<u>MEASURING WELFARE</u>

Calculating the "gains from trade" in our numerical comparative advantage examples was simple. But how do we decipher the gains in a demand-and-supply diagram? I have suggested the areas of the 'little triangles' was the way to measure it. Here, we elaborate upon that more fully.

To understand how, we need to introduce two measures of welfare gain/loss: the "**consumers' surplus**" and the "**producers' surplus**". They are both used to measure **welfare**, or more precisely the benefits and costs of exchanging goods. It is a major building block of **cost-benefit analysis** of public projects.

Contents

CONSUMERS' SURPLUS

If you said "thank you" the last time you purchased a product, did you mean it? If you were willing to pay more than the price you actually paid, you received what's called 'consumer surplus'. Consumer surplus is a measure of how well-off a household or society is when buying a product.

Defined:

> "**Consumer surplus**" is the difference between the maximum price a consumer is *willing* to pay and the price he or she *actually* pays.

Intuitively, suppose you are prepared to pay up to $15 for a shirt. You go to a store and find that the price of shirts is actually $10. You buy it. What is your gain from trading with the store? The $5 that you *don't* spend is the **gain** to your personal welfare. You were willing to put down $15 for it, but only had to put down $10, pocketing the $5 difference for yourself to spend on whatever else you want.

In a general market situation, figuring out the total gains consumers make is a simple matter of adding up the consumer's surpluses for all the buyers. Diagrammatically, the **total consumers' surplus** is measured by the area of the triangle formed by the demand curve and the market price, e.g. the shaded triangle in the diagram below.

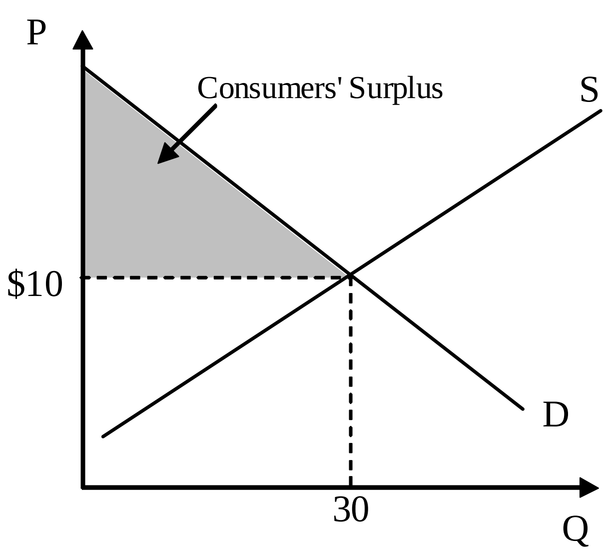

Fig. 1 - Consumers' Surplus

The usefulness of 'consumers's surplus' is that it is a great way to track and measure precisely how changes in the market affect the welfare of consumers. Take fluctuations in prices. As you can see immediately from Fig. 2, if the market price rises from $10 to $12, then the area of the triangle gets smaller, that is, the consumer surplus contracts. That is a way of saying that consumers are worse off as a whole.

How much worse off exactly? By the difference in the areas of the large triangle and the smaller one. The area of the darkly-shaded shaded polygon gives us the precise measure of the consumers' **welfare loss** from a price increase.

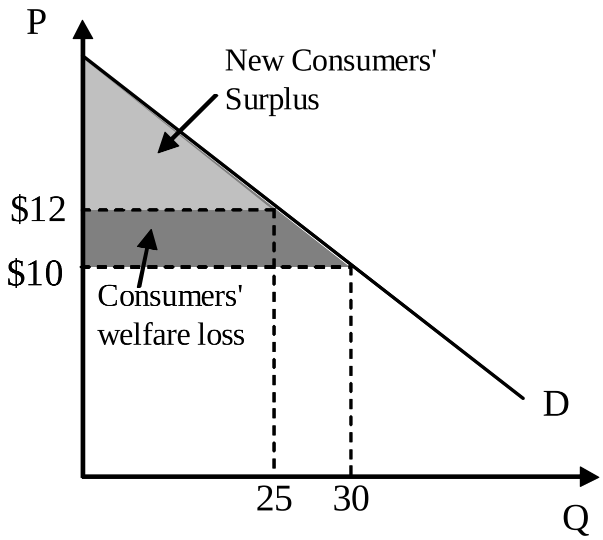

Fig. 2 - Rise in market price – consumers' welfare loss

If, contrarily, price falls, the area of the triangle gets larger, that is consumer surplus increases. Consumers are better off as a whole. e.g. If we started off at $12 and the price fell to $10, then the darkly-shaded polygon would be the exact measure of the consumers' '**welfare gain**' from the price drop.

[For the record, the 'consumer's surplus' measure of welfare was popularized in 1890 by the English economist Alfred Marshall in his *Principles of Economics* – at about the same time he arbitrarily decided (boo, hiss!) to flip the x & y-axes of the demand-and-supply diagram into the unnatural manner we use it today.]

Constructing Consumers' Surplus

I have given you definition of consumers's surplus, and asserted the area under the demand curve is the measure of its total amount. You may not immediately see the connection between the two, so it is useful to take a step back and see how we get from one to the other.

Look at a very simple example of the market demand for shirts in Figure 3. Suppose for the moment that market price of shirts was $20. In that case, then 1 shirt would be bought. That means there is someone out there *willing* to pay $20 for it. Call him Mr. Alpha.

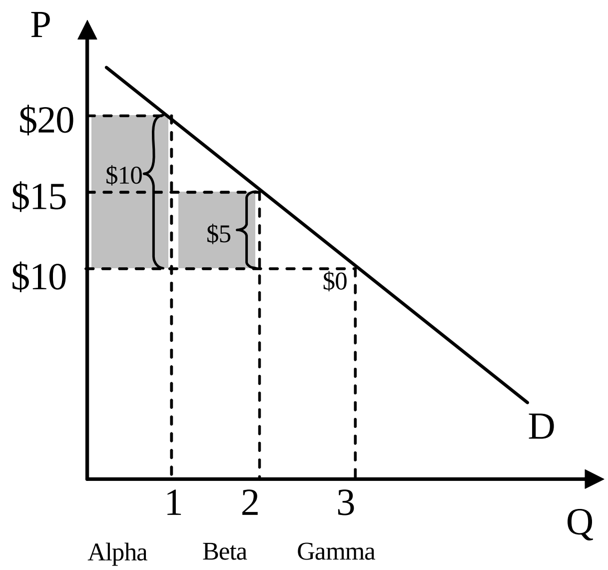

Fig. 3 – Consumer's Surplus, broken down

At the market price of $20, Mr. Alpha pays *exactly* what he is willing to pay for it. In other words, it ain't no bargain for him. He has no consumer's surplus.

But suppose instead that the market price is $15. As we see, at this price, 2 shirts will be sold. Say, Mr. Alpha buys the first shirt and another buyer, call her Ms. Beta, buys the second. Now, because market price is the *same* across all consumers – Mr. Alpha pays $15 for his, Ms. Beta pays $15 for hers – then Mr. Alpha is making a welfare gain on his purchase. Specifically, we know from before Mr. Alpha is *willing* to pay $20 for his shirt, but now he only *actually* pays $15. In other words, he is saving $5. That is *his* consumer's surplus.

In contrast, Ms. Beta makes no consumers' surplus at $15. She wasn't willing to buy a shirt when the price was $20, and is only willing to buy when the price fell to $15. So *she* is paying the maximum she is willing to pay for her shirt. She can live with the $15 shirt, but it ain't as much a bargain for her as it is for Mr. Alpha.

Now, suppose the market price was $10 instead. A this market price, 3 shirts are sold – the first to Mr. Alpha, the second to Ms. Beta and the third to Ms. Gamma. At this price, Mr. Alpha, who, as we saw, was willing to pay $20 for his shirt, only pays $10, and thus is enjoying a welfare gain of $10. Ms. Beta, who we know was willing to pay $15 for hers, is making a welfare gain of $5, while Ms. Gamma, who only now, at $10, is willing to buy a shirt, is paying the maximum she is willing – thus makes no welfare gain.

So what is the total welfare gain across all consumers when the price is $10? The total consumers' surplus at this price of $10 is the sum of all the individual consumer's surpluses:

> Total Consumers' Surplus ($10 price) = Alpha's gain + Beta's gain + Gamma's gain
> = $10 + $5 + $0

= $15.

Going back a step, for comparison's purpose, the total consumers' surplus at the higher price of $15 is:

Consumers' Surplus ($15 price) = Alpha's gain + Beta's gain + Gamma's gain
= $5 + $0 + $0
= $5.

So, when we *lower* the market price from $15 to $10, total consumers surplus increases by $10 (from $5 to $15). That $10 increase is the measure of the welfare *gain* from the proposed price decrease.

Now, to get the final step from the measure of consumer's surplus to area of the triangle under the curve, presume that price is $10 again. At this price, we know the total consumers' surplus is $15. That happens to be precisely the total area of the shaded rectangles in Fig. 3. The shaded rectangle on the left (Mr. Alpha's) has area (base × height) = 1 × $10 = $10. The shaded rectangle in the middle (Ms. Beta's) has area 1 × $5 = $5. So the sum of the areas of the two rectangles (total shaded area) = $10 + $5 = $15. Exactly the value of the consumer's surplus we calculated.

Now, the two little rectangles in Fig. 3, taken together, former a triangular-like shape. As you will notice, $15, is the area of the triangle under the curve (or, equivalently in our discrete case, the sum of the shaded rectangles). Remembering elementary school math:

Area of a triangle = (1/2) × base × height

Since, in our figure, base = 3 (shirts bought) and height = $10 (price difference between $20 and $10), then:

Consumer's Surplus = Area of triangle = (1/2) × 3 × 10 = (1/2) × 30 = 15.

Of course, two little rectangles don't make a very convincing-looking triangle. But if there were more customers and more prices, there would be more rectangles filled. And the more rectangles, the more "triangular"-looking the shape. Look at Fig. 4 below. The total consumers' surplus at $10 (shaded rectangles) begins to look nearly like a proper triangle.

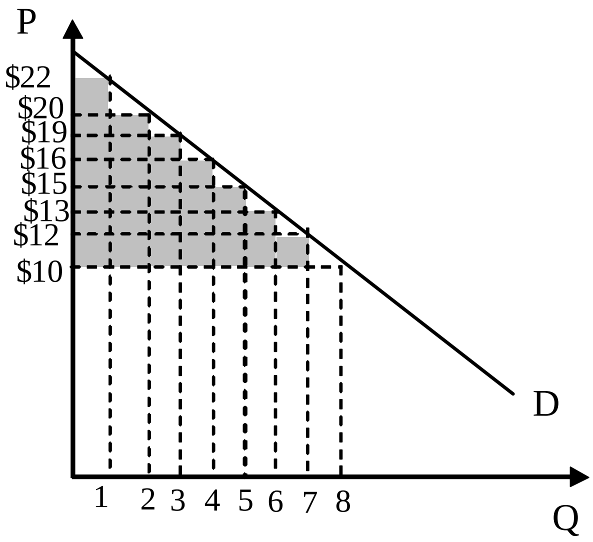

Fig. 4 - Lots of customers & prices

What is the consumers's surplus in Fig. 4? Well, the base = 8, height = 12, so:

Total Consumers' surplus = (1/2) × 8 × 12 = $48.

With enough customers the whole area under the demand curve above the $10 price would be perfectly filled in and we'd have a perfect triangle we had in our very first Fig. 1

[**Warning!**: we can only use our triangle area formula *if* our curves are straight lines and the numbers actually make the area under the curve a triangle rather than some other polygon. There is *no* economic requirement that demand and supply curves are straight lines. They can be curved, wiggly, misshapen, etc. We have forced these examples to yield us exact triangles so as to allow us to use the simple triangle area formula to calculate the consumers' surplus. But in *most* cases, the area is usually not a perfect triangle and the consumers' surplus area must be calculated by the more complicated methods of integral calculus.]

Aside: Water-Diamond Paradox

Philosophers in ancient times were puzzled by the famous "water-diamond" paradox. The paradox was that things which have great value to people – like water, which is essential for living and people would not want to do without – are often very cheap to acquire, even free. Whereas things which are largely useless – like diamonds, wanted for vanity, but not essential for living – are very expensive.

To sort it out, Adam Smith differentiated between two kinds of value – "value-in-use" (usefulness) and "value-in-exchange" (price on market). If a thing's *usefulness* to a person is reflected in his *willingness* to pay for it, then we can say that consumer's surplus is precisely the difference between value-in-use and value-in-exchange.

[Note: The solution to the water-diamond puzzle was simple: scarcity. Diamonds are extremely rare, water is abundant, so the market price for diamonds will always be higher. Except, of course, if you're in the middle of a desert.]

<u>An Aside: Price Discrimination</u>

The consumer's surplus measures welfare in the sense of the total *savings* consumers make by buying shirts at the market price – adding up the amounts they were *willing* to pay, but didn't end up having to pay. That is, the gains from trading in the market.

What if a firm could *force* them to pay the prices they were willing to pay? Then they couldn't charge the same price for every shirt sold, but would charge different customers different prices.

This phenomenon is known as "**price discrimination**". That is, charge people different prices for the same good – say, charge Mr. Alpha $20 for his shirt, charge Ms. Beta $15 for hers and charge Ms. Gamma $10 for hers. If the store does that, then there is *no* consumer's surplus since each consumer pays *exactly* what they are willing to pay and no less. It would look like the diagram below: no savings (shaded areas) for anybody, no consumer surplus at all.

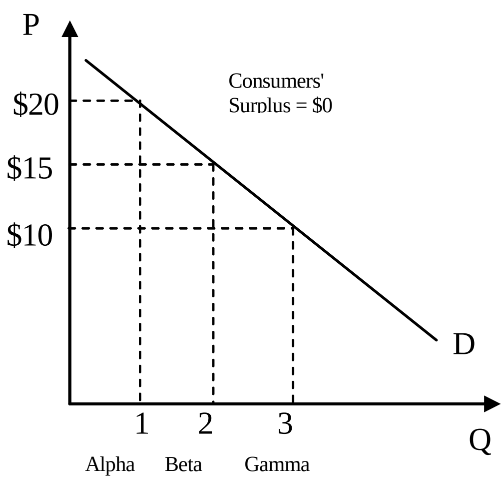

Fig. 5 - Perfect Price discrimination

Naturally, the store selling shirts would *love* to be implement perfect price discrimination. It gets every penny it can out of its customers. But that rarely happens in real life. That is because stores usually have no idea *who* is willing to pay what. Faced with this lack of personal information, stores have no choice but to treat all their buyers the same and charge them all the same price.

That doesn't mean that they don't try to discriminate and use all sorts of ruses to figure it out. But, by and large, this is difficult to do in most regular retail situations. However, it is not altogether impossible and, in certain areas, like bank loans, insurance and education, where a lot of personal information is provided, they succeed.

Sometimes, even without requiring information, it is relatively easy to induce a customer to reveal his willingness. For instance, a pharmaceutical company that sells a drug that can be used for both humans and animals, knows it can charge a higher price to customers who intend to use it on themselves but must give a lower price for customers who intend to use it on their pets. Price discrimination is just a matter of different packaging – labelling one bottle "People Medicine" and the other "Dog Medicine", with different price tags. Then just sit back and let the innocent, ignorant customers "discriminate" themselves.

(Note: firms attempting to price-discriminate must be careful not to fall foul of the law. The Robinson-Patman Act of 1936 makes price discrimination illegal under certain circumstances. And attempts to implement price discrimination sometimes overlaps into thoroughly illegal areas of social discrimination – like racial profiling, which frequently happens in bank lending.)

Perhaps the most familiar example of price discrimination is the internal "**financial aid**" offered by colleges. Everyone knows college tuition is exorbitant, but in most cases, colleges also offer "need-based" financial aid packages to students. That is usually lauded, but it can also be seen as a stealthy form of price discrimination. Financial aid is a "discount" on the sticker price of college tuition, and since aid varies by student, then the university is effectively charging students different prices for the same college education.

Universities manage to do this because students normally provide an immense amount of personal financial information – information other retailers can only dream of. They use this information to guess the price the student is willing to pay, and then adjust the price – sorry, I mean, the "aid" - accordingly. By doing so, the university is nabbing the consumers' surplus for itself.

Now, before you howl in indignation, keep in mind that the university's defense is that it is doing so in the name of egalitarianism. In other words, that it is exactly by price discrimination that they can *equalize* consumers' surpluses across students.

To see what they're getting at, suppose there was a single, flat tuition price charged across all students, no financial aid "discounts". Then rich students, willing and able to pay a lot more, will be getting a heck of a bargain, while poor students are paying as much as they possibly can. In consumer surplus terms, rich students would be like Mr. Alpha above, reaping an enormous welfare gain, while poor students would be like Ms. Gamma, paying the maximum she's willing to pay, reaping no consumer surplus at all.

So, a single, flat tuition price with no discrimination may *sound* egalitarian. But when you look at the resulting consumer surplus, they are actually very different for different students. By discriminating with "aid", the universities argue, they are trying to equalize consumer surpluses across students: the rich student who is charged more will be getting about the same amount of consumer's surplus as the poor student who is charged less. (e.g. in Fig. 6, high, medium & low tuition-paying students all reap the same-sized consumer surplus, x).

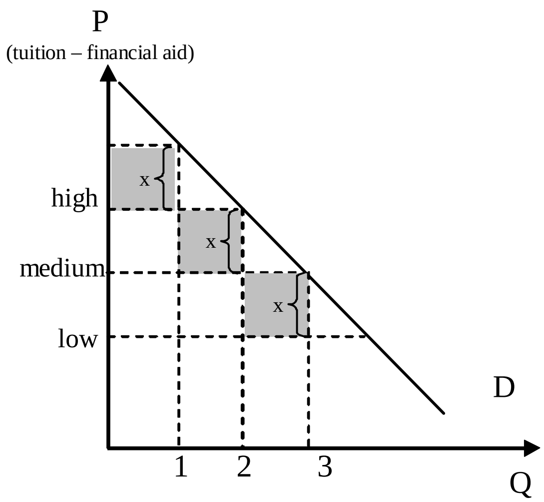

Fig. 6 – Equalizing Consumer Surplus: College financial aid?

Of course, if college tuition *with* aid is tailored in a perfectly extractive fashion (as in Figure 5) that it ends up leaving all students with *no* consumer's surplus at all, then the argument goes down the toilet. In that case, all have equal consumer surplus, yes – equally zero for all.

PRODUCERS' SURPLUS

The "producer's surplus" is a measure of the welfare gain of firms. It is the mirror image of the consumers question. Basically,

> "**Producer surplus**" is value of the difference between the minimum price a producer is *willing* to accept for their goods and the price they *actually* receive.

Intuitively, a store may be willing to sell a shirt for $6. But if the market price is $10, then they will receive $10 for it, and pocket the difference. This difference - the $4 over and above what they were willing to sell at – is the producers' surplus.

In diagrammatic terms, the **total producers' surplus** is the area of the triangle formed by the supply curve and the market price, e.g. the shaded triangle in the diagram below.

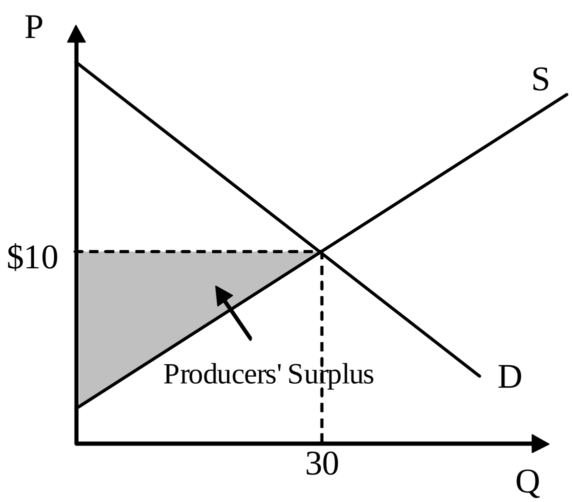

Fig. 7 - Producers' Surplus

The construction method of the producers' is analogous. It is basically the same as the consumers' surplus and demand curve, only with firms and the supply curve.

For instance, in Fig. 8 (imperfectly drawn), we see that if the market price was $2, then firms would only be willing to produce and sell one shirt. But if the market price was $3, then firms would be willing to produce & sell 2 shirts. The gain they make if the market price is $3 is that they were willing to sell the first shirt at $2 but now can sell it at $3. So firms are making a $1 welfare gain (producers' surplus) – in other words, they receive an extra $1 more for that shirt than they'd be willing to accept. That is a windfall for the firm.

Similarly as we proceed up the price scale. If market price was $4, they'd be willing to produce 3 shirts – making a $2 welfare gain on the first, and a $1 gain on the second. And so on.

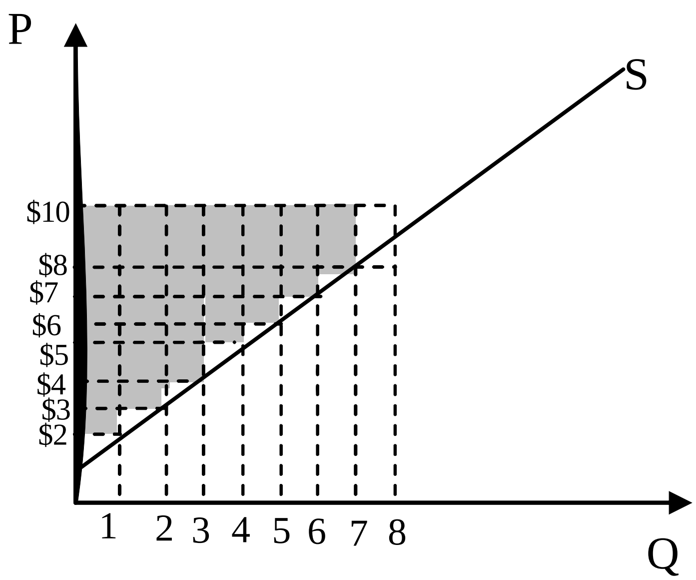

Fig. 8 - Producers' surplus construction

Going all the way up to the market price of $10, then firms will be producing and selling 8 shirts – seven of which they were willing to sell at a lower price, some at a much lower price – but they nonetheless receive the same $10 for each. The total windfall, the producers' surplus, is the shaded rectangles forming the triangular-like shape in Fig. 8.

MARKET WELFARE

Having got our concepts of consumer's and producer's surpluses down, it is natural to put them together in a single market diagram and represent them both. It would look something like this:

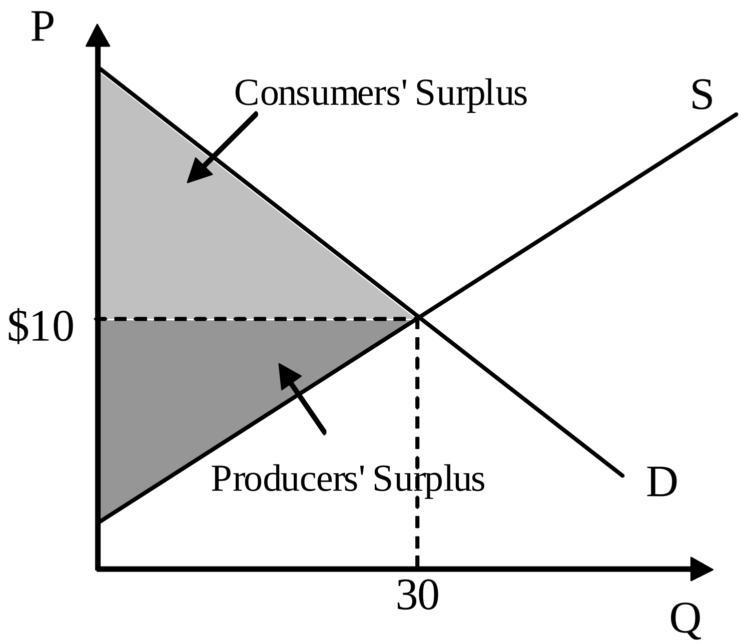

Fig. 9 – Market welfare

Very straightforward. And, as we shall see, very informative. There are a few welfare results we can analyze directly from this.

Aside: "Gains from Trade" Revisited

Before we continue, it is worth remarking on a useful way of interpreting the surpluses.

Remember in our discussion of comparative advantage, etc. we talked about the "sellers minimum" and "buyers maximum" and that the price would be somewhere in between, with both making gains from trade? Well, the concept of Consumers Surplus and Producers Surplus helps us measure that.

The way to see it is to remember that demand is the amount consumers are willing and able to pay. Consequently we can think of the demand curve as tracing the "buyers maximum" prices for a whole bunch different consumers. Alpha is willing to pay up to $20. He will accept any price below, but no price above that. Similarly, Beta willing to pay up to $15 - and any price below. Gamma is willing to pay up to $10 and so on.

So we can think of the demand curve as tracing the "buyers maximum prices", and the *entire area* below the demand curve then represents prices that are *acceptable* to those buyers. Areas *above* the demand curve are unacceptable to buyers.

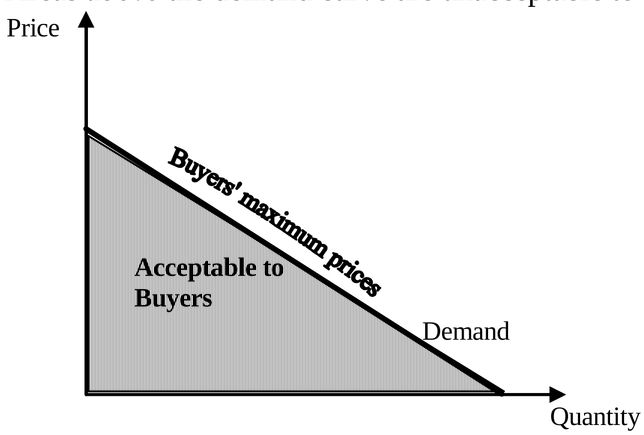

Similarly, the supply curve can be thought of as a locus of "sellers' minimum prices", and the entire area *above* the supply curve represents the prices that are acceptable to those sellers, and areas below the supply curve are the prices *unacceptable* to sellers.

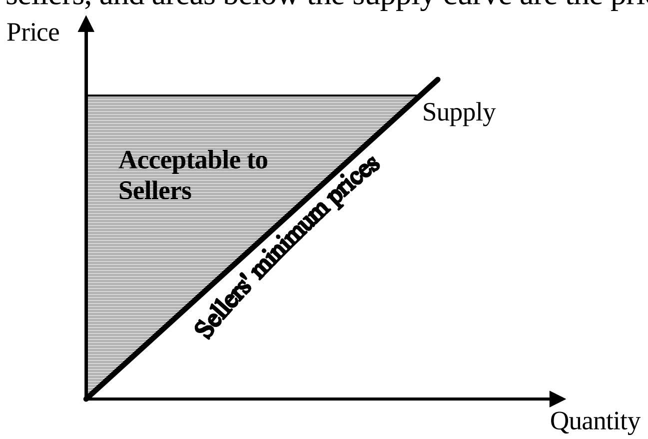

When we take demand and supply together, we have three areas. The area of interest is the middle triangle, where the acceptable areas to buyers and acceptable areas to sellers overlap. This represents the prices - or exchanges - that are acceptable to *both* buyers and sellers.

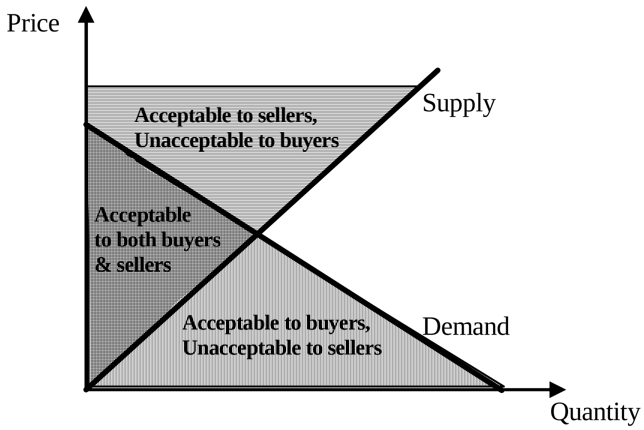

Now, just like in our comparative advantage examples, anything *within* that triangle is *possible*. But we now know that market forces will take us to the intersection point between the demand and supply curve. So rather just saying price is "somewhere" between buyers maximum and seller's minimum, we can pin it down exactly as the price and quantity where the curves intersect.

Once price is pinned down exactly, that triangle is automatically partitioned between Consumers' Surplus and Producers' Surplus. That is, the surpluses measures the "gains from trade" accruing to buyers and sellers respectively.

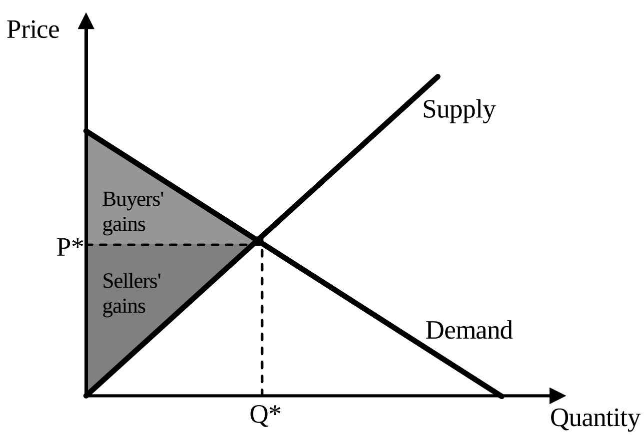

EXCISE TAXES

Since we are now playing with consumer & producer welfare, it is worthwhile to take a look at another topic of interest: **excise taxes** (i.e. sales taxes). A sales tax is added onto the price a consumer pays for a good. As a result, it necessarily creates a "wedge" between the price of a good the consumer pays and the price the producer receives. For instance, if there is a $2 tax on the sale of shirts, then if firms are willing to sell shirts at $10, then consumers *must* pay $12 for it (to cover the sales tax). The government takes the $2 difference.

However, simply tacking on $2 to the equilibrium price is not enough. The market may not clear at those prices, as we see below in Fig. 10. At the equilibrium price of $10, firms produce 30 shirts, but at $12 (the minimum price consumers can pay to cover the tax), demanders only want to buy 20 shirts. Demand does not equal supply. There is an excess supply of 10 shirts. The market is not cleared. We are not in equilibrium anymore.

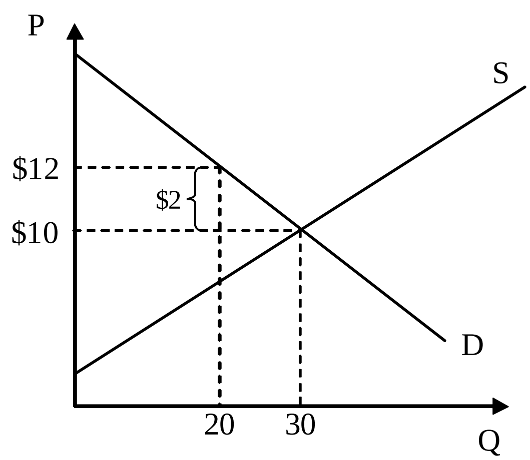

Fig. 10 – First attempt at sales tax - disequilibrium

You can guess what the "solution" to this is. Let the Law of Markets work, unleash the magic of the price mechanism and prices will adjust so that a *difference* is created between the sellers & buyers prices that exactly covers the tax *and* markets clear.. Visually, this amounts to fitting a tax-sized "wedge" between the curves. The result will look like this (Fig. 11):

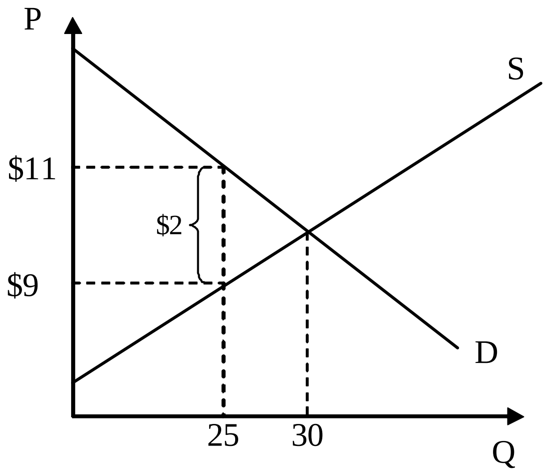

Fig. 11 – Impact of a sales tax - equilibrium

So, the market will settle where the price for suppliers is $9 and the price for demanders is $11. That exactly covers the $2 tax *and* the demand for shirts (which, at $11, is 25) is *exactly* equal to the supply of shirts (which, at $9, is 25). The markets clear.

What is the welfare impact of the sales tax? Obviously, we see that *less* shirts are produced and consumed, so obviously welfare suffers. But *whose* welfare exactly? And does the government compensate? To answer this, examine Fig.12.

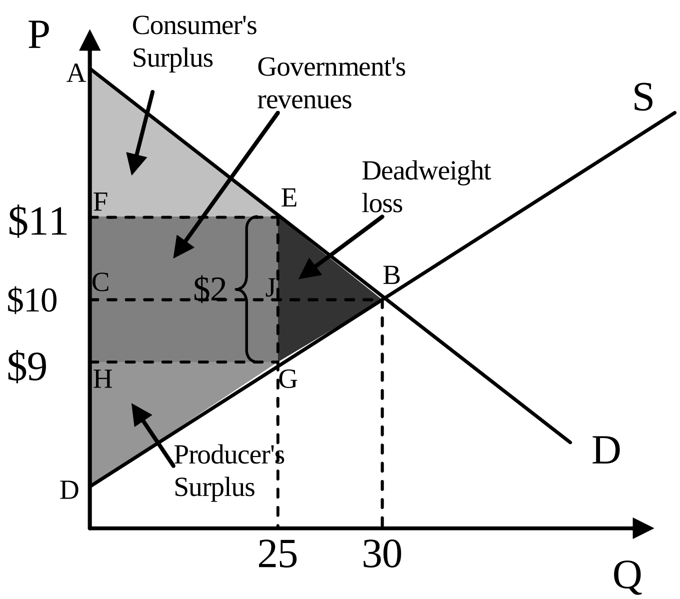

Fig. 12 – Impact of a sales tax – welfare loss

Before the sales tax, when price was $10 and quantity was 30, the consumers' surplus was the triangle ABC and the producers' surplus the triangle DBC.

But once the sales tax is instituted and the wedge fitted in, the consumers' surplus falls to the smaller triangle AEF, while the producers' surplus declines to the smaller triangle DGH. Obviously, *both* consumers and producers are *worse* off. There is no transfer of surplus from one to the other, but surplus is loss from both.

Sure, you may, but the government gets the lost surpluses, right? Wrong. What the government gets is the tax revenues, that is tax × amount sold = $2 × 25 shirts = $50. That is the area of the dark *rectangle* FEGH between the two surpluses. But the government's revenues (FEGH area) is *less* than the sum of the lost surpluses.

There is what we call a "**deadweight loss**" that goes to nobody – that is the dark triangle FEB. That represents the parts of lost consumer and producer surpluses that *don't* go to the government. It is a loss to society, it completely evaporates, a precise measure of the inefficiency caused by the excise tax.

Now, as you can see intuitively, the deadweight loss from sales tax is composed partly of consumer surplus lost and producers's surplus lost. Granted that both lose, but who loses more? This depends on the relative slope of the cures – or what is known as the **elasticity** of the demand and supply curves.

<u>ELASTICITY</u>

Elasticity is a rather simple concept, a measure of the degree of responsiveness, or sensitiveness, of the quantity demanded to a change in price.

> **Elastic:** We say demand for a good is very **elastic** if a change in price produces a *more* than proportional change in quantity demanded.
>
> **Inelastic:** We say demand for a good is very **inelastic** if a change in price produces a *less* than proportional change in quantity demanded.

So, if a small decline in the price of shirts creates a huge increase in the quantity of shirts demanded we say the demand for shirts is very "elastic". Or, what amounts to the same thing, a small increase in the price of shirts reduces the quantity demanded by a lot.

If instead, a large decline in the price of shirts generates only a very small increase in the quantity of shirts demanded, then we say the demand for shirts is very "inelastic". Or, equivalently, a large increase in the price of shirts reduces the quantity demanded by only a little.

Diagrammatically, a demand curve is **elastic** if it is relatively flat, and **inelastic** if it is relatively steep (this is not precisely correct, but close enough).

Goods can vary in elasticity. Goods that have very good substitutes for them are frequently very elastic. e.g. Coca-Cola is a very elastic good, since we usually find that a small rise in price of Coke will lead to a huge decline in Coke consumption (as everybody switches to Pepsi).

Goods which have poor substitutes, like gasoline, are relatively inelastic. If the price of gasoline rises, you may reduce demand for it somewhat, but it is not bound to be very much of decrease since you can't easily find substitutes for it (you can't switch to pouring soybean oil into your gas tank!)

Elasticity and Revenues

When the price of shirts rises, demand for shirts falls. That is the Law of Demand. But what happens to the revenues of the shirt industry? Fewer people are buying shirts, yes, but those that continue buying are paying a higher price. Is the shirt industry making *more* money or *less* money when it raises prices? This is what elasticity helps us answer.

By definition:

> **Total revenues** = price × quantity sold

Notice that the formula for total revenue is exactly the same as the formula for the area of the rectangle formed by the equilibrium price and the equilibrium quantity. Diagramatically, the shaded area below (p = 10, q = 30, so revenues = 300):

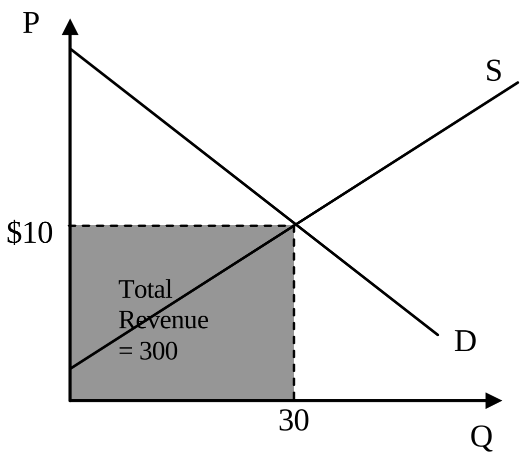

Fig. 13 – Total Revenues

So, let us go back to our original question. Suppose the price of shirts rises to $12. We know quantity sold is going to be less. But are total revenues going to rise or fall? To answer this, we need to know the elasticity of the good.

If a good is highly elastic, that means that a 5% increase in price leads to, say, a 10% drop in sales. As percentage, quantity falls more than prices rise. Thus, the total revenue number, price × quantity, will necessarily fall. (area of the rectangle falls).

If a good is rather inelastic, that means a 5% increase in price will lead to, say a 3% drop in sales. As a percentage, quantity falls less than prices rise. So higher prices more than compensate the drop the sales, so total revenues will necessarily rise. (area of the rectangle increases).

If the good is what is called "unitary elastic", that means a 5% increase in price will lead to a 5% drop in sales. So total revenues will be unchanged. The price rise perfectly matches the fall in sales.

[Conversely, lowering prices for elastic goods will *increase* total revenues (lower prices more than compensated by much higher sales), while lowering prices for inelastic goods will *decrease* total revenues (the relatively small increase in quantity sold won't make up for the decline in price].

Example: starting from p = $10 and Q = 30. Suppose prices rise by $2, and quantity sold falls by 4. Have revenues risen or fallen? Well, since $2 is 20% of $10, that means prices rose by 20%. But 4 is 13.3% of 30. So the rise in price led to a *less* than proportional change decline in quantity demanded. The good is inelastic. We should expect total revenues to fall.

Do they? Well, in the Fig. 14, notice that the there will be a new total revenue rectangle. (old rectangle = ABCD, new rectangle = EFGD). The area of ABCD = 300. Area of EFGD = $12 × 26 = 312. So the rise in price has lead to an *increase* in total revenues. As expected.

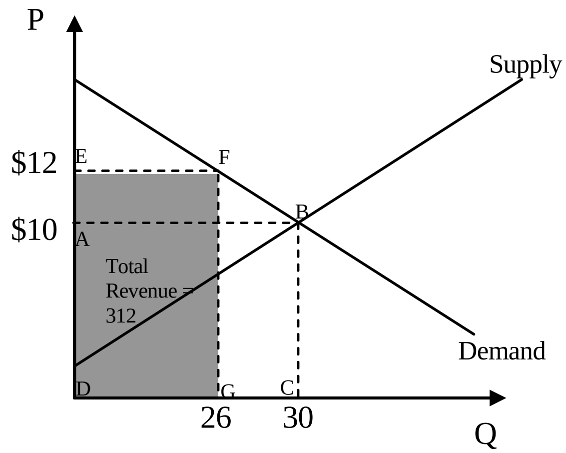

Fig. 14 – Inelastic

Suppose instead that quantity sold fell by 8 in response to a $2 increase. Since 8 is 26.6% of 30, that means that the rise in price has led to a *more* than proportional decrease in quantity demanded. The good is elastic and expect revenues to *fall*. To check, notice that the area of the new rectangle = $12 × 22 = 264. Total revenues have indeed declined.

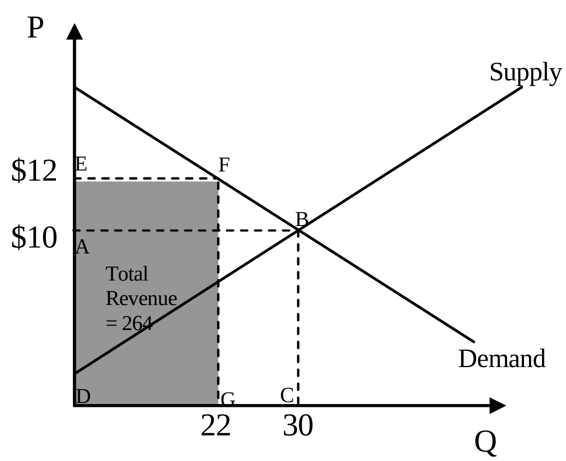

Fig. 15 – Elastic

Elasticity & the Impact of Taxation

Whether demand for a good is elastic or inelastic gives us a clue as to how an excise tax will differentially impact welfare.

For instance, take a highly elastic good like Coca-Cola. Here we have a relatively flat demand curve (Fig. 16). We see that if we impose a sales tax on it, the flatness of the demand curve's shape implies that the bulk of the burden will be taken by the producer – of the $2 tax, he will absorb $1.50, the consumer absorb $0.50. The burden won't be equally shared because the producer cannot risk passing on much of the tax to the consumer. It is an elastic good. They will flee *en masse* from the product at the slightest increase.

Notice also that the consumer's share of the deadweight loss is minimal. So, in markets where demand is elastic, the welfare loss falls mostly on the producer.

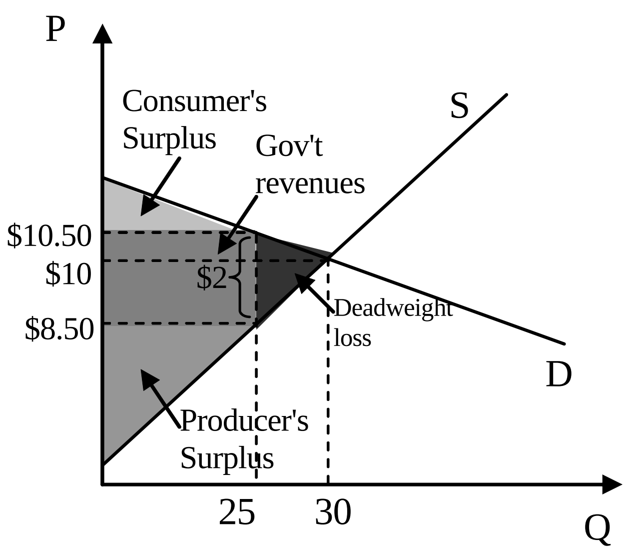

Fig. 16 – Impact of a sales tax – elastic case

In cases where the market demand is inelastic (i.e. steep demand curve, Fig. 18), the effect is the exact opposite. The burden of the sale tax will fall mostly on the consumer rather than the producer. In the sharing of the deadweight loss, the consumer takes the bulk of it.

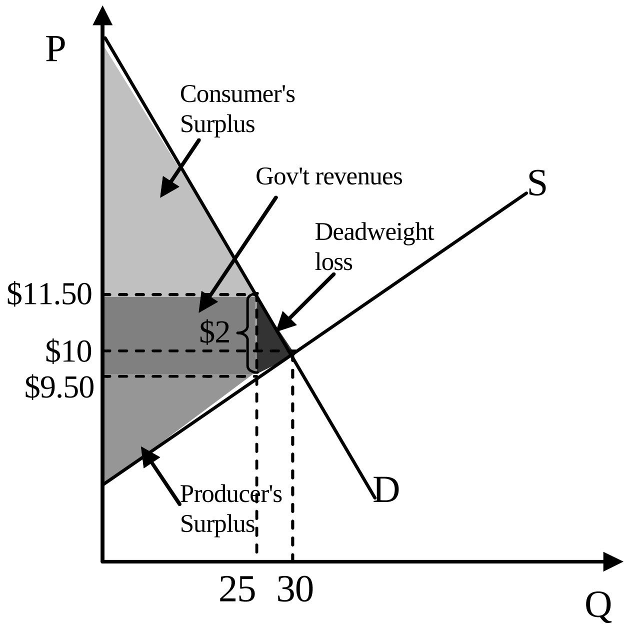

Fig. 17– Impact of a sales tax – inelastic case

A WORD ON MONOPOLY

Finally, a word must be said about the welfare impact of a monopoly – a single seller - although a full analysis of that phenomenon would take us quite beyond here (we'll deal with that in our discussion of production). But we can begin to get a clue of its impact already here.

Remember that we constructed our supply curve by *assuming* that firms took prices as "given" by the market and reacted accordingly. This reflects the *assumption* that firms are operating in a **competitive market**, i.e. a market with lots of competing firms. Competition means individual firms don't have the liberty to set prices at any level they want. No single firm is "big enough" to set prices much above their cost of production and try to make extraordinary profits. If it tried, their competitors could – and would - undercut them and steal all their customers.

The result is that, in a competitive environment, individual firms have no choice but to assume prices are beyond their control. Prices are "given" by the market and all a single firm can do is decide how much to produce *at that price*. We built the supply curve that way – at a *given* price, a firm reacts with a quantity decision.

But in a monopolistic situation, there are no competitors. As the only producer, you *can* fiddle with prices as well as quantities, without fear of competitors undercutting you. You are the only producer and the only seller. You are no longer compelled to respect the supply curve. In a monopoly situation, the supply curve merely shows the *minimum* price the monopolistic firm will consider – the one that just covers cost of production - but they are free to actually pick a different price – indeed, a much higher price - to charge their customers.

In monopolistic situations, firms are free to choose their **mark-up**, the gap between the price consumers face and the cost of production. What mark-up will they pick? Monopolistic firms, like all firms, are profit maximizers, so will try to choose the mark-up where profits are maximized. That doesn't mean it is the highest possible mark-up. They have to take demand into consideration. If they price things *too* high demand will collapse and they won't make any sales and thus no profits. So they will search around for the *optimal* mark-up where sales and costs are such that profits are at their highest.

The impact of a monopoly is pretty much like the impact of an excise tax. The optimal mark-up in a monopolistic situation will create a "wedge" between the high prices consumers face (demand curve) and the firm's own costs of production (supply curve). The result will look something like Fig. 18.

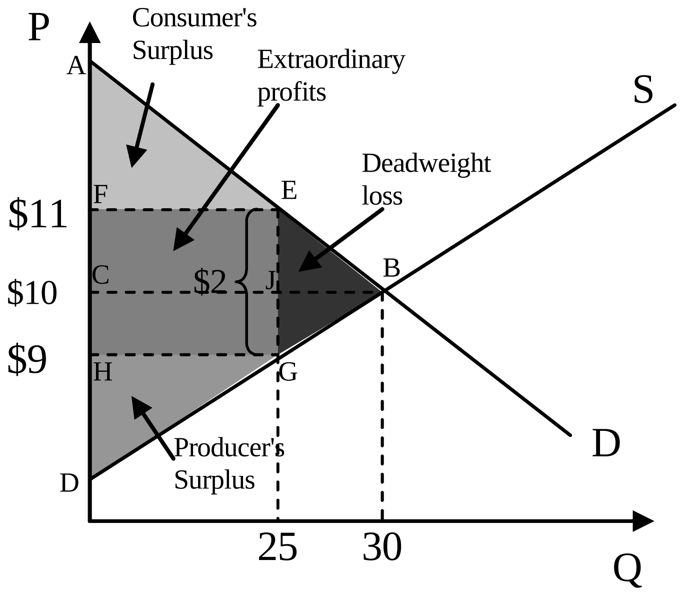

Fig. 18 – Impact of monopoly

Under competitive conditions, the p = $10 and quantity produced and sold = 30. But under monopoly conditions, assuming $2 is the profit-maximizing "mark-up" on costs, then the monopoly firm will cut back production to 25 (implying a cost of $9) and charge customers $11. In this situation, the monopoly firms is enjoying extraordinary profits that amount to $2 × $25 = $50 (the size of the dark rectangle FEGH).

The consumers' surplus declines from triangle ABC to AEF. The firm nabs only *part* of the consumers' loss for itself, but not all of it. Just like in the case of the excise tax, there is a deadweight loss corresponding to the darkly-shaded triangle EBG.

This is not an exhaustive treatment of the problems created by monopoly, but it illustrated one of its important aspects – its inefficiency (and there's nothing economists hate more than inefficiency). The deadweight loss that results from monopoly is *the* principal reason for the economic case against monopolies and the major reason behind strong anti-trust legislation.

(For the welfare impact of foreign trade and protectionist policies (tariffs, quotas, export subsidies), see our notes on "The Invisible Hand".)