Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Health Economics 22 (2003) 1085–1104
JOURNAL OF HEALTH ECONOMICS
www.elsevier.com/locate/econbase

Why did employee health insurance contributions rise?

Jonathan Gruber*, Robin McKnight

Department of Economics, Massachusetts Institute of Technology, MIT E52-355, 50 Memorial Drive, Cambridge, MA 02142-1347, USA

Received 1 June 2002; received in revised form 1 February 2003; accepted 24 June 2003

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 four 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, rising medical costs, and increased managed care penetration. Overall, this set of factors can explain more than one-half of the rise in employee premiums over the 1982–1996 period.

JEL classification: I1; H2; J3

Keywords: Employee; Health insurance; Medicaid

1. Introduction

The dominant feature of the health insurance market in the US 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]

* Corresponding author. Tel.: +1-617-253-8892; fax: +1-617-253-1330.
E-mail address: gruberj@mit.edu (J. Gruber).

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

0167-6296/$ – see front matter
doi:10.1016/j.jhealeco.2003.06.001

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, 1999). 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 most 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: managed care penetration (which is correlated with the range of choices offered by the employer); the expansion of eligibility for the public Medicaid insurance program for women and children (since the imperfect sorting model predicts that such outside options should lead to increased employee contributions); health insurance costs (since the imperfect sorting model predicts that rising medical costs will increase the pressure on firms to shift the cost of insurance to their workers); and marginal tax rates (since employer contributions are tax subsidized, but employee contributions are often not tax subsidized, higher tax rates increase the incentive for employers to finance insurance costs).

We investigate the role of these four 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 noisy 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 data capture both variation across job/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 each of these factors is strongly related to employer contribution decisions. We find that the time trend in these influences corresponds quite strikingly to that of employee contributions, and that overall these factors can explain more than one-half of the rise in employee contributions over the entire 1982–1996 period.

Our paper proceeds as follows. We begin, in Section 2, 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. Section 3 lays out our data and empirical framework for testing hypotheses about this shift. Section 4 presents our results, and assesses the extent to which the factors we investigate can explain this time series trend. Section 5 concludes.

2. Background

2.1. Employee contributions and health insurance coverage

Group health insurance provided through the workplace has been the dominant source of private health insurance coverage in the US 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 Section 1, employer-provided coverage has been declining precipitously over the past two decades. 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 (1999) 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 over this period. As noted in Section 1, employer contributions fell significantly during this time, with the share of employers paying all of the cost of insurance falling from 44 to 28%.

Micro-data evidence on the impact of employer contribution policy on employee insurance decisions has yielded only mixed evidence in support of this time series correlation. Chernew et al. (1997) used data from a sample of small firms in seven cities to model employee takeup of insurance as a function of employee premiums. They find significant but modest effects, with an implied elasticity of takeup with respect to employee premiums of −0.066. Blumberg et al. (2001) pursued a similar approach for a nationally representative sample of firms, and found somewhat smaller effects, with an elasticity of takeup with respect to premiums of −0.04. Gruber and Washington (2003) study the impact of the tax subsidization of federal employee health benefits on insurance takeup decisions, and find an elasticity of only −0.02. Cutler (2002) instruments for the employee premium share with the state tax rate, under the argument that employee premiums will be smaller in states with higher tax rates, and estimates an elasticity of takeup with respect to employee premiums of −0.12.

These estimates are all quite small, but Cutler (2002) makes the important observation that it does not take a very large elasticity to explain the relationship between the rise in contributions in the 1980s and 1990s and the decline in employer-sponsored insurance takeup over this same period. He estimates that it requires an elasticity of only about −0.06 to explain the time series trend in employer-sponsored insurance takeup. So these

past estimates suggest that increased employee premium contributions can explain between one-third and two times the existing trend in employer-provided insurance coverage. These findings highlight the importance of understanding why it is that employee contributions rose so much over this period.

2.2. Analytical framework

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 et al. (2000) and Levy (1997).

As noted by Pauly (1986), the presence of any employee contributions suggests imperfect worker sorting across firms, because employer contributions for health insurance are generally 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 (2000), suggests that half of all workers are in firms that offered such flexible benefit plans. Earlier surveys by the Bureau of Labor Statistics suggest a much lower prevalence. Data on large firms from the Bureau of Labor Statistics' Employee Benefit Survey (EBS) show that only 32% of workers in large firms had tax free employee contributions in 1993 and only 20% of workers in small firms had such arrangements in 1992.[2]

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 et al. (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

[2] In principle, every firm should set up a Section 125 plan to 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 it may be related to 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.

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.

These models have a number of interesting predictions; we follow here Dranove, Spier and Baker's discussion of comparative statistics. 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, since 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 et al. (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 in which 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.

3. Data and empirical framework

3.1. Data

Ideally, for this analysis we would use firm level data on the distribution of premiums between employers and employees. But such data are not available in samples which are large, which provide information across a variety of locations, and which provide data over a number of years. Thus, it is impossible to carry out our tests with existing firm level data sets. We therefore turn to the CPS, which has large, nationally representative samples over many years.

As noted in Section 1, 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 about 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 that have more complete information on premium sharing. The first is the Bureau of Labor Statistics Employee Benefits Survey (EBS). The EBS surveys were carried out sporadically since the early 1980s, alternating in recent years between small private firms, medium/large private firms, and government workplaces. They 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 are 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 statistics. We focus on the EBS statistics 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 over all years and to 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

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

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. Fig. 1 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.

Of course, a limitation of the CPS data is that they do not contain any information on health insurance plan characteristics. We attempt to control for this by including both state and year fixed effects in the models below, to capture both secular differences across states in the nature of insurance plans offered and national time series trends in plan characteristics.

Fig. 1. Premium sharing in CPS vs. KPMG.

Of particular concern are changes in the use of family versus single coverage over time. We will address this point by separately estimating our models for single and married workers.

Our CPS sample for this analysis consists of all adult workers (ages 21–64). We exclude the self-employed and 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 variables indicate whether the firm pays some, all, or none of the costs of health insurance. Each variable is a dummy that is equal to one if the employer pays all–some–none, and zero otherwise, not conditional on whether the individual has insurance. To interpret our findings, it is also important to measure what is happening to overall insurance coverage. We therefore also examine the impact of these factors on a fourth dependent variable, a dummy variable that indicates whether the worker has insurance on their job at all.

3.2. Measurement of key independent variables

As noted earlier, we consider the role of four 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 about 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 discussed above, the "fixed subsidy" model of employee premium sharing suggests that such premium sharing arises as a mechanism for ensuring 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 rate 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 variable is, of course, a somewhat crude proxy, but it should capture the introduction of low cost

[3] The 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.

options that would cause employers to want to induce price sensitivity in plan choice among their employees.

In order to test this proposition, we have matched these penetration data to information on the number of plan choices available in firms in the 1999–2001 Kaiser-HRET survey of firm health insurance plans.[4] We modeled the log of the number of plan choices at films as a function of the HMO penetration rate in the state, and found that each 10% point increase in HMO penetration was associated with a (highly statistically significant) increase of 5% in the number of plans offered at firms. While these results come from a cross-sectional framework, they certainly provide suggestive evidence for this central assumption of our model.

In terms of the imperfect worker sorting model, we test three predictions. The first prediction is that premium sharing should rise with the outside insurance options available to workers. We proxy for outside insurance options by using 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. Following Cutler and Gruber (1996), we 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 Medicaid eligible share (MES). This share is calculated according to

$$\text{MES} = \frac{\sum_k \text{SPEND}_k \times \text{ELIG}_k \times \text{NUM}_k}{\sum_k \text{SPEND}_k \times \text{NUM}_k} \tag{1}$$

where k indexes single year age groups of children, and broader age groups for adults.[5] 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 price is computed as

$$\text{TP} = \frac{1 - \tau_\text{f} - \tau_\text{s} - \tau_\text{ss} - \tau_\text{MC}}{1 + \tau_\text{ss} + \tau_\text{MC}} \tag{2}$$

where τ_f is the federal income tax marginal rate; τ_s the state income tax marginal rate; τ_ss the marginal payroll tax rate for the OASDI program; and τ_MC is the marginal payroll tax rate for the Medicare HI program.[6] 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

[4] We are grateful to Gary Claxton and Erin Holve for permitting us to use the Kaiser data and assisting us with this exercise.

[5] We 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.

[6] The 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 $1/(1 + \tau_\text{ss} + \tau_\text{MC})$ since each dollar of wages requires an employer payroll tax payment as well.

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 NBERs TAXSIM model, which inputs information on the major elements of taxable income and computes both a federal and state marginal tax rate.[7] Virtually all of the elements of taxable income that we need are reported in the CPS, with the major exception of information on the itemization behavior of the household. We therefore use 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 odds 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 and year, from the Health Care Financing Administration. For our measure of medical costs, we include expenditures on hospital and physician services, drugs and other medical non-durables, and vision care and other medical durables. We exclude several important categories of expenditures, including expenditures for nursing home care, home health care, and dental services, that are unlikely to affect the insurance costs of the non-elderly workers in our sample. The resulting medical spending measure is only a very noisy proxy for firm-specific insurance spending. But it should reflect the overall pressures on firms, in terms of rising health care costs, to raise employee premiums.

3.3. Identification concerns

While the measures described earlier capture 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, and to a rise in Medicaid eligibility, biasing both of these coefficients in favor of finding the expected explanatory role for contribution shifts.

To address these concerns, we use instrumental variables for two of our measures. For Medicaid eligibility, we follow Cutler and Gruber (1996) in instrumenting the Medicaid eligible share with a "simulated" MES. This variable is computed by using a measure of "simulated" Medicaid eligibility. To create this measure, we first select a random sample of

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

250 married families and 250 single persons in each decile of their marital-status-specific income distributions in each year's CPS. These same 500 observations are then assigned to each state, and the relevant odds of Medicaid eligibility are computed for each family in the sample. We then compute the average MES for each income decile/marital status/state/year cell, and use this average 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 bias arising 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.

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 for each income decile/marital status/state/year cell as an instrument.

For our remaining measures, managed care penetration and medical spending, we do not have readily available instruments. 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. As a result, our findings for medical spending and managed care are more tentative than our more well-identified findings for tax subsidies and Medicaid entitlements.

3.4. Regression framework

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

$$Y_{kjt} = \alpha + \beta_1 \mathrm{HMO}_{jt} + \beta_2 \mathrm{MES}_{kjt} + \beta_2 \mathrm{TP}_{kjt} + \beta_4 \mathrm{SPEND}_{jt} + \beta_5 X_{kjt} + \beta_6 \eta_j + \beta_7 \tau_t + \varepsilon \quad (3)$$

where k indexes individuals, j indexes states, and t indexes years; Y one of our insurance measures; HMO our managed care penetration measure; MES the Medicaid eligible share measure; TP the tax price measure; SPEND the state/year per capita medical spending; X a set of covariates; and η_j and τ_t are sets of fixed effects for state and year, respectively. The individual covariates in the model include 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; and interactions of these 20 income by marital status dummies with year dummies. We also control for the state/year unemployment rate to capture cyclical influences on contribution policies.

Table 2
Means

Variable	Mean	S.D.
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 (US$ 1000)	2.45	0.63
Unemployment rate	0.07	0.02
Number of observation	850335	850335

Our key regressor is whether the individual's employer pays all of the cost of their health insurance. For each of the coefficients β_1 through β_4, the hypothesis is that the coefficient of interest will be negative; each of these factors is hypothesized to raise premium sharing with employees. The impact on whether the employer pays some of the cost of insurance is 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. 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. Sixty-two percent of our sample of workers has health insurance coverage through their own employer. For roughly 2/5 of these covered 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 also rising rapidly. The tax subsidy to insurance is about one-third of the price of insurance on average (a tax price of insurance, relative to wages, of 0.65). Medical spending per capita in the states averages US$ 2450.

4. Results

4.1. 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.

Table 3
Basic results

Variable	Employer pays all	Employer pays some	Employer pay none	Own group coverage
Managed care penetration	−0.168** (0.039)	0.146** (0.045)	−0.005 (0.011)	−0.026 (0.037)
Medicaid eligible share	−0.140** (0.037)	0.060 (0.047)	0.023* (0.012)	−0.057 (0.045)
Tax price	−0.304** (0.099)	−0.114 (0.078)	0.054** (0.021)	−0.543** (0.086)
Medical spending	0.027** (0.008)	0.034** (0.009)	−0.001 (0.002)	0.006 (0.006)
Number of observation	850335	850335	850335	850335

Dependent variable listed in top row. Standard errors in parentheses. Regressions estimated using OLS. Regressions also include controls for: 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 state and year fixed effects. Standard errors are clustered on state and year.

* Significance at the 10% level.

** Significance at the 5% level.

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 the individual's employer. We estimate separate linear probability models for each dependent variable. Our results are very similar if we estimate probit models instead, or ordered probits on the all–some–none variable. But we prefer the individual OLS results because of the more straightforward inference with the instrumental variables estimates, and because the coefficients on the all, some, and none regressions do add up to the coefficient on the coverage regression. Our standard errors are corrected for correlations within state/year clusters.[8]

Strikingly, we find negative and significant coefficients on all four of our key predictors in this regression framework. The first coefficient of interest is that on HMO penetration. There is a very significant negative impact of HMO penetration on premium sharing, indicating that for each 10% point rise in HMO penetration, the share of employers paying all of the cost of health insurance falls by 1.68% points. There is a corresponding rise in the share of employers paying some of the cost, with little effect on overall coverage. This is a fairly sizeable impact, suggesting that the fixed subsidy model can perhaps explain some of the time series trend in premium sharing; we return to time series predictions below.

The next coefficient of interest is that on the Medicaid eligible share. We find a highly significant effect 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. The results indicate that for each 10% points increase in the Medicaid eligible share, the share of employers paying all of the cost of insurance falls by 1.4% points. There is a corresponding 0.6% point rise in the share of

[8] Another potential concern about our standard errors is that we have multiple observations from two-worker families, whose insurance decisions are surely correlated. We have assessed the importance of this factor by reestimating our model using only one randomly selected worker from each family; our results are unchanged.

employers paying some of the cost and a 0.2% point rise in the share paying none of the cost. In addition, there is a 0.57% 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), but is not significant).

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 cannot infer the premium sharing arrangement that existed for those losing (or dropping) coverage.[9] 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 assumption 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 conservative 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 0.22% of the 1.4% reduction in employers paying all comes from reduced coverage, so that on net a 10% rise in Medicaid entitlement led to a 1.18% shift from employers paying all of the cost of insurance to employers paying some or none.

The third row shows the impact of the tax price. 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, resulting in a 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 negative impact on overall coverage 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 3.8% per each 10% change in tax price.

The fourth row shows the effect of state/year medical costs; the reported coefficient is on the level of costs divided by 1000. Once again, we find a highly significant negative effect on the odds that employers pay all of the costs of insurance, and an offsetting increase in the odds that employers pay some of the costs, with no impact on the odds that employers pay none of the cost. There is no net effect on the odds of being covered by insurance. These results imply that for each US$ 1000 increase in medical costs, there is a 2.7% reduction in the odds that employers pay all of the cost of health insurance.

Thus, these results appear to confirm both the fixed subsidy and imperfect sorting models of the determination of employee premium sharing. There are effects of both managed care penetration, which is correlated with the number of choices available to employees, and tax rates, outside options (Medicaid), and medical costs, which are determinants of employee contributions in the imperfect sorting models.

[9] In 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.

Table 4
Results by marital status

Variable	Employer pays all	Employer pays some	Employer pays none	Own group coverage
Single				
Managed care penetration	−0.148** (0.047)	0.098 (0.060)	−0.011 (0.014)	−0.061 (0.043)
Medicaid eligible share	−0.169** (0.068)	0.034 (0.080)	0.024 (0.024)	−0.111 (0.080)
Tax price	−0.304** (0.118)	−0.132 (0.098)	−0.021 (0.029)	−0.457** (0.095)
Medical spending	−0.019** (0.009)	0.025 (0.013)	0.004** (0.002)	0.002 (0.008)
Number of observation	299602	299602	299602	299602
Married				
Managed care penetration	−0.177** (0.044)	0.164** (0.046)	−0.002 (0.014)	−0.015 (0.041)
Medicaid eligible share	−0.127** (0.038)	0.070 (0.047)	0.022 (0.014)	−0.035 (0.049)
Tax price	−0.453** (0.111)	−0.046 (0.091)**	−0.090** (0.28)	−0.589** (0.111)
Medical spending	0.033** (0.008)	0.040** (0.008)	0.001 (0.002)	0.008 (0.007)
Number of observation	550733	550733	550733	550733

Dependent variable listed in top row. Standard errors in parentheses. Regressions estimated using OLS. Regressions also include controls for: 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 state and year fixed effects. Top panel presents separate estimates for singles; bottom panel for married. Standard errors are clustered on state and year.

** Significance at the 5% level.

4.2. Sensitivity testing

We have also assessed the sensitivity of our findings to several variations in the sample used. The first test is to separately examine single and married individuals. There may be distinct trends in the insurance environment facing those purchasing single and family coverage. In particular, employers may be shifting costs to employees differentially for single and family plans.

In Table 4, we show the results of estimating our model separately for single and married workers. In fact, the results are quite robust across these samples. The only significant difference is a larger impact of medical costs on premium shifting to employees for married than for single workers. This finding is sensible, since a given rise in underlying medical costs will have a larger dollar impact on the cost of family than of single plans, raising the pressure on employee contributions for family plans.

The second variation that we considered is firm size. We might expect our findings to be stronger for smaller firms, for two reasons. First, larger firms are likely to have their insurance decisions made at a central headquarters, and not at the local establishment where employees work, so state or employee-specific measures may be less relevant. Second, larger firms are more likely to have Section 125 plans, so that the tax incentives to pay insurance costs at the employer level are reduced. We therefore estimated the model divided into small (less than 100 employees), medium (100–999 employees) and large (more than 1000 employees) samples, for the 1988 and onwards period for which firm size data are available in the CPS. Our findings were mixed. For HMO penetration and for Medicaid, the effects are indeed largest for the smallest firms. But, for tax prices and insurance expenditures, the effects

are largest for the medium sized firms. Moreover, none of these estimates are significantly different from each other, given the reductions in size when we split the sample. Thus, we cannot draw strong conclusions about relative impacts by firm size.

Finally, we consider another potential instrumenting strategy. 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, the change 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 addressed this endogeneity concern by recreating our instruments using *predicted* income, rather than actual 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 this alternative IV approach confirm our findings from Table 3. There is little effect on our medical spending or HMO penetration coefficients. The estimated tax effect is considerably strengthened, with the coefficient rising to −0.733 (0.186), while the Medicaid effect is weakened, with the coefficient falling to −0.089 (0.079). The change in the tax effect is largely due to a larger overall effect on insurance coverage; 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 4.4% reduction in the odds that employers pay all of the costs of health insurance. But neither the tax or Medicaid coefficient changes significantly from the results in Table 3, as the standard errors have also increased.

4.3. 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 Figs. 2 and 3. Fig. 2 shows the time trend in employer contributions, the percentage of employers paying all of the costs of health insurance over time. This series is slowly declining until 1985, more rapidly declining from 1985 to 1992, and then fairly flat. Fig. 3 shows the predicted time series, based on the time trend in our four key independent variables, times the coefficients of each in our basic Table 3 regression (normalized to the same starting point). This time series moves in a strikingly similar fashion: a slow decline until 1986, a sharp decline through 1992, and then

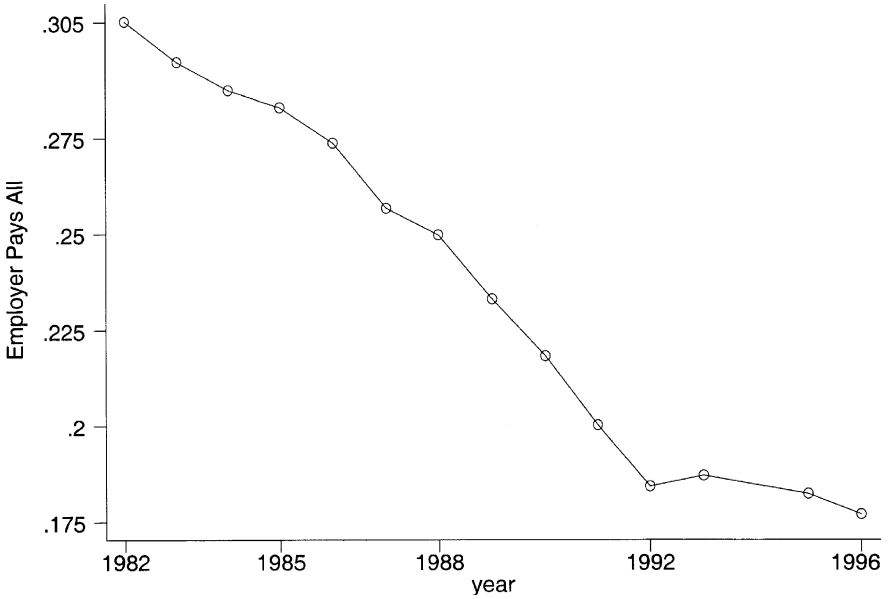

Fig. 2. Time trend in employer pays all.

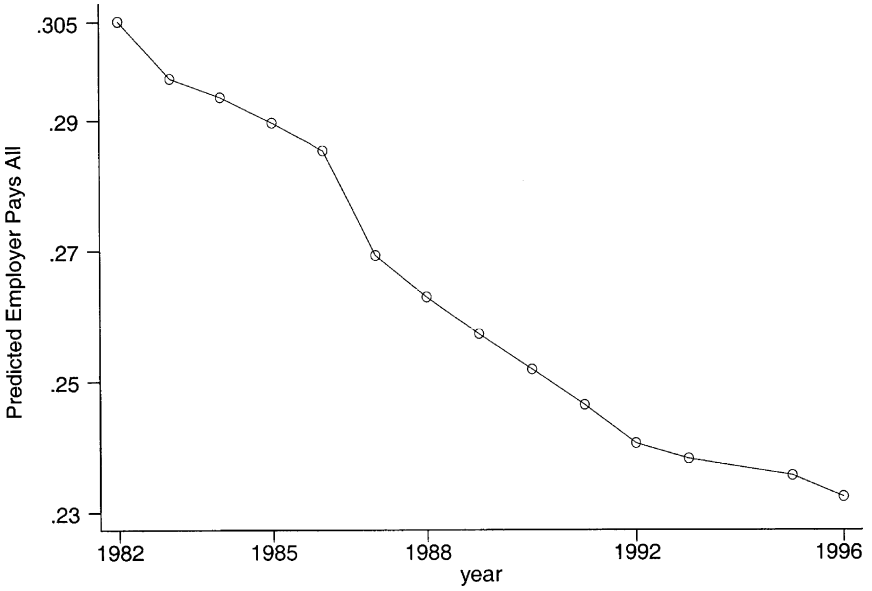

Fig. 3. Time trend in predicted employer pays all.

Table 5
Contribution of time series change

Factor	Change from 1982 to 1996	Contribution (change times coefficient)	Percentage contribution
HMO penetration	0.138	0.023	32
Medicaid eligible share	0.047	0.0065	9
Tax price	0.025	0.010	14
Medical spending	1.38	0.034	47
Total		0.073	100

a very slow decline after 1992. Indeed, even after controlling for linear time trends in both series, the partial time series correlation between them is 0.74.

Over this entire period, the share of employers paying all of the costs of insurance fell from 30.5 to 17.7%, a 13% point decline. The decline in the predicted series is 7.3%, which is more than half as large. Thus, we conclude that the factors in our model match fairly well the time series pattern of employer contributions, and that they can explain over half of the overall movements over this period.

Table 5 decomposes the time series trends in the four factors which are the focus of this study, to show which factors are driving the results. In the first column, we show the change in the value of each measure from 1982 to 1996; this is proportions in the first three rows, and in thousands of dollars per capita in the final row. In the second column, we show the implied contribution to the time series change from 1982 to 1996, which is equal to the number in the first column times the coefficient estimate from Table 3, and in the third column we show the percentage contribution to the time series trend that this represents. We find here that almost half of our estimated time series effect represents the influence of increased medical costs, and another third represents the influence of HMO penetration. The remainder is divided roughly evenly between tax price and Medicaid effects.

5. Conclusions

The large and growing literature on the determinants of health insurance coverage of the US 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 takeup 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 four 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. We find that, for all four factors, we obtain the expected relationship with employee financing, although the relationship between Medicaid expansions and employee financing is somewhat sensitive to specification. The time trends

in these predictors match the time trend in overall employer contributions quite well, and they explain more than one-half of the overall time series trend.

These findings, particularly the strong effect for tax incentives, suggest that premium financing is a price sensitive decision for firms. Thus, 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 finding provides an additional factor that must be included in cost–benefit analysis of employer versus individual subsidies as a means of expanding insurance coverage.

Acknowledgements

We are grateful to Cindy Perry, Kokkeong Puah, and Nannette Yang for research assistance, to Dan Feenberg for assistance with the NBERs TAXSIM model, to Larry Baker for sharing his HMO market share data, and to David Autor, David Cutler, Kathy Swartz, two anonymous referees, and seminar participants at the NBER, the University of Michigan, and the University of Minnesota for helpful comments. Gruber gratefully acknowledges financial support from the Economic Research Initiative on the Uninsured and from the California Health Care Foundation. McKnight acknowledges support from the National Institute on Aging, through Grant #K12-AG00983 to the National Bureau of Economic Research.

References

Blumberg, L., Nichols, L., Banthin, J., 2001. Worker Decisions to Purchase Health Insurance. Mimeo, Urban Institute.

Chernew, M., Frick, K., McLaughline, C.G., 1997. The demand for health insurance coverage by low-income workers: can reduced premiums achieve full coverage? Health Services Research 32, 453–470.

Cooper, P., Schone, B., 1997. More offers, fewer takers for employment-based health insurance: 1987 and 1996. Health Affairs 16, 142–149.

Currie, J., Gruber, J., 1996a. Saving babies: the efficacy and cost of recent expansions of Medicaid eligibility for pregnant women. Journal of Political Economy 104, 1263–1296.

Currie, J., Gruber, J., 1996b. Health insurance eligibility, utilization of medical care, and child health. Quarterly Journal of Economics 111, 431–466.

Cutler, D., 2002. Employee costs and the decline in health insurance coverage. NBER Working Paper #9036.

Cutler, D., Gruber, J., 1996. Does public insurance crowd out private insurance? Quarterly Journal of Economics 111, 391–430.

Dranove, D., Spier, K., Baker, L., 2000. Competition' among employers offering health insurance. Journal of Health Economics 19, 121–140.

EBRI (Employee Benefits Research Institute), 2000. EBRI Databook on Employee Benefits. Employee Benefits Research Institute, Washington, DC.

Farber, H., Levy, H., 1999. Recent trends in employer-sponsored health insurance coverage: are bad jobs getting worse? Journal of Health Economics 19, 93–119.

Feenberg, D., Coutts, E., 1993. An introduction to the TAXSIM model. Journal of Policy Analysis and Management 12, 189–194.

Gruber, J., 2000. Medicaid. In: Robert, M. (Ed.), Means Tested Transfer Programs in the US, submitted for publication.

Gruber, J., 2002. The impact of the tax system on health insurance coverage. International Journal of Health Care Finance and Economics 1, 293–304.

Gruber, J., Lettau, M., 2000. How elastic is the firm's demand for health insurance? NBER Working Paper #8021, November.

Gruher, J., Washington, E., 2003. Subsidies to employee health insurance premiums and the health insurance market. Mimeo, MIT.

Kaiser Family Foundation, 2000. Employer Health Benefits Survey. Kaiser Family Foundation, Menlo Park, CA.

Levy, H., 1997. Who pays for health insurance? Employee contributions to health insurance premiums. Mimeo, Princeton University.

Pauly, M., 1986. Taxation, health insurance and market failure in the medical economy. Journal of Economic Literature 24, 629–675.