In a 2–person exchange economy, or in an economy with production, there are a lot of efficient allocations. Moving from one efficient allocation to another makes one person better off while making some other person worse off. As long as we can simply redistribute endowments — of food, or clothing, or inputs to production — from one person to another, we can move from one efficient allocation to another.

Pareto optimal is not exactly the same thing as "wonderful". All it means for an allocation to be Pareto optimal is that there is no other allocation which everyone prefers. Inefficient allocations (that is, allocations which are not Pareto optimal) are "bad" because there is some other allocation which is better for everyone. So everyone would agree to move from an inefficient allocation to an allocation which is Pareto–preferred.

That's not the same thing as saying that everyone should like **any** Pareto optimal allocation more than **any** inefficient allocation. In Figure 5 allocation $A$ (the yellow square at [about] (2,3.5) ) is not Pareto optimal, since it is not on the contract curve. Allocation $B$ (the white square at (8,5)) is Pareto optimal. But person 2 certainly prefers $A$ to $B$ : he gets more food, and more clothing, at $A$. So we certainly do not expect person 2 to want to see a move from allocation $A$ to allocation $B$. However, there are allocations on the contract curve which are Pareto–preferred to $A$ (for example, points on the contract curve just to the left of $(4, 2)$, where 1's dashed green indifference curve is tangent to 2's dotted purple indifference curve). So it's not true that **any** efficient allocation is better than **any** inefficient allocation ; it's true that we can't find anything which is better (for everyone) than an efficient allocation and for any inefficient allocation we **can** find something which is better (for everyone).

When it is possible to move along the contract curve simply by redistributing endowments, and then letting competitive markets work, the economy is often described by economists as a **first–best** economy. In a **second–best** economy, which may be the more realistic case, it is not so easy to redistribute endowments. In a second–best economy, there is a bigger role for government to play, and the conflict between equity and efficiency emerges.

But back to the (imaginary) first–best world. How would a benevolent dictator, or a government, or a bunch of people voting perhaps, choose among the many efficient allocations? Microeconomics does not really offer any deep insight here. But we can talk about how this choice might be made by talking about a **social welfare function**. A social welfare function is simply a way of describing how a society's decision makers (politicians? planners? dictators?) might trade off the gains to one person against the losses of another. In a simple 2–person, 2–good exchange economy, as we move up the contract curve, person 1 is made better off, and person 2 worse off. If I could somehow measure the well–being of each person using a utility function, I could then talk about a utility combination $(u_1, u_2)$ for the two people at some allocation. Moving up the contract curve means moving to a new allocation, in which person 1 gets higher utility $u_1' > u_1$ and person 2 gets lower utility $u_2' < u_2$.

A social welfare function is some function defined over these utility pairs. So $W(u_1, u_2)$ is the measure of how well off society is when person 1 has utility $u_1$ and person 2 has utility $u_2$. One example of a social welfare function is $W(u_1, u_2) = u_1 + u_2$. A decision maker with such a welfare function would want to maximize the **sum** of people's utilities. This particular example of a social welfare function is often called a "Benthamite" utility function, after Jeremy Bentham [3], who advocated this as the goal for society. But $u_1 + u_2$ is not the only example of a social welfare function. Any function which depends ( only ) on each person's utility, and which increases when some person's utility goes up, could be a social welfare function. So

$$W(u_1, u_2) = \sqrt{u_1} + \sqrt{u_2}$$

$$W(u_1, u_2) = \ln(u_1) + \ln(u2)$$

$$W(u_1, u_2) = \min(u_1, u_2)$$

are all examples of social welfare functions.

Notice that to talk of social welfare, I have to talk about utility. To decide on which efficient allocation to choose, I have to compare the well–being of different people here. Just using people's incomes is not good enough. We may have to compare allocations when people face different prices, when some people have to work much harder than others, when some people are healthy and others are not, and so on. Income is not an adequate measure of how well off people are. I might feel better off with a moderate income, if I have to do very little work, than with a higher income, if I have to work long hours at an unpleasant job.

So we have to use utility, which in some sense assumes that we can not only measure how happy people are, but can compare the happiness of different people.

Once we have a social welfare function, then the allocation to choose is the one which maximizes the value of social welfare among all feasible allocations. One way of representing this is to use a **utility possibility frontier**. Each point on the contract curve gives rise to some utility combination. As we move up the contract curve, $u_1$ increases and $u_2$ decreases. [ Why must this always be the case? ] That means we can graph the different ($u_1, u_2$ ) combinations. This creates a downward–sloping curve, the "utility possibility frontier". Points along this curve correspond to efficient allocations. Moving down this curve to the right corresponds to moving up the contract curve in the Edgeworth box. Points inside the utility possibility frontier correspond to allocations which are inefficient. [ Why? ]

Using this picture, a society can be viewed as picking the $(u_1, u_2)$ combination on the utility possibility frontier which maximizes $W(u_1, u_2)$. Through any utility combination, we can draw an **iso–welfare** curve. This is the curve of all utility combinations which give rise to the same value of social welfare. It must slope down [ why? ], with slope

$$-\frac{\partial W}{\partial u_1} / \frac{\partial W}{\partial u_2}$$

[3] an eighteenth century British philosopher

So it's a straight line with slope $-1$ for a Benthamite social welfare function. In Figure 6, I have chosen a social welfare function

$$W = u_1 u_2$$

and drawn in three iso-welfare curves. The overall best allocation, then, is situated where one of these iso–welfare curves is tangent to the utility possibility frontier.

The second fundamental theorem of welfare economics makes sense only in a first–best world. That theorem said that any Pareto efficient allocation could be achieved as a competitive equilibrium, after suitable redistribution of endowments. So that is, by definition, a first–best result. That result means that there is no conflict between equity and efficiency in a "first–best" world. If endowments of goods ( or purchasing power ) could be redistributed costlessly from one group of people to another, then there is no reason why you would not want to have a Pareto efficient allocation, somewhere on the contract curve. That does not mean that there is no conflict between people in a first–best world. Certainly moving up the contract curve makes person 1 better off and person 2 worse off. There still is a difficult choice as to what is the best allocation. But in this first–best world, the argument in favour of efficiency is a strong one. Making the economy less efficient ( moving away from the contract curve ) does not seem to be a good way of achieving a more equitable allocation.

A second–best world is a world in which redistribution is not so easy. So — without seriously considering why or why not— suppose that it is impossible to redistribute people's endowments at all. That means that there is some initial endowment point in the Edgeworth box. If the economy is a perfectly competitive one, then people may choose to trade with each other at market prices, moving the economy to some competitive equilibrium. That competitive equilibrium is Pareto efficient. But it's the only one that we can get to, since we are now taking the initial ownership of the commodities as unchangeable. So this competitive equilibrium is still Pareto optimal. But we can't get to the other points on the contract curve, at least not in a perfectly competitive economy.

That means, if I switch to my utility possibility frontier diagram, that there is some point on the utility possibility frontier which can be achieved by a perfectly competitive economy. But the economy cannot get to the other points on the frontier as competitive equilibria.

Maybe we cannot redistribute endowments, but it seems realistic to assume that it is possible for governments to levy taxes. After all, they do levy taxes in Canada (and elsewhere). Unfortunately, taxing people is **not** the same as redistributing their endowments. So suppose that in my equilibrium, person 1 sells some of her clothing endowment to person 2, and person 2 sells some of his food endowment to person 1. ( That is, my competitive equilibrium is southeast of the endowment point. ) One thing my government might do, especially if it wanted to make person 2 better off, would be to levy a tax at the rate $t$ on each kilogram of food bought, and then to give the revenue generated by this tax to person 2.

Now if markets are competitive, it means that all people are price takers : they assume that they do not have enough impact to influence market prices. If there really were only two people in the economy, then they would each realize that they can influence prices : what person 1 buys, she must be buying from person 2.

So to be more accurate, assume now that there are **2 million** people, one million of "type 1" (all identical to each other in tastes and endowments), and one million of "type 2". There are still only two types of people, and there are equal numbers of each type. So in equilibrium, the amount of food bought by each type–1 person must equal the amount of food sold by each type–2 person (if total quantity deamnded of food is to equal quantity supplied of food).

The only difference is : now each person does not have much influence on prices.

Here the government now is taxing all food purchases, $t$ per kilogram, and dividing the tax revenue equally among the one million type–2 people.

So what is the price of food that a type–2 person gets, when he sells food to a type–1 person? Call it $p_X$. The type–1 person, who is buying food, is now paying a price of $p_X + t$, since some of the gross price of the food she is buying is actually the tax which goes to the government. There is no tax on clothing, so that the type–1 person sells clothes at a price of $p_Y$ per unit, and each type–2 person buys clothes at $p_Y$ per unit.

The type–1 person, like a good competitive price taker, chooses a point on her budget line where her indifference curve is tangent to the budget line. the slope of her budget line is, as usual, the price ratio. For her that's $(p_X + t)/p_Y$, so that her choice of consumption plan obeys

$$MRS_1 = \frac{p_X + t}{p_Y}$$

The type–2 person is getting $p_X$ for every unit of food he sells. Of course, he is also getting some of the revenue from that tax. (Remember, the government is taking the tax revenue, and giving it back to be shared equally among all the million type–2 people.)

But each type–2 person is only getting a tiny share of the tax revenue from each sale (1 millionth of the revenue). That means that he can ignore this tax revenue when making her buying and selling decisions : he is getting only 1/million of the tax revenue from the food that she sells. Every unit of food he sells gets him $p_X$ more in revenue (if we ignore his tiny share of the tax revenue).

The type–2 person chooses a consumption plan where his indifference curve is tangent to a budget line, which, in his case, means

$$MRS_2 = \frac{p_X}{p_Y}$$

His food price does not include the tax, since he's selling it for $p_X$, and the government, not he, is getting the tax revenue. Therefore, the tax imposes a **distortion**. At the competitive equilibrium with the food tax in place, it is no longer true that $MRS_1 = MRS_2$. The distortionary food tax moves the competitive equilibrium off the contract curve.

So with a food tax, and with type–2 people getting the revenue from the food tax, there will be an equilibrium that probably is southwest of the old ( no tax ) competitive equilibrium. The allocation in the presence of the distortionary tax is not Pareto optimal. There are other feasible allocations which are better for both people. But, unfortunately, I don't know how to achieve those allocations in a second best world. I have to choose between an allocation which is efficient, but not too equitable, and one which is inefficient but more equitable. The inability to redistribute endowments means that there really is a conflict possible between efficiency and equity.

In Figure 7, the outcomes possible from a distortionary food tax are graphed in a utility diagram. These possible utility combinations are inside the original utility possibility frontier — which is another way of saying the allocations they represent are not on the contract curve. There is a second–best utility possibility frontier inside the first–best.

The point $O$ in the diagram represents the utilities people would get in a "laissez–faire" economy, one with no taxes at all. The point $O$ is on the second–best upf, as well as on the first–best upf, since it can be achieved without any redistribution of endowments. A point such as $W$, which is on the second–best upf, can be achieved by a food tax. This point is inside the first–best upf. It represents an allocation which is **not** on the contract curve in my original Edgeworth box diagram. It is Pareto–dominated by some allocation on the contract curve. In the Edgeworth box diagram of Figure 5, the utility combination $O$ might arise from the efficient allocation $B$, whereas the utility combination $W$ might arise from the allocation $A$, which is not efficient in a first–best world. The utility combination $O$ can be achieved in a second–best world, simply by not levying any tax at all. $O$ does represent an efficient allocation : there is no other combination possible which can make both people better off.

But does that make the utility combination $O$ better than the combination $W$? No. If the only instrument that the government can use is a food tax, then the possible utility combinations are those on the second–best upf. There is no **possible** utility combination which both people prefer to $W$. It is true that $O$ is first–best efficient, and that $W$ is not : but person 2 prefers $W$ to $O$, since it gives him higher utility. If we could abolish the food tax, and then redistribute some of the gains to person 2, then we could make both people better off, and get to a combination on the contract curve. But if these gains cannot be redistributed costlessly, then we should look at the second–best upf, and try to get on the highest iso–welfare curve on that second–best upf. Just because $O$ is first–best efficient does not mean that it is the best point on the second–best upf.