The Incidence of Payroll Taxation: Evidence from Chile

Jonathan Gruber, *Massachusetts Institute of Technology and National Bureau of Economic Research*

I provide new evidence on the incidence of payroll taxation by examining the experience of Chile before and after the privatization of its Social Security system. This policy change led to a sharp exogenous reduction in the payroll tax burden on Chilean firms; on average, payroll tax rates fell by 25% over 6 years. Using data from a census of manufacturing firms, I estimate that the incidence of payroll taxation is fully on wages, with no effect on employment. This finding is robust to a variety of empirical approaches to the problem of measurement error in firm-level measures of taxes/worker.

Payroll taxation is a large and growing source of public finance in the United States and the rest of the world. In 1993, 38% of U.S. federal revenues were raised by payroll taxation; in 1960, this figure was only 12.4% (*Economic Report of the President* 1992). This corresponds to a similar increase in the reliance on payroll taxation in other developed countries. Among the organization for Economic Cooperation and Development (OECD) nations over the 1965–88 period, payroll taxation grew from 19% to 25% of national tax revenues and from 5.2% to 9.9%

S72

of gross domestic product (GDP), on average;[1] the payroll tax rate in Sweden, for example, grew from 6% in 1950 to 40% by the late 1970s (Holmlund 1983).

This increasing tax burden has been criticized along a number of dimensions. Critics contend that payroll taxation raises the cost of labor, lowering the competitiveness of a country's producers and leading to underemployment. This contention has found casual support in the persistent high level of unemployment in Europe, where payroll tax rates have risen rapidly since 1960. Payroll taxes will only lower employment, however, if the cost of these taxes cannot be passed on to workers in the form of lower wages. In fact, if there is full shifting to wages, employers will see no net rise in their compensation costs, and there will be no resulting disemployment. Thus, measuring the incidence of payroll taxation becomes critical for assessing the efficiency implications of this form of revenue raising.[2]

Past evidence on the incidence of payroll taxation is quite mixed. Early studies, which relied on time-series or cross-country variation in national payroll tax rates to identify shifting to wages, produced incidence estimates that varied widely. An important problem with such approaches, however, is that of omitted variables bias: there may be contemporaneous time-series changes in other variables which determine wages in a nation, or other cross-country differences in wage-setting institutions, that are correlated with tax rate differences and are not controlled for in the estimation. More recent work in the United States has examined the effects of changes in the costs of government mandated employer benefits within different states over time, controlling for correlated time-series and fixed location effects. These studies have found that the incidence of mandated employer benefits is fully on wages, with little to no disemployment effect.

In this article, I consider the incidence of a dramatic change in payroll taxation in Chile in 1981. Prior to this time, most social insurance programs in Chile were financed by a substantial payroll tax. In 1980, the average payroll tax rate for manufacturing firms was 30%, while the tax rate on workers averaged 12%. Then, in May 1981, Chile privatized its Social Security and Disability Insurance programs, as well as shifting the financing of most other social insurance programs from employer payroll taxes to general revenues. As a result, there was a drop in the average

---

[1] Payroll taxes include social security contributions and other payroll taxes; data from OECD (1993). See Summers, Gruber, and Vergara (1993) for more international comparisons of tax structures.

[2] Also at issue are the efficiency implications of the programs financed by this payroll taxation. I will not take up that issue in this article; see Gruber (1994*b*) for a review of the U.S. case.

payroll tax rate for manufacturing firms to 8.5% by 1982. This large drop in payroll taxation of firms in such a short period potentially provides a fruitful ground for uncovering further evidence on the incidence of payroll taxation.

To study this sharp change in policy, I use data from a survey of manufacturing plants in Chile over the 1979–86 period. This survey provides data on total wages and payroll taxes paid, for two classes of workers (blue and white collar), at several thousand plants in each year. By using data on taxes and wages paid at the firm level, rather than national averages, I am potentially able to surmount the problems of previous time-series studies. I create payroll tax rates for each firm by dividing total tax payments by wages. I then estimate the incidence of taxation by modeling the change in wages at a given plant from before to after the policy change as a function of the change in its average tax rate. Furthermore, if there is extensive shifting to wages, then there should be little employment effect of the policy change; I test this hypothesis as well by using plant-level data on employment.

This methodology, however, faces one important econometric problem: the change in wages is the dependent variable, and the change in the calculated tax rate is a firm-specific function of the change in wages. Thus, any measurement error in wages will lead to a systematic spurious correlation between the dependent and independent variables. The resulting bias will be in favor of finding that wages are falling where tax rates are rising. Furthermore, if there is measurement error in wages, this will also bias the coefficient on the tax rate in the employment equations toward zero. Thus, this spurious variation will lead toward a consistent, but potentially false, finding of shifting to wages.

In the empirical work below, I suggest several different estimators, each requiring different identifying restrictions, as a means of surmounting this problem. These different estimators paint a fairly consistent picture: the incidence of this change in payroll taxation was fully on wages, with little effect on employment.

The article proceeds as follows. In Section I, I discuss the theory of payroll tax financing and review past evidence on the incidence of payroll taxation. In Section II, I provide an overview of the Chilean policy change. In Section III, I discuss the data and empirical framework, and the results are presented in Section IV. Section V concludes.

I. Background on Payroll Taxation

Graphical Exposition

Figure 1 presents a simple demand and supply model of the labor market which depicts the potential effects of payroll taxation. The horizontal axis measures the level of employment; the vertical axis measures

WAGE

FIG. 1.—The incidence of payroll taxation

the wage. The upward-sloping relationship S_0^L represents the supply of labor by workers in a world without taxation; the downward-sloping relationship D_0^L represents the demand for labor by firms in the no-taxation world. The no-tax equilibrium is achieved at (E_0, W_0).

In standard tax incidence analysis, a payroll tax levied on the firm reduces the demand for labor by raising the after-tax cost of employees. The demand curve shifts to D_1^L, reducing the wage that workers are paid to W_1, and reducing employment to E_1; this is the disemployment cost highlighted by opponents of payroll taxation. The magnitude of the disemployment effect will be a function of the elasticities of labor demand and supply.

However, Summers (1989) noted that this analysis missed an important feature of payroll taxation: payroll tax revenues are often used to finance programs which benefit workers only, such as retirement benefits under Social Security or compensation for workplace injuries. This restriction

of benefits to workers creates an important *tax/benefit linkage* which must be modeled as well. In the presence of such a linkage, workers are now receiving higher net compensation than in the pure tax model, since the tax is buying them some benefits. Workers are therefore willing to work harder for a given money wage, shifting labor supply outward to S_1^l. As a result, employment falls only to E_2, while the wage falls further to W_2; there is more shifting to wages. That is, since workers value the benefits that they are buying with their payroll taxes, they will accept lower wages, and this leads to a smaller net rise in compensation costs and thus less disemployment.[3]

This insight implies that, besides the elasticities of labor demand and supply, there are two additional parameters which will determine the efficiency cost of taxation. The first is the extent to which benefits are restricted to workers. If payroll taxation finances benefits available to all, then there will be no supply shift, as work brings no additional benefits beyond the wage. The second is the magnitude of the tax/benefit linkages. Even if payroll taxes are providing some benefits to workers, if they value these benefits at below the cost to the employer, the supply shift will not fully undo the disemployment effects of taxation.

Mathematical Derivation

It is also useful for interpreting the results below to express this intuition in the context of the mathematical model that underlies figure 1.[4] Consider labor demand and supply of the forms

$$D = D(w^*(1 + t_f)), \tag{1}$$

$$S = S(w^*(1 - at_e) + qwt_f), \tag{2}$$

where

w = the pretax wage;
t_f = the payroll tax rate on firms;
t_e = the payroll tax rate on workers;

[3] This general point about tax/benefit linkages goes back much earlier than Summers (1989); see, e.g., Musgrave (1959) and the references therein. Note that Summers made the argument that employers provide particular benefits to their employees with reference to government mandating. But this argument applies equally well to any case where benefits are restricted to workers, regardless of the means of finance (mandate or tax).

[4] This exposition is a variation on the derivation in Gruber and Krueger (1991) for the case of payroll taxation. I am grateful to a referee for suggesting some important refinements for this context.

a = the extent to which employees discount their payroll tax payments relative to cash income; and

q = the extent to which employees value employer payments relative to cash income.

This is a standard model where payroll taxes are levied on both employers and employees, as was the case in Chile. In the case where the social insurance benefits financed by taxation are valued at their tax cost by employees, $a = 0$ and $q = 1$; that is, tax payments by employees are ignored (since they are in essence just paying themselves), and tax payments by employers are treated as cash income.

Solving this model, we obtain the equilibrium condition:[5]

$$(dw/w)/dt_f = \frac{h_s q - h_d}{h_d - h_s^*(1 - at_e)},\qquad (3)$$

where h_d and h_s are the elasticities of labor demand and supply, respectively. This formulation is useful because it lays out the three conditions under which full shifting, a 1% drop in wages for each one percentage point rise in the payroll tax rate or $(dw/w)/dt_f = -1$, can be expected. The first is full valuation of the benefits financed by payroll taxation ($q = 1$ and $a = 0$). The second is an elasticity of labor supply of zero. The third is an elasticity of labor demand of infinity. That is, full shifting can either be due to elastic demand, inelastic supply or due to full tax/benefit linkages. In the empirical work below, I will be estimating reduced form models of the extent of shifting, but to the extent that I find full shifting I will be unable to disentangle the structural cause of this finding; this has important implications for the generalizability of the results, a point to which I return in the conclusion.

Complication—the Minimum Wage

Both the graphical and mathematical exposition above have considered a perfectly competitive labor market with no other imperfections. But this is not true in either the United States or Chile: a potentially important imperfection in either case is the presence of a minimum wage. The analysis has assumed that firms could readily pass on their costs of taxation

---

[5] Following standard tax incidence modeling, this solution is derived from the point of $\tau_f = 0$, which must hold either before or after taxes change. This is approximately true for the Chilean case, as employer contribution rates were driven close to zero by the "after" period in the analysis below. If $\tau_f > 0$ both before and after, then the benchmark for full shifting is below -1 (in absolute value).

to workers in the form of lower wages. However, if workers are already earning the minimum wage, such shifting to wages is not possible.[6]

It is difficult with data on average plant-level wages to assess the importance of the minimum wage in Chile, but it does not appear to be much more of a barrier to shifting than in the United States. In 1980, the year before privatization, the minimum wage was approximately 48,000 pesos per year. The plant-level data show that the average blue collar manufacturing worker in Chile earned 98,000 pesos in that year; the average white collar worker earned 242,000 pesos. In contrast, in the United States, the minimum wage was $4.25 in 1992, and the average blue-collar manufacturing worker earned $10.30 per hour and the average white-collar worker earned $16.10.[7] Thus, the ratio of the minimum wage to average earnings is similar in the two countries. Precise inferences on the relative importance of the minimum wage, however, would require information on both the distribution of wages and the enforcement of the minimum wage floor across the two countries.

Evidence on the Incidence of Payroll Taxes

There is a large literature that models the labor market effects of payroll taxation. Early studies, which relied on time-series or cross-country variation in payroll tax rates, produced a wide range of incidence estimates. In a classic cross-country study, Brittain (1972) found that capital's share did not decline with the imposition of a payroll tax, suggesting full shifting to wages; but Feldstein (1972) compellingly critiqued this conclusion. More recently, Holmlund (1983) used the sharp increase in payroll taxation in Sweden, noted above, to estimate that 50% of the increased costs were shifted backward to wages. This estimate would imply nontrivial disemployment effects of payroll tax increases (depending on the elasticity of labor demand). However, time-series and cross-country studies

[6] Recent evidence on the employment effects of the minimum wage (Card 1992a, 1992b; Card and Krueger 1993) has demonstrated that, contrary to the conventional wisdom, increases in the minimum wage in recent years in the United States have led to *increases* in employment. This suggests that the simple competitive supply and demand model, depicted above, may not be the appropriate one; instead, minimum wage labor markets may be better represented by a monopsony-type model. Under this model, the efficiency cost of employer payroll taxation in the presence of a minimum wage is the same (if the minimum is above the competitive equilibrium wage) or lower (if the minimum is below the competitive equilibrium wage but above the monopsony equilibrium wage, and the tax is small) than in the competitive model (see Gruber 1994b).

[7] These figures tabulated from the Current Population Survey (CPS) (March 1993). Hourly earnings are total earnings last year over total hours worked. Blue collar is defined as craftsperson, operator, or laborer; white collar is defined as executive, manager, salesperson, or technician.

suffer from a fundamental problem: there may be other important omitted variables, which are correlated with both tax rates and wages, but are unobserved to the econometrician. Given slowly rising tax rates over a long period, it may be difficult to disentangle the effect of taxes on wages from the influences of other time-series wage trends. Similarly, high wage countries may choose high tax rates in order to finance more generous government programs. Summers, Gruber, and Vergara (1993) find that the structure of a nation's labor market is strongly correlated with the structure of its tax system, suggesting that taxes may be correlated with wages through a mechanism outside of the incidence model.

An alternative approach is to move beyond time-series and cross-country variation to find payroll tax changes which had differential effects within a country over time. This is difficult to do with most payroll taxes (such as that which finances Social Security in the United States), since they are applied nationally. This approach was used by two recent studies of government mandated increases in employer-provided benefits, where the mandated cost increase differed across states at a point in time and within each state over time; this allowed the authors to control for both time-series trends and fixed location differences in wages. Gruber and Krueger (1991) examined the incidence of increases in the cost of workers' compensation (insurance for workplace injuries) for several high cost industries in the 1980s, and they found that 85% of these costs were shifted to wages, with little disemployment effect. Gruber (1994a) studied mandated comprehensive coverage for childbirth in health insurance plans, which exogenously increased the costs of employing women of child-bearing age, and found that all of these costs were shifted to wages, with no effect on net labor supply.

The applicability of these recent studies to other types of payroll taxation and to other countries, however, is uncertain. As noted above, a key parameter driving the extent to which costs are shifted to wages is employee valuation of the benefits that they are receiving. The valuation of workers' compensation for on-the-job accidents and maternity health insurance may be quite different than the valuation of forced savings for retirement. Furthermore, the ability of firms to shift these relatively small costs to wages may not be relevant for the ability of firms to pass on to workers' payroll taxes which amount to 30%–40% of wages.[8] Finally, the ability of firms to shift costs to wages may differ in other countries; shifting could be constrained or enhanced by the differing structure of labor market institutions across nations.

---

[8] The maternity benefits studied in Gruber (1994a) amounted to approximately 5% of wages for the most costly group. The costs of workers' compensation studied in Gruber and Krueger (1991) averaged 7.5%, with a maximum of 26%.

II. Institutional Background

The dramatic change in the financing of social insurance in Chile in May 1981 provides a unique opportunity to assess the incidence of payroll taxation. Chile's Social Security system began in 1924. By the 1970s, it had evolved into a standard pay-as-you-go defined benefit plan. Social Security "contributions" were paid to one of many "Social Security Institutions" (SSIs) in Chile. The three major institutions provided benefits for most white-collar, blue-collar, and public sector workers. There were also a host of other smaller SSIs for particular groups of workers, but, in 1980, only 5% of private sector workers were in these smaller SSIs.[9] Coverage of the Social Security system was fairly high by developing bp3country standards: 61% of the labor force was covered in 1980. In the manufacturing sector, approximately 84% of the labor force was covered.[10]

The SSIs collected contributions to finance not only old age pensions but a host of other social insurance benefits as well: benefits for disability and maternity leave, family allowances (lump sum annual benefits paid for each child), unemployment compensation, and compensation for work injuries. These benefits were financed by payroll taxes levied on firms and workers. The structure of financing over time is described in table 1. For both blue- and white-collar workers, employers paid payroll taxes on the order of 30% of wages in 1979. The employee portion of the payroll tax was higher for white-collar workers than for blue-collar workers. Both the employer and employee portion of the tax was paid to a maximum level of taxable earnings. In 1980, this maximum was approximately 528,000 pesos per year; this is over 5 times the mean earnings of blue-collar workers in my sample and over twice the mean earnings of white-collar workers.

In terms of the theory laid out above, the benefits financed by employer payroll taxes were all restricted to workers, which raises the possibility of important tax/benefit linkages. However, the level of many of the benefits (i.e., family allowances) were based not on earnings but on legislated standards. Furthermore, while Social Security benefits above the minimum level were related to earnings, approximately 70% of pensions

---

[9] My calculations using data from Superintendencia de Seguridad Social (1980). Foxley (1979) notes that these smaller systems were mostly set up for well-to-do groups of workers in order to avoid the redistributive benefits structure of the larger institutions.

[10] My calculations using data from Superintendencia de Seguridad Social (1980) and Banco Central de Chile (1989). Number of workers covered by sector are only provided for the two largest plans (white and blue collar); I grossed this up to a total coverage figure for manufacturing by assuming that 5% of manufacturing workers were in the smaller plans (following the national ratio).

Table 1
Details of Chilean Social Insurance Financing

	1979		1981		1985	
	Employer	Employee	Employer	Employee	Employer	Employee
Blue-collar:						
Social security	15.95	7.25	0	13/18.9	0	13.5/18.9
Sickness/maternity	0	0	0	3.74/4	0	5.74/6
Work injury	1 + 0–4	0	1 + 0–4	0	.85 + 3.4–3.8	0
Unemployment	2	0	0	0	0	0
Family allowances	10	0	0	0	0	0
Total	28.95–32.95	7.25	1–5	16.74–22.9	4.25–4.65	19.24–24.9
White-collar:						
Social security	10.83	16.67	0	13/19.9	0	13.5/19.9
Sickness/maternity	2.54	1	0	4.55/4	0	6.55/6
Work injury	1 + 0–4	0	1 + 0–4	0	.85 + 3.4–3.8	0
Unemployment	2	0	0	0	0	0
Family allowances	10	0	0	0	0	0
Total	26 4–30.4	16.67	1–5	17.55–23.9	4.25–4.65	20.05–25.9

SOURCE.—U S Social Security Administration (various years)
NOTE —Figures are percentage of payroll Data for employee contributions to social security after privatization give contribution rates for those joining the new system and those staying with the old system

paid in 1979 were minimum pensions (Diamond and Valdes-Prieto 1993), which are a function only of having worked and not of earnings history. The result is that there may have been little marginal tax/benefit linkage for many workers. That is, while employment per se was valuable because it brought the worker under the purview of the Social Security system, additional work did not provide any marginal benefits in terms of increased social insurance entitlement. Nevertheless, since the focus below is on employment and on average wages, there may be extensive shifting on average even if there was little marginal linkage.[11]

Social insurance policy shifted dramatically in May 1981. The major policy change at this time was the privatization of Social Security, with individuals' contributions being saved for their own retirement, rather than being used to finance the consumption of current retirees. Individual retirement savings were invested with one of several competing private pension funds. Upon retirement, savings could be converted to a lump sum payment or an annuity. The details of the privatized system, and a comparison to the previous system, are provided by Diamond (1993) and Diamond and Valdes-Prieto (1993).

Privatization brought with it major changes in the financing of social insurance as well. Privatized pensions and disability insurance were funded by a mandatory contribution of 13% of earnings from both blue- and white-collar workers, with no contribution from employers. The system retained a minimum benefit, which is financed from general revenues. Joining this new system was optional for existing workers and mandatory for new labor force entrants; of existing workers, 85%–90% switched to the new system (Diamond and Valdes-Prieto 1993). Those who had accumulated an obligation under the old unfunded system were given "recognition bonds," which were financed by a large fiscal deficit.

At the same time that Social Security was privatized, the financing of the other social insurance programs was largely shifted from employers to general revenues. As table 1 shows, the only remaining employer contribution after privatization was for work injuries. The financing of sickness benefits (national health insurance) was somewhat shifted from general revenues to payroll taxes on workers at this same time.

Time-series analysis of the effects of this change in financing structure are confounded by a further policy change at the time of privatization: employers were mandated to give an 18% nominal wage increase to workers to finance this shift in the statutory tax burden. However, inflation was fairly high during this period, averaging 25% per year from

---

[11] One implication of this lack of marginal linkage is that the reduction in payroll taxation may have caused a rise in average hours worked. Unfortunately, I cannot explore this in my data.

1979 through 1982. As a result, the large nominal wage rise was not a very binding constraint.

III. Empirical Strategy

Sources of Variation

The data set used for this analysis derives from a census of all manufacturing plants in Chile with more than 10 employees. This size restriction only excludes 15% of Chilean manufacturing employment in 1979, although there is no data on size variation over time.[12] The data contain information on employment, wages, and payroll taxes paid by white/blue collar categories. All amounts are deflated by an industry-specific wholesale price index.[13]

The data do not report payroll tax rates but instead total state mandated "contributions" paid on behalf of each group of workers. I therefore define the payroll tax rate as

$$t_{yt} = \text{TAX}_{yt}/W_{yt}, \tag{4}$$

where

t_{yt} = the tax rate for group i (white/blue collar) in plant j in year t;

TAX_{yt} = total tax payments; and

W_{yt} = total wages.

Relative to the individual-level data used for recent studies of tax incidence, these data have some disadvantages and some advantages. On the one hand, the major disadvantage is that there is only information on the average wage in the firm. This is particularly problematic because the dependent variable for the wage equations is the log of the average wage, and the log of the average wage is greater than the average of log wages. I also cannot control for individual demographic characteristics.

On the other hand, the data have one potentially important advantage: I have information on the actual payroll tax cost of the firm. Most studies of tax incidence have not been able to use firm-level data, and they have

---

[12] My tabulations.

[13] This could be problematic if industry-specific tax changes were shifted forward to prices. There is not a significant relationship, however, between the industry price deflator and the average industry tax rate change; this is consistent with the findings below that taxes were shifted backward to wages.

relied on imputing the average cost to the firm based on, for example, the state of residence. This imputation can be quite noisy.[14]

In fact, there is substantial variation in tax rates across firms. Figures 2 and 3 show the distribution of blue collar tax rates across firms before (1979 and 1980) and after (1984 and 1985) privatization. While privatization led to both a noticeable fall in the mean tax rate and a compression of the distribution, there is wide variation in tax rates in both years. The nonzero tax rate in the later years may reflect continuing contributions for worker injuries. Furthermore, some firms continued to pay family allowances to their workers, being reimbursed by the government ex post.[15]

Of course, the central question for the empirical work is, Is such variation legitimate? That is, does the variation in tax rates, which I observe in the data used below, reflect true differences across firms in the underlying legislated costs of social insurance? Or is it simply capturing other spurious variation in my measure of the firm-specific tax rate? There are seven possible sources of this variation in tax rates:

1. Variation in industry risk for workers' injury compensation costs. Workers' compensation tax rates are a function of both an industry-specific payment and a plant-level payment/discount based on accident rates.
2. Membership in different Social Security institutions, with different tax rates.[16]
3. Noncoverage by Social Security of different shares of the wage bill. While all of the plants in these data are in the covered sector, not all compensation paid incurs a payroll tax obligation. In particular, firms may contract for temporary services, for which a lump sum is paid but for which there is no formal (taxable) service contract. Since these payments are included in the reported wage data, this will lead to variation in the calculated tax rate.[17]
4. Differential child allowance payments across firms. This arises from the fact that firms may lay out their allowance payments and be

[14] For example, Gruber and Krueger (1991) use the average workers' compensation cost in the state and ignore the fact that some firms may receive discounts from insurers.

[15] Francisco Bernasconi of the Chilean Ministry of Finance, personal communication with author, April 1994.

[16] Unfortunately, there is little information available on the contribution rates for the smaller SSIs.

[17] Once again, there are no data which enable me to separate taxed and untaxed wage payments. There is also little data on the use of untaxed temporary workers in either the pre-privatization or post-privatization labor market.

FIG. 2.—Average blue-collar tax rates, 1979–80

reimbursed ex post, and they may only report gross payments (not payments net of reimbursement).

5. Measurement error in tax payments.
6. Measurement error in the wage bill.
7. The maximum taxable earnings may bind differentially across firms.

The first three sources of variation can be broadly classified as "legitimate," that is, they provide plausibly exogenous variation in the tax rate which can be used to identify incidence. Of course, these may not be ideal means of identification, because these factors may be correlated with other omitted variables which drive wages. That is, it may be that the industries that have the greatest risk of injury also pay the highest wages, or those firms which are high wage are the ones which contract most of their services to untaxed temporary workers. Ideally, the rich empirical specification used below, and in particular the use of within-plant wage changes, will be sufficient to control for these potential omitted variables.

The second four sources of variation are "illegitimate," that is, they provide no exogenous information which can be used to identify the incidence of taxation. Noisy information on the true net cost of child allowances and classical measurement error in the tax bill will bias the coefficient on the tax rate toward zero in both the wage and employment equations. The final two sources of variation, however, will induce biases in favor of finding shifting to wages.

FIG. 3 —Average blue-collar tax rates, 1984–85

As wages rise in a firm, a higher fraction of wages will be above the taxable maximum, mathematically generating a negative correlation between average wages and the calculated tax rate. I cannot disentangle this mathematical effect from the behavioral effect of lower tax rates generating higher wages through the incidence mechanism. It is difficult to assess the relevance of this confounding factor without data on the distribution of wages. But a comparison with the United States suggests that this may not be an important problem. In the United States, the taxable maximum for Social Security is much more binding, on average, than that of Chile; it is less than three times as high as the average blue-collar wage and less than twice as high as the average white-collar wage. Yet only 1.4% of blue-collar workers and 13.8% of white collar workers earn more than the maximum.[18] So if the distribution of wages among manufacturing workers is similar between the two countries, this point is of limited importance.

Classical measurement error in the wage bill will cause a similar, but potentially more important, bias toward a finding of shifting to wages. Suppose that reported total wages can be represented by

$$W_{ijt} = W_{ijt}^{*} + h_{ijt},\qquad(5)$$

[18] My calculations using the CPS.

where, as before, W_{yt} is the firm's reported wage bill, W_{yt}^* is the firm's true wage bill, and h_{yt} is the (white noise) error with which that wage bill is reported. The basic regression specification for wages is

$$\log(W_{yt}/E_{yt}) = a + bT_{yt}/W_{yt} + e_{yt}, \qquad (6)$$

where E_{yt} is employment at the plant. In this regression, full shifting to wages would be represented by a coefficient $b = -1$. Substituting in using (5) (and dropping subscripts for convenience), we find:

$$\log[W^*/E + h/E] = a + b(T/W^*)\cdot[1 - h/(W^* + h)] + e. \quad (7)$$

In this model, even if there is no correlation between the true average wage (W^*/E) and tax rate (T/W^*), there will be a spurious negative coefficient on b, as positive innovations to h raise the left-hand side and lower the right-hand side.

One means of addressing this problem of correlated spurious variation in wages and tax rates is to use the employment results as a "specification check" on the wage findings. That is, in the simple labor market model used earlier, if there is full shifting to wages, then there should be no effect on employment. Thus, a finding of both full shifting to wages and a significant employment fall would suggest that the former result is spurious.

This solution has the problem, however, that measurement error in the wage bill will bias the tax coefficient in the employment regression toward zero. The basic regression specification for employment is:

$$\log(E) = a + b(T/W^*)\cdot[1 - h/(W^* + h)] + s. \qquad (8)$$

In this case, the measurement error does not induce a spurious correlation between the X and Y variables. Instead, it has the usual effect of biasing the coefficient on the tax rate toward zero. This assumes, of course, that any measurement error in the data on employment is uncorrelated with the measurement error in wages. If there is a positive correlation between the two types of measurement error, then there will be a systematic negative bias to the tax rate coefficient in the employment equation. Thus, the employment findings do not provide definitive evidence on the extent of bias to the wage equations.

Dealing with Spurious Variation

I will follow two alternative approaches in the empirical work below for dealing with this potential spurious variation from measurement error and the taxable maximum. The first is to assume that there is a

dimension of variation along which the true tax rate varies, but the spurious components of the measured tax rate do not. The true and spurious components of variation in the tax rate can be decomposed along seven dimensions:

$$t_{ijt} = \{a_i + b_j + c_t + d_{ij} + e_{jt} + f_{it} + g_{ijt}\} \text{ true}$$
$$+ \{a'_i + b'_j + c'_t + d'_{ij} + e'_{jt} + f'_{it} + g'_{ijt}\} \text{ spurious.}$$

(9)

That is, both the true and spurious components of the variation in the tax rate for employee group i in plant j in year t can be decomposed into main effects for group, plant, and year; second-level interactions of group and plant, plant and year, and group and year; and a third-level interaction of group, plant, and year. Both the components may vary along each or all of these dimensions. For example, in the absence of the legitimate sources of firm-specific variation, documented above (i.e., industry risk or use of temporary workers), the true tax rate would vary purely along the time-series and worker group dimensions; b_j and c_t would be nonzero, and the remaining true effects would be zero. In that case, all of the variation that is observed across firms in the tax rate would be spurious.

Alternatively, due to the legitimate sources of variation, the true tax rate may vary along all of the above dimensions, and the spurious variation may be more limited. Consider the case where firms systematically overstate or understate their wage bill, but there is no year-to-year or group-to-group variation in that overstatement. Then the fixed firm-specific component of variation ($a_i + a'_i$) in the tax rate will be unidentified, but within firm changes in the tax rate will represent true variation only. This approach could also control for systematically high wage workplaces exceeding the taxable maximum, if the fact that a plant is high wage is largely fixed over time.

Unfortunately, there are no overidentifying restrictions that can be used to identify which of these dimensions of variation are true and which are spurious. Thus, the primary strategy used below will be to estimate the model under different assumptions about the structure of the measurement error. This strategy is useless if the spurious component of the tax rate varies somewhat along *all* of the dimensions of variation listed above. But so long as it does not, if the results are robust to alternative estimators, it suggests that they are not due to measurement error alone.

The second methodology, which is potentially more powerful, is to find an instrument for the tax rate which is uncorrelated with both the (conditional on included regressors) wage determination process and the measurement error in the tax rate. Unfortunately, there are not many

natural candidates for such an instrument. I suggest one candidate, the location of the plant, below.[19]

The Confounding Influence of Employee Contributions

There is one other potentially important misspecification in the models used below. The reduction in the employer tax burden in the early 1980s was funded from two sources: an increase in the payroll tax rate on employees and general revenues. While the latter is likely to be exogenous to a particular firm's tax rate, the former may not be. That is, the firms that saw the greatest reduction in their tax rate, for example, because all of the wage bill in the plant was taxed by the Social Security system, may be exactly those firms where the average worker is seeing the greatest rise in his or her tax rate.

Such a firm-specific correlation between the change in employer and employee contribution rates could induce an upward bias to the estimated extent of shifting to wages. This is readily illustrated by returning to the model above, and by allowing the tax rate on employees, as well as the tax rate on employers, to vary. Solving this model,

$$(dw/w)/dt_f = \frac{h_s^*(q - a) - h_d}{h_d - h_s^*(1 - at_e)} - (dw/w)/dt_e. \qquad (10)$$

The first term in equation (10) is of the same form as equation (3). But there is an additional (nonpositive) term which is a function of the change in wages in response to changes in the tax rate on employees. If any one of the conditions for full shifting is met, for example, if there is full valuation of the tax payments by employers and employees, then this term will be zero; in that case, wages won't change with employee tax payments since, as discussed above, employees will in essence be paying themselves. But if none of the conditions for full shifting is met, then this term will be nonzero. This will lead to an upward bias to the estimated extent of shifting, if there is a correlated change in t_e when t_f changes.

In practice, even if one or more of the conditions for full shifting is met, then there may still be a bias from the second term in equation

---

[19] A third strategy, which was attempted in an earlier draft, was to create a tax rate which was, by construction, spuriously *positively* correlated with wages. This tax rate was created by dividing actual tax payments by a predicted wage bill for the firm; this predicted tax rate would be too high for (residually) high-wage firms, leading to a spurious positive correlation between wages and tax rates. By using this predicted tax rate as an instrument, I could then bound the true shifting estimate. Unfortunately, the instrument was only weakly correlated with the calculated tax rates; as a result, this strategy yielded extremely imprecise estimates which did not permit any useful inferences (although the pattern of coefficients was the same as that reported below).

(10). The reason is that, while the wage data do not include employer contributions, it is unclear whether these data do or do not include employee contributions. If they do not include employee contributions, then the second term will equal zero, as noted above. But if they do include employee contributions, then measured wages will automatically rise as employer contributions fall; but the rise will be an artificial one, arising from increased employee contributions, not higher pretax wages.

To summarize, then, there will be a bias to the estimates if there is a firm-specific correlation between the reduction in taxes on employers and the increase in taxes on employees, and if either (a) there is less than full valuation, somewhat elastic labor supply, and less than perfectly elastic labor demand (i.e., if none of the conditions for full shifting is met) or (b) measured wages include employee contributions. Unfortunately, there is little data on the key parameter here, the firm/employee correlation in the change in contribution rates.[20]

I address this potential problem in two ways in the empirical work below. First, the various estimation strategies used will purge this bias if the correlation between employer and employee contributions is either entirely within firm and constant across worker groups, or entirely within group with no residual common firm component.

Second, I note that this is a much less important problem for white-collar than for blue-collar workers. That is, for blue-collar workers, employer contributions fell by 23% on average over this period, while employee contributions rose by roughly 14.8% (see table 1). In the worst case scenario of full valuation ($a = 0$, $q = 1$), employee contributions being included in measured wages and perfect within-firm negative correlation between employer and employee contributions, this would introduce a negative bias to the shifting estimate of 64% on average. That is, even if there were no response of cash wages for blue-collar workers, the rise in employee contributions would lead to a shifting estimate of 64%

[20] Alexandra Cox-Edwards has suggested to me a reason that the firm-level correlation between employer and employee contributions will be -1: firms may have differed in the extent to which employers themselves paid the employee contributions in the early years of privatization. If some employers paid the employee contributions even after privatization, and if employee contributions are included in wages, then this will lead to some firms appearing to have high taxes and low wages, and others appearing to have low taxes and high wages, but the variation will be solely from this accounting practice. In addition to the solutions suggested below, this problem will be addressed by the use of firm differences from a period before privatization to a period several years after, by which time all firms should have adjusted to the policy change. Figure 3 indicates that firms have fully shifted their contributions to employees by 1984/85: no firms have tax rates in those years that are as high as the minimal level of employee contribution for blue collar workers (as shown in table 1).

in this worst case. However, for white-collar workers employer contributions fell by 24% on average, while employee contributions only rose by approximately 6.3%: in the worst case scenario for white-collar workers, the spurious measured shifting to wages would only be 26%. Thus, to the extent that there is evidence of full shifting for white-collar workers, it suggests that the shifting is not solely the result of a negative correlation between employer and employee contributions.

IV. Results

Time-Series Evidence

The most straightforward means of assessing the incidence of this striking policy change is to examine the trend in wages around 1981. The advantage of this approach is that, since the measurement error in wages almost certainly averages to zero within a year, by grouping all firms into annual averages we can surmount the biases discussed above.[21] There are two major problems with this approach, however. First, the mandated 18% rise in wages confounds the natural employer reaction to this reduced tax burden. Second, Chile went into a massive recession in 1982 (Corbo and Fischer 1993). This might lead to a fall in wages which would offset any natural rise from lower employer payroll tax costs.

The time-series evidence is presented in table 2. In fact, average real wages per worker rose dramatically in 1981: real wages rose 27% for blue-collar workers and 29% for white-collar workers. This is beyond what would be expected from the mandated 18% nominal (9% real) wage increase. From 1982 onward, wages fell despite the continuing fall in contributions, as the recession took hold of the Chilean economy. Contributions fell steadily until 1985, and then rose again somewhat in 1986, despite no increase in mandated social insurance contributions. This may be due to the introduction of a new survey instrument in that year.

These time-series results offer two lessons for the cross-sectional analysis below. The first is that there may be other sources of variation in the dependent variable, wages, which are spuriously correlated with true variation in the tax rate. An obvious example is the time-series variation used in this subsection: wages fell in the early 1980s, despite falling contribution rates, due to the recession. This suggests that, in using different models below, it is important to always control for those elements of variation which are spuriously correlated with wages, such as time-series variation, even if the true tax rate varies along these dimensions as well.[22]

[21] See Angrist (1993) for a discussion of grouping estimators as a class of solutions to measurement error problems.

[22] This is simply a restatement of the omitted variables bias critique of the previous literature noted in Section I.

Table 2
Time-Series Data on Labor Market

Year	N	Blue-Collar Contributions	White-Collar Contributions	Log Blue-Collar Real Wages	Log White-Collar Real Wages	GDP Growth (%)
1979	5,802	.303 (.090)	.304 (.100)	4.36 (.53)	5.07 (.88)	8.3
1980	5,301	.281 (.066)	.278 (.069)	4.39 (.50)	5.06 (.84)	7.8
1981	4,870	.160 (.067)	.159 (.067)	4.66 (.49)	5.35 (.81)	5.5
1982	4,467	.088 (.061)	.084 (.060)	4.67 (.50)	5.39 (.79)	−14.1
1983	4,196	.062 (.044)	.060 (.043)	4.36 (.55)	5.10 (.86)	−0.7
1984	4,373	.060 (.031)	.057 (.030)	4.24 (.55)	5.01 (.87)	6.3
1985	4,329	.054 (.027)	.052 (.027)	4.07 (.55)	4.83 (.87)	2.4
1986	4,170	.075 (.056)	.065 (.055)	4.07 (.58)	4.80 (.91)	5.7

SOURCE.—Columns 2–5 tabulated from Chilean manufacturing data, column 6 from Corbo and Fischer (1993)
NOTE.—Blue/white collar-contributions to social insurance are expressed as fraction of wages Real wages are total wages over number of workers Standard deviations in parentheses Results similar if sample of firms with data on both white- and blue-collar workers in all years is used ·Number of observations for blue-collar workers Number for white-collar workers is about 10% lower

As a result, all of the models below will include time dummies. Similarly, the pooled models will always include a dummy for worker group, since white- and blue-collar workers will have both systematically different tax rates and systematically different wages. Finally, all models will be estimated in differences to control for fixed-firm effects in wages. This is because, for example, it may be that workplaces with the highest fraction of uncovered payroll are also the systematically lowest wage workplaces. In this case, not controlling for fixed-firm characteristics will lead to a spurious positive correlation between tax rates and wage levels.

The second lesson is that it may be useful to restrict the set of years used in the analysis. In the regressions below, I will use only data from 1979, 1980, 1984, and 1985. In this way, I avoid effects of the recession which may not be captured by the time dummies, that is, high injury industries may be more procyclical than low injury industries. I also avoid any problems associated with the change in the survey instrument in 1986. To do so, I create a subsample of firms which report data for all four of these years; the sample size is 3,305 firms per year.[23]

Cross-Sectional Evidence: Differences Regression

The first strategy for dealing with the spurious variation in the tax rate is to present a variety of estimators, which are identified under differing assumptions about the measurement error process. I do so with a sample that contains two observations per plant/year: one for blue-collar and one for white-collar. I run regressions using both this "pooled" sample and samples for white- and blue-collar workers separately.

The first estimator is a differences estimator of the form

$$D \log(W_{yt}/E_{yt}) = a + b_1 Dt_{yt} + b_2 l_{yt} + e_{yt}, \qquad (11)$$

where l_{yt} is a dummy for the group of workers in plant j in year t (blue- or white-collar); this dummy is included in the pooled regressions only. Thus, this regression controls for main effects for plant, group, and year, and the interaction of group*year and group*plant (effects a, b, c, d, and e in eq. [9]). That is, so long as the spurious variation is only along these dimensions, and the true correlation of taxes and wages varies along other dimensions (plant*year or plant*group*year), then this model is

---

[23] Some of these firms are missing data on either white- or blue-collar workers (but not both) in a given year; this leads to the variation in sample sizes in table 3. Note that there will be a problem of "reverse attrition," whereby I miss the employment response from firms that start up in response to any net lowering in labor costs. If my finding of full shifting is correct, however, this is not an important consideration.

identified. This would be the case if, for example, measurement error only arose from firms systematically overstating or understating their wage bills through time.

In order to mitigate the measurement error in the tax rate, I use an "average" difference estimator, where I take the difference of average employment, wages, and tax rates over the 1979–80 period and the 1984–85 period. The results are reported in the first row of table 3. The wage coefficient in the pooled sample is very close to one in absolute value, and the employment coefficient is insignificant and quite small. Both results therefore indicate full shifting of payroll tax changes to wages.

The next two sets of columns examine the results separately for white- and blue-collar workers. For both groups of workers, there is evidence of full shifting to wages, although the shifting coefficient is larger for white collar workers. In neither case is there a significant employment effect; for blue-collar workers, the employment effect is "wrong-signed." The consistency of the full shifting finding across both groups of workers, and across both the wage and employment regressions, is striking.[24]

As noted above, the relative coefficients across the blue- and white-collar groups provide some indication of the importance of any bias from correlated changes in employee contributions: if these correlated changes are important, we would expect a larger shifting coefficient for the blue-collar group. In fact, the shifting coefficient is roughly 50% larger for the white-collar group. This suggests that any bias from changes in employee contributions is small. At the same time, it suggests that there may be heterogeneity in the behavioral parameters underlying equation (3) across these two groups of workers; for example, white-collar workers may perceive stronger linkages between tax payments and benefits. Indeed, the only economic explanation for a wage coefficient that is greater than one in absolute value is that white-collar workers value government social insurance benefits at more than their tax cost, that is, $q > 1$ or $a < 0$. This is possible given the failures in private markets for providing these types of insurance, for example, in the absence of a private real annuity

---

[24] I have reestimated these models using a variety of different "windows", such as 1979–85 and 1980–84. The results are insensitive to the window chosen. In the context of Griliches and Hausman (1986), this suggests that measurement error is not an important problem in the employment equation, since if there was measurement error (that was not highly serially correlated) the signal to noise ratio would rise as longer differences were used. This offers no insight about measurement error in the wage equation, however, since the Griliches and Hausman framework does not apply to the case of correlated left-hand-side and right-hand-side measurement error.

Table 3
Coefficient on Contributions/Wages in Cross-Sectional Regressions

	Pooled		Blue-Collar		White-Collar	
	Wages	Employment	Wages	Employment	Wages	Employment
Basic differences regression	-1 120	008	- 899	190	-1,350	-.183
	(.099)	(.106)	(.108)	(130)	(.172)	(.170)
DDD	-1.022	-.113				
	(.180)	(.165)				
Instrument by other group	-1.412	.131	-.901	031	-1.924	.231
	(.245)	(.260)	(.278)	(.335)	(.397)	(.394)
Instrumental variables by area	-1.561	-.260	-2.005	.127	-1 374	-.574
	(557)	(.593)	(.576)	(.679)	(971)	(.965)
	[40.64]	[20.62]	[27.37]	[32.65]	[26.30]	[8.03]
N	6,066	6,066	3,298	3,298	2,768	2,768

NOTES — All regressions are run as first differences, using data from 1979/1980 and 1984/1985 Pooled regressions include both white- and blue-collar workers, there is an indicator variable for blue-collar included in the regression Difference-in-difference-in-difference (DDD) (Gruber 1994a) regression includes plant effects in the differenced regression. In row 3, tax rate for blue-collar workers is instrumented by that for white-collar workers, and vice-versa. In row 4, tax rate is instrumented by a set of 13 area dummies Standard errors in parentheses, χ^2 test for overidentification in square brackets

market, individuals will be willing to pay more than one dollar for one dollar of Social Security benefits.[25] But it may also be indicative of spurious correlation between wages and tax rates, which is especially pernicious for white collar workers. This highlights the value of using alternative estimators to deal with this potential spurious correlation.

Cross-Sectional Evidence: Differences-in-Differences-in-Differences

In the regression framework (eq. [7]), the effect of the tax rate is identified by two elements of variation: plant*year and plant*group*year. It is plausible that the spurious component of the tax rate varies along the former dimension but not along the latter. This would be the case, for example, if overreporting of wages rose over time, but it did so equally for both white- and blue-collar workers. It would also be the case if there is a within-firm correlation between reduced employer contributions and increased employee contributions (which are included in measured wages), but that correlation is equal for both blue- and white-collar workers.

Under this assumption, then the true component of tax variation is identified in a model of the form

$$D \log(W_{ij}/E_{ij}) = a + bD \log(T_{ij}/W_{ij}) + l_{ijt} + d_j + De_{ijt}, \quad (12)$$

where, as before, i indexes blue/white-collar workers and j indexes plants; d_j is a full set of plant dummies. The regression can only be run pooled across blue- and white-collar workers.[26] In this model, shifting to wages is identified by the relative change in wages within a plant for white- and blue-collar workers, in response to the relative change in contribution rates for those groups of workers. This is the spirit of the "difference-in-difference-in-difference" estimator used by Gruber (1994a).

This estimate is presented in the second row of table 3. While the standard error approximately doubles, both the extent of shifting to wages and the employment effect are virtually unchanged, once again supporting full shifting.

Cross-Sectional Evidence: Within-Plant Instrumental Variables

A likely form of the measurement error in wages is that it is pure white noise coding or reporting error. In this case, the spurious variation in the

---

[25] Indeed, in his study of the incidence of mandated maternity benefits, Gruber (1994a) estimates shifting coefficients in excess of 1 as well (although they were not significantly different from 1).

[26] Chow tests marginally reject the pooling of blue and white collar samples in the earlier regressions, but the coefficients are economically similar.

tax rate would be along the plant＊group＊year dimension, invalidating the assumptions underlying the previous estimators. But so long as such reporting or coding error is not correlated across groups within a plant, it can be corrected by instrumenting each group's tax rate by the tax rate of the other group of workers in the plant/year. Similarly, if any correlation between employer and employee contributions is purely within group, as seems likely given the different nature of employee contributions for white- and blue-collar workers, then instrumenting by the other group's tax rate will purge the model of this correlation.

This strategy will be able to identify true variation from factors such as changes in plantwide workers' compensation costs or the plantwide coverage of the wage bill by Social Security. The model to be estimated is equation (11), but with the tax rate for blue-collar workers instrumented by the tax rate for white-collar workers, and vice versa. The first stage of the model fits quite well; the F-statistic is over 1,000.

The estimates from this instrumental variables procedure are presented in the third row of table 3. For the pooled sample, the shifting coefficient is now greater than one, but it is insignificantly different from one. The employment coefficient is once again small and insignificant. The results vary somewhat across the subsamples of blue- and white-collar workers. For blue-collar workers, the wage coefficient is very close to −1, and the employment coefficient is very close to zero. For white-collar workers, the wage coefficient is almost −2, and it is significantly larger (in absolute value) than −1. Once again, the employment coefficient is insignificant, although it is wrong-signed and fairly large.

To summarize, the results from the previous subsections presented three different approaches to estimating the extent of shifting to wages. Each was identified under different assumptions about the measurement error process. Yet each estimator yielded almost exactly the same inference: payroll tax reductions were fully shifted to wages, with no effect on employment. This effect was stronger for white-collar than for blue-collar workers, suggesting that it is not simply the artifact of within-firm or within-worker group correlation in decreased employer contributions and increased employee contributions. And the fact that, even after instrumenting, the coefficient for white-collar workers remains greater than one suggests further that there may be "overvaluation" of government-provided benefits, such as real annuities, under Social Security.

Instrumenting by Plant Location

The ideal solution to the problem of spurious variation in the tax rate is to find an instrument which is correlated with the true tax rate and uncorrelated with the measurement error or any nonlinearity in W that enters the average tax rate. Such an instrument could derive its power from the legitimate sources of variation described above.

Unfortunately, there is little direct data on the legitimate sources of variation, such as which firms are of the highest risk or coverage of the wage bill by firm. Nevertheless, following the discussion in Section III, there is reason to believe that there will be systematic true variation in the tax rate according to factors such as industry and area. Thus, an alternative approach is to simply instrument by indicators for industry or area, as a grouping instrumental variables strategy (see Angrist 1993). It seems likely that the measurement error in wages should average to zero within broad industry and area groups. Thus, so long as there is some true variation within these groups, and the instrument is uncorrelated with wage setting other than through differential tax changes, then this will provide a means of separating true and spurious variation. In using this approach, one must be careful that the first stage fit if sufficient to overcome the problem of small sample bias discussed in Bound, Jaeger, and Baker (1993) and Staiger and Stock (1993).

I have explored two candidates for instruments: a set of 29 three-digit industry dummies and a set of 13 area indicators. The former set of instruments did not provide a good first-stage fit, but the latter set did; the first stage F-statistic was 16 for the pooled regression, 10.4 for blue-collar workers, and 7.4 for white-collar workers. Note that this identification strategy is pursued within the differences regression framework, so that any fixed effects of area on wages has been differenced away. Thus, the identifying assumption is that there is no reason for there to be location differences in wage changes, besides the influence of area-specific changes in the tax rate. This is similar to the identifying assumptions in earlier studies of mandates in the United States (Gruber and Krueger 1991; Gruber 1994a).

The regressions using area of plant location as an instrument are presented in the final row of table 3. In the pooled sample, the wage coefficient is once again significant and indicates roughly full shifting to wages. The employment coefficients is insignificant, which is consistent with full shifting, but the standard error is quite large. The term in brackets below these estimates is the test of overidentifying restrictions; the critical value (given 12 df) is 21. So the overidentifying restrictions are rejected in the wage equation and are accepted in the employment equation. In the separate white- and blue-collar samples, there is once again evidence of full shifting to wages, although this is mitigated by both the large standard errors and the rejection of the overidentifying restrictions in three of four regressions.

It is interesting to note that in both this and the previous section, instrumental variables (IV) estimation has yielded more negative coefficients than ordinary least squares (OLS). If measurement error were an important problem, one would expect the coefficients to become more positive when instrumented. This result therefore further suggests that measurement error is not causing a spurious finding of full shifting. As opposed to the earlier

estimates, the blue-collar shifting coefficient is higher here, but the white-collar coefficient is still greater than one in absolute value; and the large standard errors make cross-group comparisons difficult in this case.

V. Conclusions

The findings in this article suggest that the shift in financing of social insurance in Chile in the early 1980s did not have important consequences for labor market efficiency. The reduced costs of payroll taxation to firms appear to have been fully passed on to workers in the form of higher wages, with little effect on employment levels. This finding is robust to a variety of different estimators designed to deal with the data problems in the Chilean plant survey. Of course, I have not been able to definitely solve the pernicious measurement error problem which motivated much of the methodological discussion above. But the fact that the IV estimators did not lower the shifting coefficient, as would have been the case if the full shifting finding was due to correlated measurement error, suggests that this problem may not be an important one.

It is tempting to conclude from these findings that there is generally little efficiency cost from financing employee benefits through payroll taxation. But there are at least two factors which may limit the applicability of the results beyond the setting considered, a payroll tax reduction in Chile in the early 1980s. First, the incidence of taxation in a very inflationary environment may differ considerably from incidence in a less inflationary environment. Second, if wages are rigid downward, they may react more flexibly to tax cuts than to tax increases. Nevertheless, the consistency of these findings with those for tax rises in a less inflationary period in the United States (Gruber and Krueger 1991; Gruber 1994a) are suggestive of a general incidence pattern for workplace financing of employee benefits.

Another important limitation in applying these findings more broadly is that the reduced-form approach used cannot disentangle the structural sources of wage and employment changes. In particular, I noted earlier three possible explanations for full shifting: full employee valuation of benefits, so that an outward shift in labor supply is "undoing" the inward shift in labor demand; inelastic labor supply, so that employers are able to pass the full costs of taxation onto workers regardless of their valuation of benefits; and perfectly elastic labor demand. Distinguishing these explanations is crucial for assessing the applicability of the findings to the incidence programs that extend to *nonworkers* also, such as President Clinton's recent health care reform proposal.[27] Doing so would require

---

[27] Since these programs will not have tax/benefit linkages, so their costs will only be shifted to wages if my results arise from inelastic labor supply or from elastic labor demand.

assessing the incidence of similar taxes with different degrees of tax/ benefit linkage (along with some means of measuring employee valuation). This would be a fruitful direction for future work.

Finally, I have only considered social insurance financing in this discussion. More important than the change in financing in Chile in this era was the radical change in the nature of the social insurance system. The literature on Social Security in the United States generates a broad set of predictions for the effects of the program on savings and labor supply, but these predictions have not been convincingly tested in the United States due to limited variation in the parameters of the Social Security system. Chile's privatization provides an ideal source of variation for examining the behavioral effects of social security more broadly.[28]

References

Angrist, Joshua. "Grouped Data Estimation and Testing in Simple Labor Supply Models." *Journal of Econometrics* 47 (1993): 243–66.

Banco Central de Chile. *Indicadores economicos y sociales, 1960–1988.* Santiago: Banco Central de Chile, 1989.

Bound, John; Jaeger, David A; and Baker, Regina. "The Cure Can Be Worse than the Disease: A Cautionary Tale Regarding Instrumental Variables." Technical Working Paper no. 137. Cambridge, MA: National Bureau of Economic Research, 1993.

Brittain, John A. *The Payroll Tax for Social Security.* Washington, DC: Brookings Institution, 1972.

Card, David. "Do Minimum Wages Reduce Employment? A Case Study of California." *Industrial and Labor Relations Review* 46 (1992): 38–54. (*a*)

———. "Using Regional Variation in Wages to Measure the Effect of the Federal Minimum Wage." *Industrial and Labor Relations Review* 46 (1992): 22–37. (*b*)

Card, David, and Krueger, Alan. "Minimum Wages and Employment: A Case Study of the Fast Food Industry in New Jersey and Pennsylvania." Photocopied. Princeton, NJ: Princeton University, 1993.

Carbo, Vittorio, and Fischer, Stanley. "Lessons from the Chilean Stabilization and Recovery." Paper presented at the Brookings Conference on the Chilean Economy, Washington, DC, 1993.

Diamond, Peter. "Privatization of Social Security: Lessons from Chile." Working Paper no. 4510. Cambridge, MA: National Bureau of Economic Research.

Diamond, Peter, and Valdes-Prieto, Salvador. "Social Security Reforms." Paper presented at the Brookings Conference on the Chilean Economy, Washington, DC, 1993.

[28] For an example of work that uses this type of "natural experiment" variation to assess the behavioral effects of social security, see Krueger and Pischke (1992).

Economic Report of the President. Washington, DC: Council of Economic Advisors, 1992.

Feldstein, Martin. "Comment on Brittain." *American Economic Review* 62 (1972): 735–38.

Foxley, Alejandro. *Redistributive Effects of Government Programmes: The Chilean Case.* Oxford: Pergamon, 1979.

Griliches, Zvi, and Hausman, Jerry A. "Errors in Variables in Panel Data." *Journal of Econometrics* 31 (1986): 93–118.

Gruber, Jonathan. "The Incidence of Mandated Maternity Benefits." *American Economic Review* 84 (1994): 622–41. (*a*)

———. "Payroll Taxation, Employer Mandates, and the Labor Market: Theory, Evidence, and Unanswered Questions." Photocopied. Cambridge: Massachusetts Institute of Technology, 1994 (*b*).

Gruber, Jonathan, and Krueger, Alan B. "The Incidence of Mandated Employer-Provided Insurance: Lessons from Workers' Compensation Insurance." In *Tax Policy and the Economy,* vol. 5, edited by David Bradford. Cambridge, MA: MIT Press, 1991.

Holmlund, Bertil. "Payroll Taxes and Wage Inflation: The Swedish Experience." *Scandinavian Journal of Economics* 85 (1983): 1–15.

Krueger, Alan, and Pischke, Jorn-Steffen. "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation." *Journal of Labor Economics* 10 (1992): 412–37.

Musgrave, Richard A. *The Theory of Public Finance.* New York: McGraw-Hill, 1959.

Organization for Economic Cooperation and Development (OECD). *Revenue Statistics of OECD Member Countries, 1965–1992.* Paris: OECD, 1993.

Staiger, Douglas, and Stock, James H. "Asymptotics for Instrumental Variables Regressions with Weakly Correlated Instruments." Photocopied. Cambridge, MA: Harvard University, Kennedy School of Government, 1993.

Summers, Lawrence H. "Some Simple Economics of Mandated Benefits." *American Economic Association Papers and Proceedings* 79 (1989): 177–83.

Summers, Lawrence; Gruber, Jonathan; and Vergara, Rodrigo. "Taxation and the Structure of Labor Markets: The Case of Corporatism." *Quarterly Journal of Economics* 108 (1993): 385–412.

Superintendencia de Seguridad Social (SSS). *Seguridad Social Estadísticas.* Santiago, Chile: SSS, 1980.

U.S. Social Security Administration (USSSA). *Social Security Programs throughout the World.* Washington, DC: USSSA, various years.