the well known Employment Compensation Index (ECI) measure of inflation. We match to these data, from the

1983 through 1995 period, information on the tax subsidy to insurance faced by these workers. This tax subsidy across workers of different income levels, over time, and across states, providing substantial exogenous variation in the after-tax price of insurance. And we exploit the heterogeneity in the tax subsidy across workers within firms to assess how insurance offering and spending varies with both the characteristics of the median worker in the firm, and the distribution of those characteristics.

We have three findings of interest. First, we estimate that there is a modest response of firm offering, and a more elastic response of firm insurance spending, to after-tax prices. Our central estimates indicate that the elasticity of firm offering with respect to the tax price of the median worker is -0.31 to -0.41 , and that the elasticity of firm spending conditional on offering is -0.66 to -0.99 . Second, we also find that the responsiveness of firm offering is much larger for small firms, while the responsiveness of conditional spending is larger for medium and large firms.

Finally, we provide some pieces of evidence that can help inform a theory of how firms incorporate worker preferences, in this case tax prices, in making benefits decisions. We find in particular that the median voter model provides a more accurate description of benefits provision than does a model using average tax prices. But we also find that the data reject the pure median voter model, in that the standard deviation of tax prices matters as well as the median, suggesting that intensity of preferences are weighted to some extent in making insurance decisions. We find some evidence that this reflects extra weight on the preferences of the highest paid (lowest tax price) workers, particularly in decisions on conditional insurance spending.

Our paper proceeds as follows. In Part I, we review the literature on the elasticity of employer-provided health insurance, highlighting the value added from our approach. Part II introduces our data source, and discusses the empirical strategy. Part III presents our results using the median tax price. Part IV then explores how the heterogeneity of tax rates matters for firm health insurance decision-making. Part V presents simulations of the impact of tax reform on insurance offering and spending, given our results. Part VI concludes.

Part I: Previous Literature

The previous literature on the elasticity of employer-provided insurance has proceeded along two parallel tracks, one focusing on the elasticity of insurance offering, and the other focused on the elasticity of insurance spending. We review each in turn.

Elasticity of Insurance Offering

There are several approaches to estimating the elasticity of insurance offering by firms. The first is to use variation in the premiums faced by firms to identify their price sensitivity. Two examples of this work are Feldman, Dowd, Leitz, and Blewett (1997), who use information from 1993 for a sample of small firms in Minnesota to estimate price elasticities of -3.9 (single coverage) to -5.8 (family coverage); and Marquis and Long (1999), who use data from 1993 for 10 states to estimate a much smaller price elasticity of only -0.14. A key problem with this approach, however, is that one only observes premiums for the firms that do offer insurance, and they must be imputed to firms that do not. Thus, instruments must be found that are correlated

with the price of insurance but not firm demand, and previous articles have not used firm characteristics that are likely to meet this criterion (e.g. whether the firm is unionized).

The second approach is to use variation in taxation to identify the price elasticity of offering, in essence asking whether those firms with higher tax-related subsidies to insurance purchase are more likely to offer insurance. Leibowitz and Chernew (1992) use variation in tax rates across states to examine the impact of after-tax prices on insurance offering of small firms, as well as using variation in premium quotes across locations obtained from small group insurers. They separately estimate the response to premiums and subsidies, and obtain an elasticity of between -0.8 (premiums) and -2.9 (subsidies). Royalty (1999) also uses cross-state variation in marginal tax rates to estimate an elasticity of firm insurance offering across all employers of -0.63. Gentry and Peress (1994) study cross-city differences in the average share of workers offered health insurance benefits, as a function of cross-state differences in after-tax prices of insurance. They find that for each percentage point increase in the price of health insurance, the percentage of workers covered by employer-provided insurance declines by 1.8 percent, which implies a price elasticity of demand of -1.8.

These types of studies have the advantage that differences across cities and states in tax rates should be independent of insurance offering decisions. But they may not be entirely independent: cities and states with substantial taste for insurance may be the ones that offer the largest tax breaks, which would lead to a strong relationship between price and offering. This criticism is addressed in recent work by Finkelstein (1999), who studies the removal of the large (25%) tax subsidy to supplemental private health insurance in Quebec in 1993, and finds an

elasticity of -0.42 to -0.54 for employer offering. But it is somewhat unclear how to apply the elasticity of offering of supplemental insurance for a national health insurance scheme to the decision of U.S. firms to offer full private health insurance plans.

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

A final approach is to use responses of firms to hypothetical questions about changes in the price of insurance. Morrissey, Jensen, and Merlock (1994) use the response to such hypothetical questions to estimate a price elasticity of insurance offering among small firms of -0.92. But it is unclear whether firms respond in the same way to hypothetical question as they do when faced with an actual insurance purchase decisions.

A final issue with this literature is that, with firm level data, one does not observe the characteristics of the employees to which the firm is responding in making its benefits decisions. This is important because the after-tax price of insurance for these employees depends on their characteristics. Consider two firms of identical size in different states, one of which has only low wage workers, and one of which has only high wage workers. By using the state tax rate on the average worker, an analysis using cross-state tax differences will misstate the price impacts due to progressive state tax structures. Similarly, analyses that use premium differences across firms will misstate the effective price differences by using the tax rate for some representative worker. For example, the price difference between a low wage firm with low premiums and a high wage firm with high premiums is much smaller than would be captured by assuming constant tax rates in both firms. This is not surmounted by the approach taken in Royalty (1999) of looking at the worker level. In this case, the characteristics of a randomly selected worker are well known, but this worker may not well represent the set of relevant workers for the firm's fringe benefit decision.

This concern raises a deeper limitation of the previous literature: it does not assess the importance of worker heterogeneity within the firm. A central question for understanding firm benefits provision is the appropriate underlying model of how worker preferences are aggregated to inform the firm's decision; who is (are) the marginal worker(s)? We address this issue further below.

One article which attempts to address both the identification and "marginal worker" issues is Gruber and Poterba (1994). They study how the self-employed responding to the Tax

Reform Act of 1986, which introduced a subsidy to the insurance purchases of the self-employed.

This “natural experiment” provides exogenous variation in the after-tax price of insurance.

Moreover, for the self-employed, there is no issue of deciding who is the marginal worker. They find significant increases in the insurance coverage of the self-employed relative to the employed over this period, with an implied price elasticity of as large as -1.8. Unfortunately, however, it is unclear how generalizable these results are to firms, who must aggregate the preferences of all their workers in making benefits decisions.

Price Elasticity of Demand for Insurance Spending

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

A second set of studies, including Phelps (1973), Taylor and Wilensky (1983), Woodbury (1983), Holmer (1984), and Sloan and Adameche (1986), analyze cross-sectional data on individuals or firms and ask whether those with higher tax-related subsidies to insurance

purchase spend more on insurance coverage. But a potential problem with these studies is that differences across individuals in their tax rates arise in part from differences in the underlying behavior of individuals or firms, such as differences in labor supply, family structure, or the nature of the workforce. It is impossible to tell whether differences in observed insurance coverage are due to taxes or these behavioral differences. A wide range of estimates emerge from these studies; Pauly (1986) summarizes the consensus range as -0.2 to in excess of -1.0.¹

Moreover, once again, all of these studies suffer from the major problem noted above: they either use the characteristics of some average or representative worker in a firm, or they select random workers and assume that they represent the firms they work for. We address these concerns below.

Part II: Data and Empirical Strategy

Data

The data used for the empirical analysis are from the Bureau of Labor Statistics' Employment Cost Index (ECI). The ECI measures three-month changes in total compensation per hour for U.S. workers. Thus, the ECI micro data contain detailed information on employers' costs for nonwage as well as wage compensation, including the employer's cost for health insurance. All private workers

¹Woodbury and Hamermesh (1992) attempt to overcome these types of problems by analyzing all fringe benefit expenditures around the Tax Reform Act of 1986 in a panel data set of colleges and universities, using the variation in federal and state taxes in this era. They conclude that tax reform substantially reduced the demand for fringes, with an estimated elasticity in excess of -2. But this is not focused on health insurance spending per se, so it is difficult to disentangle the impact on health insurance.

outside of agriculture are in the ECI universe, as are workers in state and local government. The data used here are for June of each calendar year from 1983 through 1995, with values multiplied by annual hours when necessary to represent annual measures.

The ECI is a survey of jobs within establishments. For each establishment, the ECI collects data for multiple jobs, usually four, six, or eight depending on the establishment's size. A job refers to the most detailed level of job recognized by the establishment. Within each establishment, the jobs sampled are chosen randomly with the probability of selection proportionate to the job's employment. The ECI micro data are average data for the group of employees who hold the sampled job; for example, earnings are the average earnings for all workers in that job. Our data set contains 203,836 jobs from 48,605 establishments. By year, the sample size ranges from about 10,000 jobs in 2,500 establishments for the earlier years to nearly 20,000 jobs in 5,000 establishments for the later years. Although the ECI includes part-time jobs, we restrict our analysis to full-time jobs, which are defined as jobs with annual hours worked of at least 1,500. We do this in an attempt to restrict our analysis to eligible workers, because if the sampled worker for a given job is part time and thus ineligible for insurance, the job will be coded as not offered insurance.

Although the ECI data have a number of advantages for our purposes, they have one important weakness: we do not have any data on worker demographics, non-earned taxable income, or other features of tax responsiveness (e.g. whether the worker itemizes). Therefore, we augment the ECI data with data on individual workers from the Current Population Survey (CPS), and with data on family taxes from the Statistics of Income data (SOI). Effectively, we treat the average data for each job in the ECI as if they represent the earnings and employer's cost

for health insurance of an individual worker. We then impute tax rates to the “workers” in our sample in five steps.

First, for each ECI observation, we create four states of the world: married, non-itemizer; married, itemizer; single, non-itemizer; and single, itemizer. Second, we impute to each observation in each state of the world the other elements of taxable income using demographic and income information from the CPS and itemization information from the SOI. For the CPS we use data on industry, occupation, wages, and state to impute spousal earnings (for the two cases where the “worker” is married) and non-earned taxable income (using separate models for those married and single). For the SOI, which is a dataset with information on all elements of taxable income and taxes paid for a nationally representative sample of workers in each year, we use information on earnings and state of residence to impute the average amount itemized, conditional on itemizing.²

An important consideration for such a match is that wages will reflect firm spending on health insurance. Gruber (2000) reviews the literature on the incidence of health insurance costs and concludes that there is strong evidence for full, or close to full, shifting to wages. Thus, by using observed wages in the CPS and ECI data, we would systematically mismatch other income as the cost of health insurance varied.³ At the same time, we don’t know the cost of health

²We actually divide the SOI data by total earnings, so we first impute other earnings on to the ECI observations from the CPS, and then do the matching to SOI information.

³For example, suppose that non-labor income is positively correlated with true marginal product, and that reported wages equal true marginal product less the cost of health insurance for those covered by insurance. In this case, for any wage cell the mean of non-labor income will average the non-labor income that corresponds to that level of marginal product (for those workers who don’t have health insurance) with the higher non-labor income that corresponds to

insurance facing each worker in the CPS data. Therefore, to workers in the CPS covered by health insurance from their job, we match the average cost of health insurance in their industry/region/year cell from the ECI. We then do a “full earnings match” to the ECI data, using reported wages plus imputed health insurance costs in the CPS to create compensation categories, and creating corresponding compensation categories in the ECI using wages plus actual health insurance costs. While mitigating measurement error, this solution of course introduces a natural endogeneity bias into our measure of tax rates, since health insurance costs is our dependent variable; this is one reason for the instrumental variables strategy that we pursue below.

For each observation in each state of the world, we then compute the marginal tax rate. We do this using the NBER’s TAXSIM model, which inputs information on the major elements of taxable income and computes both a federal and state marginal tax rate.⁴ We then use the CPS and SOI data to impute the odds of being married (the CPS) and of itemizing (SOI). Finally, we compute a weighted average marginal tax rate by taking the odds of being married and itemizing as the weights to average the tax rates from these four states of the world.

For the marginal decision on insurance spending, it is the worker’s actual marginal rate on the next dollar of earnings which is relevant. But, for the decision to offer insurance, it is the

that level of marginal product plus the cost of insurance (for those workers who do have insurance). Thus, there will be systematic measurement error in the assignment of non-labor income: those with health insurance will systematically be assigned non-labor income which is too low, and those without health insurance will be systematically assigned non-labor income which is too high.

⁴For more information about TAXSIM, see Feenberg and Coutts (1993). A public use version of TAXSIM is available at www.nber.org/taxsim.

marginal rate over the next X dollars which is relevant, where X is the amount that the firm will spend on insurance. With a single bracket linear tax, there would be no distinction between these concepts. But, with a progressive tax, they may differ, particularly before TRA 86 when the federal income tax system had 14 brackets. We therefore compute the marginal rate for the firm's spending decision as the rate over the next \$1,000 dollars. However, the marginal rate for the offering decision equals the rate over the next X dollars, where X is the average employer cost for health insurance among jobs with a plan in the same year and census division.

Our key regressor is the tax-price of insurance for each worker, which is computed as:

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

where τ_f is the federal income tax marginal rate; τ_s is the state income tax marginal rate; τ_{ss} is the marginal payroll tax rate for the OASDI program; and τ_{MC} is the marginal payroll tax rate for the Medicare HI program.⁵ We differentiate the latter two programs because, beginning in the early 1990s, the taxable maximum for the HI program was increased above that for the OASDI program (and was eventually removed altogether); the marginal rate is zero above the taxable maximum for payroll taxation.

Our dependent variables are insurance offering or spending. For the former, we use a dummy which measures whether the firm offers health insurance to any of its workers. A small share of firms in our data set report offering insurance to some full time workers and not to

⁵The reason that the payroll tax rate is additive in the denominator is that the employer is indifferent between purchasing one dollar of benefits or paying wages of $1/(1 + \tau_{ss} + \tau_{mc})$, since each dollar of wages requires a payroll tax payment as well.

others, but the reason for this occurring is unclear, since (due to anti-discrimination statute) health insurance is always offered to all eligible workers, and non-eligible workers are not supposed to be coded as not offered. But our results for insurance offering are relatively insensitive to how we treat firms that offer to only some full time workers and not others. For spending, we use the log of annual spending on health insurance, averaged over all jobs in the firm.

Table 1 presents the means for our sample.⁶ On average, 91% of the workers in our sample are offered health insurance, with an annual average cost of \$2,037 (in \$1993 and including zeros for no plan). The average median tax price in our sample is 0.644. The typical worker has annual earnings of \$23,215 (in \$1993). More than two-fifths of the firms in our sample have fewer than 100 employees, and about one-fifth have 1000 or more employees. Although the ECI intends to collect data for at least four jobs per establishment, about 42 percent of our sample are from establishments with three or fewer observations. This is due in part because we exclude part-time jobs, but it's also because the ECI does not resample if it selects the job more than once. Following the ECI, we weight these multiple hits by their frequency in calculating firm averages.

The high insurance offering rate in our data is notable. But this reflects the facts that (a) our sample is restricted to full time workers, (b) our sample is implicitly employee-weighted, rather than establishment-weighted, and (c) we assume by construction that if a worker

⁶Although the data are already self-weighted to represent the universe of workers, our means are weighted by ECI weights designed to replicate benchmarked totals. We do not weight the regressions below since the weights are simply a function of industry and firm size, which are

participates in health insurance on the job, all workers participate. Employee-weighting is appropriate if the goal is to assess the impact of tax changes on the overall health insurance coverage of the U.S. workforce, and we are unable to incorporate employee takeup due to data limitations.

Accounting for these factors, our estimates are very similar to the widely cited estimates from the BLS' Employee Benefits Survey (used, for example, in EBRI (2000)). The EBS for 1993 reported that, among medium and large firms (with 100 or more employees), 81.9% of workers are covered by health insurance. If these estimates are corrected to assume full takeup by employees, the resulting estimate is 94.5%, which is close to the ECI estimate for medium and large firms in 1993 of 95.9%. For small firms (fewer than 100 employees) for 1994, the EBS adjusted estimate is 71.4%, and our corresponding estimate from the ECI is 73%, once again very close.

Our results for insurance offering are also very similar to those obtained from a worker-based data set, the 1988 and 1993 Employee Benefits Supplements to the Current Population Survey. In 1988, the CPS estimates that 81.6% of workers in small firms worked in firms where insurance was offered; our comparable estimate for 1988 is 83.6% (which is somewhat larger than above because the concept is now whether one's firm offers insurance at all, and not necessarily whether that offer covers one's job). In medium and large firms, the CPS reports an offering rate of 97.9%, and we estimate 98.4% in the ECI data. Similarly, in 1993, the CPS figure for small firms is 79.5% compared to our ECI estimate of 82.3%; for medium and large

controlled for in the regressions.

firms in that year, the figure in the CPS is 97.2% and in the ECI is 96.8%.

Empirical Strategy

We use the data described above to estimate models of the form:

$$INS_{ijt} = \alpha + \beta_1 TP_{ijt} + \beta_2 X_{ijt} + \beta_3 \delta_j + \beta_4 \tau_t + \varepsilon$$

where i indexes firms, j indexes states, and t indexes years

INS is either insurance offering or log of insurance spending

TP is the tax price, as defined above

X is a set of firm-level controls (see below)

δ_j and τ_t are a full set of state and year dummies, respectively

This model controls in a non-parameteric way for the factors that enter into the construction of the tax rate, which may also be determinants of the demand for insurance by the firm, by including a full set of dummies for industry, occupation, firm size, number of sampled workers, state, year, and earnings; in each case, we use dummies for each category that was used to match the data to the CPS for imputing taxable income.⁷ Of course, with earnings and occupation, the controls must be made appropriate to the tax measure used. So, for example, if we use the median tax rate in the firm, we control for the earnings and occupation categories of the median “worker”. In addition, in all models, we also control for the averages of the categorical variables for earnings and occupation, to capture the average characteristics of the firm. Because the mean of the offering variable is very close to one, we estimate this model as a probit.

⁷In particular, we used twelve industry categories, nine occupation categories, 10 earnings categories, a full/part time (greater than or equal to 35 hours vs. less than 35 hours), and 51 states (including Washington, D.C.).

Even in this framework, however, one might be concerned that interactions of the variables used to create the tax measure (hours, industry, occupation, earnings, state, and year) could be simultaneously correlated with the demand for insurance, biasing our estimates in an indeterminate manner. Moreover, as is clear from the above description, there is considerable measurement error in our computed tax prices, which could bias our estimates towards zero. And there is direct endogeneity bias to our tax measure due to matching CPS characteristics based on health insurance coverage and spending levels.

To address these issues, we instrument the tax price with a “simulated” tax price, which is computed as a function only of state, year, and ten earnings groups (corresponding to the ten deciles in each year’s earnings distribution). In particular, we first divide our sample into deciles of earnings, and then compute the average for each decile of earnings and other elements of taxable income. We then assign these ten averages to each state, and compute marginal income and payroll tax rate for each of the 510 earnings group*state cells in each year.⁸ These tax rates are then used to compute a simulated tax price which instruments the actual tax price in our regressions. This instrument varies only by state, year, and earnings group, as well as interactions of these variables. Thus, by including full sets of dummies for the earnings categories, states, and years, we are clear that the source of identification of our model is only interactions of states, earnings groups, and years.

These interactions provide substantial variation in tax rates, due to three sources. The first is differences across states at any point in time in the progressivity of their tax structures.

⁸We also use the earnings group*state average marriage probabilities and itemization

The second is changes over time in the structure of federal taxation which has differential impacts along the income distribution. And the third is changes in state taxation, both on average and relatively across the income distribution. Of course, it is possible that some of these changes over time are correlated with insurance demand. We explore below the sensitivity of our findings to controls for changes in the income distribution or in average state characteristics over time.

Finally, once again, a potential identification issue arises from the fact that health insurance costs are shifted to wages. By using observed wages to create our instrument, we potentially bias our estimates against finding a responsiveness of health insurance to taxation: as health insurance spending rises, wages fall, lowering tax rates. At the same time, in creating the instrument, we can't add the firm's cost of health insurance directly back into wages to address this problem, since this would make the instrument a direct function of the dependent variable, and might therefore worsen the bias unless shifting is exactly 100%. We therefore explore below the sensitivity of our findings to an alternative instrument which addresses this identification concern, and find similar results, suggesting little bias from wage shifting.

Part III: Results for Median Tax Price

Basic Tax Price Regression

We begin our analysis by considering a single parameterization of the tax price in the firm, using the median tax price across the sampled workers in the firm. That is, we take as the

probabilities to average our four cases.

null hypothesis the “median voter” approach to modeling the demand for fringes in the firm. If there are an even number of observations, we take the average of the two workers around the median. Once again, in this setting we control for the values of the earnings, occupation, and the dummies associated with that median worker, as well as firm averages of these variables. The models are estimated by instrumental variables, as described above; our first stage fit is, as to be expected, very precise, with an F statistic on the excluded instrument of 5,000 - 11,000.⁹

We show the results both for the firm’s offering decision and for the log of annual spending in Table 2. The first row of the table shows the median tax price coefficient, the standard error, the associated marginal probability effect from the probit, and the implied elasticity. Due to the large number of control variables, we split each regression into two columns; the first and third columns show the coefficients from the probit regression on insurance offering, while the second and fourth columns show the coefficients from the log spending regression.

In fact, we find significant negative coefficients for both the firm’s decision to offer insurance, and the amount of spending conditional on offering. The coefficient in the offering equation indicates that, for every percentage point increase in the tax price, the offering rate drops by 0.44 percentage points. At the means of our sample, this implies an elasticity of offering of -0.31 . Although significant, this estimate is towards the bottom of the previous range

⁹For the offering probit, we estimate the model in two steps, first predicting the tax price, and then estimating the probit as a function of the predicted tax price. This results in somewhat understated standard errors as we don’t account for the imprecision in the first stage in our second stage estimate. The bias to our standard errors is trivially small, however, as the first stage fit is so tight; estimates for the OLS case suggest that correcting such bias would raise the

of findings for firm offering elasticities.

We also find that each percentage point increase in the tax price lowers insurance spending by one log point, for an implied elasticity at the mean of -0.66 , which puts our estimate towards the middle of the range estimated by the previous literature. Of course, given the fact that offering is changing simultaneously, this estimate cannot be interpreted purely as the change in spending among those firms offering insurance; it also confounds any changes in the set of firms offering. But if, as seems likely, those firms who drop insurance as the tax price rises are the ones who are spending the least on insurance, then this will bias our estimates in a positive fashion, suggesting that we are understating the response of spending among those firms offering. In any case, such a bias is likely to be small.¹⁰

The table also shows the coefficients on some of the control variables in the regression. We find that insurance offering rises with firm size, as does insurance spending (except for somewhat higher levels in the very smallest firm sizes, perhaps reflecting high loading factors on insurance for the smallest firms). The probability of offering and the level of spending on health insurance increase fairly steadily as the jobs from the higher average earnings categories become more prevalent. Conditional on the distribution of average earnings, however, the distribution of median earnings has a smaller and negative impact on offering. This finding suggests the

standard errors by less than one percent of their value.

¹⁰For example, in 1995, the last year of our sample, the mean spending on insurance among those offering was \$2,773 per year (in \$1993). Suppose that the firms that drop insurance were spending \$2000 per year, considerably less than the average. And suppose further that we were evaluating a 10% change in the tax price, a considerable change for our sample period. Then our dropping effect would only impart a positive bias of 0.09 to our elasticity estimate of -0.64 .

importance of worker heterogeneity within the firm, an issue we will explore more below.

Examining the pattern of industry dummies, we see that health insurance offering is most likely in Durable Goods manufacturing, Finance, Insurance, and Real Estate, Professional Services, and Public Administration; spending is highest in mining (the omitted category) and durable manufacturing. Even conditional on our firm size categories, insurance offering rises by number of jobs sampled, which is likely picking up additional firm size gradations.

Identification Concerns

In Table 3, we address two sources of identification concerns with our findings. The first is any potential bias arising from wage shifting; as noted earlier, if health insurance costs are shifted to wages, then tax rates calculated based on gross earnings (as our instruments are) will be automatically negatively related to health insurance costs, reducing our estimated elasticity; since we control for gross wage categories in the regression, the net bias is unclear. To address this concern, we have reestimated our models using an alternative instrument, which creates tax rates based on earnings plus imputed health insurance costs (“full compensation instrument”). These costs are imputed by taking the average costs in that worker’s industry/region/year cell, averaging over all firms (including those that don’t offer insurance), and matched on to all workers in that cell. In this way, we allow wages in general to reflect variation in the costs of insurance, without building in a direct dependence with the individual firm’s decision to offer insurance or how much to spend on that insurance.

This is an imperfect adjustment, but, as the second row of Table 3 shows, it makes

relatively little difference to our findings; the first row of the table repeats our basic results from Table 2 for comparison purposes. The offering elasticity is essentially unchanged with this new instrument, and the log spending elasticity is 25% larger. The relatively small impact of this adjustment to the instrument should not be surprising given the relatively large income brackets used by state and federal income taxes in most years in our sample. As a result, changing earnings by health insurance costs changes tax rates for a relatively small share of our sample. This finding therefore suggests little bias to our results from shifting to wages.

The second concern raised above is that part of the identification of this model will be coming from changes over time in both the tax rates of workers at different earnings levels, and in state tax systems. But there may be reasons why there are other time series changes that are affecting firm insurance spending and offering differentially by earnings levels and state of residence; for example, the general widening of the earnings distribution over our sample period may have caused differential changes in taste for insurance in high vs. low wage firms. Fortunately, we can deal with this quite straightforwardly in our instrumental variables framework by also including in the model interactions of both our set of state dummies and our (median, since we are using median tax price) earnings controls with linear time trends. This will control for any differential trends in tastes for insurance over time by state or earnings group.

The results of including these additional time series controls are shown in the final row of Table 3. In fact, adding these controls raises both our estimated offering and spending effects. Our offering elasticity rises to -0.41 , and our spending elasticity rises to -0.99 . At the same time, however, our estimates become somewhat less precise, as is to expected since we are

removing a source of identifying variation, so that we cannot reject our basic estimates. Thus, we conclude based on Table 3 that the true offering elasticity is in the range of -0.31 to -0.41 , and the true conditional spending elasticity is in the range of -0.66 to -0.99 .

Interpreting the Spending Results

The data from the ECI on spending provide information only on employer spending on health insurance, on average across all employees. This raises two important issues of interpretation of these spending results. First, some of the reaction of employer spending to tax changes may be shifting the costs of insurance to their employees through premium sharing. As discussed in Gruber and McKnight (2000), there has been a significant rise in employee premium-sharing over this sample period, with the share of employers paying all of the cost of health insurance falling from almost half to about one-third.

In some firms, employee premiums are also paid on a pre-tax basis, and in this case there is no reason why tax changes should induce increased premium sharing, since the decision over who pays premiums is tax neutral. But, in the majority of firms where employee contributions are not made on a pre-tax basis, tax reductions will level the playing field between employee and employer payments, and could lead to a shifting of premium payments to employees. We cannot distinguish in our data true reductions in the generosity of insurance provided to employees from this type of premium shift.

It seems implausible, however, that spending responsiveness of the magnitude that we uncover could be driven solely, or even primarily, by premium shifting to employees. In 1996,

one year after the last year of our sample, firms paid roughly 75% of the costs of insurance on average (Kaiser Family Foundation, 2000). If employees bear only 25% of the costs, then it is difficult to explain an overall spending elasticity of -0.66 completely from shifting of costs to employees. Given that during our sample period about a third of employees made contributions on a pre-tax basis, this would require an elasticity of employee contributions with respect to taxes on the order of -3.8 . In fact, Gruber and McKnight (2000) find no evidence of a responsiveness of employee premium shares to tax rates, supporting the notion that our finding arises from real reductions in firm spending.

The second issue is employee takeup. Our spending data represent average spending on health insurance for firms that offer insurance to any workers. As a result, if some workers decline to takeup coverage, that will lower measured spending per worker (our dependent variable in the spending regressions). We cannot assess directly the takeup issue because we do not have data spending per covered worker, nor statistics on employee takeup.

But, once again, this explanation also seems very unlikely to capture much of the cause of reduced employer spending, since takeup of employer-provided health insurance is roughly 80%. So it would again require an enormous elasticity of takeup for changes in takeup to explain our results. Thus, we conclude that, while these are interesting issues, the bulk of the spending response that we uncover must be real responses of the generosity of insurance per worker covered.

Large vs. Small Firms

If firm offering of insurance is responsive to its tax price, this responsiveness is likely to manifest itself most strongly among small firms. Thus, in Table 4, we examine the variation in our results when we divide the sample by firm size.

In fact, unsurprisingly, we find that our results for the sensitivity of insurance offering to its tax price are largely driven by small firms, those with fewer than 100 employees. For those firms, we estimate a larger elasticity of offering of -0.62 . For medium size firms, we obtain a much smaller and very imprecisely estimated elasticity of only -0.05 . And for large firms we are unable to estimate an elasticity of offering because the mean is so close to one.

On the other hand, we find that the conditional spending on insurance is more sensitive to its price for medium and large firms than for smaller firms. For large firms, the elasticity of spending is -1.1 , while it is only -0.72 for medium size firms and an insignificant -0.15 for small firms. Of course, this finding may reflect selection; as noted above, the responsiveness of offering to price for the smallest firms may bias towards zero our conditional spending estimate. But, once again, the selection resulting from these offering elasticities cannot come close to explaining the difference in spending elasticity that we observe across firm sizes. Putting these estimates together, we find that the overall elasticity of insurance spending, incorporating offering decisions, may not differ much for larger and smaller firms; it is the margin of response which differs most.

Part IV: Exploring Within-Firm Heterogeneity

As discussed above, a key question that has eluded previous research is the role of within-firm heterogeneity in determining the firm's benefits decisions; there is little evidence on how worker preferences are aggregated to determine the provision of health insurance or other workplace amenities. The best discussion of this issue in the context of benefits provision is in Goldstein and Pauly (1976). They conclude that the equilibrium benefits determination could arise in one of two ways. One is through the collective choice of the existing set of workers, through an insider/outsider or union mechanism. In this case, through standard voting arguments, the benefits chosen will reflect the tastes of the median worker. The second is through the choices of employers, whose goal is to minimize their total labor costs, and will therefore design their benefits packages to reflect the average preferences of their workers.

If there is a perfect Tiebout equilibrium across firms, with workers sorted completely by their tastes for insurance, then the average and median tax prices will be everywhere the same and the distinction between these models will not be important. In fact, as Table 1 shows, the mean of the median and mean tax prices are quite similar across the firms in our sample. However, there remains considerable dispersion between these measures within firms; almost 10% of firms have a median and mean tax price which differ by five percentage points or more. If there is imperfect sorting, then these models can have very different implications, depending on the difference between the median and the average. Further complications arise when considering the fact that both mobility and influence within the firm differs across workers. Firms may not consider the average of all workers preferences in making benefits, but may

weight more highly the preferences of either mobile workers, or the preferences of the “most influential” workers in the workplace.

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

We cannot fully distinguish these alternative models with our ECI data. But we can provide some new evidence which sheds light on these different models, by exploring how the structure of tax prices within firms determines benefits decisions. To do so, however, we must restrict our sample to firms with at least two jobs sampled, so that we have sufficient heterogeneity within firms to allow ourselves to identify different moments of the tax price distribution.

We first, following Goldstein and Pauly, consider a “horserace” between the median and average tax prices; if all workers are equally marginal to the firm, then this test distinguishes between their insider/outsider and employer choice models. The first row of Table 5 replicates our basic median tax price regressions from Table 2, but with this restricted sample. The estimated elasticities of insurance offering and spending are somewhat higher than in our basic results. But these estimates are similar enough that our basic conclusion remains: a small

offering elasticity and a larger conditional spending elasticity with respect to the median tax price.

The second row of Table 5 shows that the impact of average tax price is similar to, but slightly smaller than, the impact of the median tax price. When these are entered together in the third row, however, the median tax price is clearly more important; the average tax price is insignificant and positive in both specifications. Although the standard errors are large, in both cases, the coefficient on the median tax price is (marginally) statistically significantly lower than the coefficient on the average tax price. This finding suggests that the insider/outsider voting model fits the data better than does the average of preferences model. It is certainly possible, however, that not all workers are “marginal”, so that we are taking the average over the wrong population.

We therefore consider a second test of the median voter model: including along with the median the standard deviation of the tax price. If the median voter model holds in its strictest form, then such higher moments of the distribution should not matter for benefits determination.

In fact, as the fourth row of Table 5 shows, such a model rejects the purest form of the median voter model. There is a positive and highly significant impact of the standard deviation for insurance offering, and a positive and marginally significant impact for log costs. The magnitude of the standard deviation effect is roughly three-fifths that of the median tax price effect in both the offering and log cost models.

The positive coefficient on the standard deviation is striking, suggesting that more within firm heterogeneity in tax prices leads to more insurance provision. This could be possible, for

example, if the highest tax rate (most highly paid) workers have excess influence on insurance decisions, so that as tax systems become more progressive, there is more insurance provision because the incentives for the highest tax rate workers to buy insurance have increased.

To assess whether this is the explanation for our finding, the final row of Table 5 includes along with the median and the standard deviation the minimum tax price for the firm, which is the tax price of the highest tax rate worker. The results do in fact suggest that our standard deviation finding is driven by the tastes of the highest wage worker. For offering, both the tax price for the median and the minimum tax price are negative and at least marginally significant, and are quite similar in magnitude; but the coefficient on the standard deviation has been driven completely to zero by the inclusion of the minimum tax price. For log cost, the minimum tax price is not itself significant, but its inclusion once again drives the standard deviation coefficient to zero.

Thus, it appears that there is full support for neither of the extremes laid out by Goldstein and Pauly (1976). Neither a pure median voter model, nor an average preferences model seems to explain the patterns in the data. Rather, it appears that a model which weights most heavily the preferences of the median worker, but also considers the dispersion of preferences, is the appropriate one to consider when assessing benefits decisions. And the role of dispersion is at least partially explained by a distinct role for the highest tax rate (lowest tax price) workers.

Part V: Policy Simulations

The findings reported here have important implications for not only how benefits decisions are made, but also for the implications of tax reform for employer provision of health insurance. In this section, we provide some simulation evidence to document the implications of our results for the effects of tax reform on health insurance.

We consider a number of possible reforms. In each case, we use the data from the last year we have available (1995, converted to \$1993) and reestimate our tax price variables under the reformed tax law regime. We then compute, for each firm in our sample, the offering rate and spending level under the old and new regimes. Finally, we take the difference between the baseline and post-reform offering rates and spending levels, and express these differences as a percent of baseline. We consider four different specifications of the tax price model: the median tax price only (from Table 2); the median price only, adding earnings and state time trends (final row of Table 3); the median and standard deviation (fourth row of Table 5), and the median, standard deviation, and minimum (final row of Table 5).

The results of this exercise are presented in Table 6. We present three sets of columns, corresponding to the three reforms we consider. The first row for each set of columns show the percentage change in the tax price induced by the reform. The remaining rows show the percentage change in offering, conditional spending, and total spending (assuming that the firms that reduce offering were spending the mean amount before dropping out), as well as the effects on dollars of total spending.

We first consider the most radical reform: complete inclusion of the health insurance

spending in the individual tax base for both income tax (both state and federal) and payroll tax purposes. This is a very significant change in the tax-price of insurance; as the first row of the table shows, the median tax price rises by over 50%. We find that this has sizeable impacts on both firm offering and insurance spending decisions. Insurance offering falls by 17 to 30%, and conditional spending falls by 30 to 42%. Assuming that those who stop offering were spending the mean amount, this results in a reduction in total spending on insurance of 48 to 67%, or \$120 to \$1669. These enormous effects follow naturally from the radical nature of this change in the price of insurance.

Of course, this simulation involves a projection that is vastly out of sample, so it is clearly possible that the responses would be smaller for such a large change. But it is also possible that the responses would be larger. After all, in our data, there is no variation in offering among large firms; it is possible in a world without a tax subsidy that some large firms would decide to forgo the provision of insurance to their employees.

We next consider a reform which does not tax employer-provided health insurance in the payroll tax base, but does include it in the state and federal income tax bases. This raises the tax price by roughly one-third. Correspondingly, there are smaller, but still quite large, responses of offering and conditional spending: offering falls by 10 to 19%, and conditional spending falls by 20 to 28%, for a fall in total spending of 30 to 43%.

Finally, we consider a reform which is within the variation in our sample, a 10% cut in all income tax rates, which is comparable to recent tax reform proposals of congressional Republicans and Presidential Candidate George Bush. This would raise the tax price by only

3.4%. Offering falls by 1.1 to 1.9%, and conditional spending by 2.2 to 2.7%; total spending falls by 3.3 to 4.7%. Thus, even this very modest tax reform could have a non-trivial impact on the spending of employers on health insurance in the U.S.

Part V: Conclusions

The central behavioral parameter for estimating the impact of fundamental changes to the tax treatment of health insurance is the firm's elasticity of demand for health insurance.

Unfortunately, the large body of past work in this area has been unable to deliver a consistent, and perhaps even a convincing, estimate of this critical parameter. This is largely because of the absence of the “right” data for this exercise, and related difficulties in deriving an empirical framework that can deliver cleanly identified estimates of the relevant elasticities. We have attempted to address this deficiency by using the ECI data base that includes information not only on firm insurance offering and spending, but also data on the distribution of workers within the firm.

Using these new data, we estimate that firms are fairly responsive to the tax price of their median worker in their benefits decisions. We obtain central estimates of the elasticity of firm offering of -0.31 to -0.41 and of conditional insurance spending of -0.66 to -0.99 . We also estimate that small firms are particularly responsive to taxes in their offering decisions, while larger and medium size firms are more sensitive in their spending decisions. We also have some interesting findings on how benefits are determined within firms that suggests that neither the average nor the median tax price is the perfectly correct one to be used for modeling benefits

decisions. The appropriate model appears to be one where the median worker plays a central role, but where the dispersion of tax prices matters as well, at least partly through the specific preferences of the highest tax rate workers.

Our estimates have important implications for tax policy towards firm benefits provision. Most importantly, they suggest that the majority of the impact of changes in tax policy towards health insurance would be on the level of spending on insurance, not on firm provision of insurance. This is an important conclusion because one of the major benefits claimed for removing the tax exclusion is that it would lead to a reduction in inefficiently high spending, while the major concern about this policy is that it could lead to sizeable reductions in insurance coverage. Our findings suggest that the impact on insurance coverage is likely to be small relative to the reduction in spending.

Nevertheless, the reductions in employer-provided insurance offering that we estimate are non-trivial. For a reform that removes the income tax subsidy to insurance (while maintaining the payroll tax subsidy), we estimate that offering would fall by 10-19%. There are currently 155 million workers and dependents covered by employer-provided health insurance (EBRI, 2000). A 10-19% reduction in coverage of this population is therefore 16 – 29 million persons, which amounts to 36 to 66% of the size of the current pool of uninsured. A critical question for future research is whether the persons who lose their employer-provided coverage will find coverage from other sources, such as the non-group market, or through spouses. In addition, one limitation of our work for analyzing policies such as removing the entire exclusion of employer-provided insurance from taxes is that there may be more substantial responses to major policy changes

such as this than there are to the minor variation in tax rates used in our analysis.

It is also important to get inside the black box of “total spending” on insurance and understand which elements of spending are responding so elastically to tax prices. Gentry and Perress (1994), in their cross-sectional study, do find that elements of insurance which are likely to be more elastically demanded, such as dental and vision care, were more elastic with respect to tax price differences. Further work of this nature, particularly focusing on the choice of form of insurance (traditional fee for service vs. managed care), is a high priority.

Finally, it would be useful to understand the corresponding impact of tax policy on the provision of other fringe benefits by the firm. Is the provision of health insurance more or less tax sensitive, for example, than the provision of pension benefits? Ultimately, data such as the ECI data used here can allow us to assess the sensitivity of the entire compensation package to taxation.

References

- Chernew, Michael, Kevin Frick, and Catherine G. McLaughlin (1997). "The Demand for Health Insurance Coverage by Low-Income Workers: Can Reduced Premiums Achieve Full Coverage?," *Health Services Research*, 32, 453-470.
- Cutler, David (1996). "Comment on Gruber and Poterba," in Henry Aaron and William Gale, eds., *Economic Effects of Fundamental Tax Reform*. Washington, D.C.: Brookings Institution, p. 162-167.
- Employee Benefits Research Institute (2000). "Sources of Health Insurance and Characteristics of the Uninsured: Analysis of the March 1999 Current Population Survey". Washington, D.C.: EBRI.
- Farber, Henry, and Helen Levy (1999). "Recent trends in Employer-Sponsored Health Insurance Coverage: Are Bad Jobs Getting Worse?," forthcoming, *Journal of Health Economics*.
- Finkelstein, Amy (1999). "The Effect of Tax Subsidies To Employer-Provided Health Insurance on Workplace Pooling: New Evidence From Canada," forthcoming, *Journal of Public Economics*.
- Feldman, Roger, Bryan Dowd, Scott Leitz, and Lynn A. Blewett (1997). "The Effect of Premiums on the Small Firm's Decision to Offer Health Insurance," *Journal of Human Resources*, 32, 635-658.
- Gentry, William and Eric Peress (1994). "Taxes and Fringe Benefits Offered by Employers," NBER Working Paper #4764.
- Gruber, Jonathan (1999a). "Tax Subsidies for Health Insurance: Evaluating the Costs and Benefits," paper prepared for the Kaiser Family Foundation.
- Gruber, Jonathan (2000). "Health Insurance and the Labor Market," in Joseph Newhouse and Anthony Culyer, eds., *Handbook of Health Economics*, p. 646-706.
- Gruber, Jonathan and Robin McKnight (2000). "Why Are Employee Health Insurance Contributions Rising?," mimeo, MIT.
- Gruber, Jonathan and James M. Poterba (1994). "Tax incentives and the demand for health insurance: Evidence from the self-employed," *Quarterly Journal of Economics* 109, 701-733.
- Gruber, Jonathan, and James Poterba (1996a). "Fundamental Tax Reform and Employer-

- Provided Health Insurance," in *Economic Effects of Fundamental Tax Reform*, Henry J. Aaron and William G. Gale, eds. Washington, D.C.: Brookings Institution, 125-170.
- Gruber, Jonathan, and James Poterba (1996b). "Tax Subsidies to Employer-Provided Health Insurance," in *Empirical Foundations of Household Taxation*, Martin Feldstein and James Poterba, eds. Chicago: University of Chicago Press, 135-164.
- Helms, W. David, Anne K. Gauthier, and Daniel M. Campion (1992). "Mending the Flaws in the Small-Group Market," *Health Affairs*, 11, 8-27.
- Holmer, Martin (1984). "Tax Policy and the Demand for Health Insurance," *Journal of Health Economics* 3, 203-221.
- Kaiser Family Foundation (2000). *Employer Health Benefits 1999*. Washington, D.C.: KFF and Health Research and Educational Trust.
- Liebowitz, Arleen and Michael Chernew (1992). "The Firm's Demand for Health Insurance," In U.S. Department of Labor, *Health Benefits and the Workforce* (Washington: US Government Printing Office).
- Long, James E., and Frank A. Scott (1982) "The Income Tax and Nonwage Compensation," *Review of Economics and Statistics* 64, 211-219.
- Manning, Willard, et al. (1987). "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment," *American Economic Review* 77, 251-277.
- Marquis, M. Susan, and Stephen H. Long (1994). "Worker Demand for Health Insurance in the Non-Group Market," mimeo, RAND.
- Marquis, M. Susan, and Stephen H. Long (1999). "To Offer or Not To Offer: The Role of Price in Employer Demand for Insurance," mimeo, RAND.
- Marquis, M. Susan, and Charles E. Phelps (1987). "Price Elasticity and Adverse Selection in the Demand for Supplementary Health Insurance," *Economic Inquiry* 25, 299-313.
- Morissey, Michael A., Gail A. Jensen, and Robert J. Morlock (1994). "Small Employers and the Health Insurance Market," *Health Affairs*, 149-161.
- Newhouse, Joseph P. (1993). *Free for All? Lessons from the RAND Health Insurance Experiment*. Cambridge, MA: Harvard University Press.
- Pauly, Mark (1986). "Taxation, Health Insurance, and Market Failure in the Medical Economy,"

Journal of Economic Literature, 24, 629-675.

- Royalty, Anne (1999). "Tax Preferences for Fringe Benefits and Workers' Eligibility for Employer Health Insurance," *Journal of Public Economics*, forthcoming.
- Sheils, John, and Paul Hogan (1999). "Cost of Tax-Exempt Health Benefits in 1998," *Health Affairs*, 18 , 176-181.
- Sloan, Frank, and Killard Adamache (1986). "Taxation and the Growth of Nonwage Benefits," *Public Finance Quarterly* 14, 115-139.
- Taylor, Amy, and Gail Wilensky (1983). "The Effect of Tax Policies on Expenditures for Private Health Insurance," In Jack Meyer, ed., *Market Reforms in Health Care: Current Issues, New Directions, Strategic Decisions*. Washington, D.C.: American Enterprise Institute.
- Thorpe, Kenneth L, et al. (1992). "Reducing the Number of Uninsured by Subsidizing Employment-Based Health Insurance: Results from a Pilot Study," *Journal of the American Medical Association* 267, 945-948.
- Turner, Robert (1987). "Are Taxes Responsible for the Growth of Fringe Benefits?," *National Tax Journal* 40, 205-220.
- Vroman, Susan, and G. Anderson (1984). "The Effect of Income Taxation on the Demand for Employer-Provided Health Insurance," *Applied Economics*, 16, 33-43.
- Woodbury, Stephen A., and Daniel S. Hammermesh (1992). "Taxes, Fringe Benefits, and Faculty," *Review of Economics and Statistics* 73, 287-296.

Table 1: Sample Descriptive Statistics

Variable	Mean	Standard Deviation
Plan Offered	.910	.286
Annual Cost	2036.9	1623.2
Log Annual Cost	7.459	.830
Median Tax Price	.644	.051
Annual Earnings	23214.9	16239.7
< 100 Employees	.4399	.4964
100-999 Employees	.3440	.4750
1000+ Employees	.2161	.4116
1 Job Sampled	.1190	.3238
2 Jobs Sampled	.1456	.3527
3 Jobs Sampled	.1589	.3656
4 Jobs Sampled	.1529	.3599
5 Jobs Sampled	.1126	.3161
6 Jobs Sampled	.1001	.3001
7 Jobs Sampled	.0786	.2691
> 7 Jobs Sampled	.1323	.3388
Number of Establishments	48,605	48,605

Table presents means of ECI sample used for insurance offering regression; see text for description of data.

Table 2: Basic Results Using Median Tax Price

	Plan Offered	Log Annual Cost		Plan Offered (con't)	Log Annual Cost (con't)
Median Tax Price	-5.272 (1.873) [-0.443] {-0.314}	-1.023 (0.383)			
			Business Services	-0.035 (0.153)	-0.379 (0.036)
10-24 Employees	0.429 (0.035)	0.005 (0.013)	Personal Services	0.078 (0.164)	-0.443 (0.044)
25-99 Employees	0.867 (0.036)	-0.054 (0.012)	Entertainment	0.482 (0.184)	-0.098 (0.041)
100-499 Employees	1.255 (0.044)	-0.029 (0.013)	Professional Services	0.729 (0.154)	-0.202 (0.034)
500-999 Employees	1.394 (0.082)	0.058 (0.015)	Public Administration	1.599 (0.197)	-0.075 (0.035)
1,000-9,999 Employees	1.745 (0.096)	0.055 (0.014)	2 nd Earnings Decile of Median	-0.053 (0.092)	0.017 (0.028)
10,000+ Employees	1.748 (0.329)	0.055 (0.020)	3 rd Earnings Decile of Median	-0.134 (0.111)	-0.001 (0.029)
Construction	0.012 (0.149)	-0.280 (0.034)	4 th Earnings Decile of Median	-0.514 (0.120)	-0.046 (0.030)
Durable	0.765 (0.162)	0.085 (0.034)	5 th Earnings Decile of Median	-0.485 (0.142)	-0.064 (0.032)
Nondurable	0.407 (0.162)	-0.070 (0.035)	6 th Earnings Decile of Median	-0.719 (0.179)	-0.126 (0.036)
TCPU	0.514 (0.156)	-0.074 (0.033)	7 th Earnings Decile of Median	-1.043 (0.207)	-0.142 (0.040)
Wholesale Trade	0.753 (0.152)	-0.099 (0.034)	8 th Earnings Decile of Median	-1.281 (0.238)	-0.157 (0.043)
Retail Trade	0.436 (0.150)	-0.340 (0.034)	9 th Earnings Decile of Median	-1.785 (0.290)	-0.200 (0.050)
FIRE	0.806 (0.157)	-0.155 (0.034)	10 th Earnings Decile of Median	-0.802 (0.228)	-0.124 (0.037)

Table 2 (continued)

Executive Occupation of Median	-0.014 (0.164)	0.041 (0.020)	Avg. of 7 th Earnings Decile	2.321 (0.191)	1.105 (0.041)
Sales Occupation of Median	-0.054 (0.169)	0.113 (0.027)	Avg. of 8 th Earnings Decile	2.522 (0.215)	1.284 (0.043)
Administrative Support of Median	-0.214 (0.152)	0.019 (0.019)	Avg. of 9 th Earnings Decile	3.080 (0.258)	1.384 (0.045)
Craft Occupations of Median	0.450 (0.166)	0.071 (0.022)	Avg. of 10 th Earnings Decile	2.422 (0.249)	1.497 (0.044)
Operators, Assemblers of Median	0.077 (0.196)	0.107 (0.027)	Average of Executive Occupation	-0.181 (0.194)	-0.008 (0.028)
Transportation Occupations of Median	0.360 (0.195)	0.113 (0.030)	Average of Sales Occupation	0.210 (0.201)	-0.062 (0.036)
Handlers, Laborers of Median	0.095 (0.174)	0.152 (0.027)	Average of Administrative Support	0.138 (0.183)	0.267 (0.028)
Service Occupations of Median	-0.299 (0.174)	0.054 (0.027)	Average of Craft Occupations	-0.866 (0.194)	-0.030 (0.030)
Avg. of 2 nd Earnings Decile	0.685 (0.107)	0.381 (0.036)	Average of Operators, Assemblers	-0.386 (0.234)	-0.064 (0.037)
Avg. of 3 rd Earnings Decile	1.071 (0.127)	0.648 (0.037)	Average of Transportation Occupation	-0.554 (0.229)	0.015 (0.041)
Avg. of 4 th Earnings Decile	1.594 (0.136)	0.801 (0.037)	Average of Handlers, Laborers	-0.494 (0.205)	-0.096 (0.036)
Avg. of 5 th Earnings Decile	1.655 (0.151)	0.908 (0.038)	Average of Service Occupations	-0.277 (0.201)	0.015 (0.033)
Avg. of 6 th Earnings Decile	1.966 (0.175)	1.046 (0.040)	Union	0.898 (0.077)	0.411 (0.009)

Table 2 (continued)

Year = 1984	0.021 (0.079)	0.108 (0.017)	Year = 1994	0.031 (0.082)	0.581 (0.021)
Year = 1985	-0.027 (0.078)	0.139 (0.017)	Year = 1995	0.019 (0.083)	0.507 (0.022)
Year = 1986	0.025 (0.078)	0.164 (0.017)	2 Jobs Sampled	0.429 (0.033)	-0.121 (0.011)
Year = 1987	0.085 (0.095)	0.170 (0.022)	3 Jobs Sampled	0.696 (0.037)	-0.077 (0.011)
Year = 1988	0.108 (0.091)	0.230 (0.022)	4 Jobs Sampled	0.686 (0.043)	-0.044 (0.011)
Year = 1989	0.131 (0.089)	0.306 (0.022)	5 Jobs Sampled	0.888 (0.062)	-0.059 (0.012)
Year = 1990	0.084 (0.084)	0.409 (0.021)	6 Jobs Sampled	0.827 (0.072)	-0.034 (0.013)
Year = 1991	0.055 (0.086)	0.477 (0.022)	7 Jobs Sampled	0.878 (0.101)	-0.024 (0.014)
Year = 1992	0.081 (0.087)	0.537 (0.022)	> 7 Jobs Sampled	1.650 (0.184)	-0.028 (0.014)
Year = 1993	0.082 (0.087)	0.574 (0.021)	Number of Establishments	48,605	44,637

First and third columns show coefficients from probit for insurance offering; second and fourth columns show coefficients from log spending regression. Standard errors are in parentheses, the implied marginal probability effects (for probit regression) are in brackets, and the implied elasticities are in braces. The omitted categories are establishments with fewer than ten employees, establishments from mining, professional occupations, the first earnings decile, year equals 1983, and one job for the establishment in the sample. The regressions also include state dummy variables.

Table 3: Coefficient for Median Tax Price

	Plan Offered	Log Annual Cost
Basic Result	-5.272	-1.023
From Table 2	(1.873)	(0.383)
	[-0.443]	
	{-0.314}	{-0.659}
Full	-5.053	-1.308
Compensation	(1.870)	(0.405)
Instrument	[-0.442]	
	{-0.313}	{-0.842}
Earnings	-6.799	-1.538
Interacted with	(2.709)	(0.485)
Trend and	[-0.577]	
State	{-0.408}	{-0.990}
Interacted with		
Trend		

Plan Offered Equation is Probit; Log Cost is OLS. Standard errors are in parentheses, the implied marginal probability effects are in brackets, and the implied elasticities are in braces. Table shows coefficients on tax price only; regressions include all controls shown in Table 2 and footnote to that Table. First row replicates basic result from Table 2; second row shows effect of using “full compensation instrument” that incorporates predicted cost of health insurance into wage categories used to create instrument; third row shows the effect of including in the regression model interactions of both state dummies and median earnings category dummies with a linear time trend.

Table 4: Results by Firm Size

	Plan Offered	Log Annual Cost
Small Firms (<100)	-4.979 (1.802) [-0.777] {-0.617}	-0.222 (0.687) {-0.145}
Medium Firms (100-999)	-2.234 (5.016) [-0.071] {-0.047}	-1.117 (0.615) {-0.717}
Large Firms (1000+)	---	-1.823 (0.571) {-1.134}

Plan Offered Equation is Probit; Log Cost is OLS. Standard errors are in parentheses, the implied marginal probability effects are in brackets, and the implied elasticities are in braces. Table shows coefficients on tax price only; regressions include all controls shown in Table 2 and footnote to that Table. Because virtually all establishments with at least 1,000 employees offer a health insurance plan, the coefficients for the offering equation are not identified.

Table 5: Exploring Within-Firm Heterogeneity

Table 5: Exploring Within-Firm Heterogeneity						
	Plan Offered			Log Annual Cost		
	Median Tax Price					
Median Tax Price Only	Median			Median		
	-7.269			-1.188		
	(2.360)			(0.398)		
	[-0.480]					
	{-0.340}			{-0.765}		
	Average Tax Price					
Average Tax Price Only	Average			Average		
	-5.849			-1.141		
	(2.357)			(0.463)		
	[-0.374]					
	{-0.263}			{-0.731}		
	Median and Average Tax Price					
Median & Average Together	Median	Average		Median	Average	
	-10.397	4.976		-1.397	0.358	
	(5.269)	(5.835)		(0.720)	(0.865)	
	[-0.738]	[0.244]				
	{-0.522}	{0.172}		{-0.899}	{0.230}	
	Median and Standard Deviation					
Median & Std. Dev Together	Median	Stan. Dev.		Median	Stan. Dev.	
	-8.889	4.778		-1.048	0.589	
	(2.445)	(1.937)		(0.413)	(0.351)	
	[-0.606]	[0.234]				
	{-0.429}	{0.008}		{-0.675}	{0.019}	
	Median, Standard Deviation & Min					
Median & Std. Dev & Minimum	Median	Std. Dev	Min	Median	Std. Dev	Min
	-6.594	0.097	-5.188	-1.115	-0.288	-0.160
	(3.141)	(2.909)	(2.856)	(0.458)	(0.479)	(0.424)
	[-0.421]	[0.005]	[-0.321]			
	{-0.298}	{0.000}	{-0.211}	{-0.718}	{0.009}	{-0.096}

Plan Offered Equation (first set of columns) is Probit; Log Cost (second set of columns) is OLS. Standard errors in parentheses; implied marginal probability effects in brackets, and implied elasticities in braces. Table shows coefficients on moments of tax price only; regressions include all controls shown in Table 2 or mentioned in the footnote to that Table. Sample is establishments with at least two different jobs sampled. Tax price measures included in regression are: first panel, median; second panel, average; third panel, median and average; fourth panel, median and standard deviation; and fifth panel, median, std. dev, and minimum.

Table 6: Policy Simulations

	No Tax Subsidy			No Income Tax Subsidy			10% Income Tax Cut		
	Change in Plan Offered	Change in Annual Cost	Total Change in Annual Cost	Change in Plan Offered	Change in Annual Cost	Total Change in Annual Cost	Change in Plan Offered	Change in Annual Cost	Total Change in Annual Cost
Change in Median Tax Price		54.9%			33.5%			3.4%	
Response of Establishments:									
Median Tax Price	-17.1%	-30.1%	-48.0% (-\$1,201)	-10.4%	-19.5%	-30.4% (-\$760)	-1.1%	-2.2%	-3.3% (-\$83)
Median Tax Price with Linear Trends	-22.4%	-41.6%	-65.0% (-\$1,625)	-13.6%	-27.7%	-42.0% (-\$1,051)	-1.4%	-3.3%	-4.7% (-\$118)
Median & Std. Dev. Together	-24.1%	-31.8%	-56.9% (-\$1,421)	-14.9%	-21.0%	-36.4% (-\$911)	-1.5%	-2.4%	-4.0% (-\$100)
Median & Std. Dev. & Minimum	-29.8%	-35.9%	-66.8% (-\$1,669)	-18.8%	-23.6%	-43.0% (-\$1,076)	-1.9%	-2.7%	-4.7% (-\$117)

Table reports results from policy simulations described in text. The policy simulations are based on the 1995 sample of 4,051 establishments, and the coefficient estimates from the previous tables. Each set of three columns (panel) shows the impacts of a different tax reform. First row in each panel shows effect on median tax price; remaining rows show impacts on plan offering, conditional spending, and total spending, for different specifications of the regression model.