Topic 8 – General equilibrium and welfare economics

3. *Why are pollution problems best understood as externalities? Can we rely on private agents to negotiate these away?*

The definition of an externality is that the consumption of a good (or use of an input in a production economy) by one agent enters directly into the utility function of another. This is precisely what pollution does. For example, suppose we have a two person (A and B), two good (x and y) pure exchange economy. Suppose initially there are no externalities, and that good x represents a specific good which is a small part of total expenditure (e.g. playing music on a hi-fi system) and good y represents everything else consumed. This means that we can use quasi-linear utility functions as an approximate model, which greatly simplifies the algebraic argument:

$$U_A(x_A,y_A)=f(x_A)+y_A$$
$$U_B(x_B,y_B)=g(x_B)+y_B$$
$$MRS_A=MU^x{}_A/MU^y{}_A=\partial f/\partial x_A$$
$$MRS_B=MU^x{}_B/MU^y{}_B=\partial g/\partial x_B$$

The First Theorem of Welfare Economics tells us that (provided we don't have a corner solution) at the general equilibrium price ratio, each consumer will set their marginal rate of substitution (MRS) equal to the price ratio. If we let good y be the numeraire good, this gives us:

$$p_x=\partial f/\partial x_A=\partial g/\partial x_B$$

The First Theorem also tells us that provided there is a market for every good which enters into each consumer's utility function (which is the same thing as saying that there must be no externalities) then the general equilibrium must be Pareto efficient. We can see this in the above example since once the MRS has been equated for consumer A and consumer B, there is no way to make either consumer better off without making the other worse off, because the two consumers each have the same marginal valuation on good x and each have the same marginal valuation on good y. Another way of putting this is that the marginal amount of good y that consumer A would be willing to pay to reallocate one unit of good x from consumer B to himself ($\partial f/\partial x_A$) is exactly equal to the amount that consumer B would have to paid in order for them to agree to this ($\partial g/\partial x_B$). This implies there are no Pareto-superior reallocations from the competitive general equilibrium.

Now suppose that person B is negatively affected when person A plays their hi-fi (e.g. they find person A's music taste abysmal, and hence to them this constitutes noise pollution). The utility functions will now be given by the following (where $h(x_A)$ will always be negative for positive x_A, and will be decreasing, i.e. becoming more negative, as x_A increases):

$$U_A(x_A,y_A)=f(x_A)+y_A$$
$$U_B(x_B,y_B)=g(x_B)+h(x_A)+y_B$$
$$MRS_A=MU^x{}_A/MU^y{}_A=\partial f/\partial x_A$$
$$MRS_B=MU^x{}_B/MU^y{}_B=\partial g/\partial x_B$$

We have assumed that person B's utility function is additively separable between x_A and x_B, which means that the marginal utility of consumer B playing their own music is independent of the marginal disutility caused to consumer B by consumer A playing their music, and vice versa. This simplifies the algebraic argument, but does not change the essential conclusions. Note that because person B cannot directly choose x_A, their only relevant choice as a price taker (and therefore their only relevant MRS) is between x_B and y_B. The presence of the externality does not, therefore, affect the condition derived from the tangency between the MRS and the price ratios, which remains:

$$p_x=\partial f/\partial x_A=\partial g/\partial x_B$$

What has changed, however, is the fact that the general equilibrium is no longer Pareto-efficient. Pareto-efficiency would require that we cannot reallocate goods between the individuals without making one of them worse off. However, the amount that consumer B would be willing to pay to have a unit of good x reallocated to him from consumer A is equal to: $(\partial g/\partial x_B - \partial h/\partial x_B)$. Since $h(x_B)$ is decreasing in x_B then $(- \partial h/\partial x_B)$ is positive and so $(\partial g/\partial x_B - \partial h/\partial x_B) > \partial g/\partial x_B$. Using the condition derived from the equality of the MRS at the general equilibrium, we get $(\partial g/\partial x_B - \partial h/\partial x_B) > \partial f/\partial x_A$. This means that the amount of good y that consumer B is willing to pay to reallocate a unit of good x to themselves from consumer A is greater than the amount of good y that consumer A would require to keep them indifferent after the reallocation. This means that there are reallocations which are Pareto superior to the general competitive equilibrium and which involve moving good x from consumer A to consumer B in exchange for movements of good y from consumer B to consumer A.

If consumers A and B could meet together, negotiate, and then enforce an agreement, we would expect them to strike a bargain involving some kind of reallocation as described in the above paragraph. Provided there are no costs to bargaining, we would expect them to refine the bargain until they have reached a Pareto-efficient allocation. This is essentially the same as the argument that two consumers in an Edgeworth box will continue making mutually beneficial bilateral trades until they reach the contract curve (where no further such trades are possible). The final bargaining point which will be reached depends on the initial endowments. The endowments will determine a competitive equilibrium, which then will determine the set of potential bargain allocations which could be agreed to by both consumers because they are Pareto-superior to the competitive equilibrium.

If we move now to a production economy, one of the most important features of the endowment is the initial allocation of the right to pollute. For example, if factories have a legal right to create unlimited pollution, then the only way they could be prevented from doing so is if those affected by the pollution bribe them to reduce it. On the other hand, we could start from the position that no pollution is allowed unless all those affected agree to it. This would mean that the factory would have to bribe everyone else in order to be allowed to pollute. However, whatever the initial allocation of endowments or property rights, we would expect the bargained allocation to be Pareto-efficient, since otherwise both parties could improve upon the agreement by further bargaining to reach a Pareto-superior point (as argued above for the bargained equilibrium in a two person pure exchange economy).

This result is known as the **Coase Theorem** (named after the economist Ronald Coase). It is illustrated in the diagram below. There are two agents A and B. Agent A does the polluting. For example, we could think of agent A as a chemical factory which pollutes a river and B as a downstream fishery whose fishing yields are negatively affected by the pollution. The MB_A curve represents the marginal benefit of each unit of pollution to agent A. The MC_B curve represents the marginal cost of each additional unit of pollution to agent B. We again assume quasi-linear utility and that y is the numeraire good, so that the utility functions are:

$$U_A = f(q_p)+y_A$$
$$U_B = g(q_p)+y_B$$
$$MB_A = \partial f/\partial q_P$$
$$MC_B = -\partial g/\partial q_P$$

Suppose first that agent A has the legal right to pollute as much as they want. This means that the competitive market price of pollution will be 0. In the competitive equilibrium, agent A will pollute until they no longer gain positive marginal benefit from increasing pollution (i.e. up to q_P'). Agent B would be willing to pay Z+Y (in units of the numeraire good y) in return for agent A reducing their pollution to the optimal level q_P^*, whereas agent A would only require a bribe of size Y in order to agree to this. We would therefore expect an agreement to be made with the size of bribe in between Y and Z+Y. Now suppose instead that no pollution is allowed unless those affected have agreed to it. This means that the competitive price will be infinite. The competitive level of pollution would now be 0 units of pollution. Agent A would be willing to pay W+X as a bribe to persuade agent B to allow them to increase pollution to the optimal level q_P^*. However, agent B would only require X units to be persuaded to do this. So, the bribe will be somewhere between X and W+X.

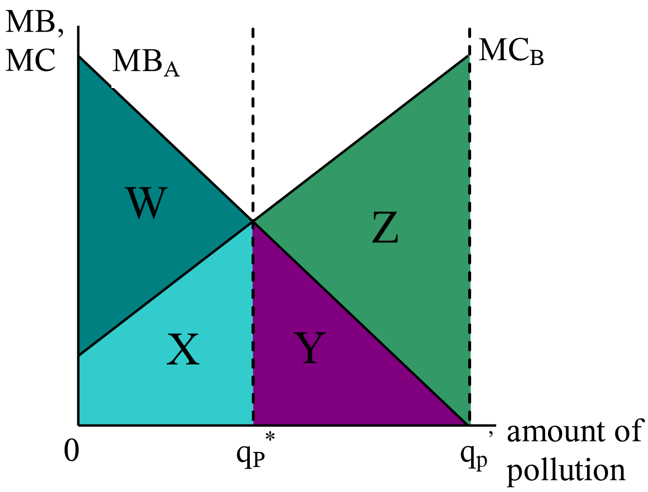

The Coase Theorem suggests that externalities will be bargained away by the interested parties, and so on the face of it need not be a concern for public economic policy. This could occur through a bargain between two autonomous firms as illustrated above, but the existence of Pareto-gains from bargaining means that one of the firms could profitably buy out the other firm. Once this had occurred, the unified profit-maximizing firm would automatically take account of the overall social costs/benefit of their pollution, since these would now be fully reflected in their private profits. Either way, once the structure of property rights leads agents to take into account the full social cost of their actions, we say that the externality has been **internalized**. So, the Coase Theorem suggests that all externalities will be internalized, regardless of the initial distribution of property rights. However, to apply in this straightforward form, it relies on a number of strong assumptions:

1. The involved parties must be able to conduct **negotiations** in which everyone's interest will be registered. This cannot occur if for example, agent A is a single firm but agent B is a large group of individuals affected, since those individuals cannot all be gathered together in the negotiations.
2. The agreement must be **monitored**. This requires that those affected can verify that the agreement is being fulfilled. Again, this will be difficult to achieve if there are a large number of affected agents because no single one of these will find it in their interest to pay for the monitoring (i.e. the monitoring framework is a public good or, to put it another way, the structure of the incentives to contribute to the monitoring framework are that of a prisoners' dilemma, so that it is a dominant strategy for each agent not to contribute). This consideration could be used to argue that the government should monitor pollution and publish the information, even if private agents will be relied upon to use this information in their bargaining.
3. The agreement must be **enforced**. It would seem that even if the government is to rely on these kinds of bilateral bargaining to solve the pollution problem, it will still need to design a legal system so as to facilitate this process. However, it is worth bearing in mind that we could think of agent A and agent B as engaging in a repeated prisoners' dilemma game, and so it is possible that they will cooperate even though they are individually rational and even if there is a detection lag, because the future cost would be higher than the current benefit from defecting (think of the bribe being paid every month in return for keeping the pollution down; if the polluter breaks the agreement, the other agent will never bargain with them again).

Whether or not we can expect agents to negotiate externalities away depends on the nature of the externality. If it is just between two firms (e.g. a single factory and a single fishery) then we could expect the two agents to make a mutually beneficial agreement which would take the economy to a Pareto-efficient outcome. However, if the agents discount the future too heavily, or if detection takes too long, then government intervention may be necessary to force agents to keep the agreement. When we move to an externality produced by a single agent but which affects many agents, we can no longer expect monitoring and enforcement to be provided privately, because they are public goods for the affected parties. However, we could get around the negotiation problem by allowing **class action litigation**, which is where a small number of individuals are able to sue a single agent for the damages caused to a large group of others. This effectively forces a polluting agent to pay the bribe for their right to pollute, even if the affected individuals fail to negotiate any kind of agreement. We have a stronger case for active government policy to facilitate this, but we do not yet have a compelling case for direct government rationing of pollution.

The most problematic pollution externalities are where a large number of agents all carry out an activity which affects a large number of others. Pollution from road vehicles is a good example. There is no-one agent we can hold legally responsible for the damages inflicted upon all. The same applies when are thinking about, say, sulphur dioxide emissions from factories. There is no easy way to solve the problem via class action against particular firms. It is difficult to see how these types of externality can be solved except by government intervention in the prices and quantities of the pollution produced. These could be done with **Pigouvian taxes** (unit taxes placed upon goods in order to bring the private cost more in line with the social cost, named after economist Pigou) or with **pollution quotas**.

Although these types of externalities require the government to directly intervene to reduce the quantity of pollution produced by each agent, there is still arguably an important role for the market mechanism. **Tradeable pollution permits** offer an appealing market based solution because the government then need only choose its pollution reduction target, and then distribute permits up to this total. These permits will then be traded among firms, ensuring through arbitrage that those firms who gain most benefit from being able to pollute do the polluting. Provided the government sets the socially optimal overall pollution level, and effectively enforces the permit system, this would result in a Pareto efficient outcome. This obviously requires information on the part of the government. Pigovian taxes arguably require more information, however, because a more pollution efficient firm would require a lower Pigovian tax (so the unit tax rate on final goods should be different for different firms), whereas if the tax is to be placed on each unit of pollution, the government will have to have a precise estimate of the pollution produced by each firm. This is a stronger requirement than simply a legal system which provides a prohibitive enough punishment for those firms which exceed their quota (e.g. with tradeable permits the government can just do a number of “surprise pollution tests” per year on random firms with fines for exceeding the quota, whereas with Pigovian taxes they would have to keep a constantly updated catalogue of data on the pollution caused by each firm). Basically, the case in favour of a tradeable permit system boils down to the idea that the individual firm will at any one time have a better idea of the costs/benefits of its own pollutive activities than the government. It is another example of the general case for decentralized decision making in the economy. An adequate system of Pigouvian taxes would be very complex to administer and would involve constant bureaucratic intervention. Tradeable pollution permit systems have, since the mid 1990s, been successfully used to reduce sulphur dioxide emissions in the US, and they are officially an important part of New Labour’s climate control policy as part of the Kyoto agreement (although further development on this front has been somewhat lukewarm).

4. *How is the level of optimal provision of a pure public good defined? Why is there a problem of providing such goods in the economy nonetheless?*

Public goods are non-rival goods. A **non-rival good** is a good which, once provided for one individual, can be provided for all other individuals at no additional cost. Some examples are: national defence (once one individual in a nation is protected from invasion, all are), television broadcasts (although television sets cannot be provided at 0 marginal cost, the actual transmission itself can), new scientific knowledge (the medium that it is stored on will have a marginal cost, but the knowledge itself can be used in an infinite number of places simultaneously without any additional cost) or street lighting (once it has been provided on a street at a certain time, any number of people can benefit). A **non-excludable good** is one which it is impossible to force people to pay for when they consume it. Excludability therefore concerns what is feasible in terms of a structure of property rights to govern the production and consumption of a good. A **pure public good** is both completely non-rival and completely non-excludable. It is a kind of ideal type which different goods will approximate to different degrees. Although the presence of either feature by itself will produce some degree of market failure, it is only pure public goods for which the case for government funding is overwhelming.

There is no necessary connection between the two qualities. Public street lighting, national defence (the nation cannot be protected from a foreign invader without all benefiting) and scientific knowledge (pure scientific progress in the broadest sense arguably has such disparate benefits that no firm can gain sufficient private benefit to make it worthwhile to fund it, whatever their property rights) are non-excludable (at least to a very strong degree), and so to all intents and purposes can be thought of as pure public goods. Television broadcasts are, however, excludable (e.g. through the use of broadcast encoding), as are specific new products and production techniques (e.g. via patents). There are also important cases of goods which are rival but non-excludable. These are the goods which are in limited supply but over which governments are unable to effectively enforce property rights so as to internalize their protection within the market. There are numerous examples: clean air, fish in the seas, rainforests. The standard model to illustrate the efficiency loss caused by the non-excludability of rival goods is “the tragedy of the commons” (see Varian p. 589). Whether or not goods are excludable often depends on the technology and institutional structure available at the time. The concepts are therefore not absolute, and there is of course controversy over them (for example, some would question whether scientific progress is really non-excludable, and therefore would argue that we could rely on private enterprise to fund all scientific research). The table below illustrates some examples from the TV industry that fall into the different categories.

	Rival	***Non-rival***
Excludable	TV sets	TV transmissions, new products and production techniques
Non-excludable	Space in river to dump old TV sets	General scientific progress, which affects the TV industry as well as many others

Non-excludable but rival goods require similar policy responses to externalities: either a better structure of property rights so that individual agents will internalize the damage they cause or some kind of overall rationing with a tradeable permit system (e.g. tradeable fish quotas). Excludable but non-rival goods provide some problems (i.e. welfare loss), but generally the market is still the best mechanism to deal with them. The fact that satellite broadcast television is excludable means that it can be provided by a private firm. However, since there is a 0 marginal cost of allowing everyone to consume the broadcast once it has been produced for one person (this is leaving aside the costs of the dish, decoder and TV, i.e. assuming people already own these), then allowing a private firm to produce and make positive revenue by excluding people who do not pay from consuming the broadcast, the market definitely produces a Pareto-inferior outcome to that which a social planner (with perfect knowledge) could achieve. The private firm producing the non-rival but excludable good will be a natural monopolist with a constant marginal cost of 0. We can illustrate the loss of allocative efficiency caused by having a natural monopolist by showing their average and marginal cost curves (the AC curve has the shape it does because of the fixed cost of providing the channel in the first instance):

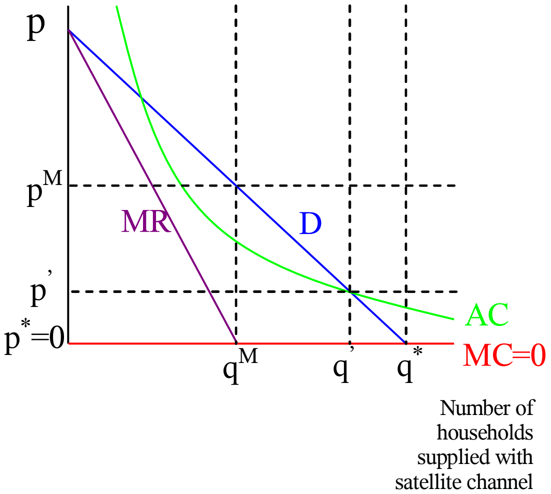

The natural monopoly will profit maximize by producing q^M units and selling at price p^M. The socially optimal outcome is that the natural monopoly produce q^* units and sell at price 0. Only when the good is free to all consumers will the Pareto-efficient outcome be reached, because otherwise the good could be provided at a cost of 0 to people who would still be willing to pay a positive price, but less than the current market price. At the socially optimal quantity, the monopolist would be making a loss and so would have to be subsidized by the government (i.e. the government would have to pay the fixed cost, which is the entire cost of supplying the channel). However, an alternative would be for the government to force the natural monopoly to price at p', where it would just be breaking even. Although this would not be socially efficient, it would have the advantage of allowing the good to continue being provided privately (thus reducing the risk of government failure with public provision) and the government would not have to raise revenue to pay for it (which might distort the economy in other areas). Another alternative would be perfect price discrimination of the part of the monopolist, which would lead to it appropriating all the consumers' surplus and producing at the optimal output level. Although perfect price discrimination is not possible, other forms of price discrimination may reduce the inefficiency of the natural monopolies producing satellite television, although as we saw in earlier weeks, imperfect (first or second degree) price discrimination can reduce overall surplus.

In the UK, we have effectively opted for government provision of terrestrial television via the BBC. Everyone who owns a TV is forced to contribute. Although the above analysis provides a part of the justification (via government coercion we are able to reach socially efficient outcome of universal provision) the argument for specific state provision probably centres on other beliefs surrounding the desirability of public service broadcasting which expands people's horizons rather than just providing them with the entertainment they want to see. Although TV is in theory excludable, the technology to encode terrestrial broadcasting is not yet in place and so, to all intents and purposes, in the UK it is a pure public good.

A pure public good x should be optimally provided at the point where the *sum* of the individual marginal rates of substitution with each other good is equal to the marginal rate of transformation with that good (for each good y $\sum_{j=1}^{N} MRS_j^{xy} = MRT^{xy}$ where N is the number of individuals). This is because taking one unit of the other good away from one individual allows the provision of the pure public good to be increased for *all* individuals. Since each individual in a competitive equilibrium only equalizes their *own* MRS with the MRT, the amount of the pure public good produced competitively will be below the socially optimal level (because MRS^{xy}, the amount of good y that is willing to be given up for an extra unit of good x, declines for each individual as more of good x is consumed). If we make good y the numeraire good and assume quasi-linear utility and a constant marginal cost of 1 for good y, then the condition for optimal provision of the public good becomes that it should be provided up to the point where the sum of the individual marginal benefits is equal to the marginal cost. This can be expressed as the following (where N is the number of people):

$$\sum_{j=1}^{N} MB_j^x = MC^x$$

The fact that we have a condition for optimal provision which looks similar to that of a private good suggests that we may be able to come up with a "market" mechanism which will allocate the pure public good efficiently. One such suggestion is the **Lindahl mechanism**. We imagine a public good which can be provided at a marginal cost of c per unit. However, once the public good has been provided for one person, it can be provided for all others at 0 marginal cost. Suppose for simplicity that there are two consumers A and B (the results fully generalize to the case where there are many consumers). The government declares tax rates t_A and t_B such that $t_A + t_B = c$. These tax rates tell each individual how much of the cost of each unit of the public good they will have to pay. So, if c=1 and $t_A=0.3$ the $t_B=0.7$, and so for each unit of the public good produced, person A pays 0.3 and person B pays 0.7. Each person then declares how much of the public good they would like the government to produce given those tax rates. These desired levels can be described using a "demand curve" for each person. If there is a difference between the desired levels of the two individuals, then the government increases the tax rate on the individual with the higher desired level, and decreases the tax rate on the other. This process continues until both individuals desire the same level of provision, at which point Pareto-efficiency has been reached. Pareto-efficiency is reached at the intersection of the two individuals demand curves because each individual optimally consumes up to the point where their marginal benefit from an additional unit is equal to the tax they are paying, and so since $c=t_A+t_B$ then the condition $\sum_{j=1}^{N} MB_j = MC$ is fulfilled. This is known as a **Lindahl equilibrium**, and is illustrated in the diagram below:

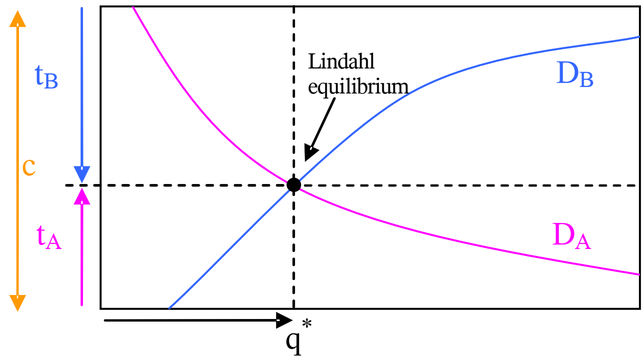

Although the Lindahl equilibrium provides an appealing way in which the government could set up a market mechanism to ensure efficient provision even of a pure public good, it is fatally flawed because it assumes that individuals will tell the truth about the amount of the public good they would like to see produced at each tax rate. It can be shown that individuals can reach a higher utility by lying, so as to cause the government to underestimate the degree to which they value the public good, and thus shift more of the tax burden onto the other individual.

The **Clarke-Groves mechanism** is similar to the Lindahl mechanism, except that it has a system of additional payments which give all agents the incentive to tell the truth about their valuations. It takes the form of a game where all agents simultaneously declare their valuations of the public good, and the structure of payments is such that it is a dominant strategy for each and every agent to tell the truth. However, it is so complicated that most people will not be able to see that it is in their interests to tell the truth. It also requires that people be forced to participate in the preference revelation game, and that none of the revenue raised from the additional payment structure be spent on any of the participants (which would be difficult if it was run on a national scale).

The bottom line is that although economic theory has come a long way towards solving the mathematical puzzle of how to induce truthful revelation of people's valuations of pure public goods, the only practicable solution to the real world problem (given real people's limited rationality, i.e. the fact that there is a "cost to thinking") remains a decision by the political authorities over how much should be produced. Hopefully democratic accountability of the authority will result in a fairly efficient outcome, although whether this is the case is the subject of an entire further sub-branch of economics!