Basic Welfare Economics and Optimal Tax Theory

There are two criteria by which economists measure the outcomes of tax policy:

1. Efficiency, which is traditionally the purview of economics, and does not involve ethical and normative judgments. Efficiency considers only how resources are allocated

2. Equity, which considers the distribution of resources. We will need to refer to social norms and value judgments to get conclusions based on this measure.

Economists can trace the implications of specific value judgments for social welfare, but economics does not give guidance on which value judgments are the correct ones.

Before discussing the efficiency implications of tax systems, we will need to lay out some basic concepts. The terms deadweight loss (DWL) and excess burden are used to describe the burden on individuals in addition to the revenue collected. DWL is a general term that can be used to refer to trade restrictions, market failures, etc., but excess burden is used specifically when discussing taxes. Optimal taxation generally refers to tax policies that minimize the DWL of the tax system. However, we may also want to consider what is optimal, or best, from the equity point of view. The goals of both efficiency and equity may not be served by the same tax system. For example, lump sum taxes, which are based purely on endowments and not behavior, are efficient but not equitable. They are optimal from the efficiency standpoint, but not from both.

Indifference curves describe the preferences of individuals or households. We generally assume that preferences are given, although this is arguable. Behavioral economics, which is an emerging field, specifically considers the possibility that preferences may change over time. For example, a procrastinator may prefer not to do something right now, but three days later, would have preferred to complete the task earlier. Such considerations can be very important for public policy issues such as saving for retirement, smoking, and the related tax policy. Forcing people to save more than they prefer makes them worse off if people know and act on their own preferences, but if not, forced saving may be welfare improving. However, we will set aside these possibilities for now and assume that preferences are stable.

An indifference curve will describe all of the combinations of apples and oranges that make an individual equally happy. In Figure 1, points A and B make the individual equally happy, but point C is preferred to both. Measures of utility are not unique, although the underlying preferences are. The numbers labeling the utility curves are convenient, but they are meaningless in and of themselves. It is only the underlying ranking of bundles that is uniquely representative of preferences.

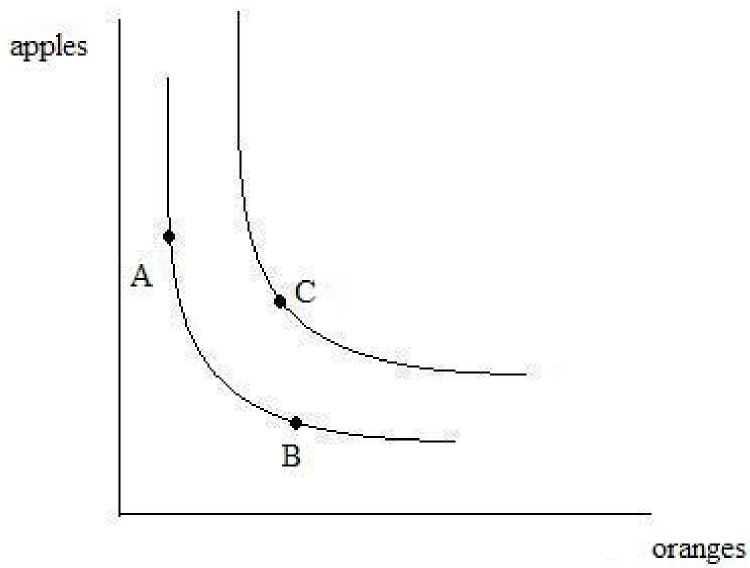

Figure 1

The slope of a line tangent to a point on the indifference curve is the marginal rate of substitution (MRS), or how many apples an individual will give up for an orange. If the slope of line at point A is -2, the individual will trade two apples for one orange and be equally well off. The steeper the curve, the more apples the individual will give up to gain an orange at that point. The convexity of the curve, that the slope increases as when moving from right to left, indicates that there are diminishing returns. If the individual has many oranges, he or she is willing to give up a lot of them to get just one apple.

If two goods are close substitutes, the indifference curves will be very flat. A common example is ten and twenty dollar bills. Generally, an individual will give up two tens for a twenty and be equally well off. If two goods are complements, like right and left shoes, adding additional units of one good alone carries no benefit. In this case, the indifference curves will look like those in Figure 2. Indifference curves need not be parallel. Non-parallel indifference curves imply that the consumption bundle changes as income rises. In figure 3, as income increases, shifting the budget curve out in parallel fashion, consumption shifts toward oranges. Preferences are stable in this case; it is just the fraction of income spent on goods that changes. Thus, the assumption that the poor and the rich have the same preferences can be consistent with their differing consumption bundles. The same tastes can, in different circumstances, lead to different choices.

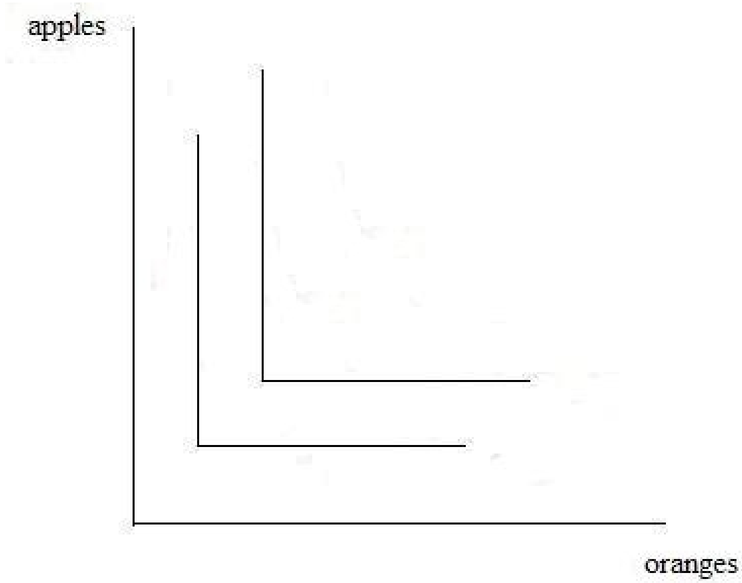

Figure 2

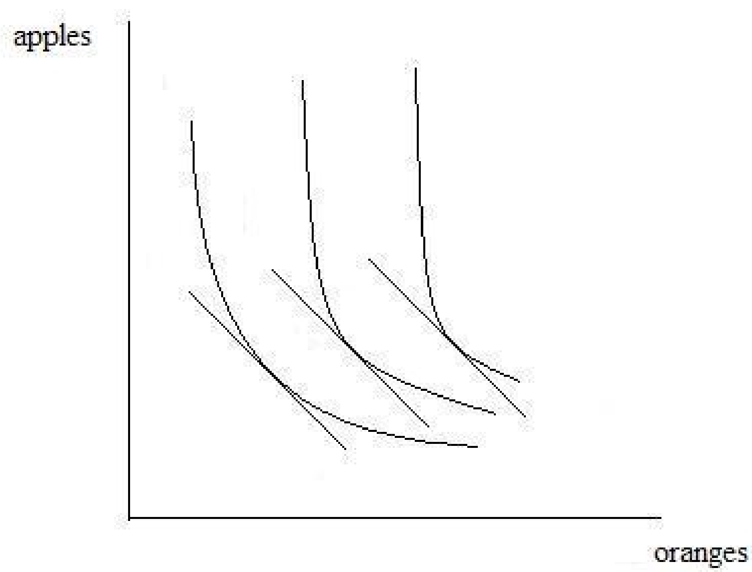

Figure 3

To consider the efficiency of an allocation, we will need to look at preferences of more than one person and decide if we can reallocate resources by trade for a more efficient outcome. A tool for combining the indifference curves of two individuals is the Edgeworth box (Figure 4). The size of the box is determined by the total amount of resources in the economy, in this case, the number of apples and oranges. We assume that the total amount of resources is given; that is, we are concerned only with how to allocate those resources between the two individuals. Any point in the box summarizes the distribution of apples and oranges between the two. At O_1, person one has neither apples nor oranges; person two has all of both. Utility is increasing for person one as the indifference curves move out from that point in a northeasterly direction. For person two, utility is increasing as the indifference curves move out from O_2 in a southwesterly direction.

In this simple economy, there will be a Pareto efficient allocation of resources when the indifference curves are tangent; that is, when they have the same slope and thus the individuals have the same marginal rate of substitution between apples and oranges. A Pareto efficient allocation of resources is one in which it is not possible to make one person better off without making the other worse off. This criterion immediately rules out allocations in which some of the resources are discarded and not distributed to the individuals; allocating the discarded resources to either individual would make that individual better off without hurting the other.

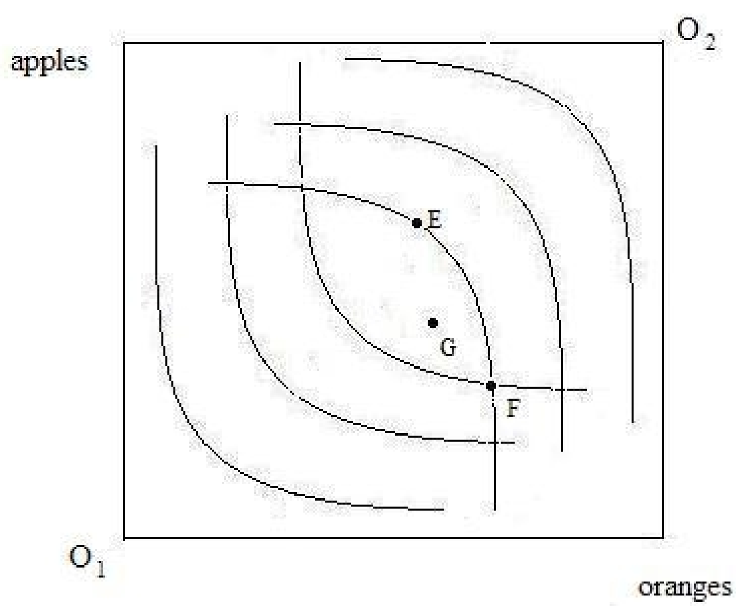

Figure 4

We can show that tangency of the indifference curves must prevail in a Pareto efficient outcome by counterexample. Suppose that an allocation is described by point F. Is this allocation efficient? Point F is feasible; it does not distribute more resources than exist in the economy. At point F, the indifference curve for person two is steep; he will give up a relatively large amount of apples for one orange. Individual one's indifference curve, however, is flat; he does not need many apples to give up one orange and be equally well-off. Thus, exchanging some of two's apples for one of one's oranges would make both better off. Such a trade could be represented by a point G, which is "above" both initial indifference curves. In fact, any point in the area between the two crossed indifference curves, by a similar argument, will be welfare improving for both individuals. This area is referred to as the core; it encompasses all points that are better than the initial allocation, or put differently, all trades that the individuals would be willing to make. Since at point F, welfare-improving trades exist and the core is non-empty, point F cannot be a Pareto efficient allocation. The core will be empty only if there is no area between the two indifference curves; that is, if the indifference curves are tangent and thus have the same slope. The line that traces out all such points of tangency in the Edgeworth box is referred to as the contract curve.

A numerical example may be helpful. If the marginal rate of substitution for person one is 1/1, he will be willing to trade one apple for one orange and be equally well-off. If the marginal rate of substitution for person two is 2/1, he will trade two apples for one orange. If two gives two apples to one and receives one orange in return, he is as well-off as he was before the trade. Person one, having received two apples and given up only one orange, is better off. The same argument can be made if person two's MRS is 1.5/1, or 1.25/1. Only if the marginal rates of substitution are equal will there be no Pareto-improving trades. The lack of any further gains from exchange is the first requirement of an efficient outcome. Note that gains to trade can arise from different preferences or different allocations.

An efficient outcome requires not only efficiency of exchange, as described above, but also efficiency of production. So far we have assumed that the amounts of apples and oranges in the economy are fixed. However, there may be a range of combinations of apples and oranges that the economy is capable of producing. This can be represented by the production possibilities frontier (Figure 5). Combinations of apples and oranges inside the frontier are feasible, but since it would be possible to produce more of one good without reducing the quantity of the other, they are not efficient. The frontier has a concave shape because we assume that there are diminishing marginal returns in production. Each additional dollar of resources devoted to producing oranges produces fewer additional oranges than the previous dollar. For example, while it is relatively easy to produce apples in New York and oranges in Florida, continuing to move resources toward orange production will eventually require using land in New York to produce oranges. Producing oranges in Florida required very little reduction in apple production because the land is better used for oranges, but shifting New York land to orange production would require losing a lot of apples to produce oranges. The slope of a line tangent to the production possibilities frontier is the marginal rate of transformation at that point. The marginal rate of transformation (MRT) is the rate at which additional apples can be produced from the resources freed up in producing one less orange.

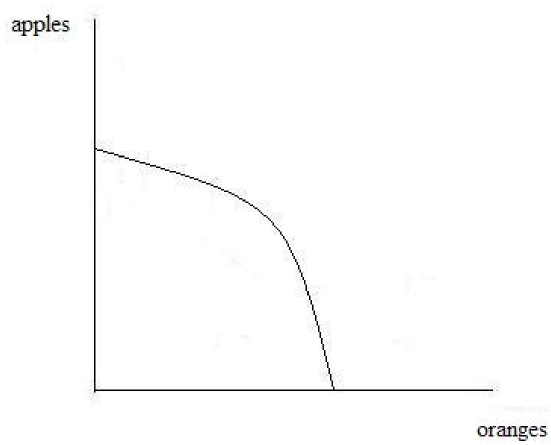

Figure 5

Efficiency of production requires both that production be on the frontier, but it also requires that the marginal rate of transformation be equal to the two individuals' marginal rates of substitution (which, as discussed above, must be equal to each other to achieve efficiency of exchange). Again, the proof of this claim is by counterexample: Suppose that there is an allocation in which the marginal rate of transformation is two, and the marginal rates of substitution are one. We could take away one orange from each individual and use those resources to produce four apples. Distributing two apples to each individual makes them better off than they were before (since trading one orange for one apple made them equally happy). Since it was possible to make an individual better off without making the other worse off, the initial allocation cannot have been Pareto efficient. It is only when the marginal rate of transformation is equal to the two individuals' marginal rates of substitution that no Pareto improving allocations can be achieved. The result can be generalized to an economy with more than two commodities; in this case, producing on the production possibilities frontier and distributing all output to individuals are still required for an efficient outcome, and the equality of marginal rates of transformation and substitution must hold for every possible pair of commodities.

Not every point on the contract curve, which by definition are efficient with respect to exchange, will also be efficient with respect to production. Moving along the production possibilities frontier will change the shape of the Edgeworth box since it changes the amount of apples and oranges available in the economy. Each box may have a fully efficient (with respect to both production and exchange) point within the contract curve. However, each point on a contract curve involves different levels of utility for the individuals. We can put together all fully efficient points and compare the utility of each using the Pareto frontier (Figure 6). Points inside this frontier are clearly not efficient, as moving to the frontier would make both individuals better off. The Pareto frontier uses the numerical representation of utility, which is arbitrary, so we cannot draw conclusions about actual welfare of the individuals from this analysis. However, changing the arbitrary assignment of numbers to utility levels would not change the efficient allocations. The Pareto efficiency criteria imply that points on the frontier are preferable to points within it, but they do not offer guidance on which of those points is the preferred outcome. Points A and B in Figure 6 are both on the frontier, but they involve very different levels of utility for the individuals.

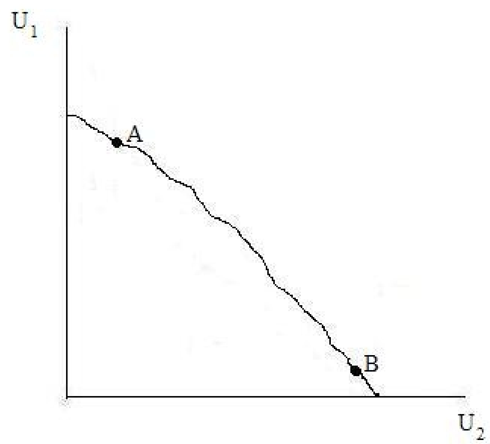

Figure 6

First Fundamental Theorem of Welfare Economics

The First Fundamental Theorem of Welfare Economics states that if both producers and consumers are price-takers, there is full information and no transaction costs, a competitive market outcome is Pareto efficient. We show this graphically using supply and demand curves for oranges, with the price measured in apples:

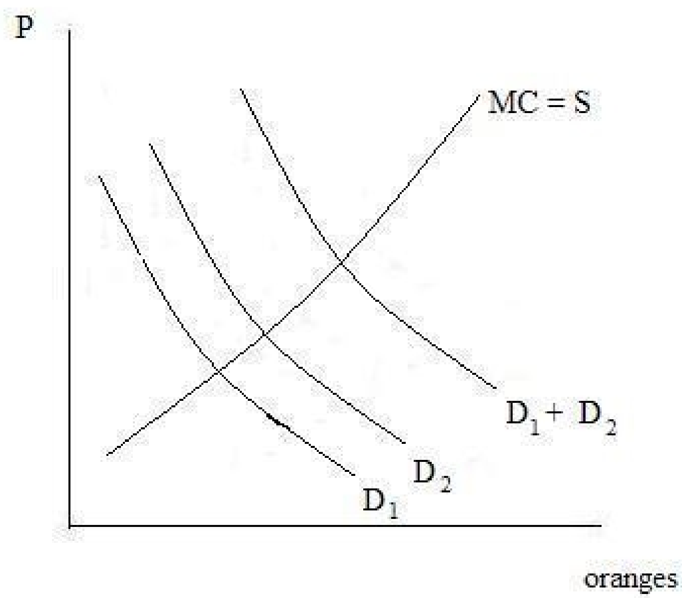

Figure 7

For each individual, we can graph the marginal rate of substitution (Figure 7), or the number of apples given up to obtain an orange. As the number of oranges increases, fewer and fewer apples will be given up. That is, as the number of oranges increases, the willingness to pay for an orange, or the price of oranges in apple terms, will decline. This willingness to pay curve is simply a demand curve. If we assume that the two individuals transact at the same price, we can add their demands at that price together. Because the price is the willingness to pay, which is in turn the marginal rate of substitution, the assumption of same price ensures that the MRS's are equal. We can similarly translate the marginal rate of transformation into a cost curve. The marginal rate of transformation is simply the number of apples given up to generate oranges, which is a price. This marginal cost rises as the quantity of oranges rises, because of diminishing returns to production. Since competitive markets have marginal cost pricing, this marginal cost curve will be equal to the supply curve in a competitive market. At the intersection of the joint demand and supply curves, the price is equal to the marginal rate of substitution for person one, which is equal to the marginal rate of substitution for person two, which is equal to the marginal rate of transformation. Since the MRS's are equal, this outcome will have exchange efficiency. Since the MRT is equal to both MRS's, there is also production efficiency. The competitive market outcome is therefore Pareto efficient.

The result seems simple, but consider the amount of information required if a social planner was required to implement a Pareto efficient outcome. The planner would need complete information about the preferences of individuals and the production processes of producers. The First Welfare Theorem states that this is unnecessary, because the market itself will create a Pareto efficient outcome.

Of course, the conclusion of the theorem holds only when the assumptions are satisfied. For the theorem to hold, both producers and consumers must take prices as given. If there are monopolies or oligopolies in which producers can restrict production to raise prices, the market outcome may not be Pareto efficient. Incomplete markets will also violate the assumptions. A common example is pollution. If a plant releases pollution into a river that is used by another producer, but there is no market in which the plant pays for the right to do so, the first welfare theorem will not hold. Violations can also occur through the structure of production. If production is declining in marginal cost, a firm cannot make a profit by producing until price is equal to marginal cost; money is lost on every unit but the last one. Public goods like national defense are a common example of production with declining marginal costs. In all three of these cases, market failure means not only that the first welfare theorem does not hold, but also that the information provided by the market is missing. Without this information, it can be difficult to correct the market failure. If there is no market for pollution damage, how can the right price to charge for pollution rights be chosen? If there is no market for national defense, how can the willingness to pay and the optimal amount of provision be known? Without the pricing and quantity information normally provided by markets, it is difficult to know the costs of the market failure and the optimal intervention. If the losses from market failure are small, intervention on the basis of imperfect intervention may actually be more costly.

The Second Fundamental Theorem of Welfare Economics

The first theorem of welfare economics states that under certain circumstances, the market will select a Pareto efficient outcome. However, it does not specify which Pareto efficient outcome. In general, there will be many Pareto efficient outcomes, involving different levels of relative utility for individuals, which may depend on their initial endowments. To choose between them, we will need to make value judgements. A common way of representing such value judgements is the social welfare function, written as W(U_1, U_2). This function describes social preferences over outcomes for each individual. In Figure 8, the outcome on the Pareto frontier represented by point B is on a higher social indifference curve and thus preferred to point A. As drawn here, the social welfare function will choose a Pareto efficient outcome. Social welfare increases if the utility of either individual increases while the other stays constant. There are many ways to write the social welfare function, each of which bring their own underlying value judgements. These will be discussed further when we consider the distributional consequences of allocations.

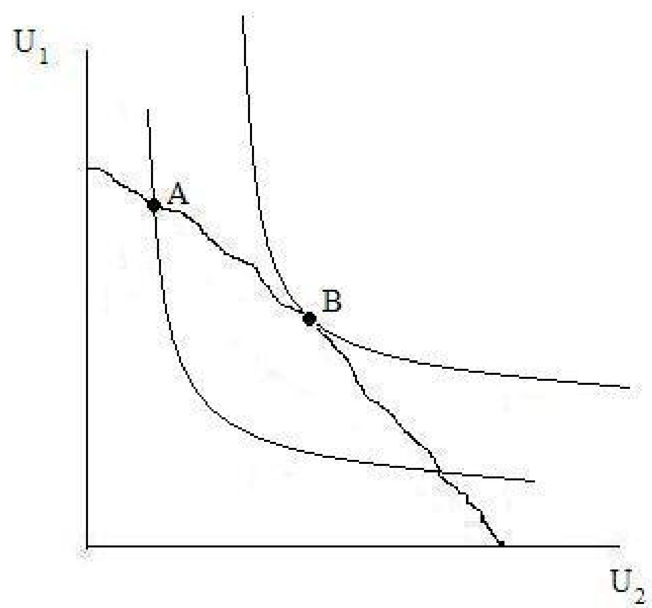

Figure 8

The second fundamental theorem of welfare economics states that any Pareto optimum (chosen by maximizing social welfare or otherwise) can be achieved through the competitive market if the government can engage in lump sum taxation. Shifting initial claims to resources, or endowments, allows any preferred

Pareto efficient outcome to come about through the market. Unfortunately, lump sum taxation requires a lot of information. True lump sum taxation would tax only innate ability, regardless of behavior, and is therefore nondistortionary. In practice, this is impossible. Ability is both hard to quantify and hard to measure. A tax on income, for example, is not a substitute for a tax on ability because it alters the incentive to work and therefore affects behavior. It is not a lump sum tax. The second welfare theorem asserts that any Pareto efficient outcome is possible through workings of the market with the right transfers, but it requires the use of a tool that is not available in practice. Thus, the study of optimal taxation is the study of other ways to transfer resources with minimal distortion, compared to the yardstick of the efficient but unavailable lump sum tax.

Theory and Measurement of Deadweight Loss

The concept of deadweight loss (DWL) measures the distortion created by non-lump sum taxes. Deadweight loss is an efficiency term, not a distributional term; it is the total cost to society of raising a certain amount of revenue. Put another way, it is what would be gained if distortionary taxation was replaced with a nondistortionary lump sum tax raising the same amount of revenue. Figure 9 is a graphical depiction of measuring deadweight loss. The initial equilibrium at P_0 and Q_0 is a Pareto optimum. The area beneath the demand curve and above the price paid is the consumer surplus. For each unit until the last, the willingness to pay given by the demand curve exceeds the price; for the last unit, the consumer surplus is zero. Similarly, the area above the supply curve and below the price is the producer surplus, the difference between the marginal cost and the price paid for the unit.

Adding a per unit tax T, paid by the consumer (we will discuss who actually bears the burden of the tax when we discuss incidence), increases the price paid by the consumer and thereby decreases the quantity traded in the market. Consumer surplus falls, for two reasons. First, for units still purchased, the consumer pays a higher price; this loss is area A. In addition, the tax raises the price higher than the willingness to pay for some units that were previously consumed, so these units are not purchased and consumer surplus falls by area B. The analysis is similar for the fall in producer surplus due to the tax. The producer gets a lower price for the units still produced (area C) and no longer produces some units because the surplus would be negative (area D). The total loss from the tax in the market, for both producers and consumers, is A + B + C + D. However, this is not the overall loss to society, since the government does receive revenue from the tax. The deadweight loss caused by the tax will be the total loss of surplus less the revenue gained from the tax. To think of it another way, the deadweight loss is the additional cost imposed by the tax's distortion of behavior, or the difference between the loss in surplus from a distortionary tax and a revenue-equivalent nondistortionary lump sum tax. To calculate the DWL, note that the revenue earned by the tax is the amount of the tax multiplied by the number of units sold in the market when the tax is in place. Graphically, it is a rectangle with height T and width Q_1, the area A + C. The deadweight loss is then (A + B + C + D) - (A + C) = B + D.

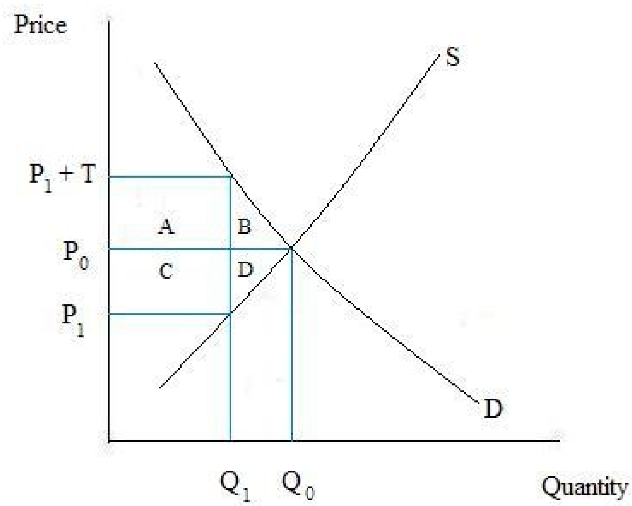

Figure 9

We can approximate the deadweight loss by assuming that the demand and supply curves are straight lines. The DWL is then a triangle with height T and base (Q_0 - Q_1):

$$DWL = -(\tfrac{1}{2})\Delta T \Delta Q$$

Note that a subsidy will also have positive deadweight loss; the sign of the change in the tax rate is negative and the change in quantity is positive. This illustrates that the loss is not from a tax's discouraging consumption, but from its distortion of consumption.

The correct demand curve for this analysis is the compensated demand curve, which accounts for the fact that changes in prices make people worse off since they are like a decrease in resources. There are two reasons that the quantity demanded changes when the tax is introduced. The quantity demanded falls partially because the government has appropriated some income, and decreased income will lead to decreased demand for all goods, not just the one being taxed. This income effect would be present even if the government could use lump sum taxation to get revenue. Quantity demanded also falls because the relative price of the taxed good has risen, and consumers will substitute away from the taxed good to other forms of consumption. This substitution effect captures the distortion of the choice between the taxed good and other goods. Using the compensated demand curve ensures that the DWL calculation captures only the substitution effect, the loss due to the distortion of consumption choices. The thought experiment here is of comparing two different tax systems, each raising $1 trillion in revenue. Both tax systems make individuals poorer by $1 trillion (the income effect), but one may be more distortionary than the other. Why does it make sense to "ignore" the income effect when calculating the effect of taxes? We assume that the income effect of the loss is offset by the income effect of the revenue, which is essentially rebated back to the individual. As long as the benefit of revenue is the same across types of taxes, there are gains to substituting an efficient tax for an inefficient one. If you always buy 50 apples, and the price of apples rises by ten cents, it's like losing $5. But even if I return $5 to you, you will probably still buy fewer apples than you did before, because the price of apples has risen relative to whatever other fruit you might buy.

There are two cases when consumption choices will not be distorted, and deadweight loss is zero. When the demand curve is vertical, demand is inelastic and consumers desire the same quantity of the good regardless of price. In this case, changes in the price do not distort relative consumption decisions. When the supply curve is vertical, supply of the good is fixed and no distortion occurs.

A retroactive tax on last year's income is an example of a nondistortionary tax. Because last year's behavior cannot be changed, the tax will cause no immediate deadweight loss. It is essentially a lump sum tax. Of course, such a tax can only be imposed once before people will begin to expect the taxation and alter their behavior accordingly, so that the tax is no longer nondistortionary.

Back to the formula:

Most taxes are given as a percentage of the purchase price, so we can rewrite the above by multiplying by P/P:

$$DWL = -(\tfrac{1}{2})\tfrac{\Delta T}{P}\Delta Q\,(P)$$

Now multiply by $\frac{\Delta T}{P}\frac{P}{\Delta T}$

$$DWL = -(\tfrac{1}{2})\left(\tfrac{\Delta T}{P}\right)^2\Delta Q\,(P)\,\tfrac{P}{\Delta T}$$

Now multiply by Q/Q

$$DWL = (\tfrac{1}{2})\left(\tfrac{\Delta T}{P}\right)^2\left(-\tfrac{\Delta Q}{\Delta T}\tfrac{P}{Q}\right)PQ$$

To simplify the formula, substitute t for the ad valorem rate and η for the elasticity to get:

$$DWL = (\tfrac{1}{2})t^2\eta\,(PQ)$$

Now we have a more useful form that tells us that the DWL due to a tax depends on three things:

1. the square of the ad valorem tax rate, t

2. the elasticity or flatness of the demand and supply curves, η
3. the size of the market for the taxed good, PQ

The elasticity in (2) is a combination of responses on the demand and supply sides. However, a common and convenient simplification is assuming that the supply curve is horizontal; that is, that supply is infinitely elastic. This assumption would be justified in the case of a small country that is a price taker for commodities. The price for those commodities is determined in the world market, and a small country cannot affect it.

Figure 10 is a graphical representation of a tax with infinitely elastic supply. In this case, the amount of the tax is exactly the change in the consumer price. The relevant elasticity for the DWL formula is the elasticity of demand for the consumer, which can be expressed as $\frac{\Delta Q}{\Delta P}\frac{P}{Q}$. If, on the other hand, the supply curve is vertical (perfectly inelastic), the elasticity of demand is irrelevant and the combined elasticity in the formula is zero. Thus, as mentioned above, the deadweight loss from a tax is zero in this case.

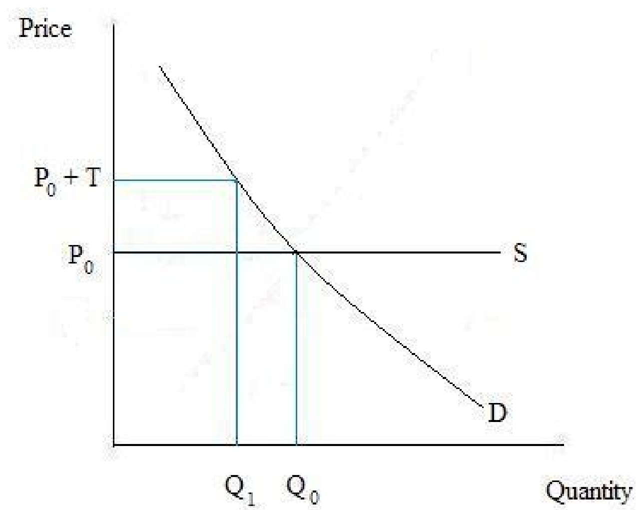

Figure 10

The third term in the deadweight loss expression, PQ, means that the DWL is larger when the market is larger. Of course, the revenue from the tax will also rise with the size of the market. Thus, a more relevant quantity is the DWL scaled by the revenue raised, or the excess burden for every dollar raised. Deadweight loss divided by revenue is what must be given up to obtain a dollar of revenue–that is, it is intuitively like the price of a dollar of revenue. In Figure 9, this is B+D/A+C. Algebraically,

$$\frac{DWL}{REV} = \frac{(\frac{1}{2})t^2\eta(PQ)}{tPQ} = (\tfrac{1}{2})t\eta$$

The key question in optimal commodity taxation is how to minimize this deadweight loss per dollar of revenue.

Basic optimal commodity tax theory

At an optimum, the marginal DWL on each commodity will be the same. We will prove this claim by contradiction: Suppose that the price of tax revenue is higher for apples than for oranges. If we lower the tax on apples and raise it on oranges to make up exactly the amount of lost revenue, the decrease in the DWL from the fall in the apple tax is larger than the increase in DWL from the increase in the orange tax. We can get the same amount of revenue with lower DWL. Therefore, the price of tax revenue being greater for apples than for oranges cannot be an optimum; at the optimum, the price of tax revenue, DWL/REV, must be equal for all taxes.

$$t_i \eta_i = k \text{ for all } i$$

The optimal tax rate for a commodity is thus inversely proportional to the elasticity:

$$t_i \propto \frac{1}{\eta_i}$$

In this simple model, taxes should differ only to the extent that elasticities differ. If elasticities are the same for all goods, uniform taxation is optimal. As η increases, the optimal tax rate falls. The intuition is that taxes that are easy to avoid have high deadweight loss relative to revenue because a higher tax rate will be necessary to offset the behavioral response and achieve the desired revenue. Since a fixed amount of revenue must be raised, avoidance behavior is just deadweight loss.

However, the simple elasticity rule holds only if the underlying assumptions about supply and cross demand elasticities are valid. If supply is not perfectly elastic as we assumed, there are profits for producers. Taxing all these profits away then applying the elasticity rule will be optimal. We have also implicitly assumed that the cross elasticities of demand are zero. That is, the simple elasticity rule requires that a tax on apples affects only the demand for apples and not the demand for oranges. Of course, such an assumption will not be valid if the two goods are close substitutes or complements. In that case, the cross elasticities will enter the expression for the optimal tax. For example, changing the corporate income tax rate may induce shifting of income between personal and corporate returns.

If the elasticity is the same for all commodities, a uniform tax is optimal. However, even if one commodity–food is a common example–is very inelastic, it would not be optimal to tax only that good, because the rate required would be very high. Since the DWL increases with the square of the tax rate, all tax rates should be as low as possible. When a commodity is very inelastic, it is optimal to tax that good at a higher rate than more elastic goods, but not tax only that good. The inverse elasticity rule is an efficiency result only. It does not take into account that taxing inelastic goods like food may have important distributional consequences. We will discuss such distributional considerations below.

It may seem at first glance that a uniform commodity tax would be a nondistortionary lump sum tax. However, this is not the case. To show this, note that an individual's budget constraint can be written as

$$P_1C_1 + P_2C_2 = wL$$

where P_i and C_i are the respective prices and quantities of commodities one and two, w is the wage, and L is labor supply in hours. If producer prices are given and an ad valorem uniform tax is placed on both commodities, the budget constraint becomes

$$(1+t)P_1C_1 + (1+t)P_2C_2 = wL$$

This can be rewritten as

$$P_1C_1 + P_2C_2 = \frac{wL}{1+t}$$

Therefore, a uniform commodity tax is equivalent to a tax on labor income. A tax on labor income is not a lump sum tax because income is not fixed; labor hours will exhibit a behavioral response to changes in the after-tax wage. While there is no distortion of the relative demand for the two commodities, there is a distortion between consumption and labor. A decline in the after-tax wage will discourage work. Thus, uniform commodity taxation will not have zero deadweight loss.

The next logical step is to look for a tax that does not distort either relative consumption decisions or the labor/leisure decision. To make this explicit in the algebra, we can express the budget constraint above as

$$P_1C_1 + P_2C_2 = w(\overline{L} - l)$$

where hours worked is now expressed as the difference between the total hours available for working, $\overline{L}$, and the amount of those hours spent at leisure, l. (We will assume that $\overline{L}$ is exogenous, although this is arguable. The maximum hours worked will actually depend on educational and occupational choices, which will in turn depend on the wage.) Rearranging the equation, there are now three "commodities" available: apples, oranges, and leisure.

$$P_1C_1 + P_2C_2 + wl = w\overline{L}$$

The price of leisure is foregone income, the opportunity cost of not working for the market wage. From this equation, it is clear that true uniform commodity taxation would include a tax on leisure:

$$(1+t)P_1C_1 + (1+t)P_2C_2 + (1+t)wl = w\overline{L}$$

This can be rewritten as

$$P_1C_1 + P_2C_2 + wl = \frac{w\overline{L}}{1+t}$$

A uniform tax on all commodities including leisure is a tax on the endowment of available working hours–essentially, a lump sum tax on underlying earning ability. Such a tax would be nondistortionary. However, taxing leisure is difficult. Leisure is defined here as any hours not spent in observable market production. To implement the tax, these hours must be accurately measured. But there will be an incentive to misstate hours if leisure is taxed. In addition, the leisure component of individual jobs is be nearly impossible to ascertain. Economists generally assume that these difficulties are large enough that taxation of leisure is infeasible.

The budget constraint can also be used to show that any tax system can be represented as a wage tax, equal to the average commodity tax, plus an additional tax or subsidy on each good.

$$(1+t_1)P_1C_1 + (1+t_2)P_2C_2 = wL$$
$$P_1C_1 + \frac{1+t_2}{1+t_1}P_2C_2 = \frac{1}{1+t_1}wL$$

Here, a non-uniform tax system is equivalent to a tax on labor income and an additional tax on commodity two.

Capital Income Taxation and Consumption in Different Periods

To consider the impact of capital income taxation, we will use a simple two-period model of the individual's lifecycle. In the first period, an individual works and earns labor income, which can be either saved or consumed. In the second period, the saved income is used for consumption. There are no transfers from the government or initial assets. Consumption in period two is then

$$C_2 = (wL - C_1)(1 + r_2)$$

where C_i is consumption in the relevant period, wL is labor income, and r_2 is the interest rate on saved earnings. The first term is saving from the first period; the second accounts for the interest on that saving. The expression can be rewritten as

$$C_1 + \frac{C_2}{1+r_2} = wL$$

From this, the household's budget constraint, it is clear that a uniform consumption tax is equivalent to a labor income tax. The reasoning is analogous to that used above regarding uniform commodity taxation. The final result in the last section also generalizes to this context; a tax on all income, both capital and labor, is equivalent to a tax on all consumption with a heavier tax rate on consumption in the second period. However, tax systems that are revenue equivalent from the household's point of view may have different implications for the timing of revenue collection by the government. A labor income tax transfers all revenue to the government in the first period, while a consumption tax transfers revenue in both the first and second periods. The government surplus and household saving will then differ across the tax systems as well.

It is simple to expand the model to allow inheritance or other initial assets. In this case, consumption in period two is

$$C_2 = (A + wL - C_1)(1 + r_2)$$

where A is the level of initial assets. Rewriting this as the household's budget constraint,

$$C_1 + \tfrac{C_2}{1+r_2} = A + wL$$

we can see that a consumption tax and a labor income tax are no longer equivalent from the household's point of view. A consumption tax will tax the value of the asset when it is spent, while a labor income tax will not. The consumption tax has a broader base and is therefore more efficient.

In a multiperiod model, there are also transitional effects of moving from one type of taxation to another. For those entering the first period, the imposition of consumption and labor income taxes are equivalent. For those entering the second period, the initial saving decision has already been made. If the imposition of a consumption tax is unanticipated, it is nondistortionary and carries no deadweight loss. The consumption tax is more efficient. Given the current tax system, a transition to pure consumption taxation would hit the old very hard. A transition to pure labor income taxation would hit the young hard.

Leaving transition issues aside, the decision on whether to tax all income or just consumption and labor income depends on the optimal relative taxation of first and second period consumption. In the two period model, a tax on all income was equivalent to a heavier tax on consumption in the second period. In a many-period model, the capital income tax represents many periods of taxation for interest income. Thus, the effective tax on consumption in each period rises as time goes on. This feature means that on efficiency grounds, it is difficult to justify high capital income taxes.

Line-Drawing (Weisbach 2000)

The issue of line-drawing arises when the tax treatment for two commodities is given, and the tax treatment for a third good, which shares characteristics with both, is in question. For concreteness, assume that good 1 is taxed and good 2 is not. There are two choices regarding good 3: tax it at the same rate as good 1, or do not tax it. Adding a tax on good 3 to the system will have two effects on deadweight loss. First, there will be an increase in DWL from adding the tax. Second, the deadweight loss from the tax on good 1 will also be affected. If the two goods are substitutes, the deadweight loss on good 1 will fall because it will be harder to avoid the tax. If they are complements, the deadweight loss will increase as the distortion of consumption away from good 1 is exacerbated. As a general rule, goods or transactions that are similar should face similar tax treatment.