

Important Conditions for Competition

Depend on Secure and Complete Property Rights

- property rights are completely specified
- all property rights are exclusive (no externalities)
- property rights are transferable and enforceable

In considering these necessities, recall Arrow's Theorem of Second Best: here, a system of property rights that satisfies most (but not all) of the conditions is not necessarily better than a system satisfying fewer conditions – counting up the satisfied assumption does not measure how near are the outcomes.

Markets

Recall from microeconomic theory that the First Welfare Theorem guarantees that a competitive market economy (with complete property rights and no transactions costs) is Pareto efficient. Economists believe that markets are generally the best way to distribute resources.

In a perfect economy people don't need to understand all the implications of their consumption on different resources; they only need to know the price. The price is the sole sufficient indicator of scarcity. So much energy is expended by modern consumers trying to balance off different criteria, even for simple choices like a lightbulb. An incandescent bulb uses 'too much' energy relative to a fluorescent, but fluorescent bulbs usually contain mercury (hazardous disposal), other types of bulb might consume particular resources in being made. How ought consumers to trade off greater electricity usage versus mercury contamination? A consumer can be left swamped with information! But in a perfect economy consumers only need to look at the price. Clearly we don't live in a perfect economy.

But many resources are already included in the price of even the most quotidian consumption item. When we choose to buy an apple we needn't worry about whether the farmer has sufficient land or uses the proper fertilizer, or if the wholesaler has a good enough inventory-control system, or if the retailer uses scarce real estate optimally. We just choose whether or not to buy it. It's only when we try to trade off between organic apples or locally-grown apples or fair-trade apples or whatever – that's difficult, because there's no single scoring system.

In a system of optimal economic competition, the price reveals relative scarcity. If supply is low relative to demand then the price will be high; if supply is great relative to demand then the price is low. Early economists often wrote about the apparent incongruity that water, necessary for life, was available for free while diamonds, not necessary for anything, were

expensive. Why this apparent paradox? Because of their relative scarcity. (And thus marginal utility, but that's for later.)

Recall supply and demand graph, plus PS, CS, DWL, so competition maximizes total surplus.

In production, supply prices in a perfectly competitive industry are determined from the minimum point of average total cost – this is the long-run industry supply curve. Firms compete to supply each commodity for the lowest price, meaning that they try to economize on inputs (use the fewest and cheapest possible).

Over a longer time period, firms will direct their Research & Development (R&D) budgets towards economizing on items which are most scarce (i.e. have high prices) – again, just because it's profitable for them to do so.

Markets are extraordinarily powerful. Recall that many countries experimented with central planning (called Communism) and that was a disaster. The best efforts by very smart people (motivated, at times, by fear for their lives) were not enough to supply even a fraction of the goods that could be provided by a market economy. Wise policy will use markets wherever possible. However markets are neither all-powerful nor omniscient. There will be cases where the simple assumptions underlying the Welfare Theorems are no longer valid, particularly where there are substantial amounts of goods with imperfect property rights (with externalities) and/or substantial transactions costs. (Bob Solow refers to the free-marketeers who see the doughnut while the interventionists see the hole (Solow 1974 AER).)

| |
|---------------|
| Externalities |
|---------------|

Externalities are cases of imperfect property rights. If my decision to consume some item has an impact on someone else, then who owns that spillover effect? This can be particularly acute in trying to resolve intertemporal or intergenerational allocations – what if my decisions affect people who will not even be born until the next century?

(Paul Krugman [blogged about Pigou](#), the English economist who first theorized about externalities.)

Examples. Smoking carries an externality: my choice to inhale smoke means that people near me will also inhale smoke. That consumption choice imposes a negative externality. Other consumption choices might impose positive externalities: economists have found significant positive externalities from education, so your decision to get more education will tend to raise the wages that your family and people around you will get. Externalities can arise from production as well as consumption. A factory belching smoke imposes negative externalities on those down-wind. A flower farm might impose positive externalities (more commonly, a beehive kept by someone who wants honey will have positive externalities because the bees can pollinate other flowers of fruits or vegetables). There can be positive or negative externalities; these externalities can arise in production or consumption.

Hanley, Shogren, & White quote Ken Arrow, that an **externality** is

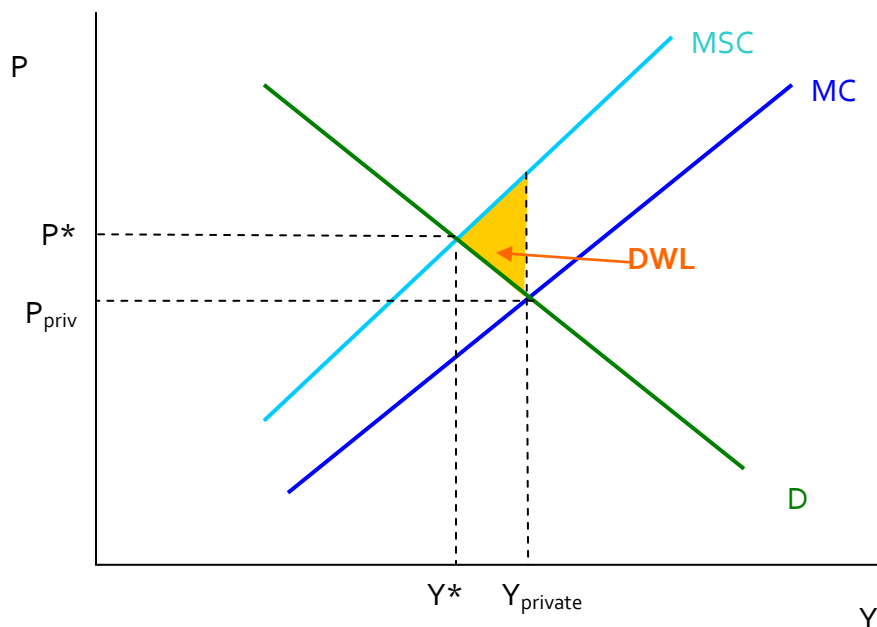
a situation in which a private economy lacks sufficient incentives to create a potential market in some good, and the nonexistence of this market results in a loss of efficiency.

Each word is essential: "lacks sufficient incentives" makes clear that it's not necessarily about technologies but organizations, "potential market" notes that even a possible market has effects (threat of entry or calls/puts), and the final phrase makes clear that not every market failure is insoluble and requires government action.

A lack of a positive externality can be considered a negative and vice versa.

Negative externalities of production produce marginal external costs (MEC) above Marginal private costs (MC, the supply curve). Since these MEC are external to the firms they do not enter into a private firm's calculations of profit maximization so the private firm produces until $P=MC$. But this creates a deadweight loss since at this level the total social costs ($MSC = MEC$ plus MC) are greater than the price, which measures the marginal benefits that people attach to this good. So it costs society more to produce than people value it, which is DWL.

Graphically,



So in this case government intervention can reduce or eliminate DWL. A tax that is just equal to the MEC, or a regulation that limits industry output to Y^* , would reduce the DWL to exactly zero. Consumers should pay more, P^* , since that is the true cost. These taxes are called Pigou taxes after the economist who proposed them originally.

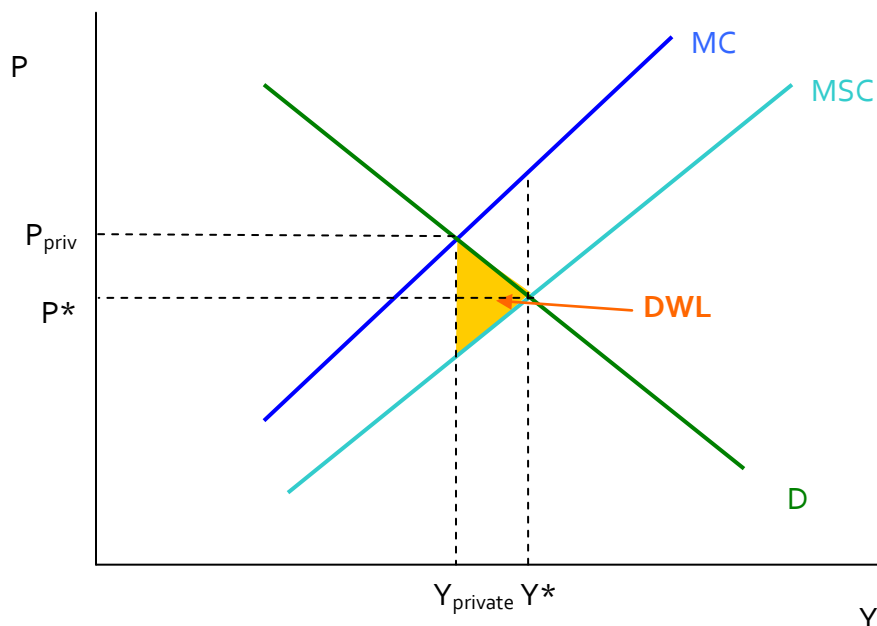
Examples of marginal social costs over and above the marginal private costs are pollution. Decades ago, a firm generating waste might simply dump it into the nearest river. This raised costs for other firms downstream if they needed clean water. (Where by 'firms' I'm including government operations for instance drinking-water treatment plants.)

Externalities loosen the case that individual maximization behavior will inevitably lead to social maximization. Consider the simple case of conversation at a party or bar: you want to talk with someone but there's so much noise that you have to speak loudly to be heard. As everyone in the bar makes this same choice, the general level of noise must rise and so everyone must, again, choose to speak even louder.

Generally externalities break down the argument that all government intervention must produce deadweight loss. Of course government actions are determined by politicians and so are often heavy-handed or even completely wrong, but this must be determined carefully and on the particular facts of each case. General statements, of the sort that politicians and newspaper editorials make, that all taxes are bad or all regulation is wrong – these statements are pure foolishness.

This is the basis for economists suggesting, for example, higher taxes on gasoline. Greg Mankiw, who advised President G W Bush, has a "Pigou Club" of economists lobbying for higher fuel taxes for just this reason (<http://gregmankiw.blogspot.com/2006/10/pigou-club-manifesto.html>). [Note: Mankiw is a clear communicator, which got him into trouble, since his views about the advisability of a gas tax, plus his views that 'outsourcing' is not really a problem, didn't mesh with that administration's overall message. I can disagree with him on many policy issues but still admire him for being intellectually honest in this case even when it was not in his best interest!]

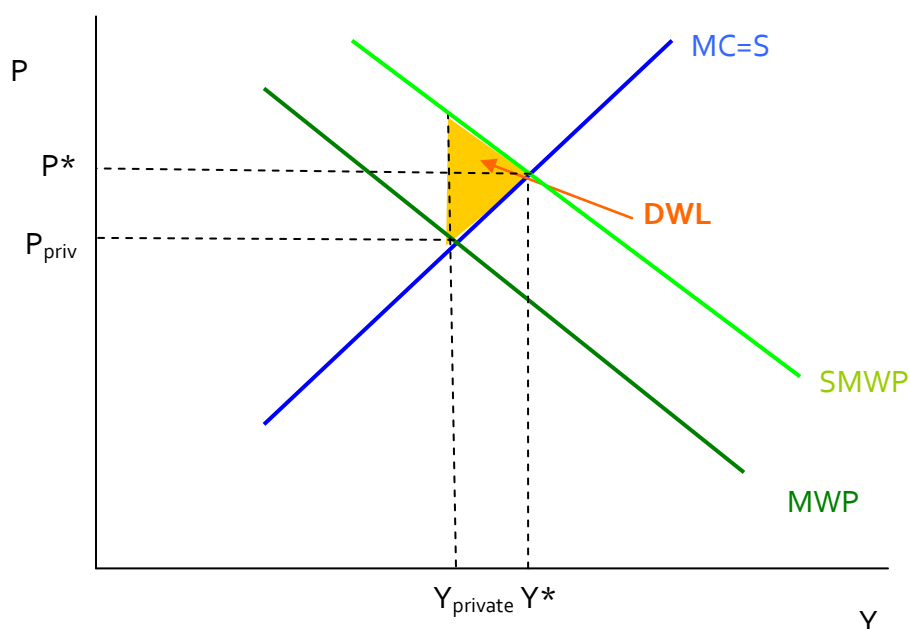
A positive externality in production would shift marginal external costs to the right of marginal cost, creating a different DWL triangle because there would now be *insufficient* production.




Sometimes government intervention in "strategic industries" or to subsidize R&D is justified by this argument. Any single firm might have relatively high costs but the total social cost is lower, so government intervention (subsidizing production) might be justified.

Research into some area, say the basic biological science behind pharmaceuticals, is expensive. There are important knowledge spillovers so a breakthrough in a particular area is likely to lower costs for the whole industry. If you've had a class in Urban Economics you know that many firms choose their location based on these sorts of knowledge spillovers. Government-sponsored research in the San Francisco Bay area led to many hi-tech firms starting up there; now Silicon Valley is a highly productive location for a wide variety of tech firms. In New York, hip design firms choose to locate in areas where there is already a density of other hip design firms (Brooklyn or Tribeca or SoHo or wherever) – there are positive externalities to these locations that are not available in Flushing, or Newark.

Externalities in demand would shift the marginal social benefit curve to the left or to the right of the marginal private benefit (demand) curve. Positive externalities of demand are "bandwagon" effects or "network" effects – Facebook is popular because 'everybody' has a FB account so MySpace is dying (and Google's effort died a while ago). Negative externalities of demand are congestion effects – when the iPhone was introduced on AT&T's network, the huge demands for bandwidth slowed down everybody's phone (will Verizon have similar troubles?). Traffic has this effect.





So in each case, a tax or price/quantity restriction can actually reduce the deadweight loss and make everybody better off.

Vertical Sum not Horizontal

Unlike the case of private demand where the market demand is the horizontal sum of the individual demands, the SMWP is the vertical sum of each individual's marginal willingness to pay (MWP). Because the nature of the externality means that the consumption is shared, we don't add up how many are demanded by each individual, at a given price. Rather we ask, if society were to consume one more unit (such consumption would be shared by many individuals), how much each individual would be willing to pay – and add up each individual's marginal willingness to pay.

These items can be positive or negative: I might be willing to pay something for public consumption of some good, or I might be willing to pay an amount to *avoid* the public consumption of that good.

Rival and/or Excludable Goods versus Pure Public Goods

A problem with providing public goods is that everybody tends to wait around for someone else to do the hard work. The idea is that, if the problem impacts somebody else, then that person might do the hard work and then I can just take the externalities – get the benefits without any of the costs. For example the global campaign to restrict carbon emissions suffers from this free rider problem: every country wants the other countries to take all the pain.

We can generally distinguish goods as either **excludable** or **non-excludable** and either **rival** or **non-rival** (in any combination).

Excludable goods mean that the technology exists to keep other people from using my stuff – kids fight in order to make their toys excludable, a mass of laws against theft and robbery help me keep my stuff excludable. Non-excludable is the opposite: I can't keep people from using it. Perhaps it's an architecturally lovely building that every passer-by can enjoy. Or the neighbor without curtains. Intellectual property law exists to try to make certain goods excludable.

Rival means that someone else's consumption of the good interferes with my own. If someone else eats my cookie then I can't eat it – cookies are rival. Non-rival is the opposite. Sometimes these distinctions are a bit arbitrary: parents don't understand why kids can't share toys, "If you're not playing with it now, why is it a problem if the other child plays with it now?" just like many people would consider their jewelry rival (even though the same argument could apply – but almost nobody, really, rents jewelry for a night out. The bling is only valuable if it's yours.).

Economists label goods that are non-rival and non-excludable "pure public goods." These are often goods that are provided by governments. Police and fire protection are difficult to exclude (both because of externalities) and, given the infrequency of occurrences, are basically non-rival. There are private security guards but these are not as common as police. National defense is non-excludable and non-rival.

But other goods, which the US government does not often produce, are also non-rival and non-excludable. Radio is – my listening to a particular station does not impact your listening (assuming the volume levels are low enough). (In other countries radio is produced by the government; in the US there is a modest subsidy to public radio.) Record companies and software companies are battling (mostly, failing) to make music and software excludable – even though any teen-ager with internet access can rip and burn music. Certainly it is non-rival, since I can copy a single mp3 file as many times as I like, without impairing my own enjoyment of it. The movie studios are terrified that their output will go the same way once users get enough bandwidth to easily swap movies peer-to-peer (since BitTorrent and similar apps haven't gone mainstream).

While these goods are not provided publicly, their peculiar character means that pricing must take different forms. Radio stations play advertisements if they broadcast for anyone; satellite radio makes their product excludable by encoding the broadcasts and selling the decoders. You can think of many more examples.

But in general, whereas we were able to prove the First and Second Welfare Theorems, in the case of no externalities and perfect property rights, to show that private markets produce Pareto-optimal outcomes, this is no longer the case when there are externalities or imperfect property rights. Markets are best wherever possible but they are not always possible.

This does not mean that every externality demands government intervention! Markets are dynamic and give participants incentives to figure out ways to exclude rivals, as the examples

above clearly show. TV stations originally broadcast over the airwaves to everyone; now cable and satellite broadcasts require de-coders. Music companies are slowly trying to figure out how to exclude copying of their products (or figure out other ways of getting revenue – right now ringtones are supporting the labels!). Internet radio like Pandora or last.fm are complicating; Apple's iTunes store crunches the music companies' margins but offers greater security.

There are also cases where private citizens will join together and voluntarily restrict their own choices. Buying a coop or condo means that you agree to be bound by the decisions of a managing board, exactly in order to keep others from imposing externalities on you. If one person doesn't maintain his unit then the board has a legal basis to force the owner to make improvements. Business Improvement Districts (BIDs) have some of this character.

Free Rider Problem

People have an incentive to 'free ride' on other people's willingness to pay. Each would want the other consumer to pay more. I might claim that, actually, my preferences are not like my neighbor's; my neighbor cares greatly about the quality of the public good while I hardly care at all – so my neighbor should pay most of the cost. My neighbor, of course, will likely make the same claim.

Consider common debates about public taxation levels. Some people want the government to levy higher taxes and provide more services; others want lower taxes and fewer services. (In defiance of the facts, the former group would more commonly be associated with Democrats and the latter with Republicans.) Sometimes lower-tax supporters will assert, "Well, if you want higher taxes, why don't you start by volunteering to pay more tax yourself?" The public good argument and marginal-willingness-to-pay argument shows why that argument is fallacious.

This problem, of consumers having an incentive to "fake" their marginal willingness to pay for an item, does not occur in the case of private goods, because for ordinary goods, if I don't pay the price I don't get to consume the item. If I go to the coffee shop and offer just 20 cents for a cup of coffee, they won't give it to me. But with a public good, I have an incentive to try to get my neighbor to pay for the public good so that I can consume it for free.


Advanced: The Consumer's Problem for the case of Externalities

Economics investigates many cases of externalities; some of these relate directly to the environment. My decision to purchase organic food might help the people who live near the farmer's fields (which no longer are sprayed with dangerous chemicals). Or externalities could relate to networks or other non-environmental issues.


But for now consider two consumers choosing between two goods, x and y , where y is a pure public good (define) that would only be provided by some external organization (like a government). How much of the public good should be provided? Or, equivalently, how much would the two people be willing to spend?

This decision can be enormously complicated if we worry too much about income effects and complementarities among goods. If the free public goods are mp3 files of top music, provided by the internet, then my marginal utility for these goods might depend quite heavily on my possession of an iPod or computer. More seriously, there has been a lengthy debate on the degree to which people demand environmental services as they get wealthier.

But for now we start simply and work our way up. For any ordinary good we can graph a consumer's demand curve: the marginal benefit gained by consuming one more unit of the good. In general this demand curve will slope downward due to diminishing marginal utility.



For a public good we can ask the same question: what is the marginal willingness to pay by the consumer for a one more unit of this public good? Again, this will generally slope downward.



This can be again caricatured as the demand curve for the public good, although it has significant differences from a typical demand curve – crucially, that payment for public goods can be difficult to arrange.

One utility function, which is easy to work with, is the quasi-linear utility, where x is typically interpreted as a composite good (a basket of ordinary consumption items) with price normalized to one and y is the public good, $U(x, y) = x + \sqrt{y}$ so that $MU_x = 1$ and

$$MU_y = \frac{1}{2\sqrt{y}}.$$

The marginal condition, that $\frac{MU_1}{p_1} = \frac{MU_2}{p_2}$, gives $\frac{1}{1} = \frac{\frac{1}{2\sqrt{y}}}{p}$ once we substitute in for each term, where p is the price that people would be willing to pay for the public good (not necessarily the price that they actually pay). Take the marginal condition and simplify to get

$$p = \frac{1}{2\sqrt{y}} \text{ or } y = \frac{1}{4p^2};$$

the graph looks like the "Marginal Willingness to Pay" above

[assuming the person has adequate budget to buy it]. There is an inverse relationship between the amount of the public good consumed and the marginal value attached to it.

Note that this is not the total value attached to the public good, just the willingness to pay for an additional unit more – that's why it's called *Marginal*. This is just the same as the case with ordinary private goods: the fact that I willingly pay \$1 for another cup of coffee does NOT imply that I would give up all of my coffee intake for \$1, only that my caffeine consumption is already high enough that I would only pay \$1 for yet another cup.

Now suppose there were two people who could consume this public good. How much would

these two people be willing to pay for this public good?

We figured out a relationship for the first person,

$$p_1 = \frac{1}{2\sqrt{y}}, \text{ where } p_1 \text{ denotes the marginal}$$


willingness to pay of the first person, and p_2 the marginal willingness to pay of the second person. So if the government provided a unit of y , which was nonrival and so could be used by both consumers, then society would be willing to pay up to $p_1 + p_2$ (as mentioned, this is the vertical sum).

If the two people had identical preferences then $p_1 + p_2 = \frac{1}{2\sqrt{y}} + \frac{1}{2\sqrt{y}} = \frac{1}{\sqrt{y}}$.

Graphically, this is a 'vertical' summation: add up the amounts that each person is willing to pay and that total price is the marginal amount that society would be willing to pay.

From micro theory, demand curves for private goods are the horizontal sum of individual demands, not the vertical sum.

So we can graph this to the right:



This basic principle applies whether the public goods have positive or negative externalities. Basically, the lack of a bad thing can be considered a good thing, for example if trash piling up is a bad then we can redefine and set trash collection as a good (last year this was a pressing concern in Naples).

Of course this assumes that there is some way to get people to reveal how much they'd be willing to pay for these public goods. This can be difficult...

Person 1 would willingly pay 0.5 in order to get 1 unit of the public good, y – which assumes that the other person is also paying 0.5. If there is not a full unit of the public good provided then Person 1 would not be optimizing. Person 2 will get utility from the public good provided by Person 1, even if Person 2 contributes nothing.

Consider now the case of two consumers with slightly different preferences: now person 2 has quasi-linear utility of the form $U_2(x, y) = x + 2\sqrt{y}$ so that $MU_y^2 = \frac{1}{\sqrt{y}}$. Now the

marginal condition gives $p_2 = \frac{1}{\sqrt{y}}$ and so the Social Marginal Willingness to Pay is

$p_1 + p_2 = \frac{1}{2\sqrt{y}} + \frac{1}{\sqrt{y}} = \frac{3}{2\sqrt{y}}$, so that, for example, if $y=1$ then person 1 is willing to pay 0.5 but person 2 would pay 1. So society overall would pay as much as 1.5.

How could these two people find this out? They have no incentive to tell the truth because they have no way of finding out the other person's true utility function.

What levels would be chosen, if the people were choosing individually? For simplicity we'll return to the case of two identical individuals with $U(x, y) = x + \sqrt{y}$, $MU_x = 1$ and

$MU_y = \frac{1}{2\sqrt{y}}$. But now we differentiate between how the individual could get this 'y' good,

since it is non-rival. Either the consumer could buy her own or she could just use what others have bought.

Notate the amount of the public good bought by an individual y ; the amount of the public good that others have already bought is Y (capital letter). Each unit purchased costs price p . So an individual consumes an amount $(Y + y)$ of the public good and $(x - py)$ of the private good (since after paying for y units of the public good she has only that much income left over for spending on x).

With the given utility function this is $U(x, y) = U(x - py, Y + y) = x - py + \sqrt{Y + y}$, where the person is choosing x and y , so we need marginal conditions for these two goods but not for Y since this is not chosen. So $MU_x = 1$ and $MU_y = \frac{1}{2\sqrt{Y + y}} - p$ (where the latter term, $-p$,

comes from taking the derivative, with respect to y , of $x - py$) and we set the marginal

conditions as $\frac{MU_x}{1} = \frac{MU_y}{p}$ (as usual we set the price of the private good equal to 1) and so

we get $\frac{1}{1} = \frac{\frac{1}{2\sqrt{Y + y}} - p}{p}$, which we simplify to get $p = \frac{1}{2\sqrt{Y + y}} - p$ or $p = \frac{1}{4\sqrt{Y + y}}$, which

we invert to find the demand curve, $y = \frac{1}{16p^2} - Y$. But how much of Y will be produced? If

we assume that all of the consumers are identical and that there are n of them, then the

other $(n-1)$ will do the same as the person under consideration, so $Y = (n-1)y$. Substituting in to find y_n , the amount chosen in private equilibrium, gets $y_n = \frac{1}{16p^2} - (n-1)y_n$ so $y_n = \frac{1}{16np^2}$, which means that the total amount chosen is $ny_n = \frac{1}{16p^2}$. (Note to those who know micro theory: yes, this is the Nash game solution.) Most worryingly, the amount chosen by each individual falls when there are more other people around, who I believe will 'pick up the slack.'

But how much would be produced, if the people could get together and agree on an optimal social amount (somehow read each others' minds to find out how much they'd be willing to pay)? Now people would maximize their utility, $U(x, Y) = x + \sqrt{Y}$, but the price of Y is $\frac{p}{n}$ since all of the population will pay for an equal part of the total amount that is consumed. So

now the marginal condition sets $\frac{MU_x}{1} = \frac{MU_Y}{\frac{p}{n}} = \frac{1}{1} = \frac{\frac{1}{2\sqrt{Y}}}{\frac{p}{n}}$ so Y^* , the optimal amount

chosen by the optimal social welfare maximization, is $Y^* = \frac{n^2}{4p^2}$.

Compare this amount with the private solution amount to see that $Y^* - ny_n = \left(n^2 - \frac{1}{4}\right) \frac{1}{4p^2}$, which will be positive whenever $n > \frac{1}{2}$ -- i.e. it will always be positive for public goods! The divergence will get bigger for larger populations, as well.

So while there will generally be some private provision of the social good, this will generally be much smaller than the amount that would be socially optimal. And the size of this divergence will grow bigger when there are more people sharing the externality.

You should be able to do this same analysis with a different utility function, such as Cobb Douglas. For this, $U(x, Y) = \sqrt{xY}$ and $MU_x = \frac{1}{2}\sqrt{\frac{Y}{x}}$, $MU_Y = \frac{1}{2}\sqrt{\frac{x}{Y}}$. For the private-

provisioning case, $U(x - py, Y + y) = \sqrt{(x - py)(Y + y)}$, $MU_x = \frac{1}{2}\left(\frac{1}{\sqrt{(x - py)(Y + y)}}\right)(Y + y)$

, $MU_y = \frac{1}{2}\left(\frac{1}{\sqrt{(x - py)(Y + y)}}\right)(px - 2y - pY)$ (this looks really ugly but many of the terms cancel so it's not quite as bad as it looks).

In our society probably the most common method of determining optimal social policies is voting, which will not in general produce optimal results but might be satisfactory. Recall the Arrow Impossibility Theorem which stated that democracy is not rational; also Churchill's "democracy is the worst form of Government except all those other forms that have been tried from time to time."

If people's preferences have some homogeneity (they're not too diverse) then voting can even be optimal.

Society has created a wide array of institutions that counteract the problems that arise from externalities. At one point these were largely based on sociological mores and traditions. Now many are contractual; in some cases governments have stepped in to formalize particular legal constructions – from the modern corporation to housing coops, condominium associations, business improvement districts, and so on.

The formal analysis mirrors the Nash game of oligopoly: although each participant would like to buy more Y (or charge a higher price), they do not do this because they assume that others would not be so 'public spirited' as to also buy more Y (or charge a high price) so they compete.

It is like a Prisoner's Dilemma. Return to the case of two identical individuals with

$U(x, y) = x + \sqrt{y}$, $MU_x = 1$ and $MU_y = \frac{1}{2\sqrt{y}}$. Their social optimum is to pay 1 and get 1

unit of public good (assuming $p=1$; this is $Y^* = \frac{n^2}{4p^2} = \frac{(2)^2}{4(1)^2}$). But if they choose individually

then they'd each choose $y_n = \frac{1}{16np^2} = \frac{1}{16(2)(1)^2} = \frac{1}{32}$ so there would be just 1/16 of this public good in total.

We could simplify this as a Prisoner's Dilemma:

| | Person 2 Cooperate | Person 2 Compete |
|--------------------|--------------------|------------------|
| Person 1 Cooperate | | |
| Person 1 Compete | | |

But we need to fill in the Utility values in each bin. We assume that each person has a budget of 1; the amount of good x that is chosen is simply the remaining budget. Setting

$Y=1$ implies that each chooses $y=1/2$ so $x=1/2$ and $U = \frac{1}{2} + \sqrt{1} = 1.5$. Setting $y=1/32$ so $Y=1/16$

means instead $U = \frac{31}{32} + \sqrt{\frac{1}{16}} = 1.22$. But if the other person buys 1/2 then I buy

$y = \frac{1}{16p^2} - Y = \frac{1}{16(1)^2} - Y = \frac{1}{16} - y - y'$ (where y' is the other person's choice of y). If the

other person sets $y'=1/2$ then I would set my own y at zero (can't be negative) so my utility would be 1.71; the other person's would be 1.21.

So this gets us this Prisoner's Dilemma table:

| | Person 2 Cooperate | Person 2 Compete |
|--------------------|--------------------|------------------|
| Person 1 Cooperate | 1.5, 1.5 | 1.21, 1.71 |
| Person 1 Compete | 1.71, 1.21 | 1.22, 1.22 |

So "Compete" is a dominant strategy. As typical with this analysis, it could be extended to multiple interactions, complete with reputational games, random strategies, etc.