

2d : Tax Incidence and Income Distribution

Often, the question of “who pays the tax?” is directed at finding how the burden of a tax varies with people’s income. A particular way of answering that question is calculating how **progressive** a tax is. A progressive tax is one in which the rich pay more, in some sense, than the poor. But the sense in which we measure the burden, and how it varies with income, is a very precise one. In looking at the progressivity of a tax, we look at the **average tax rate**, that is the tax burden as a fraction of a person’s income. That is perhaps not the only way of relating tax burden to income, but it is certainly the method most often used, and it does seem to be the best measure of how much the tax shares vary with income.

By looking at the tax burden as a fraction of a person’s income, we can compare any tax with a **benchmark**, a tax in which each person’s tax liability is exactly proportional to her income. That is exactly the same as saying that the tax burden in this benchmark tax system is constant, as a proportion of a person’s income. A tax will be said to be progressive if the tax burden, as a proportion of income, increases with income. It is said to be **regressive** if the opposite is true.

So for some tax — not just the personal income tax — the effective average tax rate for a person can be defined as the cost of the tax, divided by the person’s income. If this effective average rate increases with income, then the tax is progressive.

Now we can also define the marginal effective tax rate, as how much the tax burden increases with income — that is, the **derivative** of the tax burden with respect to the person’s income. Having the average rate increase with income is **not** exactly the same as having the marginal rate increase with income. The correct measure for looking at progressivity is the average rate, not the marginal.

For example, suppose that we did some calculations, and estimated that the **total** cost of the corporate income tax on a person, B (for burden) appeared to be a function of a person’s income, with the function having the form

$$B(y) = 1 + 0.1y + \frac{y^2}{1600}$$

where y was the person’s yearly income, in thousands of dollars. That means that the average tax rate is the tax burden B , divided by the person’s income y

$$a(y) \equiv \frac{B(y)}{y} = \frac{1}{y} + 0.1 + \frac{y}{1600}$$

whereas the marginal rate is the change in $B(y)$ as a function of y ,

$$m(y) \equiv B'(y) = 0.1 + \frac{y}{800}$$

Saying the tax is progressive in the (correct) average–rate sense is saying that $a'(y) > 0$. That is different than using marginal rates, and checking if $m'(y) > 0$.

In this example, $m'(y) > 0$ always, so you would assess the tax as progressive if you (wrongly) used the marginal rate as your measure. But

$$a'(y) = -\frac{1}{y^2} + \frac{1}{1600}$$

which is actually negative at low levels of income. In fact, this tax would be regressive at income levels below \$40,000 per year, and progressive above it. Looking at the marginal rate would give the wrong assessment about how the average burden varied by income level.

Why is average rate the proper measure? With the tax in the example above, someone whose income was \$2000 per year would have a tax burden of just over \$1100 — more than half of her income! Meanwhile someone earning \$10,000 per year would have a tax burden of about \$2060, or about 20 percent of her income. I would not call that tax progressive at low levels of income. (It is progressive at high levels of income, as that last quadratic term kicks in : someone earning a million dollars a year would wind up with a tax burden of 70 percent of her income, if the formula were accurate.)

Going the other way, a “flat tax” with a high exemption can be very progressive in the average-rate sense, even though it has a constant marginal rate by definition. The flat tax is an example that will be used a fair bit in this course, so that it is a point worth repeating : a tax can be quite progressive, even if it has a constant marginal rate.

How progressive is the actual Canadian tax system? The short answer seems to be “somewhat, but not very”. Figure 5 shows the pattern of tax incidence by income class, as estimated in three studies of the Canadian tax system. On the horizontal axis in the figure, percentiles of the income distribution are measured. (That is, the “40” on the horizontal axis represents someone whose income is greater than 40% of Canadians, and whose income is less than the other 60% of Canadians.) The vertical axis measures the tax burden of the whole Canadian tax system (federal, provincial and municipal), as a fraction of income. So the curves show different measures of $a(y)$ as a function of y . A horizontal curve would mean a perfectly proportional tax system, while a curve which sloped down would represent a regressive tax system.

In that figure (and in figures 6 – 8), “Whalley” refers to results obtained by John Whalley and France St.-Hilaire, using data from 1972, and reported in his 1984 Innis lecture cited on the reading list : Whalley, J. : “Regression or Progression : the Taxing Question of Incidence Analysis”, *Canadian Journal of Economics*, 17, 1984, 654–682 ; “Davies” refers to work on lifetime incidence, using data from 1970, done by Jim Davies, France St.-Hilaire and John Whalley, reported in : Davies, J., F. St.-Hilaire, and J. Whalley : “Some Calculations of Lifetime Tax Incidence”, *American Economic Review*, 74, 1984, 633–649 ; “Vermaeten” refers to work done by Frank and Arndt Vermaeten and the late Irwin Gillespie, using data from 1988, reported in : Vermaeten, F., W.I. Gillespie and A. Vermaeten : “Tax Incidence in Canada”, *Canadian Tax Journal*, 42, 1994, 348–416.

In figure 5, the three estimates of “annual” tax incidence look quite similar. The overall effective tax rate, as a fraction of people’s incomes, appears to increase slightly with people’s

income, but not very much. The results of Davies et al suggest that the very poorest people actually bore a higher fraction of the tax burden, relative to their incomes, than slightly higher income groups. However, Davies and his co-authors argue in their paper that this apparent regressivity at the bottom end of the income distribution is a quirk of the way we measure income. They show that using annual data can be misleading, and that if one attempts to look at the tax burden people face over their whole lifetimes, this apparent regressivity at the bottom disappears, as is shown in the curve marked “Davies (1970) [lifetime]” in the figure.

How were these estimates obtained? In brief, in each case the authors combined some model of tax incidence with data on factor ownership and expenditure patterns by income class. For example, consider the corporate income tax. The Harberger model divides the incidence of this tax between workers and owners of capital. How much of the tax is born by capital owners depends on the elasticity of substitution in production in the different industries, on capital intensity of production in different industries, and on substitutability in consumption among outputs of the different industries. So if one has some estimates of these substitution elasticities, and factor intensities, then one can calculate how much of the corporation tax is born by owners of capital. Then we need to know who owns the capital. The share of the Canadian capital stock owned by different income groups then tells how much of this part of the burden of the corporate income tax is born by different income groups.

In all three of these studies, the corporate income tax has quite a progressive incidence : the effective tax rate (of the corporate income tax), relative to people’s income, goes up with income. (For example, Vermaeten et al find an effective average tax rate of only 0.3 percent of the income of the people in the lowest income group, but is 12.3 percent of the income of people in the very highest income bracket.) This progressivity follows from two causes. First of all, all the authors estimate that capital owners do bear a large share of the cost of the corporate income tax. Secondly, ownership of capital in Canada is concentrated very much among the people in the very highest income brackets.

All three of these studies conclude that sales taxation (the provincial sales taxes, and the federal manufacturers’ sales tax, in each study ; the GST was not introduced until 1991) is quite regressive. That is because they all assume that the burden of these taxes is shifted forward onto consumers. The proportion of income spent on taxable goods and services is much higher for lower-income people, which therefore makes these taxes look quite regressive.

These studies assume that the economic incidence of the personal income tax is exactly its statutory incidence : no shifting is assumed. That makes the personal income tax look progressive, especially in the lower range of the income distribution. Vermaeten et al assume that the property tax on rental accommodation is shifted forward onto renters. That makes the tax look very regressive : the percentage of income spent on housing is much higher, on average, for people in lower income brackets. The other two studies assume that some of the property tax is shifted backward onto owners of land, and of capital. That makes the property tax look more progressive, since land and capital ownership are concentrated in the upper income brackets.

So the first broad conclusion from these empirical studies is that the Canadian tax system is close to proportional, but slightly progressive. The second is that the pattern of tax incidence does not appear to change much over time, at least between 1970 and 1990. But the third conclusion is that how progressive the tax system looks depends very much on the assumptions that are made. More important, significant variation in the degree of progressivity of the system can be obtained by making assumptions which are quite reasonable, and consistent with economic theory.

Figures 6, 7 and 8 all show this sensitivity to the assumptions. Whalley's 1984 article is devoted to showing in detail how these assumptions affect the results. Figure 6 shows his "base case" results, compared with his two most "extreme" results : "sensible" assumptions can make the tax system look extremely regressive — or extremely progressive. Figure 7 shows the same sort of variation, this time using 1988 data, and a somewhat different methodology.

Figure 8 shows the sensitivity of the incidence picture to a few other changes in assumptions. One of those changes was mentioned earlier : looking at lifetime income instead of annual income. Here is why annual income may be misleading in measuring tax incidence. For most people, annual income rises as they age, at least until their late 40s or early 50s. Then income drops sharply in their 60s, as they leave the work force. The life cycle model of saving suggests that people try to smooth their consumption compared to this varying income path : saving when employed to finance consumption when retired.

What does this life cycle path say about measured incidence of sales taxes? A large fraction of people in the lowest income bracket are old people. These people are not saving. In fact, they are running down their savings, trying to finance consumption on low incomes. Meanwhile, the higher income brackets include a lot of people in their 40s and 50s, who are saving some fraction of their income in order to finance their retirement. So annual data show sales taxes as very regressive, if they are shifted forward. Forward shifting means they are born in proportion to people's expenditure on taxed goods. For poor old people, this expenditure will be more than 100 percent of their income, since they are spending more than they earn. For well-off middle-aged people, this expenditure will be a relatively small fraction of annual income, because some of the income is saved.

But this pattern of regressivity fades when we look at lifetime income, and lifetime spending. Over their lifetime, people tend to spend about what they earn — regardless of the level of their lifetime income. (It certainly is true that the very rich tend to leave large bequests, but below that very top bracket, the ratio of lifetime expenditure to lifetime earnings is close to 1, regardless of the level of lifetime income.) That means that looking at lifetime income, rather than a single year's snapshot, makes the sales tax look less regressive. In many respects, lifetime income is a better measure of people's well-being than annual income. So here is another modification of the assumptions, which can certainly be justified as sound economics, which changes the apparent progressivity of the system.

In figure 8, there are as well some other variations. One is to assume that some markets are not competitive, which, in these data, makes the system less progressive. (Although partial

equilibrium incidence does show that, in general, a less competitive market does not necessarily mean more forward shifting.) The other is to take into account transfer incomes. These transfers (welfare, old age security, etc.) are a relatively important income source to lower-income people. In the period under analysis in these studies, these transfers were indexed to the cost of living. So increases in sales taxes would lead to increases in transfer payments, effectively insulating the recipients from the burden of sales tax increases. That makes sales taxes look less regressive, especially at the very bottom of the income distribution. The burden of forward shifting of sales taxes now falls on recipients of wage and capital income. Using a lifetime perspective, and taking account of indexing of transfers, makes the sales tax look more or less proportional, rather than regressive, in its incidence.