

**Economic Research Initiative on the Uninsured
Working Paper Series**

**WHY DID EMPLOYEE HEALTH INSURANCE CONTRIBUTIONS
RISE?**

Jonathan Gruber, Ph.D.
MIT, NBER

Robin McKnight
MIT

ERIU Working Paper 9
www.umich.edu/~eriu/pdf/wp9.pdf

Economic Research Initiative on the Uninsured
University of Michigan
555 South Forest Street, 3rd Floor
Ann Arbor, MI 49104-2531

Not to be distributed or copied without permission of the authors.

March 2002

We are grateful to Cindy Perry, Kokkeong Pua, and Nannette Yang for research assistance, to Dan Feenberg for assistance with the NBER's TAXSIM model, to Larry Baker for sharing his HMO market share data, and to David Autor and seminar participants at the NBER and the University of Minnesota for helpful comments. Gruber gratefully acknowledges financial support from the Economic Research Initiative on the Uninsured.

Abstract

We explore the causes of the dramatic rise in employee contributions to health insurance over the past two decades. In 1982, 44% of those who were covered by their employer-provided health insurance had their costs fully financed by their employer, but by 1998 this had fallen to 28%. We discuss the theory of why employers might shift premiums to their employees, and empirically model the role of six factors suggested by the theory. We find that there was a large impact of falling tax rates, rising eligibility for insurance through the Medicaid system and through spouses, and deteriorating economic conditions (in the late 1980s and early 1990s). We also find much more modest impacts of increased managed care penetration and rising health care costs. Overall, this set of factors can explain about one-quarter of the rise in employee premiums over the 1982-1996 period.

The dominant feature of the health insurance market in the U.S. is the provision of private health insurance through the workplace. But the past two decades have been a period of substantial reduction in both the scope and generosity of employer-provided health insurance. In 1982, roughly 80% of workers were covered by employer-provided health insurance. By 1998, this had fallen to 73%. Similarly, in 1982, 44% of those who were covered by their employer-provided health insurance had insurance that was fully financed by their employer. But by 1998, this had fallen to 28%.¹

There has been a voluminous literature in recent years on the causes and consequences of the decline in employer-provided health insurance coverage. But there has been virtually no work on the parallel time trend of declining employer payments for health insurance. This is a particularly glaring omission in light of recent evidence which shows that most of the time trend in private insurance coverage appears to be reductions not in employer offering of insurance, but in employee takeup of insurance conditional on offering (Cooper and Schone, 1997; Farber and Levy, 2000). Thus, the key dimension along which employers appear to be adjusting their health insurance spending is through the generosity of what they contribute. Moreover, this raises the possibility that it is reductions in employer generosity that are responsible for declining insurance coverage.

In this paper, we attempt to model the set of factors that may be driving employers to shift their health insurance costs to their employees. We begin by discussing the theory of why employers might shift premiums to their employees. There are two classes of explanations. The first is that employers are shifting premiums in order to induce employees to choose the cost effective option from the range of insurance choices offered by the employer. The second is that premium sharing results from imperfect worker sorting across firms; with heterogeneity in tastes among co-workers, premium contributions become a useful tool for separating worker types. By requiring contributions, the firm can provide insurance only to those who demand it, and can pass the savings back to employees in the form of higher wages.

We then turn to estimating the role of a number of factors which fit into these categories of explanations, most of which also have the attractive feature that they operated most strongly in the late 1980s and early 1990s, the period over which the shift from employer to employee financing was most pronounced. Our summary for the first explanation is the rise in managed care penetration. A key determinant of employer premium sharing should be the range of choices offered by the employer; as there are more managed care options available in a state, then there will

be more incentive to make employees bear insurance costs on the margin in order to motivate them to choose the low cost plan.

The second and third factors are the rise in spousal labor supply and the expansion of eligibility for the public Medicaid insurance program for women and children. A key prediction of the imperfect sorting model is that, as there are more outside insurance options available to workers, firms should increase employee contributions to insurance. Both of these factors represent a rise in such outside options: more spousal labor supply means more opportunity for spouses to be covered by insurance; and more Medicaid entitlement means more chance for coverage by public insurance.

The fourth factor is health insurance costs. In the presence of workplace heterogeneity and imperfect individual-specific wage shifting, rising medical costs will increase the pressure on firms to shift the costs of insurance to their workers. Likewise, the reductions in marginal tax rates through the tax reforms of the 1980s could be playing an important causal role. A central feature of employer-provided health insurance in the U.S. is its subsidization through the tax code. If employees are paid in wages, they must pay taxes on those wages; but, if paid in health insurance, it is tax free. Since employee contributions for health insurance are usually made in a post-tax form, higher tax rates would lead to a stronger incentive for employers to finance these costs rather than shifting premiums to employees. Thus, tax subsidies to insurance are the fifth factor.

The final factor we consider is cyclical conditions. The most dramatic rise in employee premium sharing was at the end of the 1980s and beginning of the 1990s, during which the economy went through a significant downturn. A variety of models, such as recruitment models with imperfect worker understanding of wage shifting, or rent-sharing between workers and firms that is partially through insurance premium sharing, would suggest that when the economy performs worse, there is more premium shifting to employees.

We investigate the role of these six factors using the only nationally representative annual data on premium sharing that covers this period of rapidly rising employee premium contributions: the Current Population Survey (CPS). These data provide only a crude measure of premium sharing, based on a question of covered employees as to whether their employers pay all, some, or none of premiums. Compared to more comprehensive sources available for particular years, however, these crude data capture both variation across jobs/places and over time in the propensity to share costs between employers and employees. Moreover, this disadvantage is counterbalanced by the

significant advantage that we can match to these data job and locational variation in our measures of interest. Based on these matches, we can investigate the role of these factors in driving the rise in employee premium sharing.

Our results suggest that this set of factors are all related to employer contribution decisions, but the results for some of the factors (spousal labor supply, taxes, Medicaid and unemployment rates) are much more robust than others (managed care penetration and medical costs). Interestingly, many of these factors changed significantly in the late 1980s and early 1990s in a manner that is consistent with rising employee contributions. We find that the time trend in these influences corresponds quite strikingly to that of employee contributions, but that overall these factors can only explain about a quarter of the rise in employee contributions over the entire 1982-1996 period.

Our paper proceeds as follows. We begin, in Part I, by providing background on employer and employee contributions for health insurance. We also discuss heuristically the theoretical issues involved in thinking about the tradeoff between employer and employee-financed insurance payments. Part II lays out our data and empirical framework for testing hypotheses about this shift. Part III presents our results, and assesses the extent to which the factors we investigate can explain this time series trend. Part IV concludes.

Part I: Background

Time Series Trends

Group health insurance provided through the workplace has been the dominant source of private health insurance coverage in the U.S. at least since an IRS ruling in the 1940s that health insurance costs were deductible from employer costs, but were not taxable income to employees. In 1998, over 90% of the privately insured received their coverage through employers.

But, as noted in the introduction, employer provided coverage has been declining precipitously over the past two decades. Figure 1 graphs the share of workers who have group coverage over time.² Coverage was flat until 1982, then slowly declined until 1988, when the decline was rapid before stabilizing after 1992. Over the entire period, group coverage declined by 7 percentage points, or almost 10%. This significant decline is the primary driver behind the sizeable rise in the share of the non-elderly without any health coverage, although this rise in uninsurance is smaller than the decline in private coverage due to the mitigating role of Medicaid.

Several recent studies have attempted to decompose this decline in employer coverage. Cooper and Schone

(1997) find that the decline over the 1987-1996 period is completely driven by reduced employee takeup of employer-provided coverage; they estimate that firm offering of insurance actually rose over this period. Farber and Levy (2000) also estimate that offering has risen between 1988 and 1997, and that the decline in coverage can be attributed to both reduced insurance takeup and reduced eligibility for insurance among those offered.

Why has insurance takeup declined over time? One reason may be the significant increase in required employee contributions towards employer-provided health insurance. Figure 2 superimposes the decline in the share of employers paying all of the cost of employer-provided coverage, from the CPS, on the decline in employer coverage over this period. The series are normalized to fit on the same scale, and they show remarkable consistency over time, with both series flat until the early 1980s, slowly declining until 1987, rapidly declining over the next five years, and then flattening out again. The correlation between the series is 0.98, and the residual correlation after partialing a time trend out of both series is 0.66. A time series regression of group coverage against the share of employers paying all of the costs of health insurance yields a coefficient of 0.44 (0.02), and the relationship is significant even when controlling for time trend, with a coefficient of 0.36 (0.10).

Micro-data evidence on the impact of employer contribution policy on employee insurance decisions has not, to date, yielded evidence consistent with this time series correlation. Chernew, Frick and McLaughlin (1997) and Blumberg and Nichols (2001) both model employee takeup as a function of employer premium sharing. Both papers find that there is little impact of employee premiums on insurance takeup decisions, with the highest estimated elasticities in the range of -0.1. This evidence is not fully convincing, as employer contribution policies may themselves be endogenous to tastes for employee insurance takeup. The bias from this potential endogeneity is not obvious ex ante. If employees who are likely to take up select firms with low employee premiums, this would lead to an upward bias to the estimated elasticity of takeup. But if employers cover a larger share of health insurance premiums when employees don't have tastes for insurance coverage (either through paternalism or to meet insurer requirements on employee takeup), then this could lead to a downward bias. Regardless, this striking time series correspondence is highly suggestive and highlights the value of understanding what drove the trend towards employees paying more for their health insurance.

In this section, we lay out an analytical framework for thinking about the determinants of employee premium sharing. We do not propose a new model here, but rather summarize and extend some of the insights of Dranove, Spier and Baker (2000) and Levy (1998).

As noted by Pauly (1986), the presence of any employee contributions suggests imperfect worker sorting across firms, because in general employer contributions for health insurance are excluded from taxation while employee contributions are not. In a growing share of firms with IRS Section 125 plans, employee contributions can also be excluded from taxation, but such protection of employee contributions is far from complete. The data on the prevalence of such arrangements is sketchy. The most recent available data, from a survey of employers by the Kaiser Family Foundation, suggests that half of all workers are in firms that offered such flexible benefit plans. These data also suggest that in the last year of our sample, 1996, the figure was higher (65%) (Kaiser Family Foundation, 2000). At the same time, earlier surveys by the Bureau of Labor Statistics suggest a much lower prevalence. Data on large firms from the Bureau of Labor Statistic's Employee Benefit Survey (EBS) show that in 1993 only 32% of workers in large firms had tax free employee contributions, and in 1992 only 20% of workers in small firms had such arrangements.

In principle, every firm should also set up a Section 125 plan to further maximize the size of the pie by making employee contributions pre-tax as well. The reason for less than full coverage of this generous tax benefit in practice is unclear, but some of it may have to do with extensive IRS regulation of these arrangements to ensure that they are not abused. For example, the regulations state that no more than 25% of the benefits of a plan can be attributed to any "highly compensated" employee, essentially ruling out the availability of section 125 plans for very small firms. Moreover, there are strict and complicated rules that limit the flexibility of employees to switch sources of insurance coverage during the year if they are paying their health insurance contributions on a pre-tax basis.

Levy (1997) highlights two possible explanations for the existence of employee contributions. The first is the "fixed subsidy" model, whereby employers with multiple insurance plans ask their employees to contribute funds towards insurance in order to incentivize employees to choose the lowest cost insurance plan. If this were the only motivation for employee contributions, employers would contribute the amount of the minimum cost plan, and employers with only one plan would never have employee contributions. In fact, as Levy (1997) points out, the second of these conditions does not hold in practice: more than half of firms with only one plan require an employee

contribution. Overall, she finds that only about one-sixth of employee contributions are paid by workers who have the option of a cheaper plan with no contribution required.

Alternative explanations for employee contributions rely on imperfect worker sorting across firms, and Levy (1997) and Dranove, Spier, and Baker (2000) present two different models of this imperfect sorting. The key notion behind these models is that there is not perfect worker-by-worker shifting of insurance costs to wages, so that with heterogeneity in tastes premium contributions become a useful tool for separating worker types. By requiring contributions, the firm can provide insurance only to those who demand it, and can pass the savings back to employees in the form of higher wages.

Models such as these have a number of interesting predictions; we follow here Dranove, Spier and Baker's discussion of comparative statics. First, in the absence of taxation, there should be 100% employee contributions for insurance, to maximize the ability to separate those who want and do not want insurance. As the tax rate rises, employee contributions fall, due to the tax subsidy to employer spending only. Second, as the premium rises, employee contributions rise, as the value of sorting to the firm is increasing. Third, as outside insurance options increase, employee contributions rise, since there is more possibility of shifting employees to other sources of coverage, raising the wages that can be paid to employees.

These hypotheses have been the subject of some limited testing in these previous articles. Levy (1997) shows that contributions fall with a proxy for insurance demand, worker age, and that firms where workers have higher tax rates are less likely to require a contribution. Dranove, Spier and Baker (2000) show that contributions are larger at smaller firms, are higher for firms with more female workers, are lower at firms with more older male workers, and are higher at firms with more part-time workers (a proxy for higher premium costs).

One difficulty with previous tests, however, is distinguishing the worker sorting story from a simple alternative model that high quality jobs provide higher compensation along many dimensions, including lower employee contributions. Firms with more older workers, fewer female workers, higher employee wages and thus tax rates, fewer part-time workers, and more total employees are all the type of high quality jobs that are likely to compensate their workers highly. Given imperfect controls in these models for job quality, this could easily explain the finding that such jobs require smaller employee contributions. In the empirical work presented below, we will endeavor through instrumental variables strategies to avoid such problems of interpretation in our measures of

determinants of employee contributions.

Part II: Data and Empirical Framework

Data

Our primary data for this analysis is the CPS data on premium sharing used in Table 2. As noted above, the CPS only provides information on whether the employer pays some, all, or none of the premium. An additional limitation is that this information is only provided conditional on being covered by insurance, and only for the policy through which one is covered. We cannot condition on having insurance in our regression analysis of premium sharing, since the factors that we examine may (and in fact, in some cases, do) have effects on the coverage decision itself. Therefore, our dependent variables will be unconditional, measuring (for example) the share of all workers for whom the employer pays all the costs of health insurance. This variable may change for four reasons. First, employers may shift the financing of their health insurance plans. Second, changes in employer offering may be differentially concentrated in high or low employee contribution firms. Third, changes in employee takeup may be concentrated in differentially high or low employee contribution firms. Finally, since this measure refers to the plan held by the employee, employees may be moving across plans of different contribution levels. We will address these issues in the interpretation of our results below.

These limitations raise fundamental issues of the applicability of our CPS results, however: do shifts that we observe in the share of employers paying different amounts for insurance in the CPS accurately capture shifts in employer-financing more generally? To address this concern, we have compared the CPS data with two other sources which have more complete information on premium sharing. The first is the Bureau of Labor Statistics Employee Benefits Survey (EBS). The EBS surveys were sporadically carried out since the early 1980s, alternating in recent years between small private firms, medium/large private firms, and government workplaces. They also provide data on the share of employees required to pay some of the cost of their insurance; workers required to pay all of the costs are not counted as insured for their purposes and so not included in the survey. The EBS unfortunately only provide time series data and no micro-data or cross-tabulations; we use the summary of their time series data from EBRI (2000).

Table 1 provides a comparison through time of our CPS and EBS results. We focus on the EBS results for medium and large firms, since this is the only long time series available. Since the CPS only has data on firm size

beginning with the 1988 survey (data for 1987), we compare the EBS time series both to the overall CPS patterns and the patterns over all years, and the patterns for medium and large CPS firms for 1987 onwards.

There is a rough time series correspondence between these two sources of data. Both sources show small changes in the early 1980s. The EBS shows a much larger rise from 1985 to 1988 than does the CPS. Then, from 1988 to 1993, both sources show a large rise, although it is larger in the EBS than in the CPS. The series for family premium sharing is then fairly flat in both data sources. For singles, the EBS shows a much larger rise since 1993 than does the CPS. Overall, the time series correspondence seems reasonable, particularly for family policies.

The second source is data on the share of costs of insurance for family and single plans that are borne by firms, from unpublished data tabulations purchased from the benefits consulting firm KPMG. These data have the advantage that they represent a more complete measure of premium sharing, the actual percentage of costs borne by the firm. But we were only able to obtain cross-tabulations of these data, by region, industry, and firm size, and only for years from 1991 to the present.

Therefore, to compare these data, we have collapsed our CPS data into comparable year, region, industry, and firm size cells, and examined the correlation between our CPS measure of percent of firms paying all of premiums and the KPMG measures of percent of costs borne by firms. We find a correlation for family premium sharing of 0.33, and for individual premium sharing of 0.23. Figure 3 illustrates this correlation for family premium sharing; there is a strong positive correlation with only two notable outliers. The correlations suggest that the CPS data contain real information about the degree of premium sharing.

Our CPS sample for this analysis consists of all adult workers (age 21-64). We exclude the self-employed and the federal government employees. We use CPS data from March 1983 (referring to calendar year 1982) through March 1997 (referring to calendar year 1996).³ We will focus on several dependent variables related to firms' health insurance provision. The first three are whether the firm pays some, all, or none of the costs of health insurance. As noted above, this is measured by a dummy which is equal to one if the employer pays all/some/none, and zero otherwise, not conditional on whether the individual has insurance. To interpret these findings it is also important to measure what is happening to overall insurance coverage. We therefore also examine the impact of these factors on whether the worker has insurance on their job at all.

Measurement of Key Independent Variables

As noted earlier, we consider the role of six key potential explanations for the time series trend in employer contributions. For all concepts, we would ideally measure their impact on insurance decisions at the level of the firm. But the CPS does not provide any detail on an individual's firm composition, other than their industry, location, and (from 1988 onwards) firm size.

We considered two proxies for firm-based measures of our key incentive variables. The first was to aggregate the CPS data by various combinations of state, industry, and year in order to form "synthetic firms". The alternative is to simply use the CPS respondent's information to form the measures, as a proxy for the characteristics of their firm. As part of earlier work (Gruber and Lettau, 2000), one of us has investigated both of these options using internal Bureau of Labor Statistics data, the Employment Compensation Index (ECI) data, which gathers information on both firm characteristics and the wages of workers in the firm. The data show that, for predicting the average wage of a firm, the individual worker's wage has much more predictive power than does an average wage formed by aggregating like firms into synthetic firms. We therefore create our measures at the level of the worker, as a proxy for that worker's firm characteristics.

As noted above, the "fixed subsidy" model of employee premium sharing suggests that such premium sharing arises as a mechanism to ensure efficient worker choice of health plan. This model suggests that, as new lower cost insurance alternatives become available to workers, firms should be more likely to pass premium costs to employees in order to cause them to choose these lower cost alternatives. Of course, we do not know about the insurance choices available to each of the workers in the CPS. But we can proxy for the availability of these new lower cost alternatives that might induce premium sharing by the managed care penetration rate in the worker's state. This is defined as the share of privately insured persons enrolled in HMO plans in the state, and the data come from Laurence Baker, who has compiled them for his work on HMO penetration. This is of course a somewhat crude proxy, but it should capture the introduction of low cost options that would cause employers to want to induce price sensitivity in plan choice among their employees.

In terms of the imperfect worker sorting model, we test four predictions. The first is that premium sharing should rise with the outside insurance options available to workers. We use two proxies for outside insurance

options. Our first is spousal labor force participation. Our regressor here is a dummy variable for whether the worker has a spouse who works at least 17 hours per week. We explored alternative measures that tried to use information on the quality of the spouse's job, and the results were quite similar to those reported here.

The second measure of outside options is entitlement to Medicaid. Here, we use the simulation program developed for earlier work by one of us, and described in more detail in Currie and Gruber (1996a,b), Cutler and Gruber (1996), and Gruber (2000). This program uses information on women and children in the CPS to compute their eligibility for Medicaid coverage given state eligibility rules. We then, following Cutler and Gruber (1996) use the computed eligibility for all women and children to calculate the percentage of each family's medical spending that is eligible for Medicaid, which we call MES (Medicaid eligible share). This is calculated according to:

$$(2) \quad \text{MES} = (\sum_k \text{SPEND}_k * \text{ELIG}_k * \text{NUM}_k) / (\sum_k \text{SPEND}_k * \text{NUM}_k)$$

where k indexes single year age groups of children, and broader age groups for adults.⁴ SPEND_k is the expected health spending in a year for that age group based on data from the 1987 National Medical Expenditure Survey (NMES); the appendix to Cutler and Gruber (1996) presents these figures.

The second prediction of the imperfect sorting model is that premium sharing should fall with the relative subsidy to employer spending on insurance. We test this hypothesis by computing the tax price of insurance for workers, which measures the tax subsidy to insurance purchase through the firm. This is computed as:

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

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. As the tax price of insurance rises (or as tax rates fall), there will be less pressure to pay for insurance through the firm, and therefore more premium sharing as a means of dealing with imperfect worker sorting.

To compute the marginal tax rate for each worker, we use 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.⁶

Virtually all of the elements of taxable income that we need are reported in the CPS, with the major exception of any information on the itemization behavior of the household. We therefore used data from the Statistics of Income (SOI) data to impute both the odds of itemization and the amount itemized by state and family earnings level. For each person, we compute their tax rate as a non-itemizer, and as an itemizer with average itemization equal to the imputed amount from the SOI. We then take a weighted average of the resulting tax rates, where the weights are the predicted rate of itemization based on state and earnings.

The third prediction of the imperfect sorting model is that premium sharing should rise as premium costs increase, since this raises the value of sorting. Once again, we do not know the firm's actual insurance costs. Thus, as a proxy for insurance costs, we measure average spending on medical care per capita by state, from the Health Care Financing Administration.

Finally, we consider the role of cyclical conditions. This factor is not addressed in theoretical models of premium sharing, which are full employment models. But there are a variety of rent-sharing theories which suggest that firms and workers share in the benefits of firm success (and the costs of firm failure); Budd and Slaughter (2000) provide a review of this literature and some convincing new evidence. If there is rent sharing in wages, then there may also be rent sharing through health insurance contributions as well. Alternatively, another link between health insurance contributions and economic conditions could be employee recruitment. To the extent that potential employees pay particular attention to whether they have to contribute to their health insurance plan at a prospective new firm, and do not understand that it is likely that lower contributions for insurance also generally will imply lower wages, when the labor market is tight firms may choose to pay all of the costs of insurance. But, as unemployment rises, there is less pressure on firms to use low employee contributions as a recruitment tool. Thus, we include in the model the state/year unemployment rate, from the Bureau of Labor Statistics, to capture cyclical effects on premium sharing decisions.

Identification Concerns

While each of the measures laid out above captures the influences of these factors on employer behavior, the measures suffer from two important potential limitations. The first is measurement error; these are very noisy proxies for the characteristics of a given worker's firm. The second is omitted variables bias. For each of these

measures, there are potential correlates of both the measures and the firm's insurance decisions that could bias estimated relationships. A critical omitted variable is firm-specific economic shocks. For example, if a given firm is subject to a downturn, then both wages and employer contributions for health insurance may fall. A decline in wages will also lead to a decline in tax rates and therefore the subsidy to employer-provided spending, to a rise in Medicaid eligibility, and potentially to a rise in spousal labor supply, biasing all three of these coefficients in favor of finding the expected explanatory role for contribution shifts.

To address these concerns, we use instrumental variables for the first three of our measures. For spousal labor supply, we instrument actual spousal work with predicted spousal work using the characteristics of the spouse. That is, we estimate in each year a model of labor supply (separately) for married women and married men as a function of age, race, education, and interactions of these variables. We then use the resulting coefficients to form a predicted measure of work for each spouse of each worker in our sample, and use this as our instrument. In our regression models, we control for the spouse's age, education and race directly. So this instrument is identified only by interactions of race and age, race and education, and age and education, and interactions of all of these with year of survey. All of these seem plausibly exogenous to the premium contribution of a given employee.

For Medicaid eligibility, we follow Cutler and Gruber (1996) in instrumenting the Medicaid Expenditure Share with a "simulated" MES. This is computed by using a measure of "simulated" Medicaid eligibility. To create this measure, we first select a random sample of 250 married families and 250 single persons in each decile of their marital-status specific income distributions in each year's CPS. These same 5000 observations are then assigned to each state, and the relevant odds of Medicaid eligibility are computed for each family in the sample. The average MES is then computed for each income decile/marital status/state/year cell, and this is used as an instrument for all persons in that cell.

This instrument varies only by income decile by marital status, state, and year. Each of these factors is controlled for linearly in the model, so that identification comes only from their interactions. Thus, this instrument purges any omitted variables bias other than that arising potentially from those interactions. One obvious concern with this approach is that there may be changes in employee premium sharing by income group over time. Thus, in the basic model we also include a full set of income decile by marital status by year interactions. Another concern is that there may be time trends by state that are correlated with factors such as HMO penetration, and likewise

correlated with employers' decisions on premium contributions. To control for such time trends, we include in the model a set of interactions between each state dummy and a linear time trend variable.

For our measure of the tax subsidy, we use a similar approach. We once again draw a national sample of families by income by marital status, and assign them to every state in that year. We then use that sample to compute tax prices, and use the average by income decile by marital status*state*year cell as our instrument.

For our remaining measures, managed care penetration, medical spending, and unemployment, we do not have readily available instruments. For unemployment, this is not likely to be an important issue, as the state/year unemployment rate can reasonably be taken as exogenous to the firm's decision on premium sharing. But this is a more important issue for our other measures.

For medical spending, the reverse causality may arise because rising employee contributions cause falling medical spending by making employees more sensitive to the cost of medical care. Fortunately for us, however, this biases against the hypothesis of interest, which is that higher medical costs lead to more employee contributions, so if we find the hypothesized relationship it should be convincing. For managed care penetration, the reverse causality may arise because managed care plans may choose to expand in places where employees pay a larger share of their premiums, since they will be most successful in such price sensitive environments. This bias is more problematic because it goes directly in favor of the hypothesis we are attempting to test.

Regression Framework

We will incorporate these measures of interest into a regression framework of the following form:

$$(3) \quad Y_{kjt} = \alpha + \beta_1 HMO_{jt} + \beta_2 SPLS_{kjt} + \beta_3 MES_{kjt} + \beta_4 TP_{kjt} + \beta_5 SPEND_{jt} + \beta_6 UNEM_{jt} \\ + \beta_7 X_{kjt} + \beta_8 \eta_j + \beta_9 \tau_t + \beta_{10} \eta_j * TIME + \epsilon$$

where k indexes individuals, j indexes states, and t indexes years; Y is one of our insurance measures; HMO is our managed care penetration measure; SPLS is average spousal labor supply for the cell; MES is the average Medicaid Eligible Share for the cell; TP is the average tax price for the cell; SPEND is state/year medical spending; UNEM is the state/year unemployment rate; X is a set of individual covariates; and η_j , and τ_t are sets of fixed effects for state, and year, respectively. The individual covariates in the model include own and spouse's age, race, and education; sex, marital status, and an interaction of these; occupation dummies; a set of 10 income decile dummies for married and 10 for single persons; interactions of these 20 income by marital status dummies with year dummies; and a

separate linear time trend for each state (η_j *TIME).

Our key regressor is whether your employer pays all of the cost of your health insurance.

For each of the coefficients β_1 through β_6 , the hypothesis is that the coefficient of interest will be negative; each of these factors is hypothesized to raise premium sharing with employees. The impacts on whether the employer pays some of the cost of insurance are ambiguous. On the one hand, if employers are moving from paying all of the contributions to paying some, then these coefficients should all be positive when the dependent variable is employer pays some of the cost. On the other hand, to the extent that employers react to these forces by moving from paying some of the costs to none of the costs, then the coefficient may be negative. Moreover, it is important to recall that we are using unconditional measures of premium sharing here. So if employers are reacting to these forces by simultaneously reducing insurance coverage and premium sharing, then there could be reductions in both the “employer pays all” and “employer pays some” coefficients; the reduction in the latter would reflect the net of shifting to employees and dropping insurance altogether.

The means of our data are presented in Table 2. 62% of our sample of workers has health insurance coverage through their own employer. For roughly 2/5 of these workers, the employer pays all of the cost of insurance; for the other 3/5, the employer pays some, with very few employees having employers who pay none of the costs of insurance. On average over our sample period, 12% of the privately insured are in HMOs, although this figure is rising rapidly over time. Only 3% of Medical spending for our full sample is eligible for Medicaid on average, although this figure is once again rising rapidly. Roughly half of spouses work, and on average the tax subsidy to insurance is about one-third of the price of insurance (a tax price of insurance, relative to wages, of 0.65). Medical spending per capita in the states averages \$2450, and the average unemployment rate is 7%.

Part IV: Results

Basic Results

Our basic regression results are shown in Table 3. The first three columns show the results for the odds that the employer pays all, some, or none of the cost of insurance. The coefficients across these columns need not add to zero because these are unconditional measures; rather, the coefficients add to the net change in insurance coverage induced by that factor. The final column therefore shows the impact on having coverage at all through your employer. All regressions are estimated as linear probability models for consistency of our instrumental variables estimates; results are similar if probit models are used instead. The standard errors are corrected for within state-year clustering.

The most striking feature of the first column of Table 3 is that all of our predictors have the expected (negative) sign; in every case, a stronger incentive for more premium sharing reduces the odds that employers pay all of the cost of insurance. This is a striking confirmation of the role of economic incentives in this employer decision.

But only four of the six coefficients are statistically significant. The first coefficient of interest is that on HMO penetration. There is a negative impact of HMO penetration on premium sharing, indicating that for each 10 percentage point rise in HMO penetration, the share of employers paying all of the cost of health insurance falls by 0.74 percentage points. There is a corresponding rise in the share of employers paying some of the cost, with little effect on overall coverage. But none of these coefficients are significant, and the impacts are substantively quite small; the elasticity of full employer financing of insurance with respect to managed care penetration is less than 0.04. Thus, the results here confirm the intuition from Levy's (1997) facts: the fixed subsidy model cannot explain much of the time trend in premium sharing.

The next two coefficients of interest are those on Medicaid and on spousal labor supply. Both show sizeable and highly significant negative impacts on the odds that an employer pays all of the cost of health insurance, which is consistent with the contention of the imperfect sorting model that raising outside insurance options will lead to more premium sharing. In the case of Medicaid, the results indicate that for each 10 percentage points increase in the Medicaid eligible share, the share of employers paying all of the cost of insurance falls by 1.7

percentage points. There is a corresponding 0.77 percentage point rise in the share of employers paying some of the cost and a 0.28 percentage point rise in the share paying none of the cost. In addition, there is a 0.64 percentage point decline in the odds that the individual is covered at all by employer-provided insurance (which is consistent with the “crowdout” results in Cutler and Gruber (1996)).

The fact that there is some reduction in total insurance coverage makes interpretation of the impacts on premium sharing somewhat difficult, because without longitudinal data we can not infer the premium sharing arrangement that existed for those losing (or dropping) coverage.⁷ A conservative assumption would be that those that lost or dropped coverage were distributed across the all/some/none categories in proportion to the full sample. This is conservative since it seems likely that those firms that would drop coverage in response to Medicaid expansions, or those workers that would stop taking up, would be much more likely to come from the pool of firms paying some or none of the costs of insurance, not from the pool of firms paying all of the cost.

Under this assumption, 38% of those employees losing coverage previously were in jobs where the employer paid all of the costs of insurance, 57% were in jobs where the employer paid some of the costs, and 5% were in jobs where the employer paid none of the costs. These proportions would suggest that 0.24% of the 1.77% reduction in employers paying all comes from reduced coverage, so that on net a 10% rise in Medicaid entitlement led to a 1.53% shift from employers paying all of the cost of insurance to employers paying some or none. But this is likely a lower bound, for the reasons noted above.

For spousal labor supply, there is a 1% reduction in the odds of an employer paying all of the costs of insurance for each 10% rise in the odds of having a working spouse. But there is an even larger 1.2% decline in the odds of an employer paying some of the cost, with little effect on the odds of paying none of the cost, for a total reduction in employer coverage of 2.2%. In this case, interpreting the impact on actual changes in premium sharing is more difficult. But under the conservative assumption that coverage reductions are in proportion to the initial shares of premium contributions, then there is a slight shift in financing of roughly 0.1% for each 10% rise in spousal labor supply. Thus, there is a wide range of possible impacts of spousal labor supply on premium sharing decisions of firms, but in any case the estimated impact appears fairly small.⁸

The fourth row shows the impact of the tax price. Once again, as with the case of spousal labor supply, here we find negative effects on paying all and on paying some of the cost of insurance, and even a significant

negative impact on the odds of paying none of the cost, with a resultant very sizeable decline in overall insurance coverage; this result mirrors the price sensitivity of employer-provided insurance coverage documented in Gruber and Lettau (2000) and Gruber (2002). Once again, this overall negative impact makes interpretation somewhat difficult. Under the conservative assumption used thus far, each 10% increase in the tax price leads to a 1.7% shift from employers paying all of the costs of insurance to employers paying some or none. So the impact of tax changes on premium sharing appears quite large; the effect varies from 1.7% to 3.7% per each 10% change in tax price.

The fifth row shows the effect of state/year medical costs; the coefficient is that on the level of costs divided by 1000. We find here a negative, but not significant, relationship between medical costs and premium sharing.

Finally, there is a very significant negative effect of the unemployment rate on the odds that the employer pays all of the cost of insurance. This coefficient indicates that for a 10 percentage point rise in the unemployment rate, 1.7 percent fewer firms pay all of the cost of insurance. There is a rise in the odds of a firm paying some or none of the cost of insurance by 0.5 percent, and an overall reduction in coverage of 1.2 percent. Thus, under the conservative approach pursued thus far, we say that each 10 percent rise in unemployment leads to a 1.2 to 1.7 percent reduction in the odds that an employer pays all of the costs of insurance.

Endogenous Incomes

We argued above that our instruments likely purged these models of omitted variables bias, because our instruments only vary by income group, marital status, state, and year, and we are controlling for main effects of all four factors, as well as interactions of income, marital status, and year. But there is an additional concern that is not addressed by this approach: endogeneity of income groups. The consensus in the health economics literature is that there is full or close to full shifting of health insurance costs to wages (Gruber, 2000). As a result, if firms change their insurance contributions, that should be reflected in wages, which will in turn feed back to our instruments. This creates a problematic endogenous correlation between our instruments and the dependent variables in these models.

We have addressed this endogeneity concern by recreating our instruments using not actual income but *predicted* income. That is, we predict income for each household as a function of age, sex, race, education, sex*education, sex*race, race*education, and dummies for number of children. We then use these predictions to create *predicted* income deciles, and classify households based on these predicted income deciles for the purposes of making our instruments. This approach results in instruments which are free of the potential endogeneity bias from using actual incomes, although they are also, by definition, less efficient.

The results of using this alternative IV approach are shown in Table 4. As would be expected, there is relatively little impact on the regressors where there was no change in instruments; the coefficient in the “employer pays all” regression is down somewhat for spousal labor supply, and up for the unemployment rate. There is also remarkably little impact of this instrument on the Medicaid coefficient on premium sharing, although the overall coverage coefficient is now insignificant. There is a much larger impact on the tax price coefficient, which has almost doubled in size. This is partly due to a larger overall effect on insurance coverage, and partly due to a larger concentration of the effect in the employer pays all (rather than the employer pays some) category. Using the same type of calculation that we pursued above, we now estimate that for each 10% rise in the tax price, there is a 3% reduction in the odds that employers pay all of the costs of health insurance, a quite large effect.

Implications for Time Series Trends

Our paper began with the question of what factors can explain the time series trend in rising employer contributions for health insurance. We can now return to this question by applying our estimated coefficients to the time series trends in our key independent variables, and comparing the predicted time series trends that result to the actual trend in premium contributions by employers.

The results of doing so are presented in Figure 4. The figure shows two lines, which have been rescaled so that the time patterns can be easily compared. The first line is the actual time trend in the share of employers paying all of the cost of insurance. The second line is the predicted time trend, based on the time trend in our six key independent variables, times the coefficients of each in our basic Table 3 regression.

There is a remarkably close correspondence between the time series in the actual and predicted time series. The key features of the time series are replicated here: a slow decline through 1985, a much more rapid decline

through 1992, and then a flattening in the mid-1990s. The figure is very similar if, instead, the conditional share of firms paying all of insurance costs is compared to the implied conditional effects from our regressions (e.g. using our conservative assumption above to obtain the impact on premium shifting).

While the correspondence between the series is close over time, however, the magnitudes implied by our model are not large enough to explain the overall time series shift. In Table 5, we illustrate this by dividing our data into three periods: 1982-1985; 1985-1992; and 1992-1996. From 1982-1985, the share of employers paying all of the cost of health insurance fell by 2.2 percentage points. The predicted decline from our model was 0.4 percentage points, or 18% as large a decline. From 1986-1992, the share of employers paying all of the cost fell by 9.8 percentage points. The predicted decline 2.7 percentage points, or 28% as large a decline. From 1993-1996, the share of employers paying all of the costs fell by 0.7 percentage points. The prediction over this period was actually a rise of 0.3 percentage points. Over the entire period, the actual decline was 12.8 percentage points, and the predicted decline was 2.8 percentage points, or 22% as large a decline. Thus, we conclude that the factors in our model match fairly well the time series pattern of employer contributions, but that they can only explain about a quarter of the overall movements over this period.

Part V: Conclusions

The large and growing literature on the determinants of health insurance coverage of the U.S. population has been focused primarily on the decision of employers to offer health insurance. But there is a growing recognition in health economics that employee take-up decisions may be the more important margin for explaining the large declines in coverage that we have witnessed over the past two decades. This contention is bolstered by the fact that there was such an enormous shift in premium costs from employers to employees over this time period. Yet, to date, there has been no explanation for this dramatic and potentially important trend.

In this paper, we have investigated six possible determinants of this trend, drawing on the theoretical arguments for why, in the face of tax subsidized employer premiums, employers would shift premium costs to employees. Five of these six determinants (spousal labor supply being the exception) also have the attractive feature that the incentives for employee financing grow most rapidly in exactly the time period when the shift to employee financing was most pronounced, the late 1980s and early 1990s. We find that, for all six factors, we obtain the

expected relationship with employee financing, although this relationship is only significant in four of the six cases. In terms of timing of changes over this period, these factors do an excellent job. But, in terms of the overall trend over this period, they explain less than one-quarter.

These findings, particularly the strong effect for tax incentives, suggest that premium financing is a price sensitive decision for firms. This implies that policies that subsidize the employer-provision of health insurance may not only increase insurance offering, but also reduce the burden of premium payments for employees. This provides an additional factor that must be included in cost-benefit analysis of employer versus individual subsidies as a means of expanding insurance coverage.

These results also raise two further research questions. First, what other factors explain the trend towards increased employee premium sharing over this period? Future research with data that has more continuous measures of premium sharing should be employed to understand more fully this important trend. Second, what are the implications of these rising employee contributions? As noted earlier, the existing small literature on employee takeup suggests that it is not very price elastic, suggesting that this premium shift has only distributional consequences. But further work is needed to confirm or refute this contention.

References

- Blumberg, Linda, Len Nichols, and Jessica Banthin (2001). "Worker Decisions to Purchase Health Insurance," mimeo, Urban Institute.
- Budd, John, and Matthew Slaughter (2000). "Are Profits Shared Across Borders? Evidence on International Rent Sharing," NBER Working Paper #8014.
- 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.
- Cooper, Phillip, and Barbara Schone (1997). "More offers, fewer takers for employment-based health insurance: 1987 and 1996," *Health Affairs*, 16, 142-149.
- Currie, Janet, and Jonathan Gruber (1996). "Saving Babies: The Efficacy and Cost of Recent Expansions of Medicaid Eligibility for Pregnant Women," *Journal of Political Economy*, 104, 1263-1296.
- Currie, Janet, and Jonathan Gruber (1996). "Health Insurance Eligibility, Utilization of Medical Care, and Child Health," *Quarterly Journal of Economics*, 111, 431-466.
- Cutler, David, and Jonathan Gruber (1996). "Does Public Insurance Crowd Out Private Insurance?," *Quarterly Journal of Economics*, 111, 391-430.
- Dranove, David, Kathy Spier and Laurence Baker (2000). "'Competition' Among Employers Offering Health Insurance," *Journal of Health Economics*, 19, 121-140.
- Farber, Henry, and Helen Levy (1999). "Recent trends in Employer-Sponsored Health Insurance Coverage: Are Bad Jobs Getting Worse?," *Journal of Health Economics*, 19, 93-119.
- Feenberg, Daniel, and Elizabeth Coutts (1993). "An Introduction to the TAXSIM Model," *Journal of Policy Analysis and Management*, 12, 189-194.
- Gruber, Jonathan (2000). "Medicaid," forthcoming in Robert Moffitt, ed., *Means Tested Transfer Programs in the U.S.*
- Gruber, Jonathan (2002). "The Impact of the Tax System on Health Insurance Coverage," *International Journal of Health Care Finance and Economics*, 1, 293-304.
- Gruber, Jonathan and Michael Lettau (2000). "How Elastic is the Firm's Demand for Health Insurance?," NBER Working Paper #8021, November.
- Kaiser Family Foundation (2000). *Employer Health Benefits Survey*. Menlo Park, CA: Kaiser Family Foundation.
- Levy, Helen (1997). "Who Pays for Health Insurance? Employee Contributions to Health Insurance Premiums," mimeo, Princeton University.
- Pauly, Mark (1986). "Taxation, Health Insurance and Market Failure in the Medical Economy," *Journal of Economic Literature*, 24, 629-675.

Table 1: Comparison of CPS and BLS Data on Percent of Employees Contributing to their Plans

Year	CPS All		EBS Medium/Large		CPS Medium/Large	
	Family	Single	Family	Single	Family	Single

1981	54%	48%				
1982	54%	47%				
1983	55%	48%	54%	33%		
1984	56%	49%				
1985	56%	49%	56%	36%		
1986	57%	50%				
1987	58%	49%			62%	55%
1988	58%	50%	64%	44%	63%	56%
1989	61%	52%			66%	58%
1990	63%	53%			68%	60%
1991	65%	57%	69%	51%	70%	65%
1992	67%	60%			73%	67%
1993	67%	59%	76%	61%	73%	67%
1994
1995	66%	60%	78%	67%	73%	67%
1996	70%	60%			75%	68%
1997	69%	60%	80%	69%	75%	68%
1998	70%	63%			76%	70%

Table 2: Means

Variable	Mean	Standard Deviation
Own Group Coverage	0.62	0.48
Employer Pays All	0.24	0.43
Employer Pays Some	0.36	0.48
Employer Pays None	0.03	0.17
Manage Care Penetration	0.12	0.09
Spousal Labor Supply	0.46	0.50
Medicaid Eligible Share	0.03	0.12
Tax Price	0.65	0.10
Medical Spending (\$1000)	2.45	0.63
Unemployment Rate	0.07	0.02
Number of Obs	850,541	850,541

Table 3: Basic Results

Variable	Employer Pays All	Employer Pays Some	Employer Pays None	Own Group Coverage
Managed Care Penetration	-.074 (.086)	.092 (.089)	-.049 (.028)	-.037 (.071)
Medicaid Eligible Share	-.167 (.034)	.077 (.047)	.028 (.012)	-.064 (.045)
Spousal	-.104	-.122	.008	-.215
Labor Supply	(.026)	(.030)	(.011)	(.032)
Tax Price	-.367 (.098)	-.127 (.077)	-.060 (.022)	-.554 (.091)
Medical Spending	-.005 (.008)	-.003 (.009)	-.004 (.003)	-.011 (.008)
Unemp. Rate	-.170 (.071)	.065 (.080)	-.017 (.025)	-.120 (.074)
Number of Obs	850,541	850,541	850,541	850,541

Note: Dependent variable listed in top row. Standard errors in parentheses. Regressions also include controls for: own and spouse's age, race, and education; sex, marital status, and an interaction of these; occupation dummies; a set of 10 income decile dummies for married and 10 for single persons; interactions of these 20 income by marital status dummies with year dummies; state and year fixed effects; and a separate linear time trend for each state.

Table 4: Results Using Predicted Income Instrument

Variable	Employer Pays All	Employer Pays Some	Employer Pays None	Own Group Coverage
Managed Care Penetration	-.088 (.088)	.076 (.096)	-.048 (.028)	-.066 (.081)
Medicaid Eligible Share	-.163 (.067)	.180 (.096)	.012 (.024)	.025 (.094)
Spousal	-.094	-.074	.009	-.156
Labor Supply	(.025)	(.031)	(.010)	(.033)
Tax Price	-.667 (.183)	-.066 (.192)	-.078 (.043)	-.808 (.183)
Medical Spending	.0001 (.008)	.003 (.009)	-.005 (.003)	-.001 (.009)
Unemp. Rate	-.254 (.076)	-.014 (.085)	-.009 (.025)	-.276 (.085)
Number of Obs	850,541	850,541	850,541	850,541

Note: Dependent variable listed in top row. Standard errors in parentheses. Regressions also include controls for: own and spouse's age, race, and education; sex, marital status, and an interaction of these; occupation dummies; a set of 10 income decile dummies for married and 10 for single persons; interactions of these 20 income by marital status dummies with year dummies; state and year fixed effects; and a separate linear time trend for each state.

Table 5: Comparing Predicted vs. Actual Trends

Time Period	Actual	Predicted
1982-1985	- 2.2 %	- 0.4%
1985-1992	- 9.8 %	- 2.7%
1992-1996	- 0.7 %	0.3 %
1982-1996	- 12.8 %	-2.8 %

¹Source for all figures is author's tabulations of March Current Population Survey data.

²There was a major redesign of the health insurance questions in the March 1995 CPS that results in a significant upward jump in the share of the population with group coverage. We have assumed that coverage was flat from 1993 to 1994 and used the ratio of these years to reverse benchmark the earlier figures. Our regression models below will all contain year dummies to capture such shifts in survey methodology.

³The CPS data on premium sharing in the March 1995 survey, for calendar year 1994, are not useful for our purposes since they lump together firms paying all and some of the costs of insurance.

⁴We divide adults into those age 19-29, 30-39, 40-49, 50-59, and 60-64. We further divide women into ages 40-44 and 45-49 because pregnancy is assumed to occur only in the first group.

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.

⁶For more information about TAXSIM, see Feenberg and Coutts (1993). A public use version of TAXSIM is available at www.nber.org/taxsim. Marginal rates are computed by first computing the tax bill, then adding \$1000 to earned income and recomputing the tax bill, and taking the difference divided by \$1000.

⁷In the context of Medicaid, the available evidence suggests that the overall reduction in coverage arises mostly from a reduction in insurance takeup conditional on offering, not from reduced employer offering.

⁸It is interesting to note that while we find that higher levels of spousal labor supply reduce the odds of own insurance coverage, we find no effects on the odds of having any employer-provided coverage. Thus, it appears that when spouses work, they reduce the insurance coverage on their spouse's job, but equally raise the odds of insurance through their own job.