

## 2a Tax Incidence : Preliminaries

**Tax incidence** is the description of how the burden of some tax is split among different people or groups of people. It is important to recognize that this analysis is often not as simple as it might seem. It is particularly important to distinguish between **statutory incidence** and **economic incidence**. The statutory incidence of a tax is who the law says pays the tax. That's **not** the subject of this section. The economic incidence of a tax is who *really* bears the tax.

Of course to do that, we need some concept of what it means to bear the tax. The following example may show what it means in principle — and, for the most part, “in principle” is all that is done in this section. Consider the Canadian economy, with some set of government policies in place. Look at the *equilibrium* of that economy, that is the prices of commodities, and of factors of production, in that economy, and the allocation of goods and services among the different people. Suppose now that some government policies change. The main policy change considered here is the introduction of taxes of different types. The policy changes affect the economy, and we get ( eventually ) a new equilibrium. That new equilibrium will have, perhaps, different prices for some commodities and factors of production, and it will have a new allocation of goods and services among people. In principle, tax incidence analysis compares the actual equilibrium allocation of goods and services, before the tax was introduced and the actual equilibrium after the tax was introduced, and looks at the changes in people's consumption. The people whose consumption has fallen because of the introduction of the tax are the ones who are said to bear the tax.

For instance, if the new tax raises revenue of \$500,000, and if we look at the changes in everyone's allocation, and find that for all people except one, their consumption bundle has not changed at all after the introduction of the tax, and if we look at the one remaining person (call him “Bruce”), and if Bruce's new consumption bundle is smaller than his old one, and if the value of his new consumption bundle is lower than the value of his old consumption bundle by exactly \$500,000, then we conclude that Bruce is bearing 100 percent of the new tax, and nobody else is bearing any of it.

In this very extreme example, it does not matter what the legal structure of the tax was, or who actually sent the cheques to the Canadian tax authorities. In the example, the tax enabled the government to collect \$500,000 from the people of Canada ; the government has \$500,000 more to spend ; Bruce has \$500,000 less to spend ; no-one else is affected. So, it certainly appears that Bruce is bearing the entirety of the burden of the new tax.

If instead, in the new equilibrium, the value of Bruce's consumption has fallen by \$300,000, the value of Sheila's consumption has fallen by \$200,000, and the value of everyone else's consumption has stayed the same, then we can conclude that Bruce is bearing 60% of the burden of the tax, and Sheila is bearing the remaining 40%.

So how can this sort of exercise be done in a more practical manner? One way would be to somehow actually calculate the equilibrium allocations for the economy, with and without the taxes. Over the past 40 or so years, many applied economists have worked on computer algorithms

to do exactly that <sup>1</sup>. You can see some references to early examples of that approach in the Whalley survey in the *Canadian Journal of Economics*.

That's not an easy task. We often want to look at the incidence of a *proposed* new tax, one which has not yet been introduced, and maybe one which will not actually be introduced if it turns out not to be such a good idea. So we have to somehow assess what the equilibrium allocation would be if this hypothetical proposed tax were introduced. Alternatively, suppose that we wanted to know the incidence of an existing tax, such as the federal corporate income tax. Then we would somehow have to assess what the equilibrium allocation would have been in Canada, if the federal corporate income tax were not actually in place — even though the corporate income tax has existed in Canada for the past 90 years. In either case, it is not simply a matter of collecting data on people's consumption ; it is a matter of calculating how people's consumption would change in response to some hypothetical change in policy.

Sometimes that can be a relatively easy task. What if the government introduced a \$10 per tonne tax on the import of coffee beans? Two facts about coffee might be relevant for this analysis. One, Canada does not grow any coffee : all our coffee is imported. Two, Canada is a pretty small fraction of world coffee consumption : changes in quantity demanded of coffee in Canada should have hardly any impact at all on the world price of coffee. From those two facts, it can be concluded that the impact of this tax ( or duty ) is simply to raise the price of coffee in Canada by the amount of the tax. If we make a third assumption that the coffee roasting and grinding and distribution industries in Canada are perfectly competitive, then we know that the supply curve of coffee to Canadian retail buyers is perfectly horizontal. We can conclude that the only relevant effect of this tax is to raise the price of coffee in the store. The cost of the tax will be born entirely by consumers of coffee. And then we can calculate exactly how much is the cost of this tax to any Canadian individual using the methodology from the introductory section : the whole business of compensating and equivalent variations was to measure the cost, in dollars per year, to a consumer of having the price of coffee go up by 1 cent per kilogram (or by any other amount).

So we can use what we know of microeconomic theory to calculate the incidence of a tax sometimes, without having to calculate and compare two whole general equilibrium models. To use this theory, in the coffee tax example above, I had to make some assumptions about the way the economy works. The analysis of this section is meant to demonstrate that

- i* the basic diagrams of microeconomics say a lot about tax incidence, and
- ii* we have to make a lot of assumptions.

Worse, what gets concluded about tax incidence is very sensitive to the particular assumptions made.

Another feature of the coffee example, which will carry over to all the analysis of this section, is that the statutory incidence of a tax does not matter in determining the economic incidence. In the example, it does not matter if the duty on coffee import is payable by the firms doing the

---

<sup>1</sup> starting with the work of Mantel and Scarf in the late 1960's

importing, by the Brazilian firms doing the exporting, or by the ships doing the transporting. None of these three groups will bear any of the cost of the tax : under the assumptions I made, it is the consumers of coffee in Canada who bear the cost, regardless of who has the legal responsibility of remitting the money to the government.

To a seasoned microeconomist, this irrelevance of the statutory incidence is an obvious point. But it is a very important one, and one which often gets neglected in debate about public policy. When payroll tax rates increase ( as they have done often in the past few decades ), the increase is almost always accompanied by a fierce debate about how the increase is to be split between employers' and employees' contributions. Typically, of course, the unions argue that employers are far better able to afford the increases than workers. Business groups argue that increasing employers' contributions would hamper Canada's ability to compete on world markets. However, what is being decided is the *statutory* incidence. What the unions and business groups should be debating, and what they care about, is the economic incidence. And under virtually any reasonable assumptions about the organization of labour markets in Canada, the statutory incidence of the payroll tax is **irrelevant** for its economic incidence. That is, what these groups are debating is a policy decision that does not matter at all. The overall impact on workers, employers — and anyone else in Canada — would be pretty much exactly the same, whether the government increases only the portion of the tax paid by employers, only the portion of the tax paid by workers, or any combination of the two.

So the claim I just made is that statutory incidence often has very little economic effect. The argument behind the claim is a pretty simple one. But the simplicity of the example should not be misleading. The point is relevant to the real world.

Suppose, as a simple example, that the government levies a tax on pet food stores, which is administered as follows. Each store must keep records of all the cat food they rent, and they must pay to CRA [that's the federal government agency which collects our taxes] a tax of \$1 for every can of cat food sold. Note that this is an example of a tax in which the statutory incidence of the tax is on the stores selling the cat food. The customers buying the cat food are not liable for the tax, and may not even know it exists. Now suppose that after the imposition of this \$1 per can tax, we observe that the retail price of cat food in Canada goes up by 70 cents a can. If we are fairly sure that this change in the price is due to the imposition of the tax [i.e. that not much else changed, except for the introduction of this tax], then we have the answer as to the incidence of the tax. In this case, even though the statutory incidence of the tax is on the stores, the stores are bearing only 30 percent of the tax, and the customers buying the cat food are bearing 70 percent of the tax.

Why? The net revenue per can that store owners are getting has actually fallen by 30 cents in the example : they get 70 cents more from customers per can, but then must remit the \$1 tax to CRA. Since the tax is \$1 per can, and pet food stores have lost 30 cents per can of cat food, they are bearing 30 percent of the cost of the tax, and the customers, whose cat food costs have gone up by 70 cents a can, are bearing the other 70 percent. [70 cents/\$1=70 percent]

On the other hand, suppose that the CRA administers the tax in a completely different manner, so that the statutory incidence of the tax falls on cat “owners”. That is, each customer must pay a tax of \$1 each time she buys a can of cat food. Perhaps it is the store which must help to administer the tax, but it is the customer who must pay it. For example, imagine that every Canadian were issued a card, and the pet food store required to scan the card each time a person bought a can of cat food, and then people would have to pay a tax annually based on the number of times the card was scanned : \$1 for each scan. Now imagine, if you will, that as a result of the introduction of this tax, on retail buyers (instead of the previous example, a tax on the stores), that the price of cat food at stores in Canada falls by 30 cents per can. If this were the case, then store owners would be bearing 30 percent of the cost of the tax, since they see their net revenue per can fall by 30 cents. Customers bear 70 per cent of the cost, since they are paying a \$1 tax, but are paying a price to the store which is 30 cents lower.

In the story the statutory incidence did not matter for economic incidence. However, to make the story work, the price changes had to be fixed just right in the two versions : a 70 cent increase if the stores paid the tax matched with a 30 cent decrease if customers paid the tax. That may seem a convenient coincidence. But it was not coincidence at all. What determines how much the price adjusts in response to the tax is the nature of the market for pet food : how competitive it is, what is the elasticity of demand, and so on. It is these market fundamentals which determine the changes in the price. So if we know that the market, for example, is a three-firm Cournot oligopoly, and that the own-price elasticity of demand for cat food is 0.6, and that there are constant returns to scale in the firms’ costs, then we can figure out how the price changes under the two different tax rules. And for virtually every market structure that we know of — including perfect competition most notably — if the price increase is 70 cents when stores must pay the tax, then the price decrease **must** be 30 cents when customers must pay the tax. In the first story, the tax shifts up the costs of firms, leading to an equilibrium with a higher price. In the second case, the tax on customers shifts down their demand curve, leading to a lower price in equilibrium.

Before getting down to the formal analysis, some preliminary points. Who can bear taxes? In tax incidence analysis, the question is which **people** bear the taxes. It’s only people — and groups of people — that concern us here. Firms are not people, and firms cannot bear taxes. That is, if as a result of a tax, the profits of some corporations fall, economists do not conclude by saying that these corporations bear whatever share of the tax implied by this fall in profits. We need to trace the incidence one step further, and attribute these costs to the owners of these corporations.

This is not just a minor theoretical nicety. With several fairly important taxes — such as the property tax and the corporate income tax — one impact of an increase in the tax rate is a fall in the general rate of return to capital. That is, the rate of return earned on financial assets in Canada will be affected by these taxes. To see how this fall affects the distribution of income in general, we need to know the distribution of capital ownership in Canada. To the extent that capital is held predominantly by people in the higher part of the income distribution, then this part of the costs of the taxes (the reduction in the net return to capital) is born predominantly by

upper income people. Table 2 of Whalley's 1984 article gives the distribution of capital income by income class in Canada. The table shows that people in the highest income class do derive a much higher percent of their income from the return to capital than do people in lower income classes. Thus, a tax which lowers the return on capital will be born much more by high-income people than a tax which lowers the return to labour.

In assessing tax incidence, we often we have to look at both sides of the ledger, **sources** and **uses** of income. How can a tax increase lower a person's real income? One way is on the source side, by lowering the total amount of wages or salary she earns. The other way is on the use side ; if a tax increases the overall cost of goods and services that she buys, then that too lowers her real income. Put otherwise, a 25 percent increase in my nominal income is exactly the same as a fall of 20 percent in the prices of all the goods and services which I buy.

For example, consider the manufacturers' sales tax which we used to have in Canada. It probably had significant effects on both the sources and uses side. Since it was levied on manufactured goods, one effect was probably to raise the price of manufactured goods, relative to the prices of other goods and services. This in turn caused a reduction in demand for the output of manufactured goods. This "output effect" should have led to a decrease in the returns to factors used most intensively in the production of manufactured goods. Thus, how much of the manufacturers' sales tax you bore would have depended both on the importance of manufactured goods in your consumption bundle, **and** on the fraction of your income earned from owning inputs used relatively intensively by the manufacturing sector.

An issue which often complicates tax incidence theory is the other side of the government's ledger, the side that gets covered in AP/ECON 4080. Can we ignore government expenditure when analysing tax incidence? The government cannot simply increase a tax, with nothing else adjusting in the government accounts. For the most part, since this is a microeconomic course, I assume that the government **balances** its budget : tax revenues equal government expenditure. So if the government increases a tax, and its tax revenues increase, so will government spending. In looking at the incidence of a tax increase, do we look at the effects of the tax side of the government's ledger alone, or do we look at both the direct effects of the tax, and the effect of the spending that the tax finances?

We can take either approach. The second approach — looking at both the effects of the tax, and the effects of the expenditure it finances — is called **balanced-budget incidence**. The first approach — looking just at the taxes collected — is called **differential incidence**. The reason for this latter name comes from the way we are able to ignore the expenditure side. What is done in differential incidence is to look at the **difference** in incidence between the tax change in question, and some alternative way of raising the money, an alternative which serves as a sort of reference point. For instance, we could look at how the costs of raising a certain amount of revenue through a corporate income tax increase compare with raising the same amount of money through an increase in the personal income tax. Here we are comparing two ways of raising the **same** amount of revenue. So it is all right to ignore the expenditure side : government expenditure here

will be the same with either tax, so what matters is the difference in tax burdens across income classes between the two taxes. For example, this approach could assess whether using the property tax to finance the local public sector imposed more of a burden on low-income people than using a proportional income tax. That's a classic question in differential incidence.

Another way of doing the same thing, but using balanced budget incidence, would be to examine the effects of increasing the corporate income tax, with the proceeds of the tax increase simply distributed in cash to people, say proportionally to their incomes, or perhaps independent of their incomes. Effectively, that's the same thing as doing differential tax incidence, with the other "reference" tax being either proportional to or independent of, people's incomes. In other words, we often do not know exactly what the government will be doing with the tax revenue. We know the tax revenue goes somewhere. So if we just make the simplest possible assumption about where it goes [for example, that it is simply returned to taxpayers, in proportion to their incomes], we can concentrate on the particular features of the tax itself.