Week 4 - Linear Demand and Supply Curves

November 26, 2007

1

Suppose that we have a linear demand curve defined by the expression $X_D = A - bP_X$ and a linear supply curve given by $X_S = C + dP_X$, where $b > 0$ and $d > 0$.

(a) The coefficients b and d represent the responsiveness of demand and supply respectively to changes in price. Mathematically they are the derivatives of the demand and supply curves with respect to price. (This implies, via the inverse function theorem in calculus[1], that $\frac{1}{b}$ is the derivative of the *inverse* demand function with respect to quantity and $\frac{1}{d}$ is the derivative of the *inverse* supply function with respect to quantity. By assuming that $b > 0$ and $d > 0$ we ensure a standard downward sloping demand curve and upward sloping supply curve.

(b) The model only makes economic sense if A is positive, because if A were negative there would not be positive demand at any price.

(c) For a long run supply curve, C would have to be negative because if C were positive then some of the good would be supplied when the price is 0. This could not possibly be profitable, so firms would shut down in the long run and so supply would have to be less than 0 when $p_X = 0$ (it is fine for supply to be negative because the only relevant part of the supply curve is at positive values of X and p_X). In the short run however, if there were perishable excess stocks (e.g. of meat) then some of the good might be supplied when the price is 0 (or rather, infinitely close to 0 - if suppliers have to get rid of the stocks, it will be better to sell them for an infinitely small price than just to let the stock remain unsold and be wasted).

(d) It is the convention in economics to always display a demand and supply curve with amount X on the x-axis and price P_X on the y-axis.

[1]See the maths workbook or a mathematics for economics textbook.

Thus, technically speaking, when we sketch the demand curve we are really sketching the inverse demand curve because X is the independent variable and P_X is the dependent variable. In order to sketch the demand and supply curves, we must first therefore rearrange to make P_X the subject of the expressions. These yield $P_X = -\frac{1}{b}X_D + \frac{A}{b}$ and $P_X = \frac{1}{d}X_S - \frac{C}{d}$. From this, it is clear that the y-intercepts of the inverse demand and supply curves are $\frac{A}{b}$ and $-\frac{C}{d}$ respectively and that the slopes are $-\frac{1}{b}$ and $\frac{1}{d}$ respectively.

The mathematical definition of the consumers' and producers surplus is also more intuitive when we think about the inverse demand and supply functions rather than the demand and supply functions themselves. Whereas the demand function gives quantity demanded as function of price, the inverse demand function gives the marginal reservation price as a function of quantity demanded. This implies that the area under the inverse demand curve gives the summation (integral in calculus terms) of the marginal reservation prices, and therefore that the area under the demand curve minus the revenue rectangle (quantity demanded multiplied by price paid) gives the total surplus of reservation prices above prices paid, which is the definition of (Marshallian) consumers' surplus. The intuition for producers' surplus is that the inverse supply curve is the marginal cost curve for a competitive industry. Firms produce more units until the marginal cost of the last unit produced is just equal to the price they get for it (any further production involves a loss being made on the last few units). This means that the inverse supply curve gives marginal cost as a function of quantity, and that the revenue rectangle minus the area under the inverse supply curve gives the summation of marginal revenue minus marginal cost, which is the same as total revenue minus total variable costs (in order to get industry profits, we must further subtract any fixed cost, which has the mathematical role of being the coefficient of integration when we integrate MC to get TC).

(e) Price elasticity of demand is defined as percentage change in quantity demanded divided by percentage change in price. Price elasticity of supply is defined as percentage change in quantity supplied divided by percentage change in price. The reason that elasticities are often more useful in economics than the straight derivative of the demand function is that whereas the value of the derivative will depend upon the units in which you measure quantity and price (i.e. it will be large if price is measured in large units and quantity in small units and small if price is measured in small units and quantity in large units). The concept of elasticity corrects for this so that we get the same answer no matter what units we measure in.

Imagine that we increase price by 1 unit along a linear demand curve. The percentage change in price will therefore be $\frac{1}{P_X} * 100$. Demand will in this case change by $-b$ units because this is the slope of demand curve. The percentage change in quantity demanded will therefore be $\frac{-b}{X_D} * 100$. The price elasticity of demand is therefore:

$$\epsilon_D = \frac{\frac{-b}{X_D} * 100}{\frac{1}{P_X} * 100} = -b\frac{P_X}{X_D} = -b\frac{P_X}{A - bP_X}$$

From this, it should be clear that the elasticity changes along a linear demand curve. As price P_X goes to 0, elasticity goes to 0. As P_X goes to $\frac{A}{b}$ (the reservation price for the first unit), the elasticity goes to $-\infty$. In order to find the elasticity along a non-linear demand curve, you must find the elasticity along a linear approximation to the curve at this point (i.e. find the elasticity along the *tangent* to the demand curve). The more general equation for price elasticity of demand is therefore $\epsilon_D = \frac{dX_D}{dP_X}\frac{P_X}{X_D}$.

By a similar process of reasoning, the general expression for price elasticity of supply is $\epsilon_S = \frac{dX_S}{dP_X}\frac{P_X}{X_S}$. In the linear case which here concerns us, the expression this simplifies to give is $\epsilon_S = d\frac{P_X}{C+dP_X}$. The key point is that given a particular value of X_D or X_S and P_X, the absolute value of supply and demand elasticity, $|\epsilon_D|$ or $|\epsilon_S|$ respectively, (remember that elasticity is usually negative along a demand curve but positive along a supply curve) are higher, ceteris paribus, for a demand curve with a higher value of b and a supply curve with a higher value of d.

(f) We find the equilibrium price ${P_X}^*$ and quantity X^* by solving the simultaneous equation system formed by the demand and supply curves and the $X_D = X_S$ condition. This gives us $A - b{P_X}^* = C + d{P_X}^*$, which rearranges to give $(b+d){P_X}^* = A - C$, and therefore finally that ${P_X}^* = \frac{A-C}{b+d}$. Plugging this back into the equation for quantity demanded gives us $X^* = {X_D}^* = {X_S}^* = A - b{P_X}^* = A - b\frac{A-C}{b+d} = \frac{A(b+d)}{b+d} - \frac{b(A-C)}{b+d} = \frac{dA+bC}{b+d}$.

(g) Assuming that in the short run the supply of pork meat is perfectly inelastic, this would give us a vertical short run supply curve. Since the market supply would no longer depend upon the price level d_{SR} would be 0. C_{SR} would be positive, and equal to the fixed short run supply quantity. Plugging these values into the above solution for the market equilibrium yields $({P_X}^*)_{SR} = \frac{A-C_{SR}}{b}$ and $({X_D}^*)_{SR} = ({X_S}^*)_{SR} = C_{SR}$. Such a short run supply curve could be realistic if the time scale of the short run in long enough that there are perishable stocks of pork meat which must be supplied that period but short enough that new pigs cannot be reared. Over a longer time scale, the supply of pork meat will not be perfectly inelastic because higher price will lead more pigs to be reared because it increases the marginal profitability of pork meat production. This can be seen mathematically because if the marginal cost of the final piece of pork sold was previously equal to the marginal revenue gained (as would be the case for a profit-maximizing firm) then a rise in the price of pork will raise the marginal revenue on the final unit sold without increasing marginal cost. Output would therefore be increased by profit maximizing firms until $MR = MC$ again.

(h) When A increases by 1 unit, market equilibrium price increases by $\frac{dP_X{}^*}{dA} = \frac{1}{b+d}$ in the long run. In the short run, this becomes $\frac{1}{b}$. In the short run, market equilibrium quantity does not change, whereas in the long run it decreases by $\frac{dX^*}{dA} = \frac{d}{b+d}$. The economic intuition for this is that when the demand curve shift outward, this creates upward pressure on the price. In the short run, since output is fixed, the entire effect is on price. In the long run however, suppliers increase quantity in response to the price increase, and this results in the price going up by less (because $\frac{1}{b+d} < \frac{1}{b}$).

(i) If we were not model the effect on demand from the discovery of "mad pig disease" as a reduction in A, then the effect would simply be the reverse of that described in the above part. Price would reduce, but by more in the short run than in the long run, because suppliers would reduce output in the long run, driving the price back up a bit, but not above the original price before the reduction in A.

(j) The *legal incidence* of a tax refers to who is legally required to pay the tax. The *economic incidence* of a tax refers to who actually pays what proportion of the tax. Economic theory predicts that the long run economic incidence of a tax on a good will be determined by the elasticity of supply and demand, and will be independent of the legal incidence (the economic incidence may, however, depend upon the legal incidence during the short run before the market has fully adjusted to the introduction of the tax). This is because both suppliers and consumers will attempt to pass on the tax as a lower post-tax price paid for the good in the case of consumers and as a higher pre-tax price in the case of the suppliers. The degree to which each side wins this "tug of war" depends upon the elasticities of demand and supply. The more elastic is demand, the more of the tax burden the supplier will have to take on in order to maintain sales and revenue. The more elastic the supply, the more of the tax the consumers will have to take on in order to maintain the level of output and their consumers' surplus.

(k) The introduction of a unit tax creates a "wedge" between the price paid by consumers per unit P_{X_D} and the price received by suppliers $P_{X_S} = P_{X_D} - T$. The market equilibrium condition therefore becomes: $A - bP_{X_D}{}^* = C + d\,(P_{X_D} - T)$. Rearranging this yields $P_{X_D} = \frac{A-C+dT}{b+d}$. Plugging this into the $P_{X_S} = P_{X_D} - T$ condition and rearranging yields $P_{X_S} = \frac{A-C-bT}{b+d}$. The tax paid per unit by the consumer and producer is therefore $\frac{dT}{b+d}$ and $\frac{bT}{b+d}$. Adding these two shares together gives T, the total tax paid per unit. The proportion of the tax paid by the consumer is therefore $\frac{d}{b+d}$ and by the producer $\frac{b}{b+d}$.

From these expressions it is clear that the higher is the elasticity of demand, the higher is b and thus the higher is the proportion paid by the supplier and the lower is the proportion paid by the consumer. On the other hand, the higher is the elasticity of supply, the higher is d, and the higher is the proportion paid by consumer and the lower is the proportion paid by the supplier.

(l) Substituting the formula for the equilibrium price paid by the consumer, $P^*_{X_D} = \frac{A-C+dT}{b+d}$ back into the equation for the demand curve $X_D = A - bP_X$ yields the following expression for the equilibrium output: ${X_D}^* = \frac{dA+bC-bdT}{b+d}$ (you would get the same answer by finding the price paid by the supplier and plugging into the supply curve equation). The output distortion caused by the tax is therefore $\frac{bdT}{b+d}$. From this, it is clear that the output distortion is increasing in the tax rate, the elasticity of demand and the elasticity of supply.

(m) The revenue R_X gained by the government is equal to the unit tax multiplied by the number of units sold post-tax. This will be given by the expression: $R_X = \frac{T(dA+bC-bdT)}{b+d}$.

(n) (o) The marginal revenue of the government will be found by finding the derivative of the revenue function with respect to the tax rate. As long as this is positive, the government can increase its revenue by raising the tax rate. Before differentiating, it is instructive to rearrange the expression for the tax revenue to the standard quadratic form: $R_X = \left(-\frac{bd}{b+d}\right)T^2 + \left(\frac{dA+bC}{b+d}\right)T$. From this form, it is clear that we have an "upside-down U-shaped" quadratic which goes through the origin (when tax is 0 revenue is 0) and will have a second positive root at $T = \frac{dA+bC}{bd}$ (this is where the tax has become so high, it has driven output down to 0 and so tax revenue has again gone to 0). Differentiating the revenue function yields $\frac{dR_X}{dT} = \left(-\frac{2bd}{b+d}\right)T + \left(\frac{dA+bC}{b+d}\right) = \frac{dA+bC-2bdT}{b+d}$. Setting the derivative equal to 0 to find the maximum yields $T^* = \frac{dA+bC}{2bd}$ (note this is half way between the roots, as it should be with a quadratic). Finally, we can check that this is indeed a maximum and not a minimum by noting that the second derivative of the revenue function is $\left(-\frac{2bd}{b+d}\right)$, which is unambiguously negative.

(p) In order to show the deadweight loss, you must either vertically shift the demand curve downwards by T (if you wish to have the post-tax price received by the producer on the y-axis) or the supply curve upwards by T (if you wish to have the pre-tax price paid by the consumer on the y-axis). Where the shifted curve crosses the unshifted curve gives you the new quantity. The deadweight loss is the sum of the differences between the reservation prices (the demand curve) and the marginal cost of production (the supply curve of a competitive industry) over the output distortion. In the linear demand and supply curve model it is a triangle with height equal to T and width equal to the output distortion.

(q) Since the output distortion is $\frac{bdT}{b+d}$, the size of the deadweight loss is $\frac{1}{2}\frac{bdT^2}{b+d}$. We can see that the deadweight loss is increasing in b and d (and therefore increasing with increased elasticity of demand or supply), and proportional to the *square* of the tax rate. This is important because it means that doubling the tax rate more than doubles the size of the deadweight loss (in the linear case exactly quadruples it, approximately so in the case of non-linear demand curves). The intuition for this result it that as the tax rate increases, the difference between what the consumer was willing to pay and the cost for the producer to supply for the last unit prevented from being produced by the tax rise is increasing. To put it another way, the *marginal social cost* of unit taxation on a particular good is *increasing* with the tax rate.

(r) We can find a precise expression for the marginal social cost of taxation in this case by differentiating the deadweight loss with respect to the tax to give $\frac{bdT}{b+d}$. The key thing to note is that it is increasing in T. (Only in the special case of the linear demand and supply system is this expression equal to the output distortion.) We saw in part (n) that the marginal revenue from a unit tax is decreasing in T. The intuition for this is that as the tax gets higher, output decreases and so the extra revenue gained from ramping up the tax a bit more goes down. Put more technically, the elasticity of demand with respect to a tax increase increases because the percentage change in the tax rate decreases as the tax gets high whereas the percentage change in output gets higher and higher as output gets smaller and smaller. At the point where marginal revenue is driven to 0, demand is unitary elastic with respect to a change in the tax rate.

Both these factors mean that it becomes less socially desirable to increase taxes on a good the higher they get. In general, a government would not wish to set the revenue maximizing tax rate on any particular good, because the deadweight loss would be prohibitively high. The only exception might be demerit goods such as alcohol, cigarettes or drugs, since if people are myopic, they do not take into account the long run health costs and so overconsume these. In this case, the deadweight loss is *not* the area between the demand and supply curves because the demand curve overstates the "true" value of the good (the true marginal social benefit curve should be the demand curve shifted downwards, so that taxing the good actually *reduces* the deadweight loss by bringing the effective demand curve more in line with this). For most goods, the unit tax rate will be set by the government somewhere between 0 and the revenue maximizing rate, where the social marginal benefit from the revenue is equal to the social marginal cost from the deadweight loss.

The above analysis helps to explain why in the real world, governments spread commodity taxation over many goods. This is because a given amount of tax revenue can be raised with less social cost if the taxes are spread over many goods, because then the decreasing marginal revenue and increasing social cost is minimized as much as possible. Note also that the model does not in general predict that the tax rate should be the *same* on all goods. Ceteris paribus, goods with less elastic demand should be taxed more heavily. This simple model does not, however, capture all of the considerations. Equality also matters when setting commodity tax rates. For example, demand for basic foodstuffs is very inelastic, but a society that cares about equality will probably not tax them heavily, because they form a large proportion of the income of the poor.

(s) In the case of non-linear demand curves, most of the qualitative results established above would apply. The deadweight loss will only be approximately proportional to the square of the tax rate, but this approximation will be good for a small tax because if you zoom in to the crossing point of two non-linear curves, they will look increasingly like straight lines the closer you get. The size of the output distortion and the deadweight loss will continue to be increasing in the elasticities of supply and demand, again for the same intuitive reasons as discussed above.

saying this is that for a good which makes up a small part of total expenditure, a quasi-linear utility function may be a good model (provided we are only interested in that good and the numeraire good, i.e. only have a 2 good model).

(iii) Since the amount of the discrete good that the consumer is choosing does not change, the change in CS is simply the change in CS per unit multiplied by the number of units consumed. The change in CS is therefore -£10.

4. (i) The income elasticity of the demand for pork meat products will depend on whether pork meat products are inferior, normal or luxury goods, and how extreme they are as examples of these cases. Income elasticity is defined as:

$$\frac{\textit{\% change in quantity demanded}}{\textit{\% change in income}}$$

For an inferior good it must be negative. For a normal good it must be positive. For a luxury good, it must be greater than 1 so that the proportion of the consumer's income spent on that good increases as they become richer. Note that a luxury good must be normal, but a normal good does not have to be a luxury good (i.e. I can continue to consume more of a good as I get richer even though the *proportion* of my income spent on that good is dropping). Refer to question 1 for the definition of price elasticity of demand.

How do we apply these concepts to a specific case like pork meat products. Well, we use common sense! Pork meat products are very unlikely to be luxury goods, certainly not at higher incomes because there is a biological limit to how much meat people can actually eat, however rich they are. They are also probably unlikely to be inferior goods since there is plenty of evidence that people do tend to prefer to eat more meat products rather than cereals and grain as they are able to move out of poverty. So we would conclude that pork meat products are likely to be normal, but not luxury, goods.

The price elasticity of demand for pork meat products will depend on the availability of substitutes. If we lived in a world where pork meat products are the only food permitted by the government, the price elasticity would be very close to 0 (remember that it is always negative, assuming the good is normal) because it is a necessity, so people would just have to put up with the higher price in order to fulfil their nutritional requirements. However, in the real world in which we live, there are many available substitutes for pork meat products, and so we would expect the price elasticity of demand to be fairly negative (i.e. demand to be fairly elastic).

(ii) For the supply to be perfectly inelastic, it must be the case that the amount of pork meat sold will be the same, whatever the price level. In the short run, this might be fairly close to the truth, because farmers have a certain number of animals that they must kill given the plans they have made at an earlier date. If demand were to suddenly collapse, farmers would have to sell the meat they have produced at whatever price they can get, because that is better than just letting it rot. By the same token, if demand were to suddenly increase, it would not be instantly possible for more farmers to enter the market and produce more meat, so farmers would be able to charge very high prices. However, in the long run, we would expect a sustained drop in the price of pork to lead farmers to leave the industry in search of better employment, and a sustained increase to lead to farmers entering the industry. In the long run, therefore, the supply would be much more elastic in the price.

(iii) (a) The diagram below illustrates what will happen if scientists publish research which shows that mad cow disease can be transmitted to pigs. The demand curve will shift to the left, because consumers are now willing to consume less pork meat products at each price. This results in a new equilibrium with lower price. The amount of pork meat sold/consumed does not change, because the supply is perfectly inelastic, as represented by the vertical supply curve. Note also that the demand curve has a shallow slope, representing the fact that demand is fairly elastic.

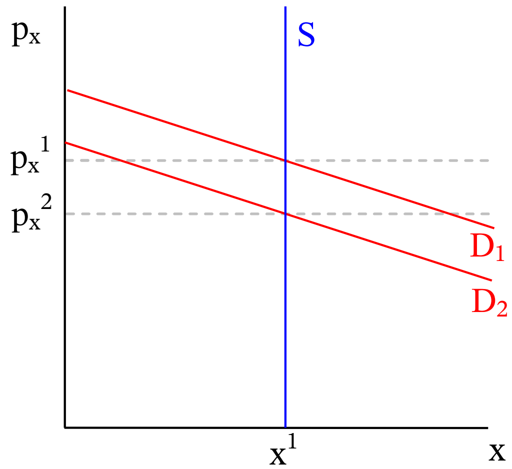

(b) If the government introduces a unit tax on pork meat, it creates a wedge of 1 unit between the price paid by the consumers of pork meat, and the price received by the suppliers of pork meat. This can be represented as a downward shift of the demand curve from D to D'. The consumers of pork meat still pay p_x per unit, whilst the suppliers now receive only p_x'. The amount sold remains the same x^1. So, the burden of taxation falls purely on the suppliers. This result is a special case which occurs because supply is perfectly inelastic. If the supply curve was not vertical, some of the burden would fall on consumer in that they would not be able to buy so much. Because the amount sold to the consumer does not change, there is no deadweight loss. Again, this is a special result due to the fact that the supply curve is perfectly inelastic (i.e. vertical).

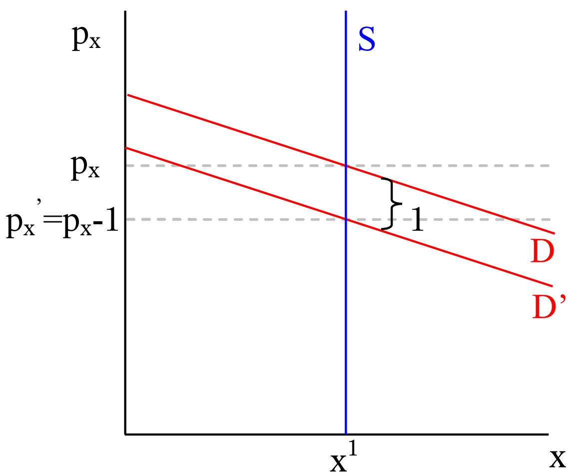

(c) In the long run, the supply curve will be elastic, and so will be upward sloping – as the price of pork meat increases, more will be supplied by farmers. By introducing a minimum price guarantee, the government **must** commit itself to buying the excess supply, because otherwise competition between suppliers to sell their excess supply will push the price down below the minimum price guarantee. Once the price has been fixed at p_x^m by the government, consumer demand will also be reduced, because no consumer will be able to buy a unit of pork meat for less than p_x^m because the government would always be committed to buying that unit for p_x^m. The cost to the government depends on the shortfall between demand and supply at p_x^m. If the supply curve and demand curve are relatively inelastic, like S_1 and D_1, then the total cost to the government would be represented by the small "double hashed" rectangle. If the supply and demand curves are relatively elastic (D_2 and S_2) then the total cost would be the double hashed rectangle plus the vertically hashed rectangles.

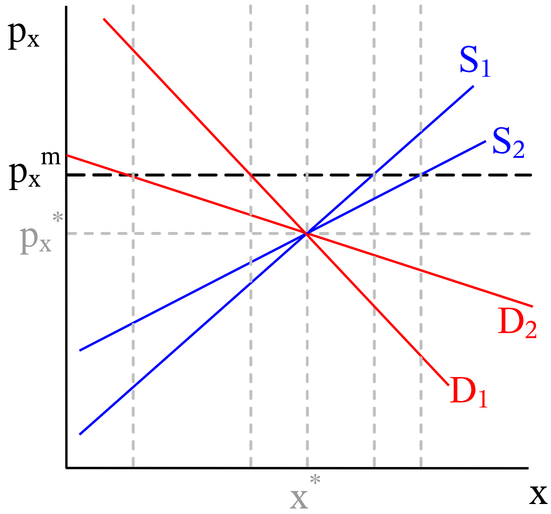

(d) We have already seen that CS is the area under the demand curve above the price, and represents (roughly) the total benefit received by the consumer from being able to buy and consume x units at price p_x. Producers' surplus PS is based on the same idea; it is the total welfare gain the producer gets from being able to sell x units of good x at price p_x. It is the area above the supply curve but below the price because the supply curve represents the price that suppliers are willing to supply each additional marginal unit for, so the area above it represents the total of what the suppliers get in terms of the numeraire good minus what they require to keep them at the same utility level from selling each marginal unit. If this seems unsatisfactory, it is because we have not yet looked properly at the theory of the firm, which we shall cover in topic 4. By the end of next week, we will be able to think about producers' surplus in terms of the relationships between revenues, profits and the cost of producing goods.

At the initial domestic market equilibrium, where foreign pork meat is not allowed into the country, the total CS is area c. The total PS is area a+b. Once pork meat is allowed into the country, domestic producers will no longer be able to supply units at prices over p_x^w. Domestic pork production will be reduced, and domestic PS will go from a+b to just a. So, domestic pork producers will be made worse off. On the other hand, domestic CS will increase from c to c+b+d. The area b represents the welfare transferred from domestic pork producers to domestic consumers. However, our welfare economic analysis allows us to say more than just that domestic producers have suffered so that domestic consumers are better off. The area d represents a pure welfare gain. We could also think of it as the deadweight loss from following a policy of ***not*** allowing foreign pork imports. The overall utility of domestic consumers and producers has increased, i.e. the total surplus TS=CS+PS from domestic pork production and consumption has increased. Does this mean that we can say society is better off allowing pork imports? It depends on how much you value domestic producers and consumers. If you value everyone's utility equally, then you might say that society has been made better off. We will look at these kinds of issues later towards the end of term when we look at welfare economics in a lot more detail.

Probably the most convincing argument is that even if we cannot make them do it in practice, in theory domestic consumers could pay domestic pork meat producers compensation for the introduction of foreign pork imports, and still be better off themselves (because of area d). When the government decides whether to introduce a new policy (e.g. build a road, lower a tax, etc.) then it employs economists (or at least, it should do!) to try to measure these changes in welfare. Usually, some people will be made better off, but others worse off, and it is often in practice impossible to make the winners fully compensate the losers. However, if in principle the losers could be compensated and everyone made better off, the policy should go ahead, since society as a whole has been made better off. By the way, what I have said in this last paragraph is still a contentious issue both within and beyond economics, and goes right to the heart of moral philosophy. There are problems from a practical (can we really measure welfare with any degree of accuracy), theoretical (sometimes we get paradoxes where the winners would be able to "bribe" the losers to accept a policy change but the losers would be able to "bribe" the winners back to the original policy; which policy would be the best then?) and moral (is it right to make some suffer for the benefit of others?) perspectives. However, I personally still find it exciting that welfare economics can shed a great deal of light on what constitutes desirable government policy.