In an economy with production, there are three sets of conditions for Pareto optimality : the marginal rate of substitution between any two goods X and Y must be the same for all people ($MRS^i_{XY} = MRS^j_{XY}$ for any two people i and j) ; the marginal rate of technical substitution between any two inputs must be the same in all industries ($MRTS^X_{KL} = MRTS^Y_{KL}$ for any two industries X and Y, and any two inputs K and L) ; the marginal rate of transformation between any two commodities must equal any consumer's marginal rate of substitution between those goods ($MRT_{XY} = MRS^i_{XY}$ for any pair of commodities X and Y, and any person i).

Any tax which prevents these efficiency conditions from holding will result in an allocation which is not Pareto optimal, and will have an excess burden. The excess burden is (sort of) a measure, in dollars, of how far away from Pareto efficiency the tax moves the economy.

Another way of describing that inefficiency is that taxes **distort** the prices faced by consumers and firms in the economy. In the examples used in Figure 3 and Figure 4, the prices without the taxes represent the "true" costs of the goods, the producer prices. Efficiency requires that the price ratio faced by consumers be the same for all consumers, and that this price ratio equal the ratio of "true" opportunity costs, which is the marginal rate of transformation. What makes the diagrammatic analysis work is that the "after" price ratio is the price ratio faced by consumers, and that this differs from the price ratio faced by producers. In the figures, the tax on good Y made the price line faced by consumers different from the price line faced by producers, violating the third efficiency condition $MRT_{XY} = MRS^i_{XY}$ above. The problem just mentioned was that the price line faced by consumers was **different** from that faced by producers', not necessarily that it was less steep. Just as there is an excess burden to excise taxation, there will be an excess burden to a subsidy on consumption of a single good, since that too will cause the MRS between goods faced by consumers to differ from the MRT.

Put otherwise, if the government were to subsidize consumption of a good, then the cost of the subsidy would be more than the value of the benefit received by the consumer. Figure 6 depicts a subsidy on consumption of good Y.

The subsidy on good Y pivots the consumer's budget line out, around its intercept on the X–axis (the point $(80, 0)$ in Figure 6). The consumer chooses a consumption bundle where her indifference curve is tangent to the new, steeper, budget line, point F (where the indifference curve is tangent to the dashed dark blue line "budget line : subsidy on clothing" in Figure 6). An alternative policy would be, instead of subsidizing consumption of good Y, just to give a cash grant to the consumer. A cash grant would shift the original budget line out parallel (to the solid gray line "budget line : equivalent cash grant" in Figure 6). From the consumer's perspective, the cash grant will be equivalent to the subsidy on food if they both move her to the same indifference curve, which is true for the grants depicted in Figure 6.

The cost of the equivalent cash grant is just the amount that the original budget line must be shifted to get the consumer to the higher indifference curve (the horizontal distance between the solid green line "budget line : no subsidy" and the solid gray line "budget line " equivalent cash grant" in Figure 6). The cost of the subsidy on consumption of good Y is the difference to the true cost of the consumer's consumption of Y, and the amount paid by the consumer. As in the case of a tax, to find the true cost of a consumption bundle, such as F in the figure, draw a line with a slope reflecting the true costs of the goods, $-p_x/p_y$, through the consumption bundle. (That's the dot–dashed light blue line "true cost of subsidized bundle F" in Figure 6.) The horizontal intercept of this line is the true cost of the subsidized bundle. The cost of the subsidy is the difference between this cost and the total paid by the consumer, which is just her income. So the cost of the subsidy is the horizontal distance between the line through F with slope p_x/p_y, and the consumer's income, which is the horizontal intercept of her budget line.

That means that the cost of the subsidy exceeds the cost of a cash grant, as long as the line with slope $-p_x/p_y$ through F is to the right of the line with slope $-p_x/p_y$ through G. The distance between those lines is the excess burden of the subsidy. But G is where the indifference curve is tangent to a line with slope $-p_x/p_y$, and F is where it is tangent to a line with slope $-p_x/(p_y - s)$, where s is the subsidy per unit of good Y. That means F must be to the left of G, so that a line with slope $-p_x/p_y$ through F must hit the horizontal axis to the right of the line through G, as in the figure. That horizontal distance is the excess burden of the subsidy (here it's \$10, which happens to be 25 percent of the value of the subsidy to the consumer).

Figure 6 used the equivalent variation to the subsidy to measure the excess burden, but the exercise could also be done using the compensating variation, which would give a slightly different value.

Again, the expenditure function indicates the same result depicted graphically in Figure 6, that there is an excess burden to a subsidy. The excess burden to the subsidy is the difference between the cost of the subsidy S, and the value of the subsidy to the consumer EV. Here

$$S \equiv s_y Y^H(p_x, p_y - s_y, u_1) \tag{28}$$

which is the subsidy s_y per unit of clothing, times the number of units of clothing chosen by the consumer (when the price of clothing is $p_y - s_y$, and when her utility level is u_1, the indifference curve corresponding to her utility after the subsidy has been introduced). And the equivalent variation to the subsidy is

$$EV = E(p_x, p_y, u_1) - E(p_x, p_y - s_y, u_1) \tag{29}$$

So the excess burden is $S - EV$, which is zero when $s_y = 0$, and where

$$\frac{\partial}{\partial s_y}[S - EV] = Y^H(p_x, p_y - s_y, u_1) - s_y \frac{\partial Y^H(p_x, P_Y, u_1)}{\partial P_Y} - \frac{\partial E(p_x, P_Y, u_1)}{\partial P_Y} \tag{30}$$

where $P_Y = p_y - s_y$ is the net–of–subsidy price of clothing. Again, Shephard's Lemma mean that the first and third terms in expression (30) cancel each other out, so that

$$\frac{\partial}{\partial s_y}[S - EV] = -s_y \frac{\partial Y^H(p_x, P_Y, u_1}{\partial P_Y} \tag{31}$$

As in the previous sub–section, the fact that Hicksian demand curves slope down means that expression (31) is zero or positive, and must be positive if the indifference curves have no kinks. So increasing the subsidy above zero must increase the excess burden, and increasing the subsidy further makes the excess burden even bigger.

[This result can also be obtained using equation (21) or (25) of the previous sub–section, and recognizing that a subsidy is just a negative tax. So $t_x < 0$ is a subsidy on good X, and equations (21) and (25) show that decreasing t_x below zero increases the excess burden above zero.]

The figure also shows that any policy must have an excess burden — unless the policy leads the person to consume at a point where the (absolute value of the) slope of her indifference curve equals the ratio of the true costs of the goods, p_x/p_y.

This statement may lead to another question : What would happen if both goods X and Y were taxed, and at the **same** proportional rate τ? Then the tax-inclusive prices faced by consumers for the two goods would be $(1+\tau)p_x$ and $(1+\tau)p_y$ respectively, so that their MRS would be p_x/p_y, since the $1+\tau$ terms would cancel — meaning that it would still be the case that $MRS = MRT$.

So this argument suggests that there would be no excess burden to a general sales tax, that taxed consumption of **all** commodities at the same rate. Going a little further, in the section on general equilibrium tax incidence, I concluded that a proportional sales tax on all goods was equivalent to a proportional income tax on all income. If a general proportional sales tax has no excess burden, then neither does a proportional income tax.

That equivalence between taxes is still valid [1] : a proportional sales tax is equivalent to a proportional income tax. However it is very misleading to suggest that neither of these taxes has an excess burden.

To see why, just look at the individual's budget constraint. That equation can be written

$$P_X X + P_Y Y = M \tag{32}$$

where M is her income, and where P_X and P_Y are the prices she actually faces for the two commodities. This equation was used to show the equivalence of a sales tax and an income tax. Now where does that income M come from? If it is really and truly **exogenous**, then a tax on that income is a lump–sum tax. Then so too is a proportional consumption tax. So if everyone's income consisted entirely of bequests from long–dead relatives, we could feasibly impose lump–sum taxes which were proportional to income. If your income consists entirely of an already–made bequest, then there is nothing you can do to affect your income. Therefore your tax liability does not depend on what you do, so that a tax on that income is a lump–sum tax.

On the other hand, if your income M depends on choices that you make (and have not yet made), your tax liability **will** depend on your choices if a proportional income tax is levied. So that is how the apparent paradox is resolved : a proportional consumption tax is equivalent to a proportional income tax, but in general neither tax is non–distortionary. These taxes would

[1] at least when intertemporal complications are ignored

be lump–sum taxes only if people's income were exogenous — and that is not a very realistic or likely characterization. People's labour income does depend on how much they choose to work, and, in the longer term, the choice of training they undertook. People's non–labour income looks exogenous at a point in time : what assets you own on October 30 2014 depends on decisions you have made prior to October 30 2014. But your future non–labour income does depend on decisions you have yet to make. So a tax on non–labour income will certainly not be lump–sum, unless it is a one–shot tax collection which comes as a complete surprise to people.

One way of representing the distortions, and excess burdens, of a tax on labour income, is to look at the labour–leisure choice of individuals. This approach is usually discussed briefly in ECON 2300 (and is discussed in much more detail in ECON 3240). It consists in viewing time not spent at work as a commodity which you consume. We usually refer to this time as "leisure", though it is meant to cover all uses of time other than paid employment in the market. If we look at a person's weekly budget constraint, she has 168 hours per week to spend, either in the workplace or elsewhere. So her available leisure is 168 hours minus the time spent earning income. Her budget constraint is then

$$P_X X + P_Y Y = w[168 - L] \tag{33}$$

where w is her net wage per hour, and L is the time spent at leisure, that is the time not spent on the job. We can re–write this equation

$$P_X X + P_Y Y + wL = 168w \tag{34}$$

which looks like the budget equation for a person with exogenous income $168w$, who can spend that income on any of **three** goods (X, Y, L), with prices of (P_X, P_Y, w) respectively.

If we write the budget constraint in this manner, then a consumption tax on goods X and Y is **not lump–sum** ; the relative price of leisure has been changed. Similarly a proportional income tax alone will have an excess burden, since it changes the price of leisure, relative to good X, from w/P_X to $w(1-\tau^i)/P_X$. We would have a lump–sum tax if we taxed consumption of good X, consumption of good Y, **and leisure**, all at the same rate. But such a tax does not seem feasible, because we would have to measure and tax the amount of leisure time a person spends. Were there such a tax, workers would have strong incentive to go to their employers, and ask them to write a phony contract, saying that they were paid \$10 per hour for an 80–hour work week, rather than the true \$20 per hour for a 40–hour work week.

Given the answer to this question, the excess burden of a tax on labour income can be derived, by noting that it is actually a subsidy on leisure. But while recalling the labour–leisure choice material from intermediate microeconomics, you may recall the possibility of the labour supply curve sloping down. Recall that the theoretical sign of the effect of an income tax on someone's labour supply is ambiguous. If we tax labour earnings, we lower the net–of–tax wage. This has a substitution effect, inducing the person to consume more leisure (because the price of leisure relative to good X has gone from w/P_X to $w(1-\tau^i)/P_X$). This substitution effect leads to a decrease in hours of work in response to an income tax increase. However there is an income effect working in the opposite direction. Increases in tax collections make the person worse off. When she is worse off, she wants to consume less of all normal goods. The evidence seems quite strong that leisure is a normal good. Therefore, when her real income falls, she will want to consume less leisure, in other words to work more.

As a practical matter, it does appear that the effect of a change in the net wage on hours worked is quite small, at least for prime–age married males who are the primary wage earner in their household. This empirical evidence — that the labour supply curve of these men is virtually vertical — can be explained by these substitution and income effects more or less cancelling each other.

Does that mean that the excess burden of the taxation of labour income of married males is trivial? The answer is a definite "no". Both the equivalent variation and the compensating variation are measured as movements along an indifference curve. The compensating variation, for example, looks at the impact of a change in the wage rate, together with an income transfer which leaves the person as well off as before. This last effect exactly neutralizes the income effect. The compensating and equivalent variation are determined by movements along a single indifference curve : a pure substitution effect. As mentioned earlier, the magnitude of the deadweight loss is determined by the shape of the indifference curve. So even if the overall effect of a wage change on labour supply is virtually zero, this overall effect may consist of a fairly large, negative substitution effect being cancelled by a fairly large, positive income effect. Figure 7 illustrates. In that figure, a tax on wage income moves the person's preferred leisure–consumption choice from A to B. So, in the figure, a tax on wage income has little or no effect on labour supply. (The person's preferred quantity of leisure hours has not changed.) But the move from point A to point B in Figure 7 represents an "uncompensated" change. To measure any excess burden, the compensated change must be used : how much would her consumption of leisure change if the tax were introduced, and if she stayed on the same indifference curve. That's the "substitution effect", and it's measured by the movement between points B and C in Figure 7. So, in the case illustrated in Figure 7, a tax on wage income has a non–trivial substitution effect — a reduction in the person's hours of work — offset by an income effect (between points A and C in Figure 7) which works in the opposite direction.

The excess burden is determined by the substitution effect alone. And it does appear that the income and substitution effects of wage changes on the labour supply of prime–age married males are non–trivial. There is a substantial deadweight loss from the taxation of labour income, even though the supply curve for labour may be close to vertical.