Hidden Order

The economics of everyday life

David D. Friedman

Second Edition
Audio

Graphics
Equations
Links

This document contains figures, equations, and links from the print copy of *Hidden Order* along with associated text, since figures and links cannot be included in the audiobook and equations are hard to follow in a spoken format.

Contents

1: Rush Hour Blues and Rational Babies

For Further Reading

A more extensive discussion of the economics of warfare can be found in my essay, "The Economics of War," in J.E. Pournelle (ed.), *Blood and Iron* (New York: Tom Doherty Associates, 1984).

3: Thinking on Paper — The Geometry of Choice

A Geometric Interlude

The simplest version of the grocery store problem is one in which each store sells only two goods and the consumer has a fixed amount to spend. Two goods are sufficient to explain the paradox and few enough to let me diagram the problem on the two-dimensional screen you are reading this on.

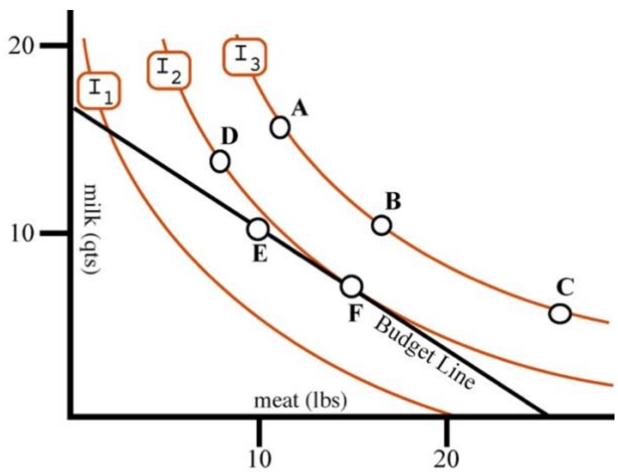

Figure 3-1a

Mrs. Smith in Kroger

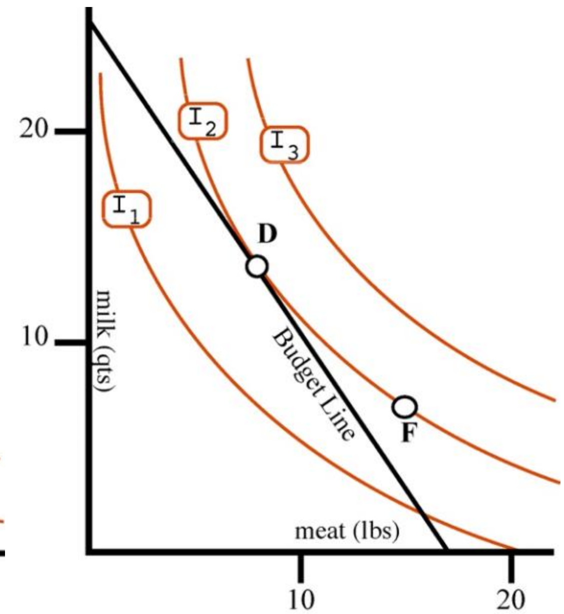

Figure 3-1b

Mrs. Smith in A&P

The logic of rational choice is simple: Out of all the available alternatives, choose the one you prefer. So our analysis of choice requires a way of representing available alternatives and a way of representing preferences. Figure 3-1a shows both.

Mrs. Smith enters Kroger with twenty-five dollars in her pocket. Milk costs $1.50 a quart at Kroger; meat (on sale) is $1/lb. Her *budget line* shows the alternative combinations of meat and milk ("bundles") she could buy with her money. Bundle **E**, for example, contains ten pounds and ten quarts, adding up to $25. Bundle **G** contains twenty-five pounds of meat and no milk at all, also adding up to $25. If Mrs. Smith decides to buy a quart less of milk she can use the money to buy a pound and a half of meat, so **B** is a straight line with a slope of -2/3.

We show Mrs. Smith's alternatives with a budget line, her preferences with a set of indifference curves. An indifference curve such as I_3 on the figure shows bundles all of which Mrs. Smith considers equally desirable. Bundle A on indifference curve I_3 is ten pounds of meat and fifteen quarts of milk. Bundle **B**, also on I_3, is fifteen pounds and ten quarts. Mrs. Smith is indifferent between them, does not care which she has.

If one bundle has less meat than another yet is equally attractive to Mrs. Smith, it must have more milk. The argument applies to any two bundles that are on the same indifference curve, so indifference curves slope down and to the right.

The more you have of a good the less you value having a little more (*the principle of declining marginal value*). As you move down and right along I_3 to bundles with less milk and more meat, additional milk becomes more valuable and meat less. Going from A to B, Mrs. Smith gives up five quarts of milk in exchange for an extra five pounds of meat. From B to C, the amount of milk drops by another five quarts and it takes an extra ten pounds of meat to make up for the loss. That is why the indifference curves all have the same general shape — with the curve getting less steep as you move right and down.

I do not actually know Mrs. Smith, nor her tastes for milk and meat. The purpose of indifference curves is not to present real information about the tastes of a real person but to help us think clearly. The arguments we construct using budget lines and indifference curves to think through the logic of rational choice will depend only on the general characteristics of indifference curves, not on the precise shape of the curves describing the tastes of a real person.

Every possible bundle is on some indifference curve: the curve showing all bundles equivalent to that one. If I drew all of those curves, the figure would be solid black. Curves I_1, I_2 and I_3 are the three I have drawn out of an infinite number I could draw.

If Mrs. Smith shifts from Point A on I_3 to point D on I_2, she gives up both milk and meat; since both are goods, she prefers A. As you move down and left, you move to less and less desirable indifference curves. The complete set of indifference curves would provide a complete description of Mrs. Smith's preferences with regard to milk and meat, since it tells us, of any two bundles, which she prefers — the one on the higher indifference curve.

Since these are the only goods available, Mrs. Smith might as well spend all of her money; there is nothing else to buy (and never will be; in our simplified world she only goes shopping once). Her choice is simple: Out of all the bundles on her budget line, pick the one she likes best. The solution is bundle F.

How do we know that F is the preferred bundle? F is on I_2, which is the highest indifference curve that touches the budget line. Mrs. Smith would prefer a bundle on I_3, but she does not have enough money to buy one. There are lots of bundles on I_1 that she could afford to buy but she prefers F.

We now know how to describe what happens when Mrs. Smith goes into Kroger graphically, but it is only when Mrs. Smith moves on to the A&P that our drawing begins to tell us things we did not already know.

Mrs. Smith is still Mrs. Smith, so the indifference curves representing her tastes are unchanged. At A&P, however, milk is on sale and meat is not; the prices are $1.50/lb for meat and $1/quart for milk. With different prices, Mrs. Smith must now choose among a different set of alternatives; her budget line on Figure 3-1b no longer runs through the point F. At the A&P's prices, Mrs. Smith cannot afford the quantities of meat and milk she bought at Kroger. Kroger's ad told the truth.

Does it follow that Kroger is really a cheaper store and that Mrs. Smith is better off doing her shopping there? No. She cannot duplicate what she bought at Kroger for the same amount at A&P. But, if she were in the A&P, she would not want to.

Point D on Figure 3-1b is what Mrs. Smith would choose to buy at A&P with her twenty-five dollars. Like point F on Figure 3-1a, it is, out of all the bundles she can afford, the one on the highest indifference curve. Faced with a different pattern of prices, Mrs. Smith chooses a different bundle of goods. Meat was cheap and milk expensive at Kroger, so she bought lots of meat and little milk; at A&P the pattern is reversed.

As it happens, D and F are on the same indifference curve: I_2. The two bundles are equally attractive to Mrs. Smith. She is equally well off whichever store she shops at.

The same pair of figures can be used for A&P's customer, Mrs. Jones, if we assume that her tastes happen to be the same as Mrs. Smith's. Mrs. Jones goes into the A&P with her $25 and buys D, the optimal bundle on her budget line. She then goes to Kroger, prices the same bundle, and finds that it costs about four dollars more. A&P too was telling the truth.

The situation is shown in Figure 3-2. The vertical axis represents housing, the horizontal axis expenditure on all other goods. The initial budget line shows the different combinations of housing and other goods you could have chosen at the initial price of housing. Point A is the optimal bundle — the amount of housing you bought.

A second budget line shows the situation after the price of housing has risen. It has a shallower slope, since more expensive housing means that you must give up more dollars to get an extra square foot of house. The new budget line must still go through point A, since one of your alternatives is to continue living in the house you already own. You can choose to move away from A along the budget line either up (sell your house and buy a bigger one, trading dollars for housing) or down (sell your house and buy a smaller one, trading housing for money)

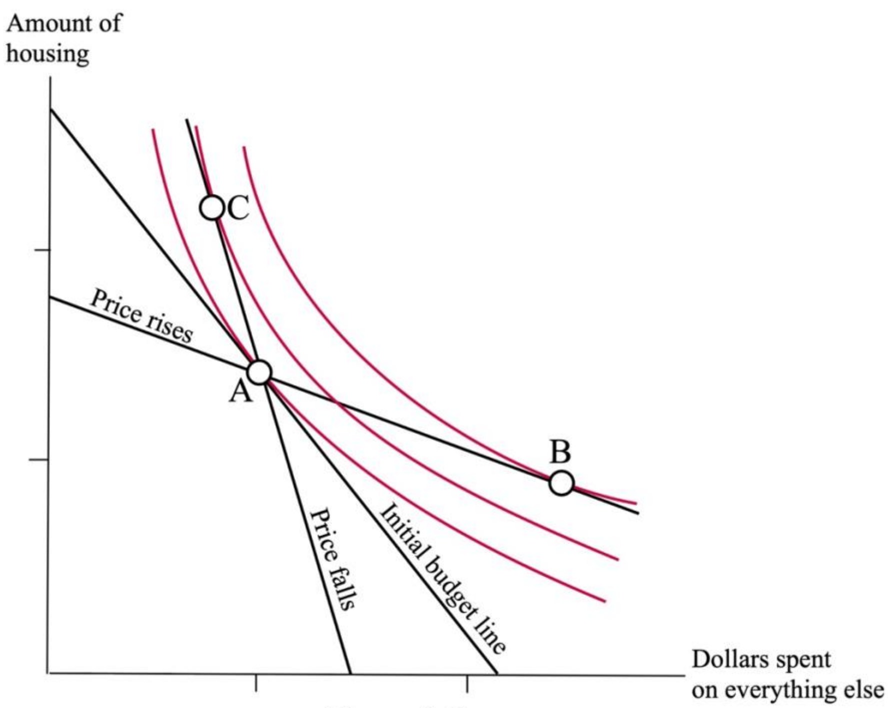

Figure 3-2

The effect on a homeowner of a change in the price of housing. The initial budget line shows the alternatives available at the original price of housing; the other two budget lines show the alternatives available if the price rises or falls. **A** shows the homeowner's bundle of housing and all other consumption after the house is built and before any change in housing prices.

The figure shows what you choose to do; your new optimal point is B. Since housing is now more expensive, you have sold your house and bought a smaller one — the gain in income is worth more to you than the loss in space. You are now on a higher indifference curve than before the price change.

A third budget line shows the situation if the price of housing goes down rather than up after you buy your house. Again you have the choice of keeping your original house, so the line has to go through A — but this time with a steeper slope, since housing is now cheaper. Your new optimal point is C; you have adjusted to the lower price of housing by selling your house and buying a bigger one. You are again on a higher indifference curve than before the price change. The drop in the price of housing has made you better off!

By looking at the figure, you should be able to convince yourself that the result is a general one; whether housing prices go up or down after you buy your house, you are better off than if they had stayed the same. The argument can be put in words as follows:

> *What matters to you is what you consume — how much housing and how much of everything else. Before the price change, the bundle you had chosen — your house plus whatever you were buying with the rest of your income — was the best of those available to you; if prices had not changed, you would have continued to consume that bundle. After prices change, you can still choose to consume the same bundle, since the house already belongs to you, so you cannot be worse off as a result of the price change.*
>
> *But since the optimal combination of housing and other goods depends on the price of housing, it is unlikely that the old bundle is still optimal. If it is not, that means there is now some more attractive alternative, so you are now better off; a new alternative exists that you prefer to the best alternative (the old bundle) that you had before.*

The advantage of the geometrical approach to the problem is that the drawing tells us the answer. All we have to do is look at Figure 3-2. The initial budget line was tangent to its indifference curve at point A, so any budget line that goes through A with a different slope must cut the indifference curve. On one side or the other of the intersection, the new budget line is above the old indifference curve — which means that you now have opportunities you prefer to bundle A.

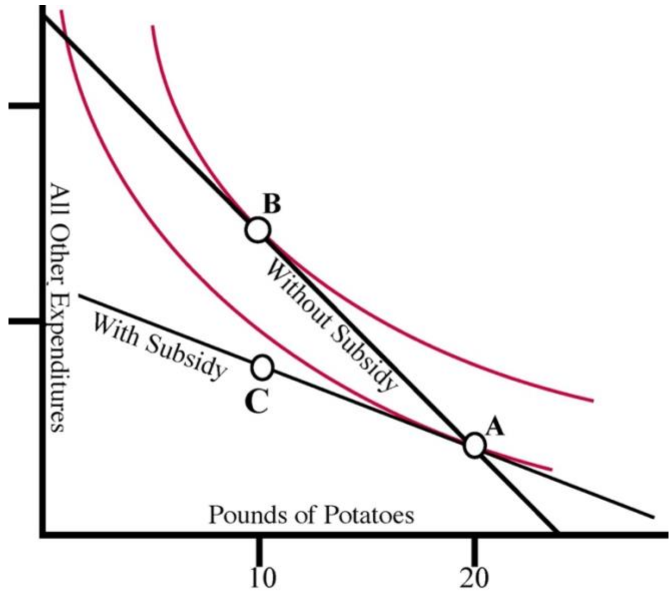

Figure 3-3

Purchases of Potatoes with (A) and without (B) a subsidy and associated tax.

Appearances are deceiving. You are paying twenty dollars a month in taxes — and so is everyone else. You are receiving twenty dollars a month in subsidy — and so is everyone else. The result is that you are worse off — and so is everyone else, with the possible exception of the potato farmers.

To see why, consider Figure 3-3, which shows the budget lines with and without the subsidy and associated tax. A is the optimal point with the subsidy, the point where the budget line just touches an indifference curve. It is the bundle — of potatoes and everything else — that you choose to consume, given the alternatives available to you.

Since potatoes are more expensive without the subsidy, the budget line showing your alternatives without the subsidy is steeper: You must give up more of everything else for each pound of potatoes you consume. It still runs through point A. Buying that bundle will cost you an extra twenty dollars, since potatoes are a dollar a pound more expensive without the subsidy — and that is exactly the amount you no longer have to pay in taxes.

You can still buy A if you want to, but you don't. As you can see from the figure, the most attractive bundle available to you, with neither tax nor subsidy, is B. You reduce your consumption of potatoes by ten pounds, spend the money you save on other goods, and shift up to a higher indifference curve.

The figure gives us the answer: We are better off at B than at A, so the combination of a potato subsidy and a tax to pay for it has made us worse off. But just as in the previous example, we need to convert the argument back into English before we can understand why.

We start by asking why I could not get from A to B without abolishing the subsidy. For the population as a whole, tax collected equals subsidy paid, and the amount of subsidy paid depends on how many pounds of potatoes people buy. If everyone cut his consumption of potatoes in half we could cut the tax in half as well, putting all of us at B.

But I do not control what everybody does; I only control what I do. If only I cut my consumption my tax remains almost the same and I am at C — worse off than if I remained at A.

We would all be better off if we all cut our consumption of potatoes in half, but each of us would be worse off if he cut his consumption of potatoes in half.

4: How much would you pay to Get Off a Desert Island?

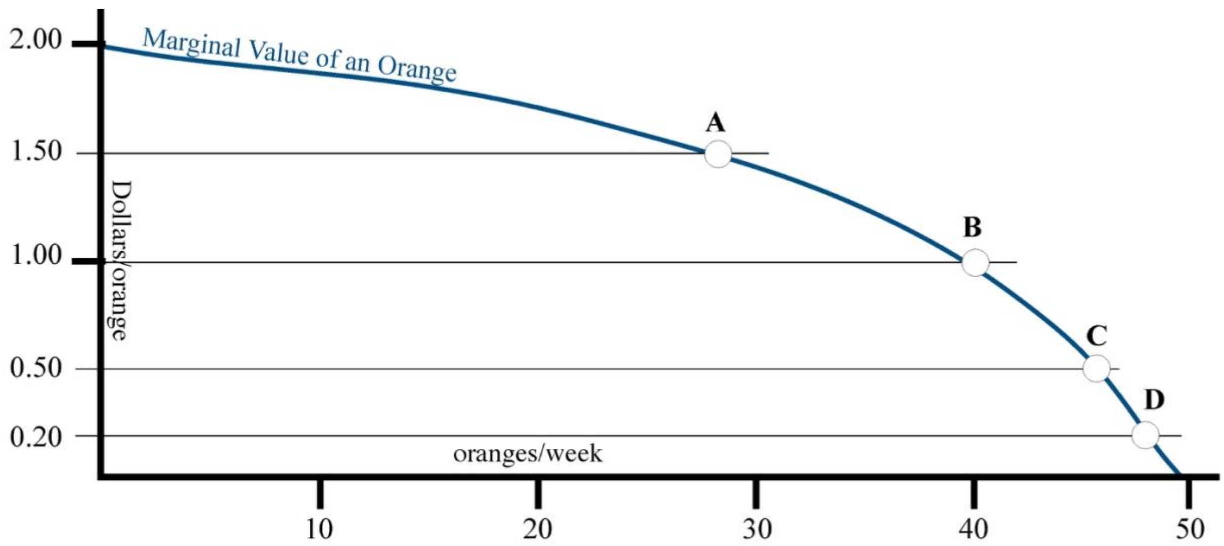

Figure 4-1

The value to you of one more orange as a function of how many oranges you are consuming.

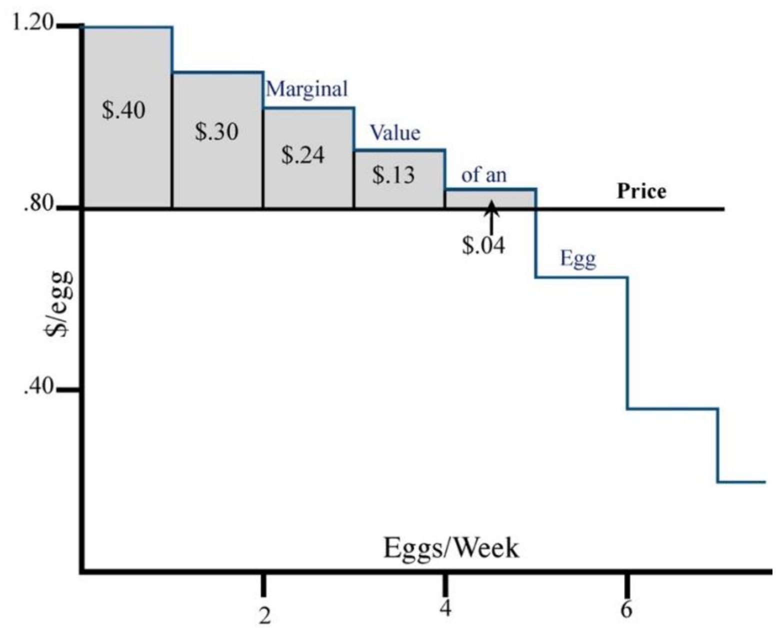

Figure 4-2

Marginal value curve and consumer surplus for a lumpy good. The shaded area under the marginal value curve and above the price is consumer surplus: the net benefit from buying that quantity at that price.

Can we make this argument more precise? Can we say how much better off you are by being able to buy as much water as you want at $0.01/gallon or as many eggs as you want at $0.80/egg? The answer is shown in Figure 4-2. By buying one egg instead of none, you receive a marginal value of $1.20 and give up $0.80; you are better off by $0.40. Buying a second egg provides a further increase in value of $1.10 at a cost of another $0.80. So buying 2 eggs instead of none makes you better off by $0.70.

This does not mean you have $0.70 more than if you bought no eggs — on the contrary, you have $1.60 less. It means that buying 2 eggs instead of none makes you as much better off as would the extra goods you would buy if your income were $0.70 higher than it is. You are indifferent between having your present income and buying 2 eggs (as well as whatever else you would buy with the income) and having $0.70 more but being unable to buy any eggs.

Up to five eggs per week, each additional egg you buy makes you better off. Your total gain from consuming 5 eggs at a price of $0.80 each instead of consuming no eggs at all is the shaded area on the figure, the sum of the little rectangles. The gain from consuming five eggs is the gain from consuming five instead of four, plus the gain from consuming four instead of three, plus ... .

Next consider Figure 4-3, where instead of a lumpy good such as eggs we show a continuous good such as wine. If we add up the gain on buying wine, drop by drop, the tiny rectangles exactly fill the region A. That is your net gain from being able to buy wine at $10/gallon.

This area is *consumer surplus*: The net gain to you from what you consume. Think of it as the value of what you buy (A+B on the figure) minus what you give up to get it (B). It is a tool of many uses. In later chapters it will help us to measure the real cost of taxes, figure out how to run Disneyland, and decide whether to legalize polygamy.

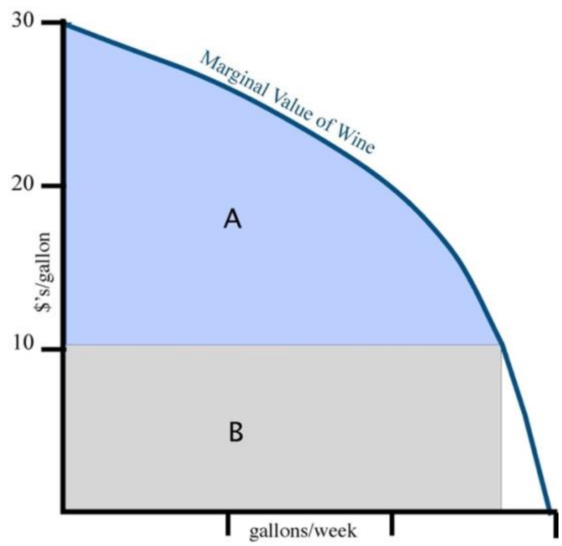

Figure 4-3

Marginal value and consumer surplus for a continuous good. A is the consumer surplus from being able to buy all the wine you want at $10/gallon. B is what you pay for it. A+B is the total value to you of 2 gallons per week of wine.

Figure 4-4 shows your demand curve for potatoes. To simplify the problem, I assume that potatoes cost $2/lb to produce and are sold at a price that just covers their cost.

Without the subsidy the price is two dollars and your consumer surplus is area A. With the subsidy, the price is one dollar and your surplus is A+B. So your gain from the subsidy is the difference: area B.

What does the subsidy cost you? Just as in chapter 3, we assume that everyone buys the same quantity of potatoes and pays the same share of taxes, so your taxes are just equal to the cost of the subsidy you are receiving: a dollar a pound times the number of pounds of potatoes you are consuming (Q_s). That is B+C on the figure. You gain B, you lose B+C, so your net loss is C.

Where does the loss come from? It comes from consuming potatoes that are worth less to you than they cost to produce. Between Q_o and Q_s, the value to you of each additional pound of potatoes is between one and two dollars, as shown by your marginal value curve, the same line as your demand curve. Because of the subsidy, you are eating potatoes that cost two dollars to produce and are worth less than two dollars to you. C is the resulting net loss.

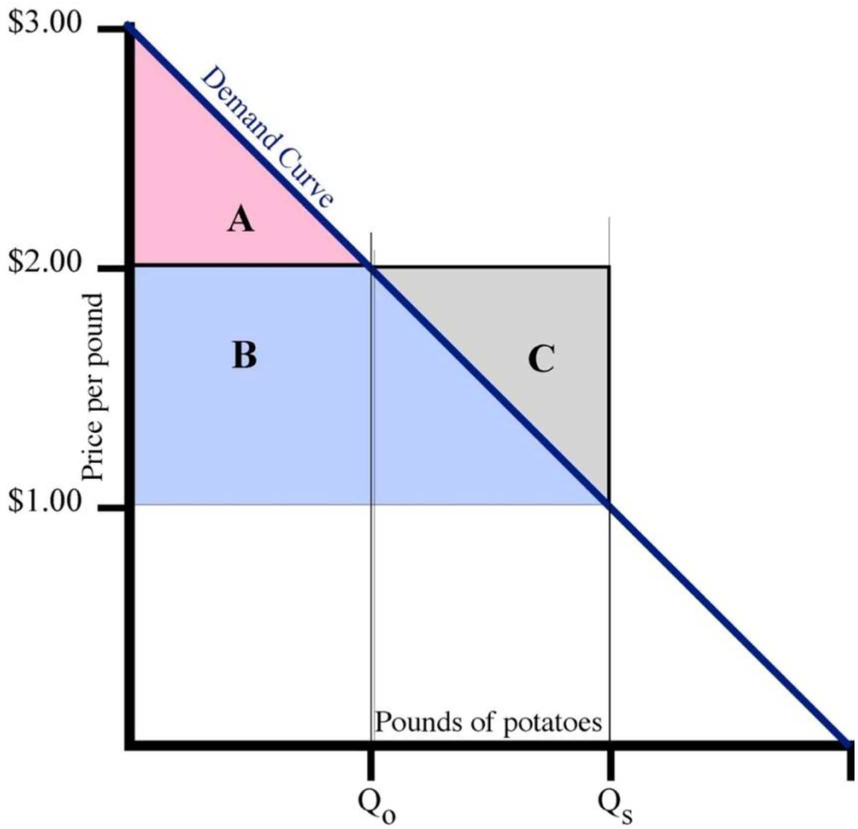

Figure 4-4

Your demand curve for potatoes. A one dollar subsidy shifts the price from \$2.00 to \$1.00, increasing your consumer surplus by B, costing you B+C in additional taxes, thus making you worse off by C.

Figure 4-5 shows your demand curve for popcorn. Suppose the theater sells it at a dollar a bag. You buy one bag for a dollar, spending area B+D; your consumer surplus is area A. If popcorn costs the theater fifty cents a bag, their cost is D, leaving them B — a profit of fifty cents.

Next suppose they cut the price to fifty cents. Your expenditure is now D+E — two bags at fifty cents apiece. Their profit is zero, since they are selling at cost. It looks as though dropping the price lost them fifty cents — area B.

We have forgotten consumer surplus. At the lower price, your consumer surplus is A+B+C. The value to you of the environment they are providing has increased by B+C, so when they cut the price of popcorn they can raise the admission price by that much without driving you off. They have lost B on popcorn but gained B+C on admission, for a net gain of C.

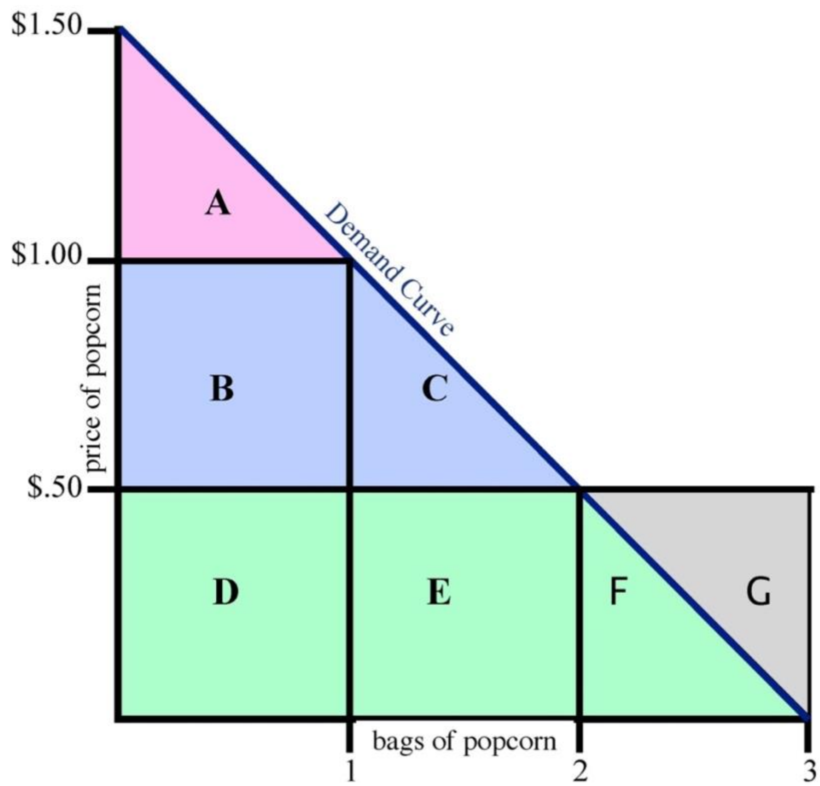

Calculating the profit maximizing price of popcorn

Suppose the theater decides to push your consumer surplus even higher by giving the popcorn away. At a price of zero, you buy three bags. Their loss from producing three bags and giving them away is their cost: D+E+F+G. The amount you are willing to pay for admission has increased by the increase in your surplus: D+E+F. They are worse off by G.

5: Bricks without Clay — production in a One-input World

Step I: How to Spend Your Life

You can produce any of three goods, as shown in Table 5-1: mowed lawns, washed dishes, or meals. The price for a mowed lawn is $10 and you can mow one lawn in an hour, so mowing pays $10/hour. Washing seventy dishes per hour at $0.10/dish yields $7/hour and cooking two meals per hour at $3 per meal yields $6 an hour. Since the only difference among the alternatives is the implicit wage, you get out the mower.

Table 5-1

	Lawn Mowing	**Dish Washing**	**Cooking**
Output	1 lawn/hour	70 dishes/hour	2 meals/hour
Price	$10/lawn	$0.10/dish	$3/meal
Wage	$10/hour	$7/hour	$6/hour

Step II: How Much of Your Life to Spend

How many lawns do you mow? Figure 5-1a shows the marginal disvalue of labor. Just as the marginal value of oranges depends on how many you have, so the marginal disvalue of working depends on how much work you are doing. If you were enjoying 24 hours a day of leisure, it would take only a small payment ($0.50 in the figure) to make you willing to work for a single hour; you would be indifferent between zero hours a day of work and 1 hour of work plus $0.50. If you were already working 10 hours a day, it would take a little over $10 to make you willing to work an additional hour.

The wage is $10/hour and you are working 5 hours per day. You would be willing to work an additional hour for an additional payment of about $3; since you can actually get $10 for it, you are better off working the extra hour. The same argument applies to the next hour; it keeps applying so long as the marginal disvalue of labor to you is less than the wage. So you end up working that number of hours for which the two are equal; the number of hours of labor you supply at a wage of $10 is the number at which your marginal disvalue for labor is equal to $10. Your marginal disvalue for labor curve is your supply curve for labor — just as, in Chapter 4, your marginal value curve was your demand curve. You work ten hours (and mow ten lawns) a day.

Producer Surplus

The wage is $10/hour. You are willing to work the first hour for $0.50; since you receive $10 for it, your net gain on that hour is $9.50. The next hour is worth a dollar to you; you receive $10 for a gain of $9. Summing these gains over all the hours you work gives us the shaded area of Figure 5-1a, the amount by which you are better off working at $10/hour than not working at all. Just as consumer surplus was the area under the demand curve (equal to the marginal value curve) and above price, so producer surplus is the area under the wage and above the supply curve (equal to the marginal disvalue curve) for labor.

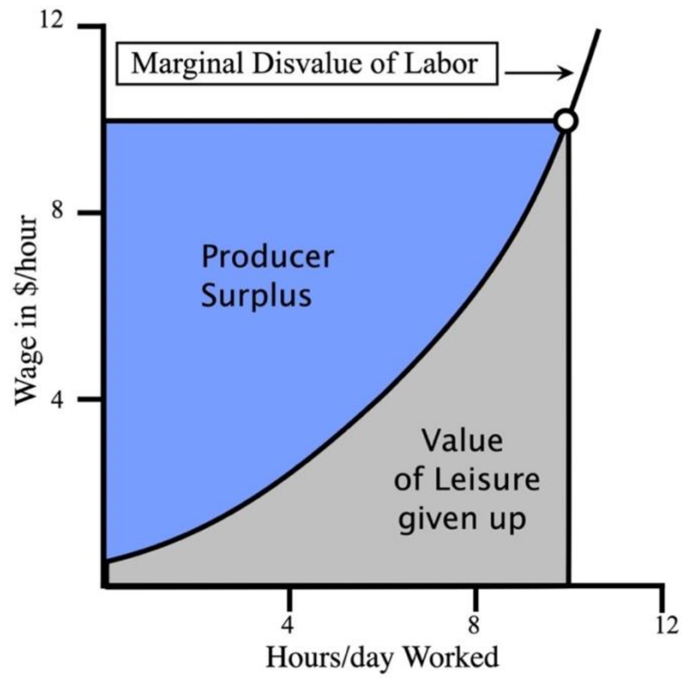

Figure 5-1a

Producer Surplus, the marginal disvalue of labor, and the supply curve for lawn mowing. The area above the marginal disvalue curve and below the $10/hr wage is the producer surplus from being able to work for $10/hr.

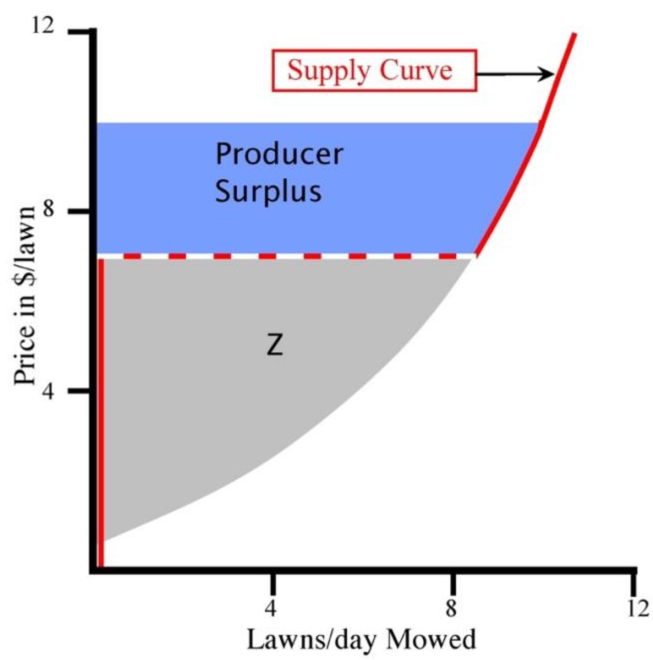

Figure 5-1b

Producer Surplus and the supply curve for lawn mowing The shaded area above the supply curve for lawns and below the price is the producer surplus from being able to mow lawns for $10/lawn. The supply curve is horizontal at the price at which you switch to your next most profitable option — washing dishes.

We now have the supply curve for labor but what we want is the supply curve for lawns. Since I can mow 1 lawn per hour, a price of $10/lawn corresponds to a wage of $10/hour and a labor supply of 10 hours per day corresponds to mowing that many lawns. It appears that the supply curve for lawns and for labor are the same; all I have to do is relabel the vertical axis "price in $/lawn" and the horizontal axis "lawns/day."

Appearances are deceiving; there is one important difference between the two supply curves. When the amount I get for mowing a lawn drops below $7, my output of mowed lawns drops to zero; I am better off washing dishes. The resulting supply curve is shown on Figure 5-1b. The shaded area is my producer surplus.

To see why it does not include Z, the area below the line at $7, consider what my surplus would be if I could get $7 for each lawn I mowed. How much better off am I being able to mow lawns at $7 than not mowing lawns? I am not better off at all; at that price, I can do just as well washing dishes.

Cost is opportunity cost: The cost to me of mowing lawns is whatever I must give up in order to do so. If the best alternative use of my time is leisure, the cost is the value of my leisure. If the best alternative use is washing dishes, the cost is the money I would have gotten by washing dishes.

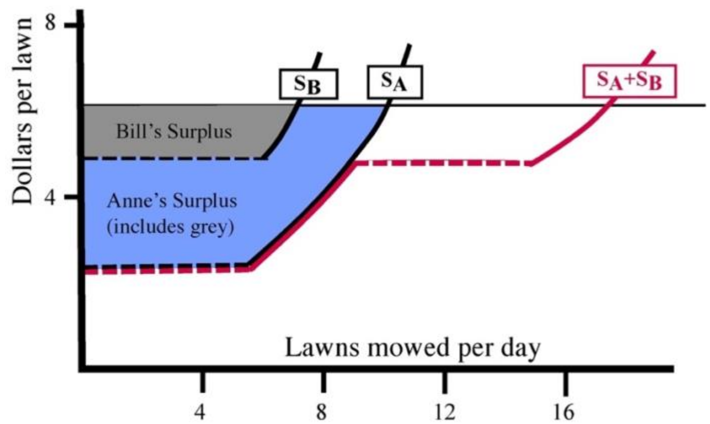

Figure 5-2a

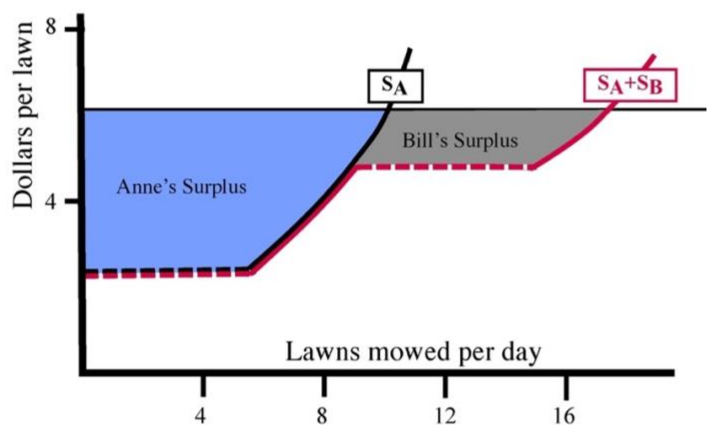

Figure 5-2b

Producer Surplus for two producers.

Step III: Summing People — The Aggregate Supply Curve

Producers differ in how good they are at producing different goods and in how willing they are to work, so different people have different supply curves. A producer who is very good at mowing lawns or very bad at doing anything else will mow lawns even at a low price; one who is bad at mowing lawns or good at something else will mow lawns only when the price is high. Figure 5-2 shows the supply curves for two such producers, A(nne) and B(ill), and their combined supply curve.

At prices below $2.50/lawn, neither Anne nor Bill produces. At prices above $2.50/lawn but below $5/lawn, only Anne produces. At a price of $5, Bill enters the market, mowing 6 lawns per day for a total output (Anne plus Bill) of 15. When the price goes from $5 to $6, Anne increases her output by another unit and so does Bill; total output increases to 17.

The combined supply curve is a horizontal sum; we are adding up quantities (shown on the horizontal axis) at each price. The same would be true if we were deriving an aggregate demand curve from two or more individual demand curves. All consumers in a market pay the same price, so total quantity demanded at a price is the quantity consumer A demands plus the quantity consumer B demands plus . . . .

As you should be able to see from the figure, the sum of the producer surplus that B receives at a price of $6 plus the producer surplus that A receives is equal to the producer surplus calculated from the combined supply curve — the area above their combined supply curve and below the horizontal line at $6. The result applies to any number of producers, as does a similar result for the consumer surplus of any number of consumers. So we can find the sum of the surpluses received by consumers or producers by calculating the surplus from their aggregate demand or supply curve just as if it were the demand or supply curve for a single individual.

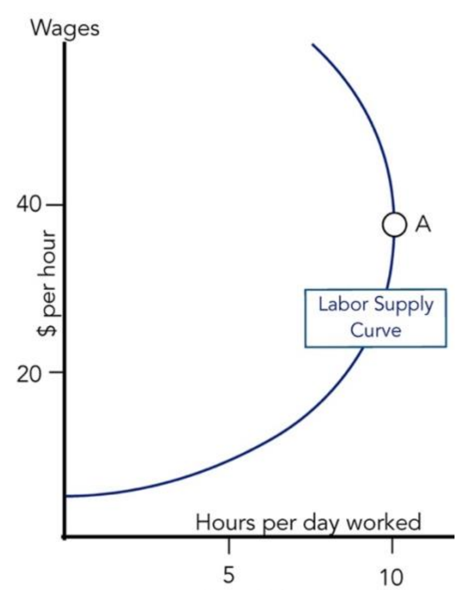

Figure 5-3

A backward-bending supply curve for labor. As the wage increases, the number of hours worked first increases (up to A) then decreases.

7: Putting It Together — Price Theory in a Simple Economy

Figure 7-1a shows supply and demand curves for widgets, an imaginary commodity consumed mostly by economics professors. The vertical axis is price, the horizontal axis is quantity; any point on the diagram represents a quantity and a price.

Suppose widgets cost ten dollars apiece. At that price, producers wish to produce and sell more widgets than consumers want to buy. Producers with widgets they cannot sell are willing to cut their price to get rid of them. Price falls — and continues to fall as long as quantity supplied is greater than quantity demanded.

What if, instead of ten dollars, the initial price was five dollars? At that price, consumers want to buy more than producers want to sell. Some consumers find that they cannot buy as many widgets as they want. Figure 7-1b shows the marginal value curve of one such consumer. At $5/widget he would like to buy six widgets but can only find four for sale. He is willing to pay anything up to nine dollars for one more widget, since that is its marginal value. He, and other consumers with the same problem, bid the price up.

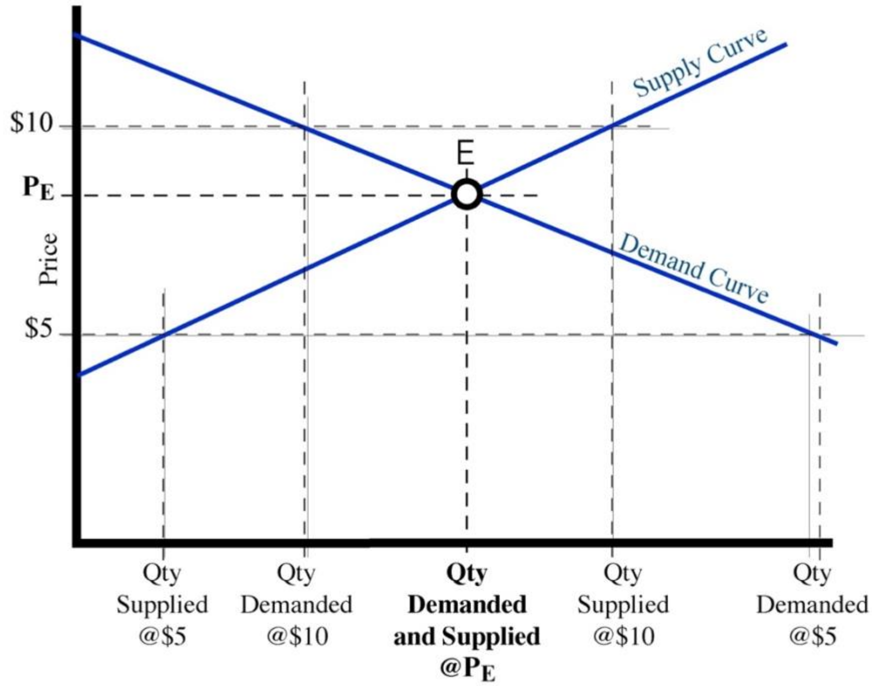

Figure 7-1a

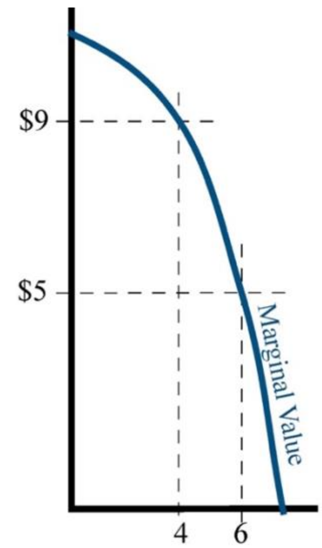

Figure 7-1b

Market equilibrium. At point E, price = P_E; quantity demanded equals quantity supplied. At lower prices, less is supplied; individuals are consuming quantities for which MV > P, as shown on Figure 7-1b, and so are willing to offer a higher price for additional quantities.

Shifting Curves

Much confusion can be avoided by distinguishing carefully between changes in demand (the demand curve shifting) and changes in quantity demanded, and similarly for supply and quantity supplied. In Figure 7-2, for example, demand changes, which changes price, which changes the quantity supplied. But supply has not changed; the supply curve is the same after the change as before.

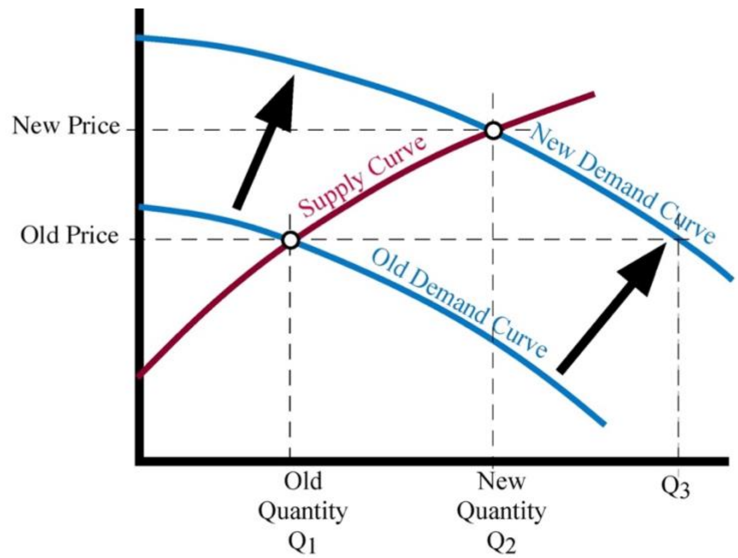

Figure 7-2

The effects of shifts in supply and demand curves.

Being careful with such distinctions can help you avoid some of the worst absurdities of newspaper economics. Consider the following:

> *"The demand for memory chips increased, which drove up the price, which drove up the supply, which brought the price back down."*

This is the change illustrated on Figure 7-2. An increase in demand (the demand curve shifts out) raises price; the increased price reduces quantity demanded below what it would have been if the demand curve had shifted but the price had remained the same (Q_3). The new quantity demanded (Q_2) is less than Q_3 but more than the old quantity demanded (Q_1). Q_2 must be greater than Q_1 because quantity demanded is equal to quantity supplied, the supply curve has not shifted, and a higher price applied to the same supply curve results in a larger quantity supplied.

Who Pays Taxes?

We are now ready to start on one of the questions sometimes asked of economists; the number of pages it has taken us to get this far may explain why answers that fit a 30-second news story are generally wrong. The question is "Who really pays taxes?" When a government imposes a tax on some good, does the money come out of the profits of those who produce it or do the producers pass it along to the consumers in higher prices?

Suppose the tax is $1/widget; for every widget sold, the producer must pay the government $1. The result is to shift the supply curve up by $1, from S_1 to S_2, as shown in Figure 7-3a.

Why? What matters to the producer is how much he gets, not how much the consumer pays. If he gets $6/widget, of which he must hand over $1 to the government, his return for each widget sold is the same as if he were selling them at $5/widget. So he produces the same quantity of widgets at \$6/widget after the tax is imposed as he would have produced at \$5 before and similarly for all other prices. Each quantity on the new supply curve corresponds to a price \$1 higher than on the old; the supply curve shifts up by \$1.

This does not mean that the market price goes up \$1. If it did, producers would produce the same amount as before the tax, consumers would consume less than before, making quantity supplied greater than quantity demanded. If, on the other hand, price did not rise at all, quantity demanded would be the same as before the tax, quantity supplied would be less, since producers would be getting a dollar less per widget, so quantity supplied would be less than quantity demanded. As you can see on Figure 7-3a, the price rises, but by less than a dollar. All of the tax is paid by the producer in the literal sense that the producer hands the government the money, but in fact the price paid by the consumer has gone up by *a* and the price received by the producer net of tax has gone down by *b*, where *a*+*b* adds up to the full amount of the tax.

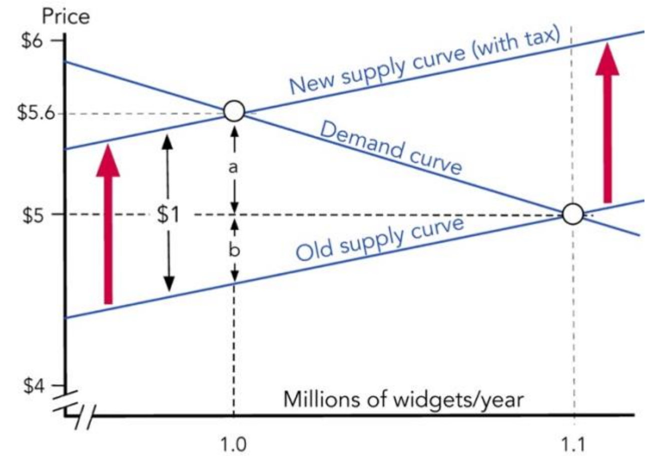

Figure 7-3a

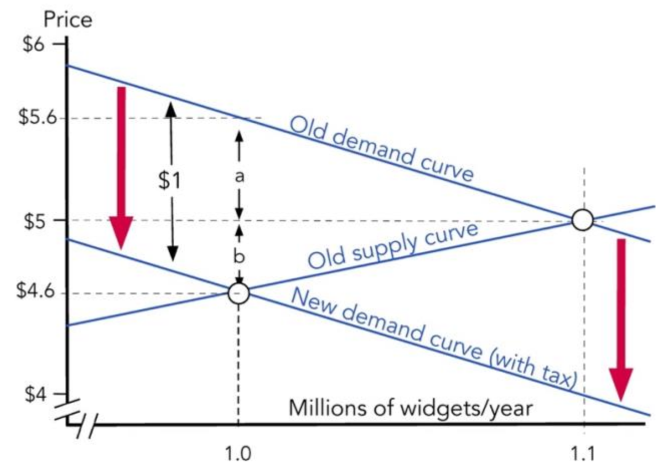

Figure 7-3b

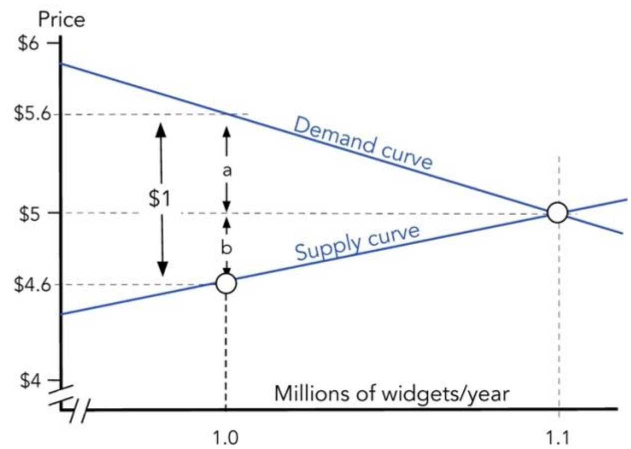

Figure 7-3c

The effect of a $1 tax on widgets. Figure 7-3a shows the effect of a tax paid by the producer; the supply curve shifts up. Figure 7-3b shows the effect of a tax paid by the consumer. Figure 7-3c shows the same situation, with the supply curve depending on price received by the producer (market price minus any tax on producers) and the demand curve on price paid by the consumer (market price plus any tax on consumers). The difference between the two prices is the tax, whichever one actually hands the money over to the government.

Suppose the government decides to tax consumers instead of producers: For every widget you buy, you must pay the government $1.The result is shown on Figure 7-3b. This time it is the demand curve that is shifted by the tax. Widgets at $5 with no tax cost you the same amount as widgets at $4 with a $1 tax, payable by the consumer; either way you give up, for each widget purchased, the opportunity to buy $5 worth of something else. Since the cost to you is the same in both cases, you buy the same quantity in both cases — and so does everyone else. So the total quantity demanded is the same at a price of $4 with the tax as it would be without the tax at a price of $5, and similarly for all other prices. The demand curve shifts down by $1, the amount of the tax.

Looking at Figure 7-3b, you can see that the tax lowers the price received by the producer by *b* and increases the cost (including tax) to the consumer by *a*, and that *a* and *b* are the same as on the previous figure. If we ignore the old supply curve on one figure and the old demand curve on the other, figure 7-3b is simply 7-3a shifted down by $1. On Figure 7-3a, the price shown on the vertical axis is price after tax, since the tax is paid by the producer. On 7-3b, it is price before tax, since the tax is paid by the consumer. The difference between price before tax and price after tax is the amount of the tax: $1.

A third way of describing the same situation is shown in Figure 7-3c. Here supply is shown as a function of price received, demand as a function of price paid. Before the tax was instituted, market equilibrium occurred at a quantity (1.1 million widgets/year) for which price received was equal to price paid. After the tax was instituted, market equilibrium occurs at a quantity (1 million widgets/year) for which price received is a dollar less than price paid, with the difference going to the government.

What we left out of our analysis of the cost of a one dollar tax on widgets was consumer (and producer) surplus, whose function is to measure the net benefit of being able to buy (sell) goods. Before the tax, the consumer could purchase and the producer sell as many widgets as he wanted at $5 apiece. Afterwards the cost to the consumer was $5.60/widget and the revenue received by the producer was $4.60/widget. The cost to producers and consumers of the tax is the difference between their surplus in the first case and their surplus in the second, shown in Figure 7-4.

The area under the demand curve and above $5 is consumer surplus before the tax. The area under the demand curve and above $5.60 is consumer surplus after the tax. The blue area above $5 is the difference between the two, the cost of the tax to consumers. It is made up of two parts: a rectangle (increased cost/widget times number of widgets purchased) plus a triangle (lost consumer surplus on widgets no longer bought because of the tax).

Similarly, the green area below $5 is the cost of the tax to producers, their loss of producer surplus. It too consists of a rectangle (lost revenue on the widgets still being produced) plus a triangle (lost producer surplus on widgets no longer sold because of the tax).

If we sum the two rectangles, we have the amount of the tax, the difference between cost per widget to consumers and revenue per widget to producers, times the number of widgets produced; that is the total revenue produced by the tax. If we sum the two triangles, we have the excess burden of the tax, a loss for producers and consumers with no corresponding gain for anyone.

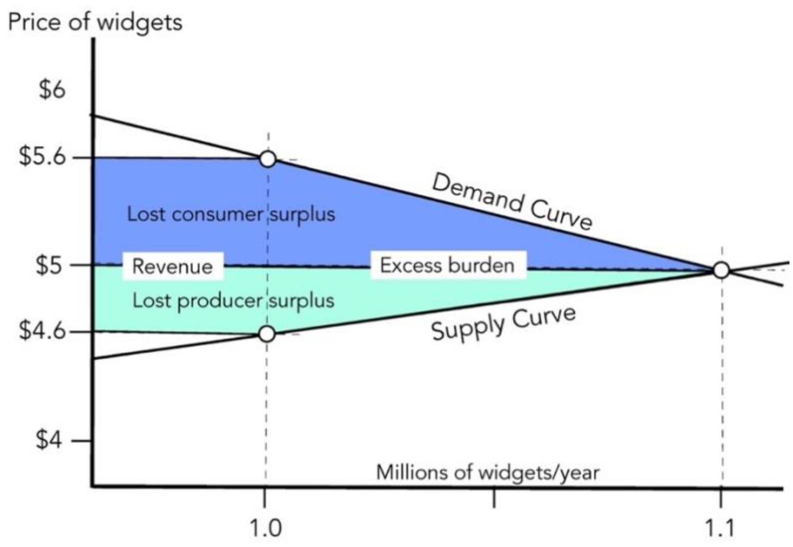

Figure 7-4

The effect on surplus of a $1 tax on widgets. The dark shaded area is lost consumer surplus, the lightly shaded area lost producer surplus. Lost surplus equals revenue collected (the two rectangles) plus excess burden (the two triangles).

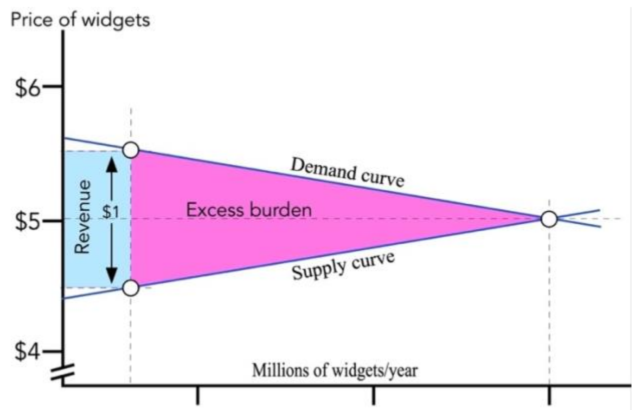

Figure 7-5a

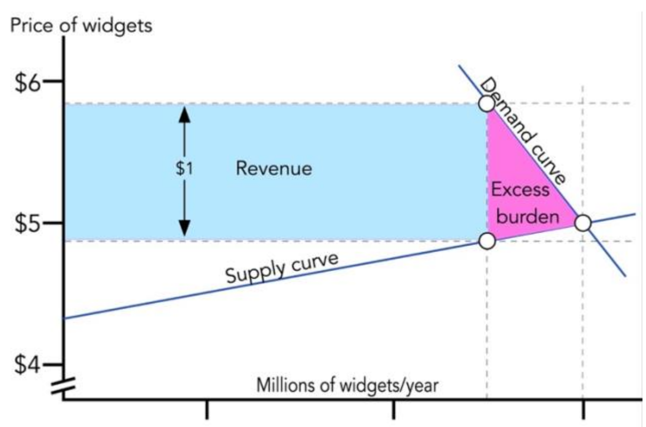

Figure 7-5b

The effect of elasticity of the demand curve on the relation between revenue and excess burden. A very elastic demand curve (Figure 7-5a) produces a high ratio of excess burden to revenue; a very inelastic demand curve (Figure 7-5b) produces a low ratio.

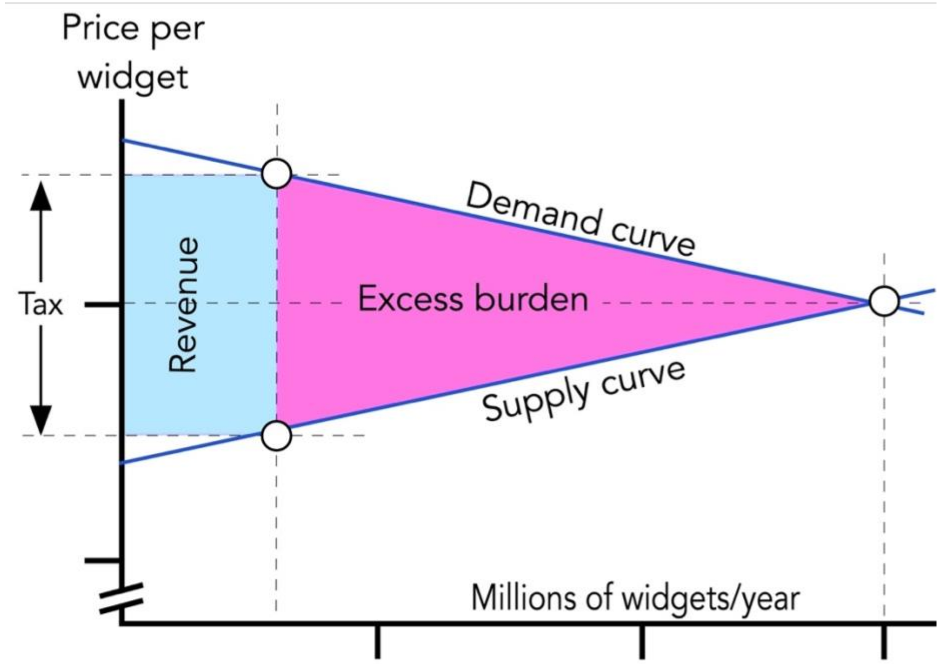

Figure 7-6a

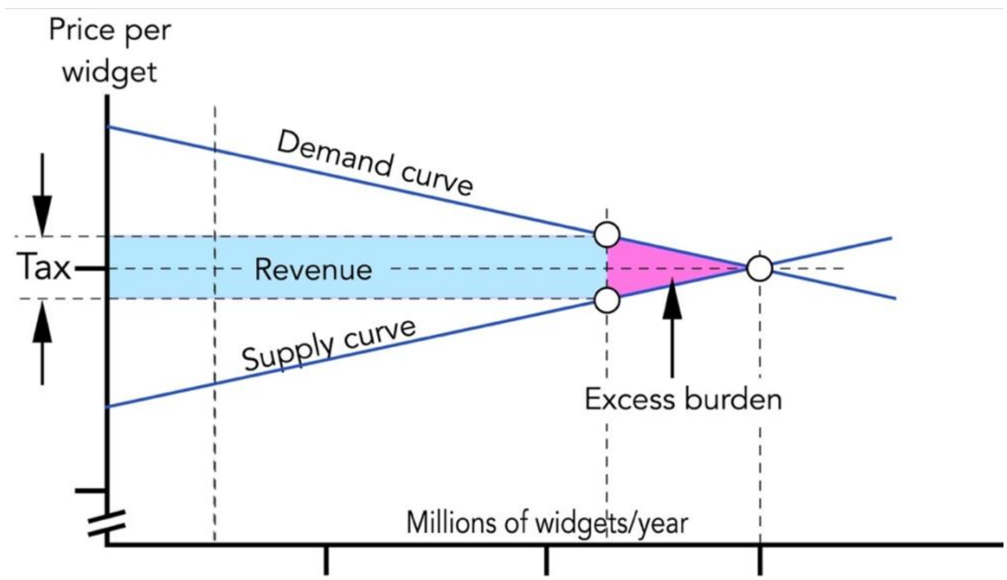

Figure 7-6b

The effect of the size of the tax. A large tax (7-6a) produces more excess burden per dollar of revenue than a small tax (7-6b).

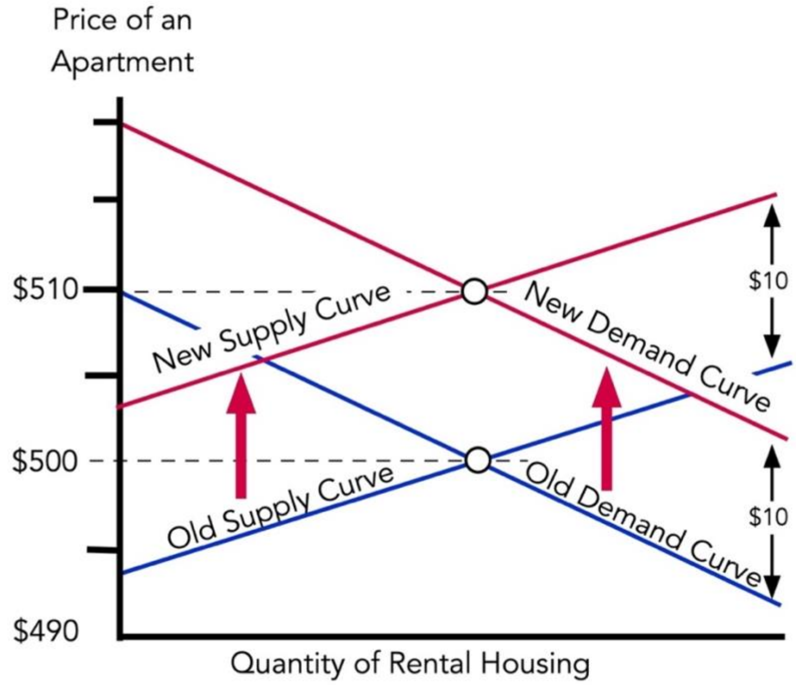

Figure 7-7a

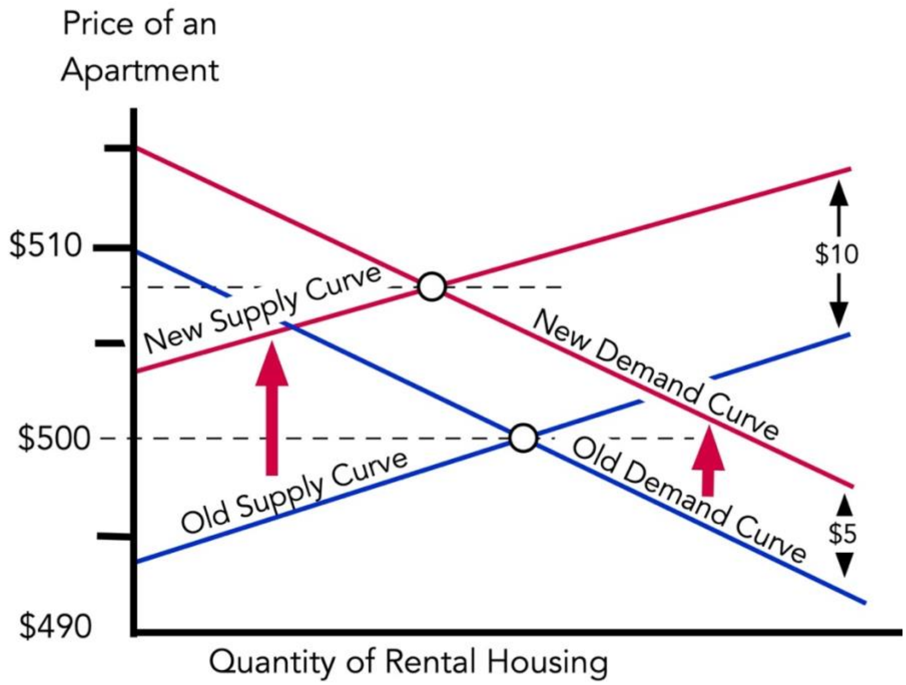

Figure 7-7b

Effect of regulations on the rental market. Figure 7-7a shows the effect of a compulsory $10 transfer from landlords to tenants. Figure 7-7b shows the effect of requiring landlords to provide tenants with six months' notice. The requirement is equivalent to a $10 tax on landlords and a $5 subsidy to tenants.

Figure 7-7a shows the result; for simplicity I am treating housing as if it were a simple continuous commodity like water and defining price and quantity in terms of some standard-sized apartment. Since both curves shift up by $10, their intersection shifts up by $10 as well. The new equilibrium rent is precisely $10 higher than the old. The law neither benefits the tenant nor hurts the landlord.

Next consider a more realistic regulation. The city council decides that the terms of some existing leases are unfair to tenants and announces that in the future landlords must give tenants six months' notice before evicting them even if the tenants have agreed in the lease to some shorter period. Again we consider the effect after enough time has passed to let rents reach their new equilibrium.

The new rule increases operating costs by making it harder to evict undesirable tenants. From the standpoint of the landlord, it is like a tax. Suppose it is equivalent to a tax of $10: Landlords are indifferent between having to provide each tenant with six months' notice and having to pay a $10/month tax on each apartment. The supply curve for apartments shifts up by $10, as shown in Figure 7-7b.

The additional security is worth something to the tenants. Suppose it is worth $5/month; a tenant who was willing to pay $500/month for an apartment without six months' tenure is willing to pay $505 for one with the additional security. The demand curve shifts up by $5, as shown in Figure 7-7b.

8: The Big Picture

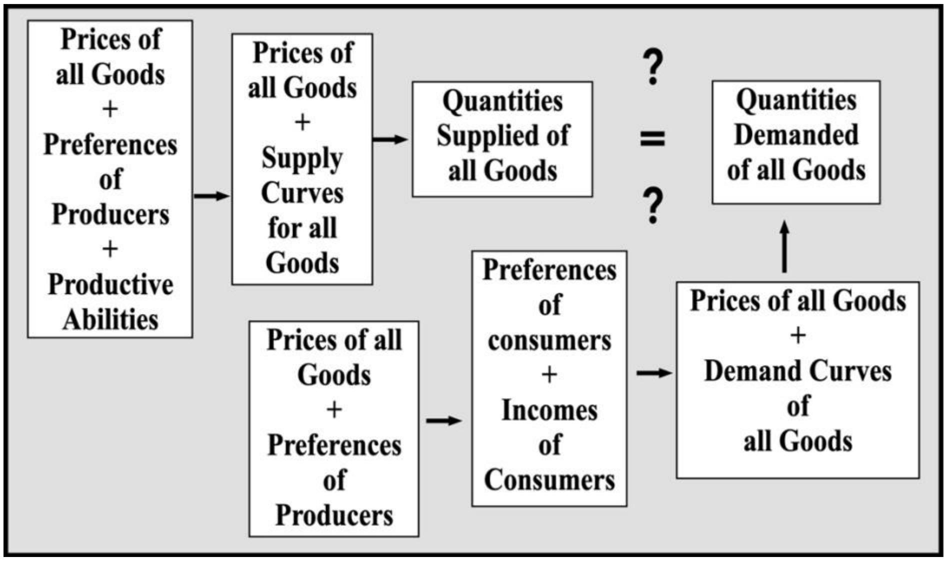

How to solve an economy. Starting with prices of all goods, productive abilities, and preferences of all consumers, derive quantities supplied and demanded. If they are equal for all goods, the initial set of prices describes a possible market equilibrium — a solution for that economy.

9: Bosses, Workers, and Other Complications

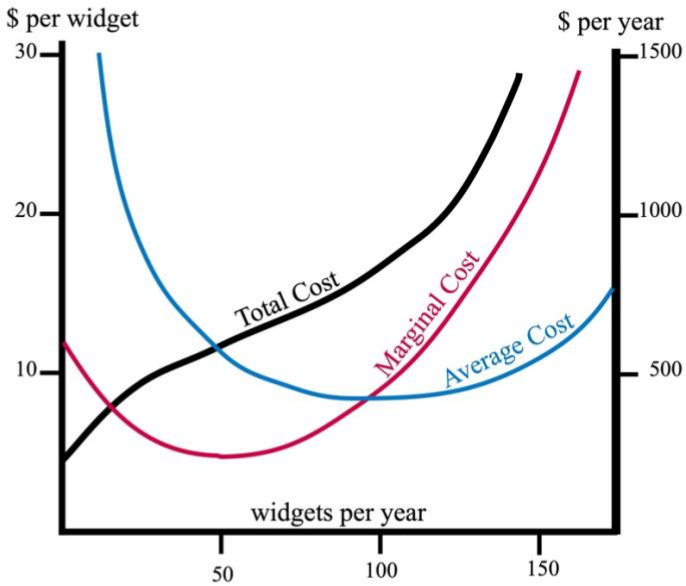

Figure 9-1

A firm's cost curves: Total cost is in different units from marginal and average cost, so the former uses the right vertical axis and the latter uses the left.

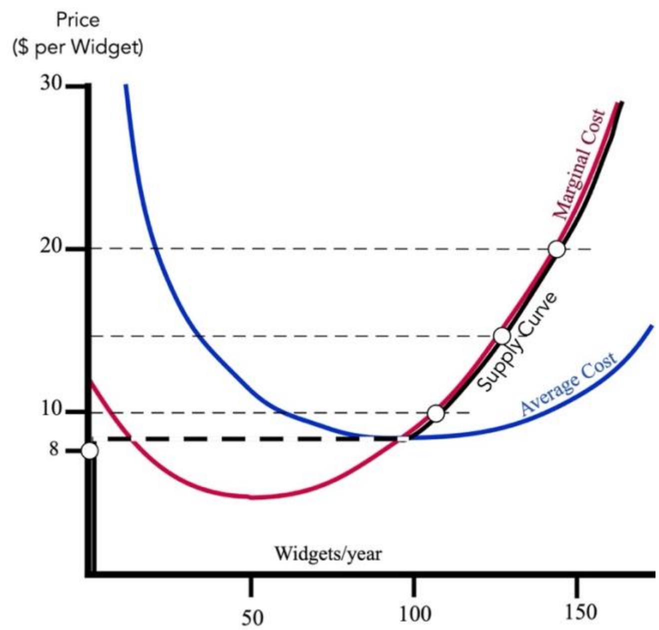

Figure 9-2

The quantity produced at each of four prices, and the resulting supply curve. As long as price is above the minimum of average cost, the firm maximizes its profit by producing a quantity for which P=MC. At lower prices, it shuts down and produces nothing. So the supply curve S is the marginal cost curve above its intersection with average cost.

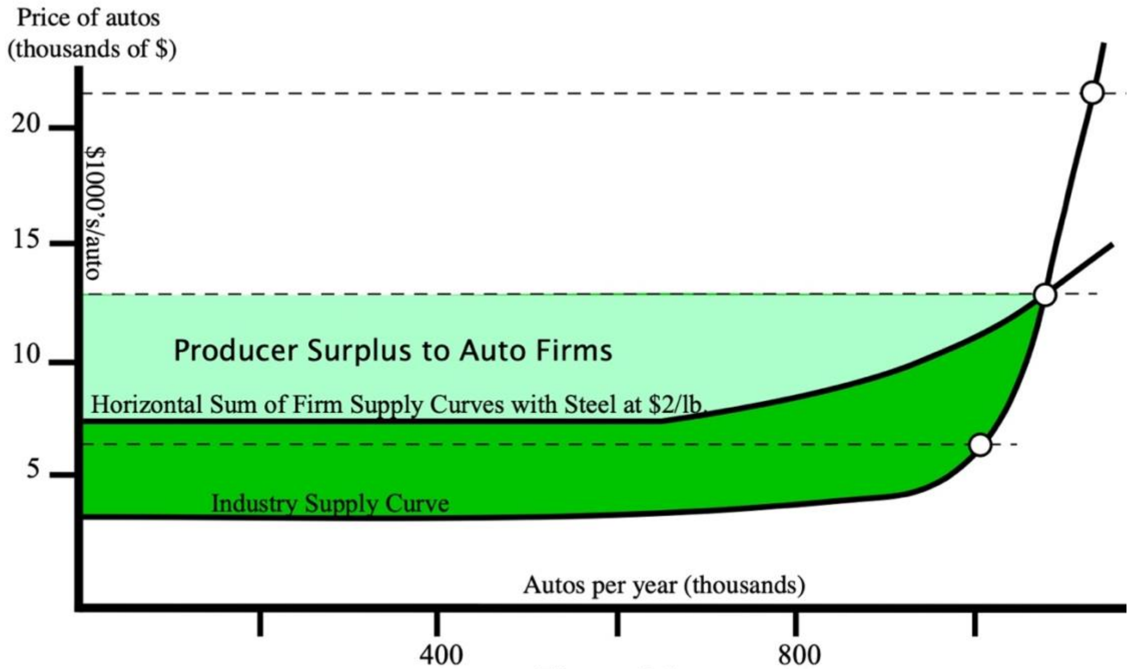

Figure 9-3

Firm and Industry Surplus: With output at 1,300,000 autos/year and steel at $2/lb, producer surplus calculated from the firm supply curves (light green) is less than producer surplus calculated from the industry supply curve (light green plus dark green).

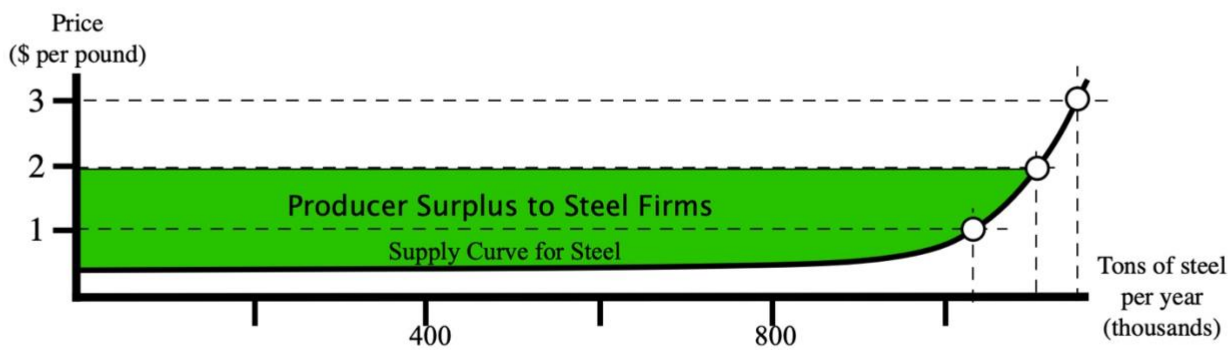

Figure 9-4

Industry Surplus Going to Producers of Inputs. As auto output increases, the industry's increased consumption of steel bids up the price, generating producer surplus for the steel producers.

10: Monopoly for Fun and Profit

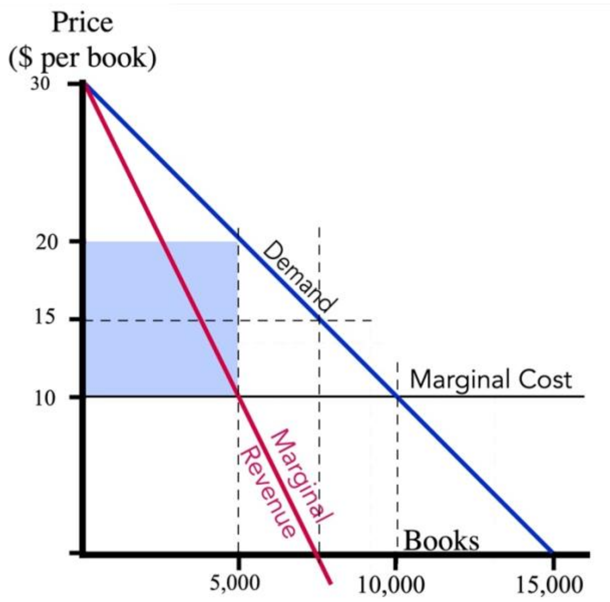

Figure 10-1

Maximizing my profit on this book: Beyond 5,000, each additional book increases revenue by less than it increases cost, so sales beyond that point would reduce profit. The grey square is my gross profit of $50,000; subtracting fixed cost gives me a net profit of $40,000.

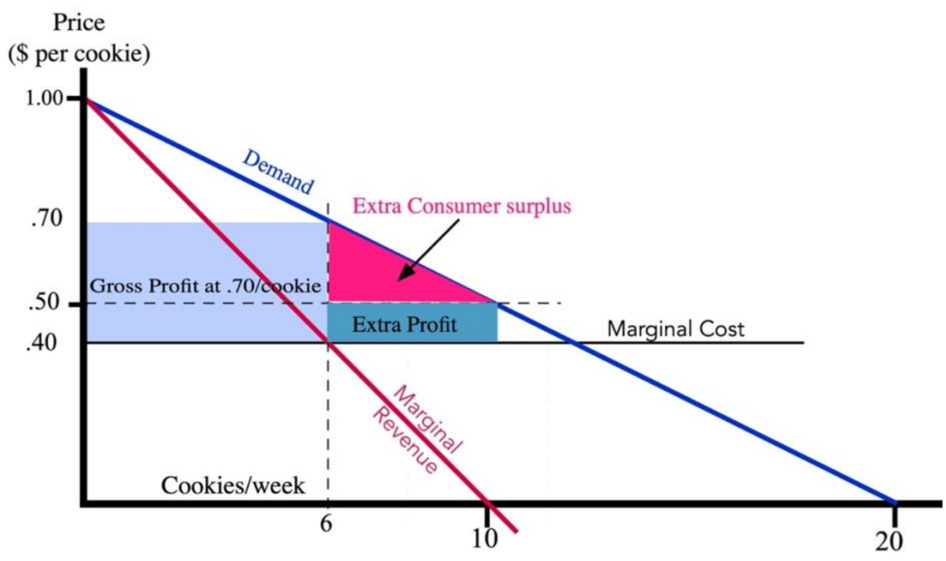

Figure 10-2

Discriminatory pricing in the cookie industry — first try. The profit-maximizing single price is $0. 70/cookie. The firm charges each customer that price for the first 6 cookies but sells additional cookies for $0.50/cookie, increasing its profit by the shaded area.

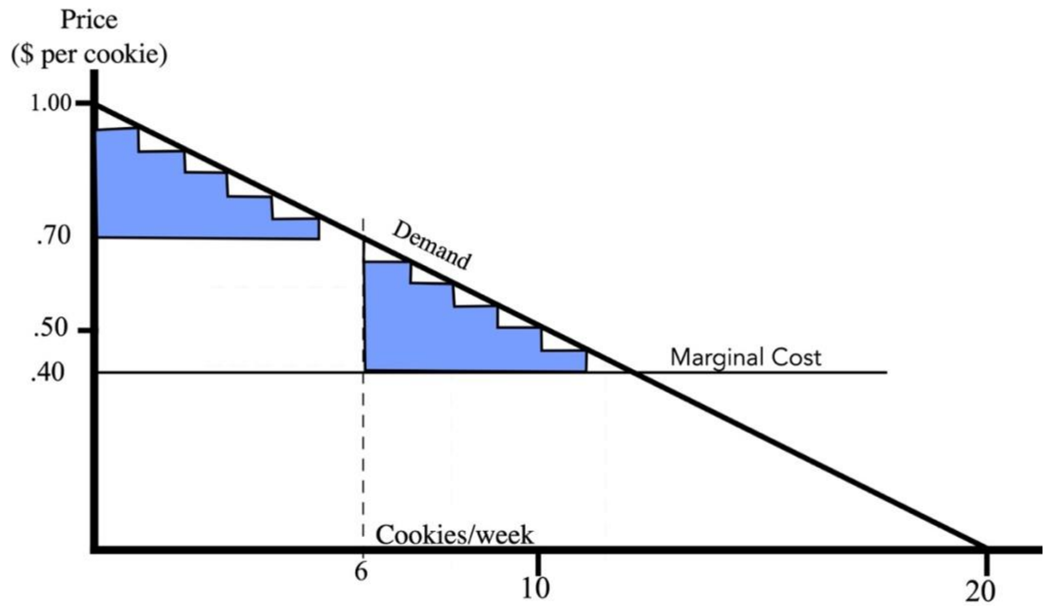

Figure 10-3a

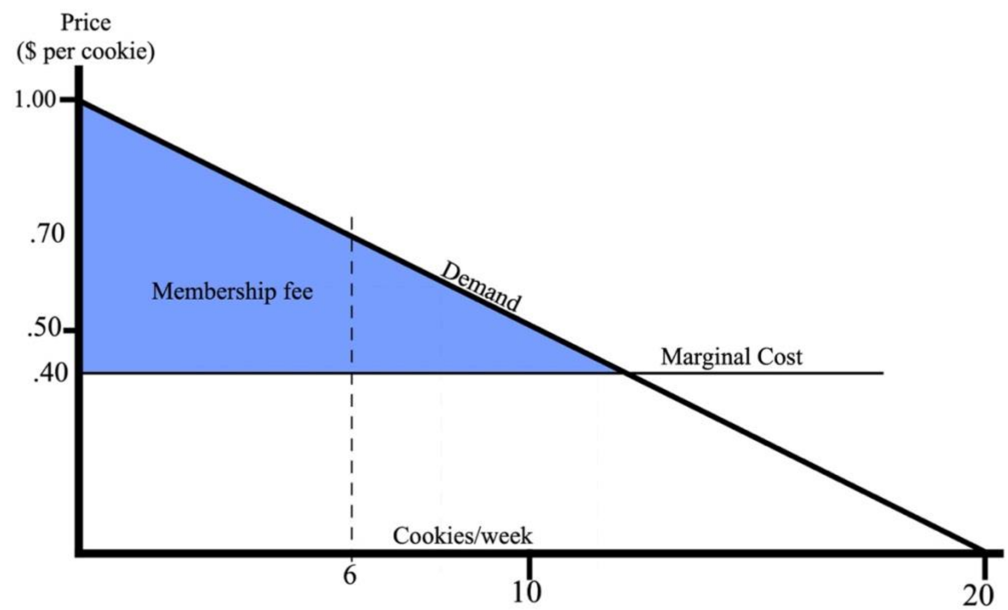

Figure 10-3b

Discriminatory pricing in the cookie industry — improved versions. On Figure 10-3a, cookies are sold on a sliding scale starting at $0.95/cookie. On Figure 10-3b, the price is $0.40/cookie, but cookies are only sold to customers who pay $3.60 for membership in the cookie club.

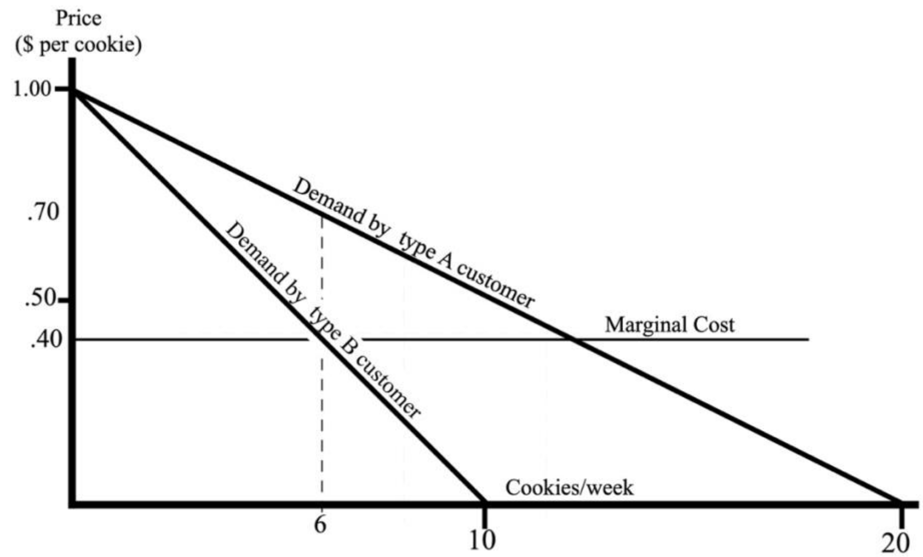

Figure 10-4

The case of nonidentical customers.

11: Hard Problems: Game Theory, Strategic Behavior, and Oligopoly

Table 11-1

		Joe: Confess	Joe: Say Nothing
Mike	Confess	2 years, 2 years	3 months, 5 years
Mike	Say Nothing	5 years, 3 months	6 months, 6 months

The payoff matrix for prisoner's dilemma: Each cell of the table shows the result of choices by the two prisoners; Joe’s sentence is first, Mike’s second

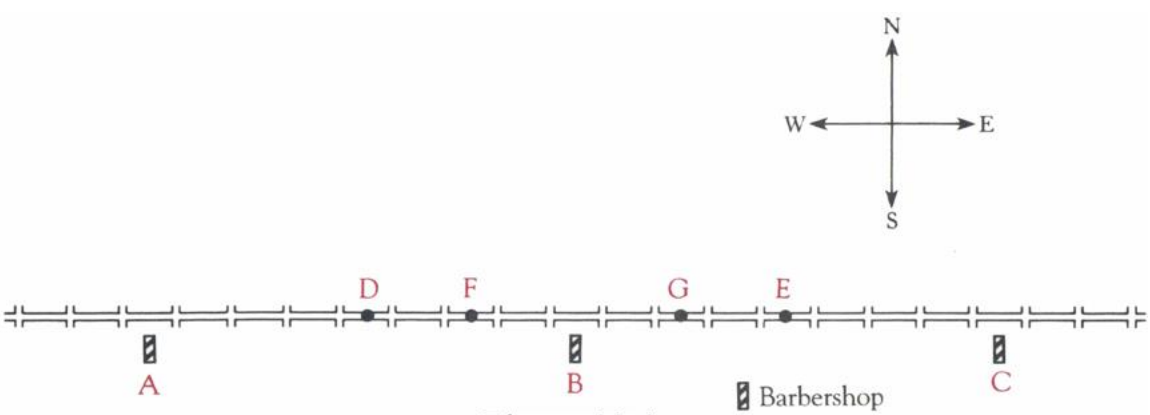

Figure 11-1

The street of barbers. There is one barbershop every eight blocks.

12: Time . . .

Table 12-1				
Year	Harvard wage	Present value of Harvard wage	Yale wage	Present value of Yale wage
1	$80,000	$80,000	$62,000	$62,000
2	$80,000	$72,727	$66,000	$60,000
3	$80,000	$66,116	$70,000	$57,851
4	$80,000	$60,105	$74,000	$55,597
5	$80,000	$54,641	$78,000	$53,275
6	$80,000	$49,674	$82,000	$50,916
7	$80,000	$45,158	$86,000	$48,545
8	$80,000	$41,053	$90,000	$46,184
9	$80,000	$37,321	$94,000	$43,852
10	$80,000	$33,928	$98,000	$41,562
Total		**$540,722**		**$519,781**

Comparing two job offers: Each is a stream of payments over time; take the present value of each payment and sum for each offer. Harvard wins.

For Further Reading

The analysis of depletable resources in this chapter is not a product of recent concerns with the problem, summarized in phrases (and book titles) such as "limits to growth" and "spaceship earth." It was produced more than eighty years ago by Harold Hotelling in "The economics of exhaustible resources," JPE 39, 137-75.

13: . . . And Chance

If the gambler gains $2 on heads but loses $1 on tails, we have:

$P_{\text{Heads}} = 0.5$; $R_{\text{Heads}} = + \$2$

$P_{\text{Tails}} = 0.5$; $R_{\text{Tails}} = - \$1$

Expected Return = $(p_{\text{Heads}} \text{ x } R_{\text{Heads}}) + (p_{\text{Tails}} \text{ x } R_{\text{Tails}}) =$

$[0.5 \text{ x } (+ \$2)] + [0.5 \text{ x } (- \$1)] = \$0.50.$

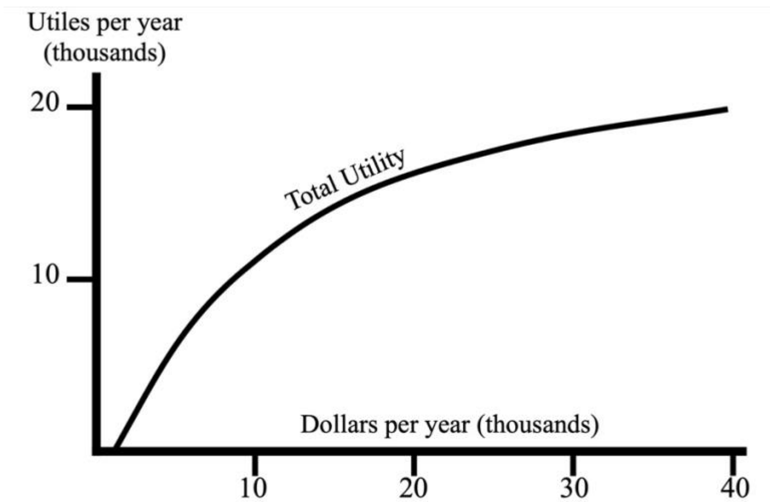

Figure 13-1a

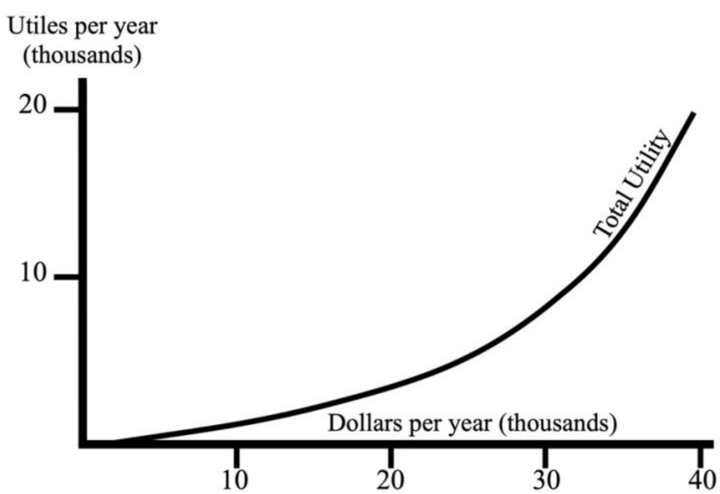

Figure 13-1b

Total utility of income for risk-averse (a) and risk preferring (b) individuals.

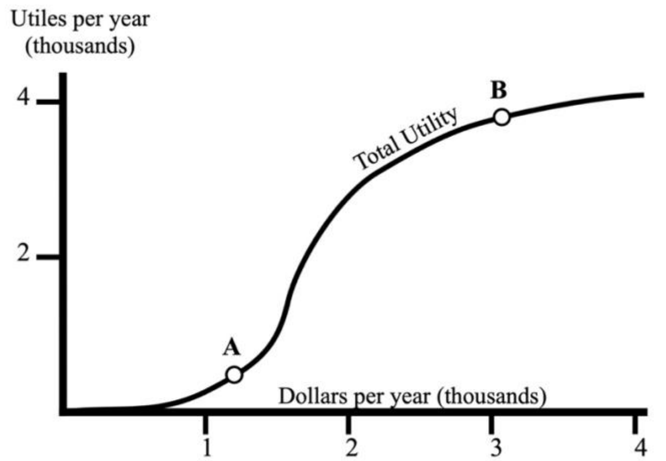

Figure 13-2

Utility function for someone whose marginal utility of income increases as his income approaches what he needs to survive then decreases beyond that.

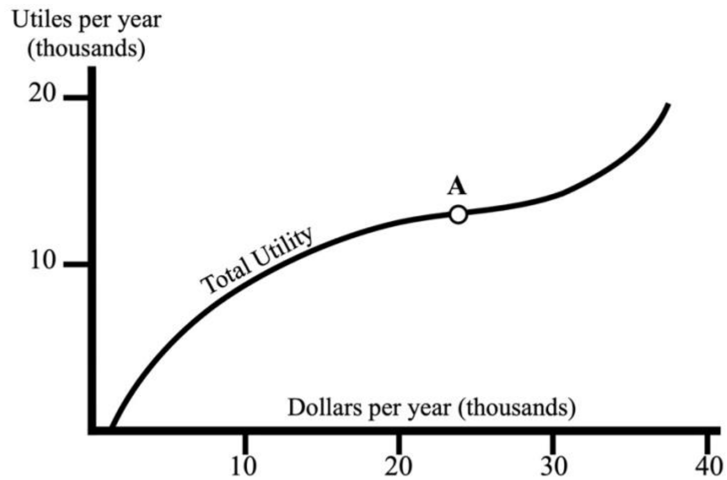

Figure 13-3

Utility function for someone whose is risk averse at low incomes, risk preferring at high. Starting at point A, he could could increase his expected utility by buying both insurance and lottery tickets.

For Further Reading

The original discussion of Von Neumann utility is in John Von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior* (Princeton: Princeton University Press, 1944), Chapter 1.

A classic discussion of the lottery-insurance paradox is Milton Friedman and Leonard J. Savage, "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy*, Vol. 56, No. 4 (August, 1948), pp. 279-304.

14: Who Gets How Much Why?

For Further Reading

The final section of this chapter is my rewrite of something first published in 1776. It can be found in Chapter X, Book I, of Adam Smith, *An Inquiry into the Nature and Causes of the Wealth of Nations*. The book is still well worth reading.

The most famous supporter of the idea of taxing the site value of land, Henry George, stated his argument in *Progress and Poverty* (New York: Robert Schalkenbach Foundation, 1984).

15: Summing People Up

For Further Reading

For an original, interesting, and readable discussion of the idea of economic improvement, see Alfred Marshall, *Principles of Economics*, (8th. ed.; London: Macmillan, 1920), Chapter VI.

Under most circumstances, the Marshallian and potential Paretian definitions of efficiency lead to the same conclusions. For circumstances under which they do not, see: David Friedman, "Does Altruism Produce Efficient Outcomes? Marshall vs Kaldor," *Journal of Legal Studies* Vol. XVII, (January 1988).

16: What is Efficient?

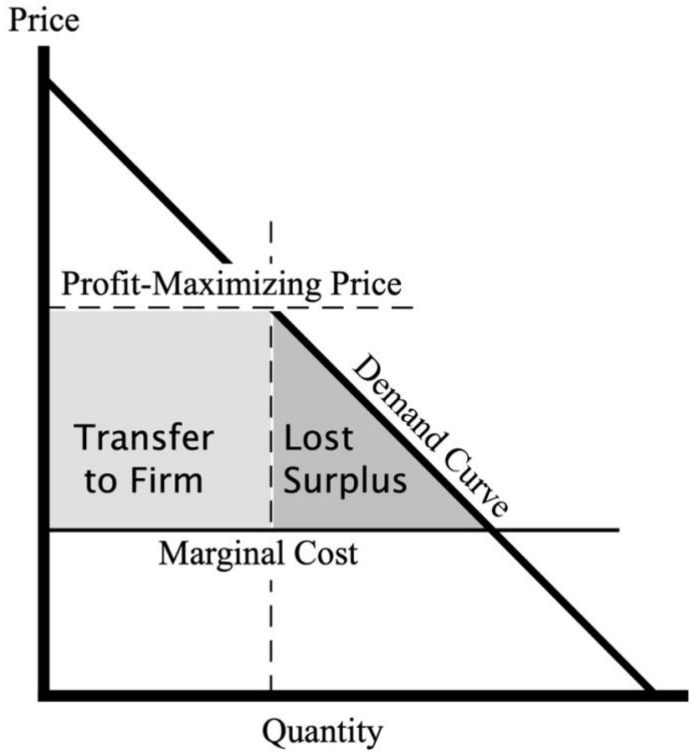

Figure 16-1

The Deadweight Cost of Monopoly: When a monopoly raises its price above marginal cost in order to maximize its profit, the increased price of the goods still sold is a transfer to the firm but the reduction in consumer surplus due to the reduction in output is a net loss.

Figure 16-1 shows the inefficiency in graphical form. The shaded region is the difference between consumer surplus at the monopoly's profit-maximizing price and consumer surplus at a price equal to marginal cost. Part of that is a transfer — a higher price paid to the firm by its customers for the goods they buy. The rest is the lost consumer surplus on units consumers would buy at marginal cost but do not buy at the price that maximizes monopoly profit; consumers lose that surplus but nobody gets it. This is the lost welfare triangle due to monopoly. It is the same triangle that the movie theater lost, back in Chapter 10, if it sold popcorn for $1.00/bag instead of $0.50/bag.

17: How to Gum Up the Works

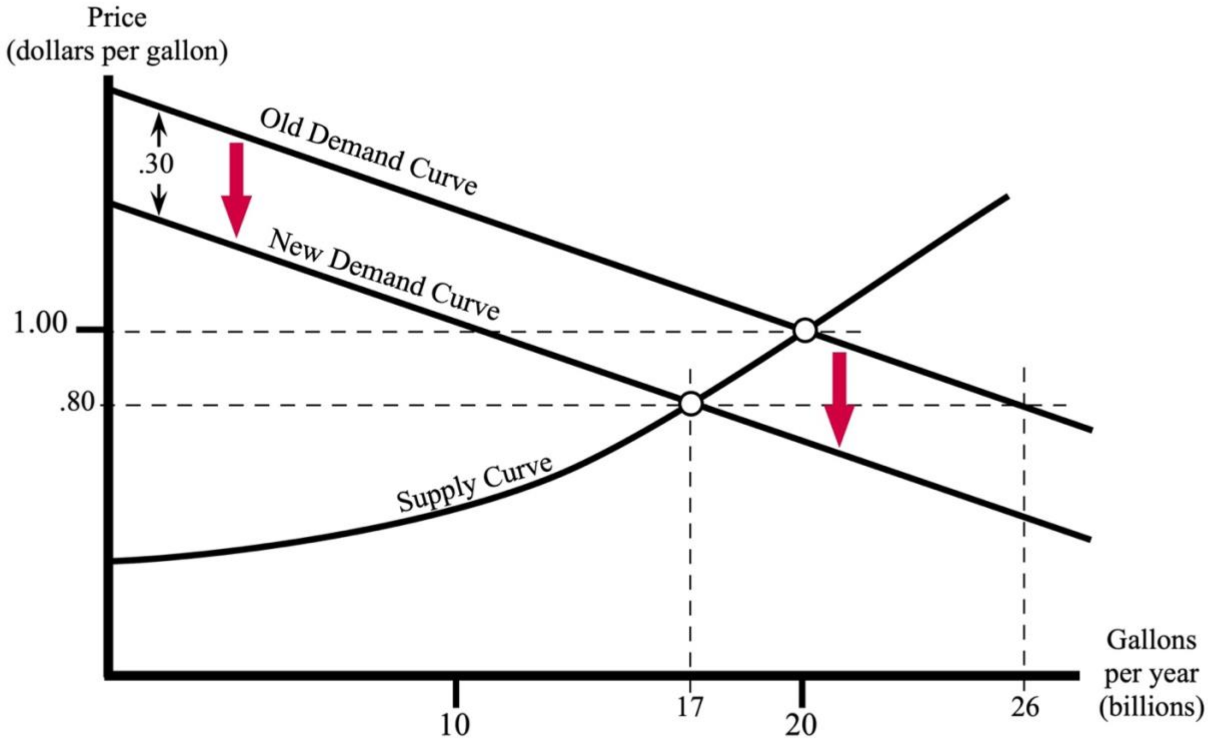

Figure 17-1

The effect of price control on gasoline. Price control at $0.80/gallon produces a shortage; quantity demanded is larger than quantity supplied. Lines grow until their cost shifts demand down far enough to make quantity demanded at the controlled price equal to quantity supplied. Consumers are paying $0.20/gallon less in money and $0.30/gallon more in time.

18: Why We Are Not All Happy, Wealthy, Wise, and Married

For Further Reading

For evidence that making cars safer increases the number of accidents, see Sam Peltzman, "The Effects of Automobile Safety Regulations," Journal of Political Economy, August 1975.

19: Law and Sausage: The Political Marketplace

The Geometric Proof. Figure 19-1 shows the supply curve for American production of automobiles and the demand curve for American consumption of automobiles, before and after the imposition of a tariff of $t per automobile. P_A is the market price before the tariff, P'_A after the tariff. Q_A is the quantity of imported cars before the tariff, Q'_A after. Figure 19-2 shows the corresponding curves, prices, and quantities for wheat.

The price at which U.S. quantity supplied equals U.S. quantity demanded is above the world market price, so the United States imports autos. Quantity demanded by U.S. consumers is equal to quantity supplied by the U.S. auto industry plus imports. The price at which quantity of wheat supplied and quantity demanded in the United States are equal is below the world price of wheat, so the United States exports wheat. Quantity produced by U.S. farmers equals quantity demanded by U.S. consumers plus exports.

Why does an auto tariff affect the price of wheat? Wheat is what we send foreigners in exchange for the autos they send us. When we impose an auto tariff, fewer dollars go abroad to buy foreign cars. Foreigners have fewer dollars with which to buy American wheat, so their demand falls and the price of wheat in America drops.

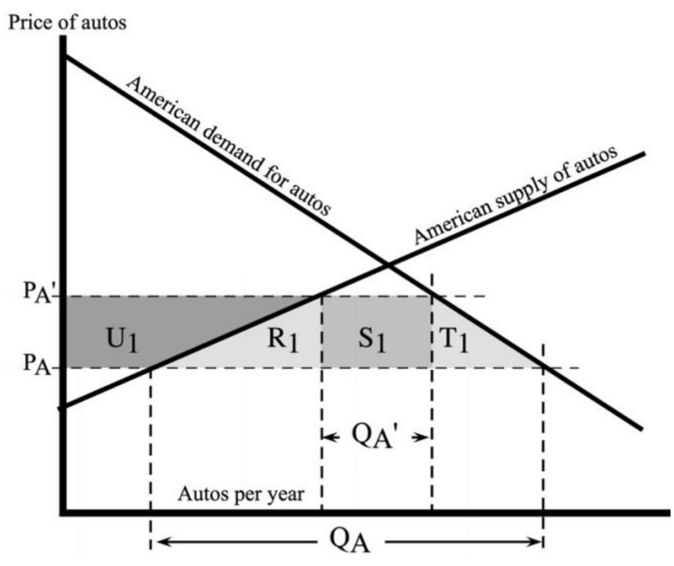

Figure 19-1

The effect on the domestic auto market of a tariff on autos. Imports of autos, equal to the difference between domestic demand and domestic supply, fall from Q_A to Q'_A.

The U.S. price rises from P_A to P'_A. U_1 on Figure 19-1 is the increase in American producer surplus as a result of the tariff; $U_1 + R_1 + S_1 + T_1$ is the reduction in American consumer surplus. The shaded area is the net loss to Americans of surplus on autos as a result of the tariff. Similarly, on Figure 19-2, U_2 is the gain in (American) consumer surplus as a result of the fall in the price of wheat produced by the tariff on automobiles; $U_2 + R_2 + S_2 + T_2$ is the loss of (American) producer surplus. The shaded area $R_2 + S_2 + T_2$ is the net loss to Americans of surplus on wheat as an indirect result of the tariff on autos.

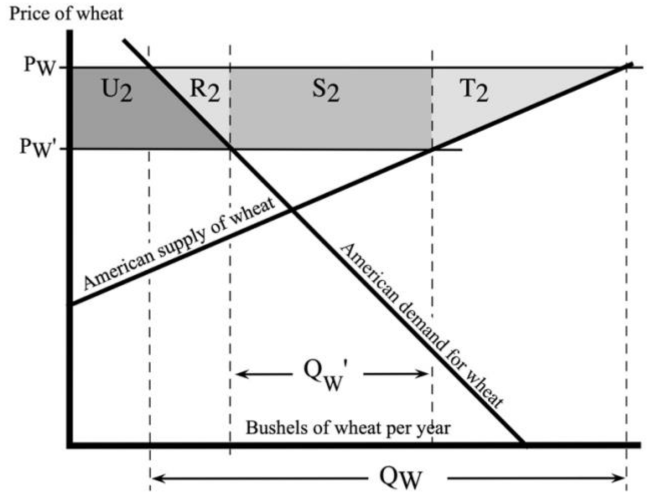

Figure 19-2

The effect on the domestic wheat market of a tariff on autos. Imports of wheat fall from Q_W to Q'_W; the U.S. price falls from P_W to P'_W.

The net loss in surplus must be weighed against the revenue from the tariff. The government collects t dollars on each of Q'_A autos imported each year, so its revenue from the tariff is t x Q'_A. If that is larger than the sum of the two shaded areas, then the tariff makes us, on net, better off — revenue collected is more than surplus lost. If it is smaller, the tariff makes us worse off.

Since America is a price taker in international markets, the tariff does not affect the relative prices of autos and wheat outside the United States. Before the tariff, the price ratio is P_A/P_W. After the tariff, the price of wheat abroad (in dollars) is P'_W, the price of autos abroad is P'_A - t, so the price ratio is (P'_A - t)/P'_W.

Why is the world price of autos P'_A - t? P'_A is the price of autos in the United States. In order to get foreign autos into the United States, you must pay their world price plus the tariff t; the price in the United States is P'_A, so the world price must be P'_A - t.

Since the price ratio outside the United States is the same before and after the tariff, it follows that:

$$P_A/P_W = (P'_A-t)/P'_W \qquad \text{(Equation 1)}$$

Autos are, by assumption, our only import and wheat our only export, so the total number of dollars foreigners get for the cars they sell to us must equal the number of dollars they spend for the wheat they buy from us. Using prices and quantities after the tariff is imposed, this gives us:

P'_W x Q'_W = \$'s spent on wheat by foreigners =

\$'s received for cars by foreigners = $(P'_A - t)Q'_A$ (Equation 2)

(We spend P'_A on each car, but since t goes to the government to pay the tariff, only P'_A - t goes to foreigners).

Finally, from Figures 19-1 and 19-2, we have:

$S_1 + S_2 = (P'_A - P_A)Q'_A + (P_W - P'_W)Q'_W$ (Equation 3)

Equations 1 and 2 imply that:

$Q'_W = Q'_A(P'_A\text{-}t)/P'_W = Q'_A(P_A/P_W)$

Substituting this into Equation 3 gives us:

$S_1 + S_2 = Q'_A(P'_A - P_A) + Q'_A (P_A/P_W)(P_W - P'_W) =$

$Q'_A\{P'_A - P_A + (P_A/P_W) (P_W - P'_W)\} = Q'_A\{P'_A - P_A + P_A - P'_W(P_A/P_W)\}$

Using Equation 1 to replace (P_A/P_W) with $(P'_A\text{-}t)/P'_W$, we get:

$S_1 + S_2 = Q'_A\{P'_A - P'_W(P'_A\text{-}t)/P'_W\} = Q'_A(P'_A - P'_A + t) = Q'_A \times t$ (Equation 4)

$S_1 + S_2$ is only part of the lost surplus due to the tariff; Q'_A x t is all of the revenue. The tariff costs us more than it brings in; on net it makes us worse off by $R_1+T_1+R_2+T_2$. And that is without taking account of the additional costs it imposes on our trading partners abroad.

For Further Reading

The article that first used the term “rent seeking” is: Anne Krueger, "The Political Economy of the Rent-seeking Society," *American Economic Review*, Vol. 64 (June, 1974), pp. 291-303, but the idea appeared earlier in: Gordon Tullock, "The Welfare Costs of Tariffs, Monopolies and Theft," *Western Economic Journal*, Vol. 5 (June, 1967), pp. 224-232 and David Friedman, *The Machinery of Freedom: Guide to a Radical Capitalism* (New York: Harper & Row, 1971; Arlington House 1978; Open Court, 1989), Chapter 38.

Gary Becker, "A Theory of Competition Among Pressure Groups for Political Influence," *Quarterly Journal of Economics*, Vol. 98 (1983), pp. 371-400 is an important source for the model of the political market described in this chapter. Other classics of public choice theory include: James Buchanan and Gordon Tullock, *The Calculus of Consent* (Ann Arbor, MI: University of Michigan Press, 1962) and Mancur Olson, *The Logic of Collective Action* (Cambridge, MA: Harvard University Press, 1965).

20: Rational Criminals and Intentional Accidents

The Economics of Law and Law Breaking

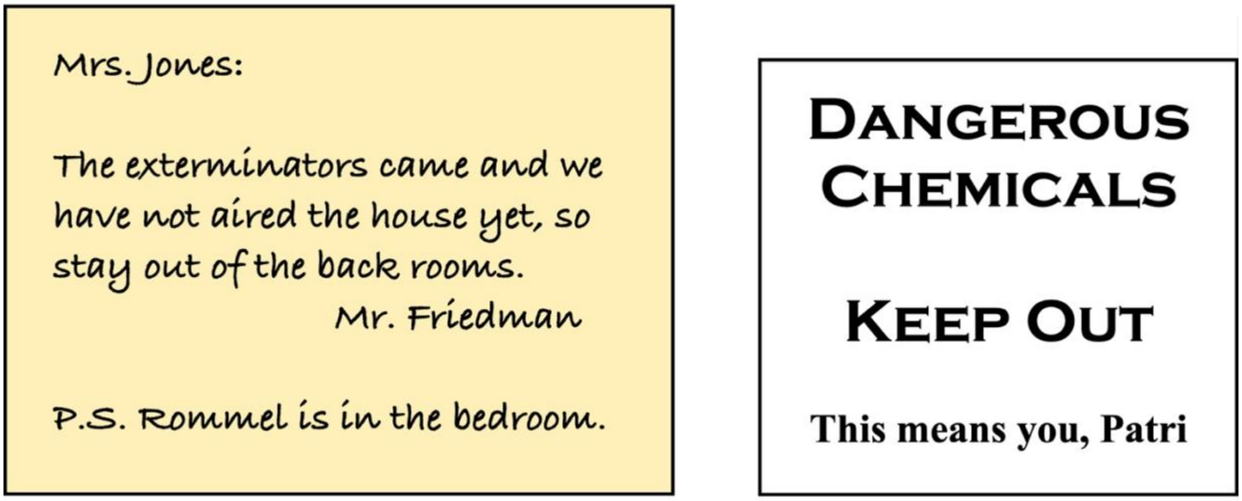

(a) (b)

Figure 20-1

Low-cost burglar repellents. Fictitious notes to a fictitious cleaning lady and a real son.

For Further Reading

My analysis of private enforcement is in "Efficient Institutions for the Private Enforcement of Law," *Journal of Legal Studies* (June, 1984). My book *The Machinery of Freedom* contains a discussion of how a fully private system of courts, police, and laws might work and a description of the Icelandic system. A later and more detailed account of the Icelandic system, along with a wide range of other legal systems, can be found in *Legal Systems Very Different from Ours*, written by me with additional chapters (on the legal systems of prison gangs and 18th Century pirates) by David Skarbek and Peter Leeson.

A more detailed account of optimal punishment is in my "Should the Characteristics of Victims and Criminals Count? *Payne v Tennessee* and Two Views of Efficient Punishment," *Boston College Law Review* XXXIV No. 4, pp.731-769 (July 1993). My book *Law's Order* and my essay "Economic Analysis of Law" in *The New Palgrave: A Dictionary of Economic Theory and Doctrine*, John Eatwell, Murray Milgate and Peter Newman, eds. (Macmillan, 1987) provide an overview of the subject and further references.

The Last Testament of Lucky Luciano, by Martin A. Gosch and Richard Hammer (Boston: Little, Brown: 1974), claims to be based on information given to Gosch by Luciano. *The Cocaine Kids: The Inside Story of a Teenage Drug Ring*, by Terry Williams (Addison Wesley 1989), provides a more recent view of an illegal market.

"Fact, fancy, and organized crime," by Peter Reuter and Jonathan B. Rubinstein, *The Public Interest* 53 (Fall 1978) pp. 45-67, provides evidence and arguments that support my view of organized crime, including the results of the study of bookmaking mentioned in this chapter.

Final Words

For Further Reading

My first economics article was "An Economic Theory of the Size and Shape of Nations," *Journal of Political Economy*, Vol. 85, No. 1 (February, 1977), pp. 59-77.

Students who would like to learn economics from its inventors should read Adam Smith, *An Inquiry into the Nature and Causes of the Wealth of Nations* (New York: Oxford University Press, 1976); David Ricardo, *The Principles of Political Economy and Taxation* (Totowa, N.J.: Biblio Distribution Centre, 1977); and Alfred Marshall, *The Principles of Economics* (8th ed., London: Macmillan, 1920).

The three books are very different. Smith's is the most far ranging and entertaining, Ricardo's the most difficult. Marshall's *Principles* is where modern economics was first put together; it is the only one of the three that could, for a sufficiently courageous reader, substitute with some advantage for a modern economics text.

Other Work by David Friedman

Non-fiction

The Machinery of Freedom: Guide to a Radical Capitalism

Price Theory: an Intermediate Text

Law's Order: What Economics Has to Do With Law and Why It Matters

Future Imperfect: Technology and Freedom in an Uncertain World

A Miscellany (with Elizabeth Cook)

How to Milk an Almond, Stuff an Egg, and Armor a Turnip: A Thousand Years of Recipes (with Elizabeth Cook)

Fiction

Harald

Salamander

Brothers

Web Page

Blog