Price Theory

Lecture 2: Supply & Demand

I. The Basic Notion of Supply & Demand

Supply-and-demand is a model for understanding the determination of the price of quantity of a good sold on the market. The explanation works by looking at two different groups – buyers and sellers – and asking how they interact.

II. Types of Competition

The supply-and-demand model relies on a high degree of competition, meaning that there are enough buyers and sellers in the market for bidding to take place. Buyers bid against each other and thereby raise the price, while sellers bid against each other and thereby lower the price. The equilibrium is a point at which all the bidding has been done; nobody has an incentive to offer higher prices or accept lower prices.

Perfect competition exists when there are so many buyers and sellers that no single buyer or seller can unilaterally affect the price on the market. Imperfect competition exists when a single buyer or seller has the power to influence the price on the market.

The supply-and-demand model applies most accurately when there is perfect competition. This is an abstraction, because no market is actually perfectly competitive, but the supply-and-demand framework still provides a good approximation for what is happening much of the time.

III. The Concept of Demand

Used in the vernacular to mean almost any kind of wish or desire or need. But to an economist, demand refers to both willingness and ability to pay.

Quantity demanded (Q_d) is the total amount of a good that buyers would choose to purchase under given conditions. The given conditions include:

- price of the good
- income and wealth
- prices of substitutes and complements
- population
- preferences (tastes)
- expectations of future prices

We refer to all of these things except the price of the good as determinants of demand. We could talk about the relationship between quantity demanded and any one of these things. But when we talk about a demand curve, we are focusing on the relationship between quantity demanded and price (while holding all the others fixed).

The Law of Demand states that when the price of a good rises, and everything else remains the same, the quantity of the good demanded will fall. In short,

$$\uparrow P \rightarrow \downarrow Q_d$$

Note 1: "everything else remains the same" is known as the "ceteris paribus" or "other things equal" assumption. In this context, it means that income, wealth, prices of other goods, population, and preferences all remain fixed.

> Of course, in the real world other things are rarely equal. Lots of things tend to change at once. But that's not a fault of the model; it's a virtue. The whole point is to try to discover the effects of something without being confused or distracted by other things.

Note 2: Is the law of demand really a "law"? Well, there may be some exceedingly rare exceptions. But by and large the law seems to hold.

Note 3: I will use the word "normal" to refer to any good for which the law of demand holds. Please note that this is different from the book's definition of normal.

A Demand Curve is a graphical representation of the relationship between price and quantity demanded (ceteris paribus). It is a curve or line, each point of which is a price-Q_d pair. That point shows the amount of the good buyers would choose to buy at that price.

Changes in demand or shifts in demand occur when one of the determinants of demand other than price changes. In other words, shifts occur "when the ceteris are not paribus."

The demand curve's current position depend on those other things being equal, so when they change, so does the demand curve's position.

> Examples:
> 1. The price of a substitute good drops. This implies a leftward shift.
> 2. The price of a complement good drops. This implies a rightward shift.
> 3. Incomes increase. This implies a rightward shift (for most goods).
> 4. Preferences change. This could cause a shift in either direction, depending on how preferences change.

Demand versus Quantity Demanded. Remember that quantity demanded is a specific amount associated with a specific price. Demand, on the other hand, is a relationship between price and quantity demanded, involving quantities demanded for a range of prices. "Change in quantity demanded" means a movement along the demand curve. "Change in demand" refers to a shift of the demand curve, caused by something other than a change in price.

IV. The Concept of Supply

Used in the vernacular to mean a fixed amount, such as the total amount of petroleum in the world. Again, economists think of it differently. Supply is not just the amount of something there, but the willingness and ability of potential sellers to produce and sell it.

Quantity supplied (Q_s) is the total amount of a good that sellers would choose to produce and sell under given conditions. The given conditions include:

- price of the good
- prices of factors of production (labor, capital)
- prices of alternative products the firm could produce
- technology
- productive capacity
- expectations of future prices

We refer to all of these, with the exception of the price of the good, as determinants of supply.

When we talk about Supply, we're talking about the relationship between quantity supplied and the price of the good, while holding everything else constant.

The Law of Supply states that "when the price of a good rises, and everything else remains the same, the quantity of the good supplied will also rise." In short,

$$\uparrow P \rightarrow \uparrow Q_s$$

A Supply Curve is a graphical representation of the relationship between price and quantity supplied (ceteris paribus). It is a curve or line, each point of which is a price-Q_s pair. That point shows the amount of the good sellers would choose to sell at that price.

Changes in supply or shifts in supply occur when one of the determinants of supply other than price changes.

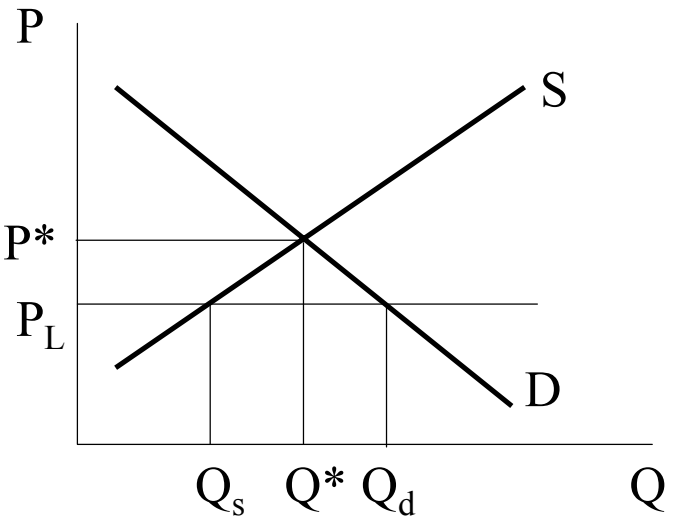
Examples:

1. The price of a factor of production rises. This would cause a leftward shift the supply curve.
2. A rise in the price of an alternative good that could be provided with the same resources. This implies a leftward shift of supply.
3. An improvement in technology. This leads to a rightward shift of supply.

Supply versus Quantity Supplied. Analogous to the demand versus quantity demanded distinction. "Change in quantity supplied" means a movement along the supply curve. "Change in supply" refers to a shift of the supply curve, caused by something other than a change in price.

V. Constructing the Market

Putting demand and supply together, we can find an equilibrium where the supply and demand curve cross. The equilibrium consists of an equilibrium price P* and an equilibrium quantity Q*. The equilibrium must satisfy the market-clearing condition, which is $Q_d = Q_s$.



Mathematical example: Suppose $P = 20 - .1Q_d$ and $P = 5 + .05Q_s$. In equilibrium, $Q_d = Q_s$, so we have a system of equations. Solve for Q like so:

$20 - .1Q = 5 + .05Q$
$15 = .15Q$
$Q^* = 100.$

Then plug Q* into either equation: $P = 20 - .1(100) = 10.$
So the market equilibrium is $P^* = 10$, $Q^* = 100$.

If price is below P*, at P_L, then we have $Q_d > Q_s$. This is called "excess demand" or "shortage." The quantity that actually occurs will be Q_s. For this quantity, buyers are willing to pay much more than P_L, so they'll start bidding against each and raising the price.

If price is below P*, at P_H, then we have $Q_s > Q_d$. This is called "excess supply" or "surplus." The suppliers will start competing against each other for customers by lowering the price.

Short-side rule: When there is a disequilibrium price, the actually quantity that gets sold is given by $Q = \min\{Q_s, Q_d\}$. This is implied by the requirement of voluntarism.

VI. Price Controls

Price floors and price ceilings are government mandated prices that attempt to control the price of a good or service.

A price ceiling is usually imposed to keep down the price of something perceived as too expensive. To have any effect, it must be imposed below the market price.

Example: Rent control on apartments.
What effect do we predict? As with any below-equilibrium price in the example above, we expect to get a shortage. But in this case, buyers can't raise the price by bidding against each other, because by law the price cannot rise.

Using the short-side rule, we discover that rent control actually reduces the amount of housing made available to the public.

Ironically, the price control may also raise the *de facto* price paid by consumers. From the demand curve, we can see that consumers would be willing to pay a very high price (much higher than the price ceiling or even the market price) for the reduced quantity (Q_s) available. They are willing to pay this money if they can just find a way to do so – and they do, in the form of bribes, key fees, rental agency fees, etc.

N.B.: If the price ceiling is imposed above the market price, it has no effect.

A price floor is usually imposed to keep up the price of something perceived as too cheap. To have any effect, it must be set above the market price.

Example: Agricultural price supports.
These are imposed, usually, because farm lobbies have convinced the legislature that they are not earning enough to stay in business.

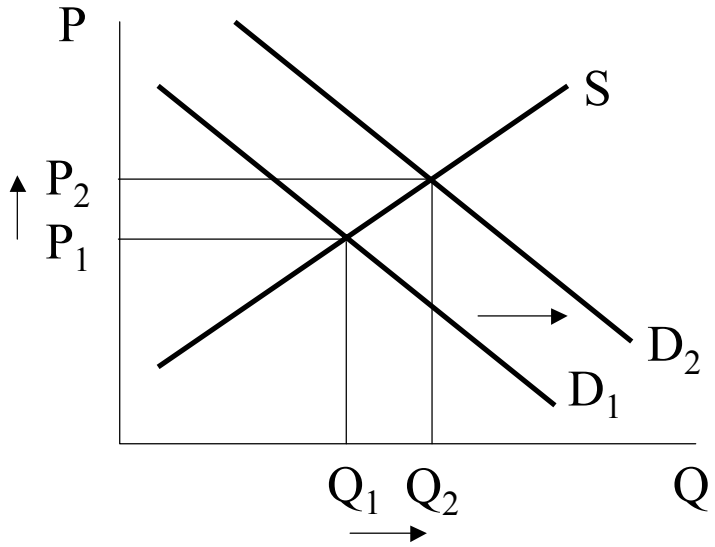
What effect do we predict? As with any above-equilibrium price, we expect to get a surplus, this time persistent because sellers can't bid down the price. And for many years, that's exactly what the U.S. had. The government usually bought up the surplus (and dumped it on 3rd World markets).

Note: If the price floor is imposed below the market price, it has no effect.

Note: It's easy to get confused if you're not thinking clearly. An effective price ceiling is below the market price, while an effective price floor is above it. (Imagine a ceiling being too low and bumping your head, or a floor rising beneath your feet.)

VII. Analyzing Changes in Market Equilibrium

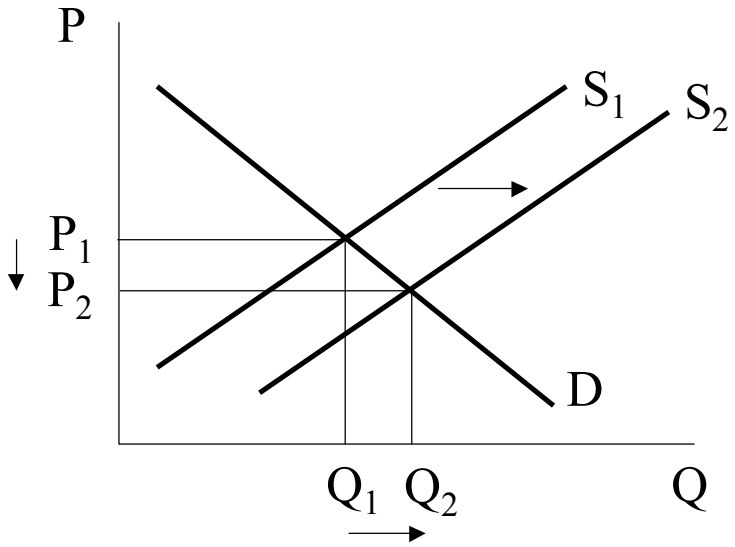
Consider first a rightward shift in Demand. This could be caused by many things: an increase in income, higher price of substitute, lower price of complements, etc. Such a shift will tend to have two effects: raising equilibrium price, and raising equilibrium quantity. [$\uparrow P^*$, $\uparrow Q^*$]



Example: Consider the market for rental housing, and suppose that a new factory or industry opens up in the city, attracting more residents. Then there will be a rightward shift in demand, driving up both price and quantity. Note: The price of housing does go up, but not by as much as you might think, because the change in demand induces suppliers to bring more housing to market. This can be seen in the movement <u>along</u> the Supply curve.

A leftward shift of demand would reverse the effects: a fall in both price and quantity. The general result is that Demand shifts cause price and quantity to move in the <u>same</u> direction.

Now consider a rightward shift of supply (caused by lower factor price, better technology, or whatever). This will tend to have two effects: raising equilibrium quantity, and lowering equilibrium price. [↓P*, ↑Q*.]



Example: A new immigration policy allows lots of low-wage labor to enter the steel business. The lower price of steel leads to a rightshift in the supply of cars, so the price of cars falls and the quantity rises.

A leftshift of supply would reverse the effects, so the general result is that supply shifts tend to cause price and quantity to move in opposite directions.

Now, what happens if both demand and supply both shift at once? In general, the two changes have reinforcing effects on either price or quantity, and offsetting effects on the other.

Example 1: The computer industry. Incredible improvements in technology, as well as the entry of many new firms into the industry, have increased supply. Simultaneously, many people have become very aware of the benefits of computers, and new software has made computers more useful for a variety of projects, thereby increasing demand as well.

The increase in demand tends to increase both P and Q.
The increase in supply tends to lower P and raise Q.

So both effects tend to raise quantity. But what happens to price? That depends on the relative size of the two changes. Observation of the computer industry shows that prices have actually fallen, so we can conclude that supply shifts have been relatively more important than demand shifts in this market. A similar pattern emerges in many high-tech products (CD players, DVD players).

A possible complication: Much technological change has been in the form of higher quality, and consumers shift demand from lower quality to higher quality machines. This is not as easy to think about in the supply-and-demand framework. The price of a new, cutting edge computer has stayed about the same over time, at about $2000.

Example 2: Higher education. (This example may not be totally accurate historically, but it demonstrates the point.) Supply has fallen because higher education is a labor-intensive business, and educated labor (which has to be attracted from other industries) has become much more expensive. Meanwhile, demand has increased because jobs for educated people have become increasingly attractive relative to other jobs.

The increase in demand tends to increase both P and Q.
The decrease in supply tends to raise P while lowering Q.

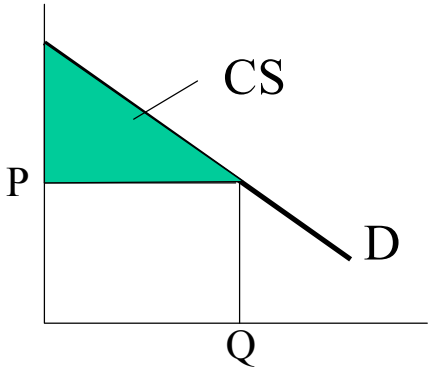
So both forces tend to raise P, and that is confirmed by observation. But they work in opposite directions on Q. Our knowledge of the market for higher education tells us that Q has actually increased, meaning that the shift in demand has been relatively more important.

When analyzing situations where both supply and demand shift at once, don't let yourself be fooled by your graph. Your graph may appear to show clear changes in both price and quantity. But we know that one variable will experience an offsetting effect. Whether that variable appears to rise or fall depends entirely on how large you've drawn the curve shifts. (Take the computer example above. If you drew the demand shift bigger than the supply shift, you would mistakenly conclude that price should rise.) Unless you are provided with additional information about the size of the shifts, you can only make a prediction about one variable.

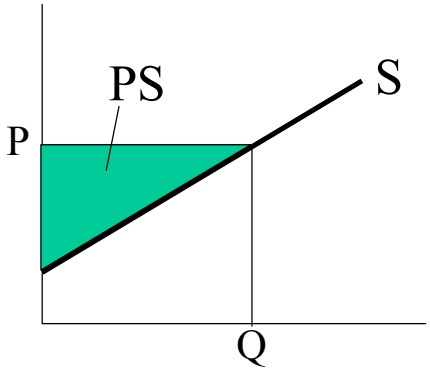
VIII. Welfare Analysis

Aside from our positive observations about the effects of price controls (and other policies), we can also do welfare analysis, which is a means of showing who gains and who loses from these policies. This is where we make use of the efficiency concepts we learned earlier in the course.

Definition of Consumer Surplus: the extra (or excess) value individuals receive from consuming a good over what they pay for it. Or, the dollar-valued benefits to buyers from all trade in a market. The CS is above the price paid, below the demand curve, and to the left of the quantity purchased.

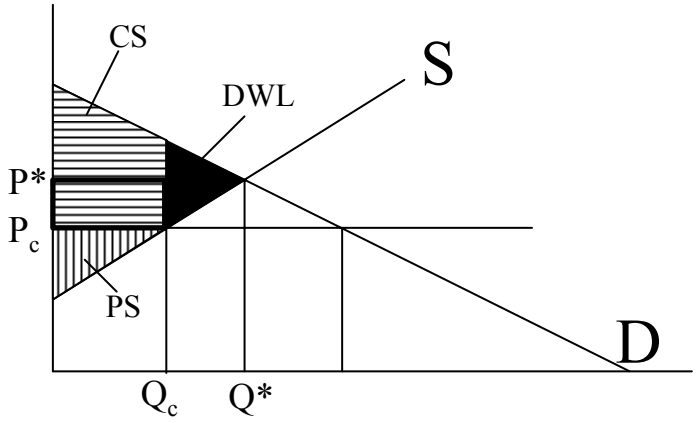


Definition of Producer Surplus: the extra value producers get for a good in excess of the opportunity costs they incur by producing it. Or, the dollar valued benefits to sellers from all trade in a market. The PS is below the price paid, above the supply curve, and to the left of the quantity sold.



Definition of Total Surplus: CS + PS. Total surplus represents the total value of all gains from trade to all parties.

Welfare analysis for a price ceiling. The quantity is reduced from Q* to Q_c. The CS extends further down, but not as far right, compared to the free market. The PS does not extend as far up, and also not as far right, compared to the free market. The transfer is the portion of CS that used to be in PS, between the equilibrium price and the price ceiling. Finally, the Dead Weight Loss (DWL) lies above supply, below demand, and between Q* to Q_c.



Dead-weight loss (DWL) is the reduction in the total surplus that results from a policy that prevents mutually beneficial trades from occurring. It is the area below demand, above supply, and between the market equilibrium quantity and the actual quantity that gets sold.

The transfer is the portion of the total surplus that moves from CS to PS, or from PS to CS, when a policy such as a price control displaces the market outcome. The transfer occurs because some consumers (those who are able to buy the reduced quantity available) are able to get a lower price; they are better off, and their gain is a loss for the producers.

In efficiency terms, the move from free market to the price-controlled market was not a Pareto improvement. The situations are Pareto-incomparable because some people (a subset of consumers) gain, while others (producers and the rest of consumers) lose. But the free market is Kaldor-Hicks superior to the price-controlled market, because the total wealth is larger. Why? Because total wealth is smaller by the amount of the DWL under the price ceiling. Note: The transfer is totally irrelevant from a (K-H) efficiency perspective, because it remains in the total. The DWL is what makes the price control inefficient.

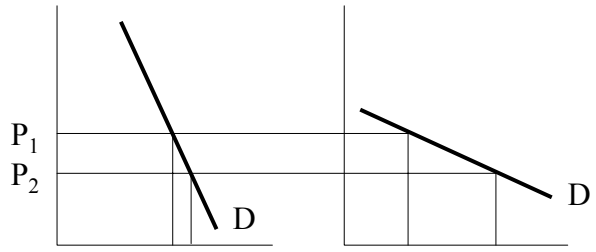
One thing this analysis leaves out is the effect of waiting, bribes, rental agency fees, and other means of rationing the reduced quantity of the good. These factors can reduce the size of CS by transferring some wealth back to the PS.

Welfare analysis for a price floor. The analysis here is similar. But this time, there is a transfer from the CS to the PS.

Consumers as a group lose from this policy. Some of the producers gain, if they are the producers who get the privilege of selling the reduced quantity. Other producers lose, because they don't get that privilege. Overall, since the price floor produces both winners and losers, the price floor regime and the free market are Pareto-incomparable. But the free market is Kaldor-Hicks superior, because it doesn't incur a DWL.

IX. Elasticity

Elasticity refers to the degree of responsiveness of one variable to another. It's not enough to say, for instance, that a rise in price leads to a fall in quantity demanded (the Law of Demand); we want to know how much quantity changes in response to price.



A simple way to see the degree of responsiveness is simply to look at the slope. A flatter demand curve represents a greater degree of responsiveness (for a supply or demand curve), as shown in the above graphs: the flatter demand curve produces a larger change in quantity for the same change in price.

Using just the slope is the quick-and-easy way to think about elasticity. The extremes are easy to remember: A perfectly elastic demand curve is horizontal, because an infinitely small change in price corresponds to an infinitely large change in quantity; the graph looks like the letter E for elastic. A perfectly inelastic demand curve is vertical, because quantity will never change regardless of the change in price; the graph looks like the letter I (for inelastic).

But using the slope can be misleading, because it doesn't tell us the significance of the quantities. Suppose a $1 dollar increase in price leads to almost everyone choosing not to buy the good. That would not surprise us for gumballs, but it would certainly surprise us for televisions. The point is that a $1 increase is not much relative to the total price of TVs, but it is huge relative to the total price of gumballs. This is why we use elasticity instead of just the slope.

Definition of price elasticity of demand: the percentage change in quantity demanded divided by the percentage change in price. That is,

$$E_d = |\%\Delta Q_d / \%\Delta P|$$

How do you find the percentage change in something? You find out how much it changed, and divide by the initial value. For example, suppose your income rises from $400 a week to $500 a week. The change is $100, so the percentage change is $100/$400 = .25 or up 25%. N.B.: The percentage change depends on the direction you're going. If your income went from $500 to $400, the percentage change would be - $100/$500 = -.2 or down 20%. [Note: For this reason, some textbooks, including ours, use a slightly different formula. When calculating the percentage change in a variable, instead of dividing by the original point, they divide by the average of the two endpoints. In this class, we will use the original point as described above.]

Thus, we can also write

$$Ed = |(\Delta Q_d / Q_d)/(\Delta P/P)|$$

Now, we can rearrange this like so:

$$E_d = |(\Delta Q_d / \Delta P) \times (P/Q_d)|$$

Look at the first term, change in Q over change in P. This is basically the slope of the curve. I say "basically" because when we talk about the slope of a line, we usually measure the rise (vertical distance) over the run (horizontal distance). In a supply and demand graph, we usually measure price vertically. So actually, the slope is change in P over change in Q. What we have here is the inverse of the slope. If we use *m* to stand for the slope, we have:

$$E_d = |(1/m) \times (P/Q_d)|$$

This is the easiest formula to use when you have a straight line for a demand curve.

> Example: You have the demand curve $P = 50 - .1Q_d$. The slope is –.1. Using our formula, the elasticity at the point (100, 40) is given by $E_d = |(-1/.1)(40/100)| = 4$. The elasticity at the point (200, 30) is given by $E_d = |(-1/.1)(30/200)| = 1.5$. Notice that the elasticity is not the same at every point on a line.

Note: In your previous classes, you may have learned "arc elasticity." With arc elasticity, you are finding the elasticity over a section of the demand curve. That's what we started with, but now I've just shown you something called "point elasticity." Point elasticity tells you the elasticity at a single point – specifically, at the (P, Q_d) point you plug in. Think of point elasticity as the elasticity for an interval (change in price) that is very, very small.

We have the following definitions:

When $E_d > 1$, we say the curve is elastic at that point.
When $E_d < 1$, we say the curve is inelastic at that point.
When $E_d = 1$, we say the curve is unit elastic at that point.

[Note: Textbooks differ on whether to take the absolute value or not. If you don't take the absolute value, you'll get a negative elasticity, which means that the demand curve is downward sloping. But to keep things positive, we'll always take the absolute value.]

In general, any demand curve will have one point that is unit elastic, which means that a one percent change in price corresponds to a one percent change in quantity.

Example: continued from above. We can find the unit elastic point by setting

$$E_d = 1$$
$$|(-1/.1)(50 - .1Q)/Q| = 1$$
$$(50 - .1Q)/Q = .1$$
$$50 - .1Q = .1Q$$
$$50 = .2Q$$
$$Q = 250$$

So the point Q = 250, P = 25 is the unit elastic point.

The price elasticity of supply is almost identical to the price elasticity of demand, except using a different curve. You can find it using the same formulas, except substituting quantity supplied for quantity demanded. Since the supply curve is upward sloping, the sign will be positive instead of negative.

X. Tax Incidence

So why is elasticity important? Many reasons, but here is one: it determines the distribution of the burden of taxes. Tax incidence is the study of how the burden of a tax is distributed over different groups. Consider the following statements:

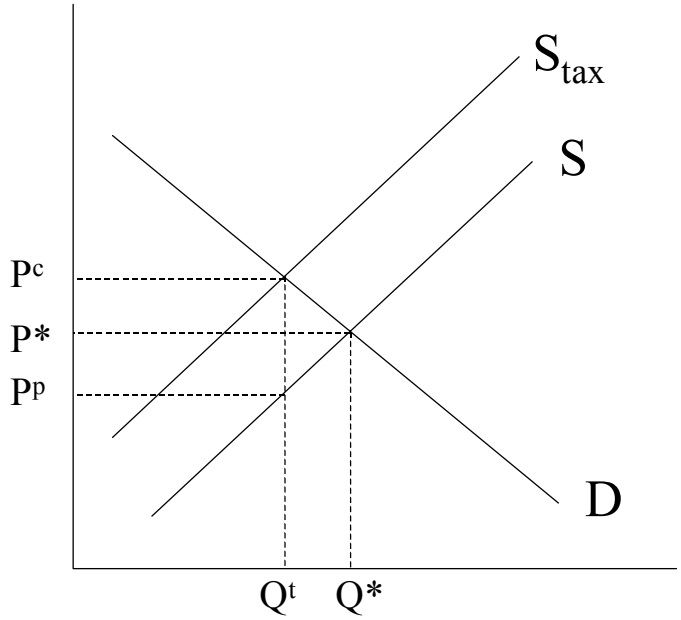
"If we raise the tax on cigarettes, tobacco companies will just pass the tax on to consumers." This statement implies that consumers bear the entire burden of a sales tax, even if the government requires firms to pay the tax.

"The Social Security tax is divided between the employer and the employee. The employer must pay half of the tax, and the employee must pay the other half." This statement implies that the government can decide how the burden of the tax will be distributed.

So which point of view is correct? As a general rule, neither. Consumers do not bear the entire burden of a tax in most cases, but neither can the government decide who pays how much. The distribution of the burden depends on the elasticity of supply and demand.

The key to understanding tax incidence is to realize that a sales tax (in fact, almost any kind of tax) is not a tax on a person -- it's a tax on a transaction. If there is a $1 tax on cigarettes, what that means is that there has to be a $1 difference between what the buyer pays and what the seller gets. It doesn't really matter who sends the check to the government.

Let's suppose, for now, that a sales tax is nominally imposed on the sellers of cigarettes. Suppose that you were previously willing to sell a pack of cigarettes for any price higher than $2.50. But with a $1 tax added to your costs, you're not willing to sell for anything less than $3.50. It's like your cost of production has gone up by $1 per pack. Overall, the effect is to raise the supply curve by exactly $1 at every quantity. Consider the following graph to see the result:



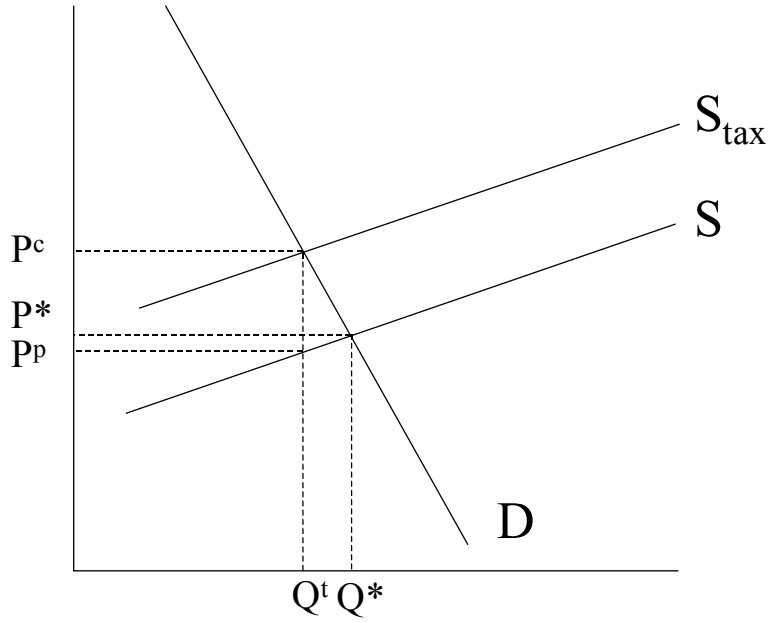
The vertical distance between S and S_{tax} is the amount of the tax ($1 in our example). The result of the tax is to reduce quantity from Q* to Q^t, and to raise the price from P* to P^c. But do the producers receive that price? No, they must receive a price that is $1 lower, at P^p. So how is the burden distributed? It should be apparent that consumers are paying more than they were before, and producers are receiving less than they were before. The burden is distributed between the two of groups.

In addition, there is also a dead-weight loss (DWL). It is the triangle between the old S and D, and between Q^t and Q*. This DWL corresponds to units that should have been sold, because they would have been mutually beneficial exchanges. But the gains from trade from selling these units, while positive, is not large enough to cover the tax, so the trades don't take place.

It might appear from the graph above that the tax is distributed evenly between consumers and producers, but that need not be true. It depends on the elasticity of supply and demand. Suppose that demand is very inelastic (consumers are unresponsive to price

changes), and supply is very elastic (producers are very responsive to price changes). Then we get a picture like the one below. Here, it should be apparent that the consumers are bearing the bulk of the tax burden, while the producers' burden is very small.

Another way of thinking about this is in terms of CS and PS. When supply and demand have about the same elasticity, the tax reduces CS and PS by about the same amount. But when demand is inelastic and supply is elastic, the CS is reduced by more and the PS is reduced by less.



On the other hand, what if the supply were very inelastic and the demand were very elastic? In that case, the producers would bear most of the burden. The general result is that when demand is more elastic than supply, producers bear the larger burden, and when supply is more elastic than demand, consumers bear the larger burden.

Consider again our examples. In the case of cigarettes, do you think the demand is relatively elastic or relatively inelastic? Given the addictive quality of cigarettes, it seems like demand is probably inelastic. If that's true, then a sales tax on cigarettes is likely to be borne mainly by the consumers. In the case of the Social Security tax, do you think the supply of labor is relatively elastic or relatively inelastic? It's probably fairly inelastic (people need to have their jobs, and almost all legal jobs are taxed), so the suppliers (i.e., the employees) probably bear most of the burden.