

Chapter 19: Compensating and Equivalent Variations

19.1: Introduction

This chapter is interesting and important. It also helps to answer a question you may well have been asking ever since we studied quasi-linear preferences right at the beginning of the book. We saw then, that with quasi-linear preferences we have an exact monetary measure of how much better off is the individual if he or she moves from one indifference curve to another (because the indifference curves are parallel). As a consequence we saw that the area between the price paid and the demand curve (or the area between the price received and the supply curve) is an *exact measure* of the surplus or profit from trade.

You may have been asking: what happens if the preferences are not quasi-linear? The answer is that there is now no unique measure of how much better off an individual is. Therefore what exactly is this area between the price paid and the demand curve (or the area between the price received and the supply curve) measuring?

We will approach an answer rather obliquely. We do this by asking the question: suppose the price of some good changes, how can we measure how much better off or worse off is the individual as a consequence? We will see that there are at least two answers to this question - unless preferences are quasi-linear – and this leads us to realise that the search for a unique measure of ‘how much better off’ is bound to be fruitless.

19.2: The Effect on Behaviour of a Price Change

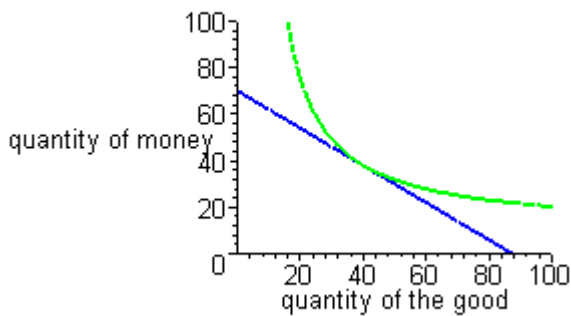
We start with simply looking to see the effect on behaviour of a price change. Then we look to see how we might get a monetary measure of the welfare effects of this price change. Here we take a price increase, but we could do the whole analysis with a price decrease (though we have to be slightly careful about the interpretations).

We return to the space we used in chapters 3 and 4. We are interested in some good – we therefore put the quantity of it on the horizontal axis. Having purchased the good the individual has money left over to spend on other goods – we put the quantity of money to spend on other goods on the vertical axis. We assume that the individual has preferences over (the quantity of the good, the quantity of money) which are representable by some indifference map. To begin with let us assume that the indifference curves are smoothly convex – though later we shall consider other cases.

We denote the price of good 1 by p . Rather naturally we take the price of the quantity of money to spend on other goods to be 1. We take the monetary income of the individual to be m . In the example that follows we take m to be 70.

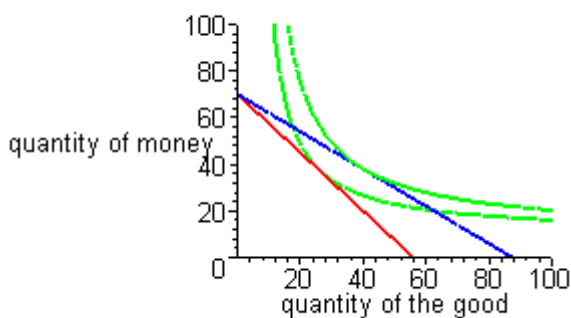
We start with the following position in which the price of the good is initially 0.8. So the budget line goes from 87.5 (m/p) on the horizontal axis to 70 (m) on the vertical axis. In figure 19.1 we draw in this budget line, the highest attainable budget line and the optimal point.

19.1: before the price rise



We now suppose that the price of the good *rises* – from 0.8 to 1.25. With the new price the budget line goes from 56 (m/p) on the horizontal axis to 70 (m) on the vertical. In figure 19.2 is shown the new budget line, the new highest attainable budget line and the new optimal point. Clearly the individual changes his or her behaviour as a consequence of the price rise – in this case he buys less of the good and has less money left over to spend on other goods.

19.2: after the price rise



Let us now look in more detail at what has happened. Look at the budget line – what has happened to it? As you can see from figure 19.2, the budget line has *rotated* around the point on the vertical axis – specifically the point m on the vertical axis. So we can say (1) that the *slope* of the budget line has changed; and (2) that the size of the budget constraint has been reduced – from the triangle $[(0, 0), (87.5, 0), (0, 70)]$ to the triangle $[(0, 0), (56, 0), (0, 70)]$. In other words:

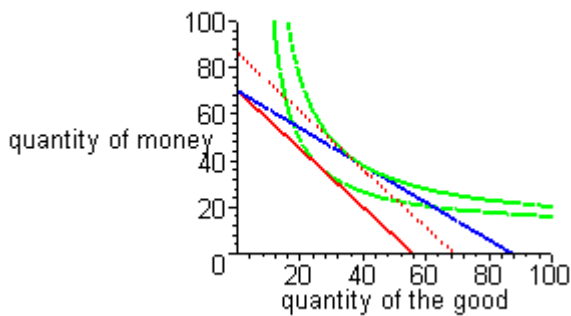
- 1) the relative price of the good has changed;
- 2) the individual is worse off – in a sense the value of his income has been reduced.

We call the first of these the *relative price effect*¹ and the second one the *income effect*.

We will find it extremely useful to separate out these two effects. At the moment they are confounded – the movement from the original budget line to the new budget line contains both effects – there is both a change in the slope and a change in the welfare of the individual. As we will see it is extremely useful if we can separate out these two effects. We do this by introducing an imaginary intermediary budget constraint – which is illustrated in figure 19.3. It is very carefully positioned.

¹ Some economists use the expression ‘the substitution effect’.

19.3: the compensating variation



This imaginary intermediate budget constraint is the dashed line in figure 19.3. Its position is determined by two considerations. First it is *parallel* to the *new* budget line and so reflects the new price. Second, it is *tangential*² to the *original* indifference curve. This implies that the individual must be indifferent between the original budget constraint and the intermediate budget constraint. Why? Because with either budget constraint the individual can reach the same indifference curve and therefore is just as well off.

We can therefore say that in moving from the original budget constraint to the intermediate budget constraint the individual is no better or worse off – his or her *real income* is the same in both situations.

We are now ready to separate out the two effects. We can consider the move from the original to the new budget constraints as equivalent (in terms of the final effect) to the following two moves:

- 1) a move from the original budget constraint to the intermediate budget constraint;
- 2) a move from the intermediate budget constraint to the new budget constraint.

Why do we do this? Well, consider the first of these moves – from the original to the intermediate budget constraint. In this the budget constraint *rotates around the original indifference curve*. The individual stays equally well off as he or she was originally – but the slope has changed from the original slope to the new slope. This move can be considered as solely the *relative price effect* – there is no change in real income involved. Now consider the second of these two moves – from the intermediate to the new budget constraint – the budget line has the same slope before and after the move, but it has shifted *parallel* to itself. This move can be considered as solely the *income effect* – there is no change in price involved.

To summarise – the move from the original to the new budget constraint can be decomposed into two moves – from the original to the intermediate and from the intermediate to the new. The first of these captures the *relative price effect* and the second the *income effect*.

We can see from figure 19.3 that the effect on the optimal behaviour of the individual of the move from the original to the intermediate budget constraint is *unambiguous* – the individual buys less of the good and has more money left over to spend on other goods. More importantly you will see that it is bound to be unambiguous – as it is the consequence of a *rotation* of the budget constraint *around the original indifference curve*. This cannot possibly increase the consumption of the good

² By which we mean that the intermediate budget constraint is a tangent to the original indifference curve.

and will decrease it if the indifference curve is smoothly convex³. So the *relative price effect* in unambiguous.

The income effect on the other hand may be ambiguous – though generally we expect that for most goods that we consider normal that as income falls so does the demand for the good. In this case it does. But there are cases when the demand for the good rises when the income falls. Such a good is referred to as an *inferior* good. Real life examples are few.

19.3: The Effect on Welfare of the Price Change

So far we have considered the effect on behaviour of the price change. What about the effect on the welfare of the individual? We can easily see that the individual is worse off than originally – but how much worse off? There is one way we can get an answer to this question – by simply asking “how much money do we need to give to the individual to *compensate* him or her for this price rise?”. You might be able to answer this question. If we give money to the individual to compensate him or her, we move the budget constraint parallel to itself (parallel to the new budget constraint) upwards and outwards. How far do we have to move it? *Until it becomes tangential to the original indifference curve and the individual is just as well off as he or she was originally.*

So, given the price rise, and given the fact that the budget constraint has shifted to the new position, to compensate the individual, to take the individual back to where he or she was originally, we need to give to the individual enough money to shift the budget constraint back parallel to itself until it is tangential to the original indifference curve. That is, we have to move the budget constraint from the new position to the intermediate position that we have so carefully constructed.

Now if you look carefully at figure 19.3 above you will see that the intermediate budget constraint has an intercept of a little over 86 on the vertical axis. Originally the income was 70. So this is telling us that if we give to the individual a little over 16 in money ($= 86 - 70$) this will compensate the individual for the effects of the price rise. This is a monetary measure of the welfare effects of the price rise⁴. It is termed the *compensating variation*. It tells us how much money should be given to the individual to compensate him or her for the price rise.

[An aside is necessary at this stage. If the price *falls* we can do all the above analysis. But note that the individual is better off. In this case the compensating variation is *negative* – the individual needs to give away money to compensate for the fact that he or she is better off than before.]

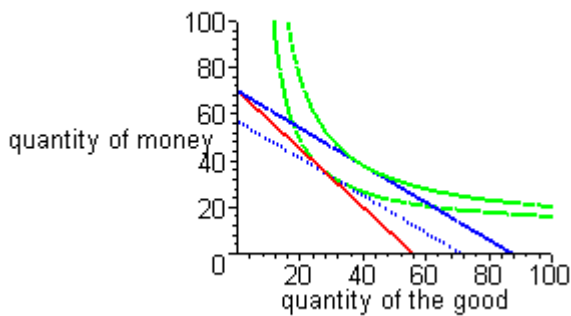
19.4: An Alternative Decomposition

If you are really alert you may have noticed that the way we decomposed the overall effect of the price change into a price effect and an income effect was a little arbitrary. You may have realised that there is another way. Consider figure 19.4 and compare it with figure 19.3.

³ It is possible that there is no change at all in the quantity purchased of the good – if, for example, the preferences are perfect complements rather than smoothly convex – but there is no way that the quantity purchased can increase.

⁴ Recall that the variable on the vertical axis is money and that its price is 1.

19.4: the equivalent variation



Again we have drawn in an imaginary intermediate budget constraint – the dashed line in figure 19.4. Notice how it is constructed. First it is *parallel* to the *original* budget constraint and therefore reflects the original price. Second, it is *tangential* to the *new* indifference curve. This implies that the individual must be indifferent between the intermediate budget constraint and the new budget constraint. Why? Because with either budget constraint the individual can reach the same indifference curve and therefore is just as well off.

We can therefore say that in moving from the intermediate budget constraint to the new budget constraint the individual is no better or worse off – his or her *real income* is the same in both situations.

We are now ready once again to separate out the two effects. We can consider the move from the original to the new budget constraints as equivalent (in terms of the final effect) to the following two moves:

- 1) a move from the original budget constraint to the intermediate budget constraint;
- 2) a move from the intermediate budget constraint to the new budget constraint.

Why do we do this? The same reason as before. Consider the first of these moves – from the original to the intermediate budget constraint. In this the budget constraint *moves parallel to the original budget constraint*. This move can be considered as solely an *income effect* – there is no change in price involved. Now consider the second of these two moves – from the intermediate to the new budget constraint – the budget line has *rotated around the new indifference curve*. The individual stays equally well off as he or she is with the new price – but the slope has changed from the original slope to the new slope. This move can be considered as solely the *relative price effect* – there is no change in real income involved.

Once again we see that the relative price effect is unambiguous – because of the rotation *around* the new indifference curve, the quantity purchased of the good must normally decrease⁵. The income effect on the other hand may be ambiguous – though generally we expect that, for most goods that we consider normal (as distinct from inferior as define above), as income falls so does the demand for the good. In this case it does.

Once again, we have considered the effect on behaviour of the price change. What about the effect on the welfare of the individual? We know that that the individual is worse off than originally – but how much worse off? There is one way we can answer this question – by simply asking “how much money would we have to take away from the individual at the original price to have the

⁵ It is possible that there is no change at all in the quantity purchased of the good – if, for example, the preferences are perfect complements rather than smoothly convex – but there is no way that the quantity purchased can increase.

equivalent effect on his or her welfare?”. You might be able to answer this question. If we take money away from the individual at the original price we move the budget constraint parallel to itself (parallel to the original budget constraint) downwards and leftwards. How far do we have to move it? Until it becomes tangential to the new indifference curve and the individual is just as well off as he or she is in the new position.

So, given the price rise, and given the fact that the budget constraint has shifted to the new position, to return to the original price and instead take enough money away from the individual so that he or she ends up with the same welfare as in the new position, we need to take away enough money so that the budget constraint ends up parallel to the original budget constraint and tangential to the new indifference curve. That is, we have to move the budget constraint from the original position to the intermediate position that we have so carefully constructed in this section.

Now if you look carefully at figure 19.4 above you will see that the intermediate budget constraint has an intercept of 57 on the vertical axis. Originally the income was 70. So this is telling us that if we take away from the individual 13 in money ($= 70 - 57$) at the original price this has the *equivalent* effect on his or her welfare as the price rise. This is a monetary measure of the welfare effects of the price rise⁶. It is termed the *equivalent variation*. It tells us how much money should be taken away from the individual at the original price to have the *equivalent* effect on his or her welfare as the price rise.

[Again an aside is necessary at this stage. If the price *falls* we can do all the above analysis. But note that the individual is better off. In this case the equivalent variation is *negative* – the individual needs to be given money at the original price to have the same effect on his or her welfare as the price fall.]

So we have two monetary measures of the welfare effect of the price rise: the compensating variation and the equivalent variation. The first tells us how much we should give to the individual at the new price to compensate him or her for the price rise; the second tells us how much money we should take away from the individual at the original price to have the same effect on his or her welfare. You will see in this example that these two measures are different: the compensating variation is 16 and the equivalent variation is 13. The reason for this is that they are measuring things from different perspectives: the compensating variation from the perspective of the new price and the equivalent variation from the perspective of the original price. The compensating variation is bigger than the equivalent variation because the price is lower with the original price – when the individual is better off. Note carefully in this example that demand increases with income – when the individual is better off he or she buys more of the good. That is why the compensating variation is larger than the equivalent variation.

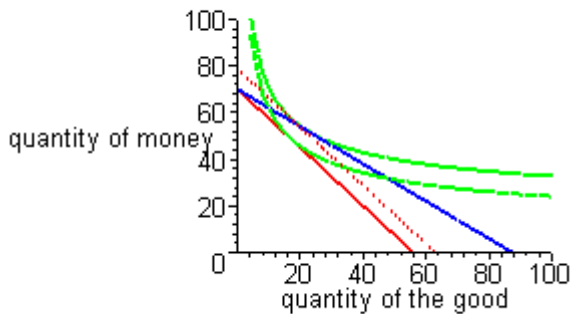
You may be able to reason that this is a property of the preferences that we have assumed. Indeed if we had taken preferences for which the demand for the good *falls* as income rises, then the compensating variation would be smaller than the equivalent variation. Moreover, if the quantity purchased is *independent* of the income we would expect that the two variations would be the same.

19.5: Quasi-Linear Preferences

But we know a case when the quantity purchased is independent of income – when the preferences are quasi-linear. In this case we have for the compensating variation figure 19.6.

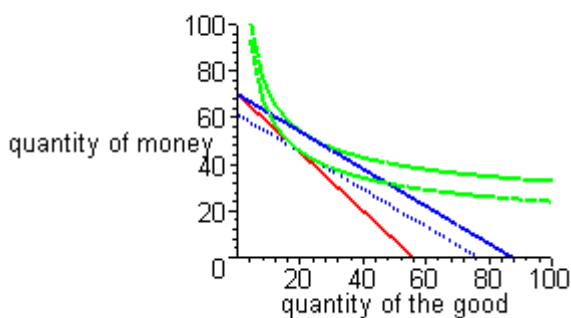
⁶ Recall again that the variable on the vertical axis is money and that its price is 1.

19.6: the compensating variation



Note that the two indifference curves are parallel in a vertical direction. We see from this figure that the compensating variation is 8.5 ($= 78.5 - 70$, the intercept on the vertical axis of the intermediate budget constraint minus the intercept of the original budget constraint). The equivalent variation is shown in figure 19.7.

19.7: the equivalent variation

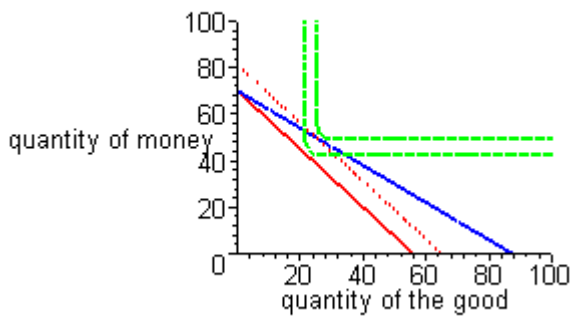


Notice once again that the indifference curves are parallel. We see from this figure that the compensating variation is 8.5 ($= 70 - 61.5$, the intercept on the vertical axis of the original budget constraint minus the intercept of the intermediate budget constraint). The compensating and equivalent variations are equal – a consequence of the fact that the indifference curves are parallel. (Study and compare figures 19.6 and 19.7 to make sure that you see this.)

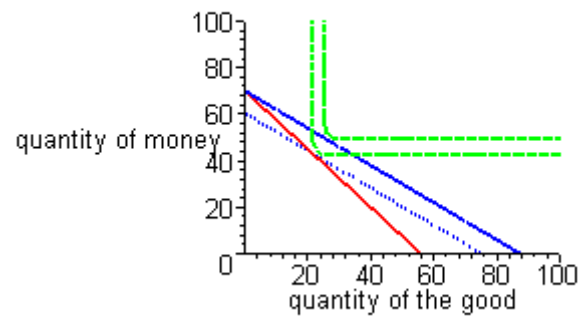
19.6: Perfect Complement and Perfect Substitute Preferences

It is of interest to look at some special cases – if only to see when the two effects are important and when they are not. Let us start with *perfect complements* – here taking the *1-with-2* case. The compensating variation is shown in figure 19.8 and the equivalent variation in figure 19.9.

19.8: the compensating variation



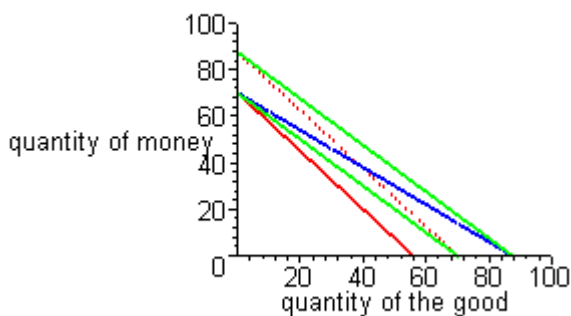
19.9: the equivalent variation



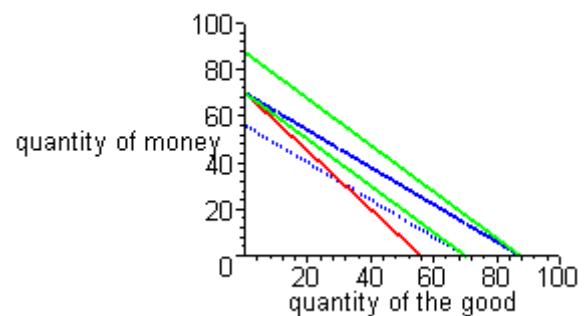
You will see that in both these cases there is no relative price effect – in the sense that the rotation of the budget line around either the original or the new indifference curve has no effect. But in a sense we could have anticipated this – with perfect complements there is no possibility for substitution. You will also notice that, once again, the compensating variation (about 11.5) is larger than the equivalent variation (about 9.5) – because the demand for the good increases with income.

In the case of *perfect 1:1 substitutes* we have the compensating variation in figure 19.10 and the equivalent variation in figure 19.11.

19.10: the compensating variation



19.11: the equivalent variation



Here the compensating variation is almost 18 and the equivalent variation 14 – once again the latter is smaller because of the income effect on the demand for the good. We should perhaps describe the two figures. In figure 19.10, from the intercept of 70 (the original income) on the horizontal axis, we have 3 lines: the lowest is the new budget constraint; the next the new indifference curve; and the top the original budget line. Then there are two further lines – the next one moving vertically is the intermediate budget line and the top one the original indifference curve. You will notice that both on the original budget line and the intermediate budget line the optimal point is on the original indifference curve. You will also notice that the intermediate budget line is parallel to the new budget line. In figure 19.11, from the intercept of 70 (the original income) on the vertical axis, we have 3 lines: the lowest is the new budget line; the next the new indifference curve; and the top the original budget line (just as in figure 19.10). Then below these three lines is the intermediate budget line, while above the three lines is the original indifference curve. You will notice that both on the intermediate budget line and the new budget line the optimal point is on the new indifference curve. You will also notice that the intermediate budget line is parallel to the original budget line. Finally note that the only thing that differs between the two figures is the position of the intermediate budget line.

19.7: Consumer Surplus and its Relationship with Compensating and Equivalent Variations

You might well be asking what is the relationship between these two measures of the effect of a price change on welfare and the measure with which we started this book – the change in the surplus. We have argued that the surplus measures the gain or surplus from trading at a particular price. So if the price changes, and hence the surplus changes, then presumably the change in the surplus measures the welfare effect on the individual? Actually we have proved this is true in the case of quasi-linear preferences but not in other cases. In this chapter we have shown for quasi-linear preferences that the compensating and equivalent variations are equal. It is not too much of a jump to realise that for the case of quasi-linear preferences these two variations are not only equal, but are also equal to the change in consumer surplus.

What about other preferences? Well we already know that the compensating variation is not equal to the equivalent variation and is generally larger than it. It can be shown (though using techniques outside this book) that in general the change in the surplus is *between* the compensating and equivalent variations. We shall show this in a particular example. But we shall start with some general definitions that will prove useful later.

We introduce a new utility function. We already have one that represents the preferences of the individual. It is denoted as follows:

$$U(q_1, q_2)$$

To distinguish this from the new utility function that we will be introducing, we shall call it the *direct* utility function, as it measures the utility that the individual gets from *directly* from consuming quantities of the two goods.

The individual is presumed to choose the optimal consumption bundle given the budget constraint. This leads to *demand functions* for the two goods which we write as $q_1 = f_1(m, p_1, p_2)$ and $q_2 = f_2(m, p_1, p_2)$. Note that these demands are functions of the income of the individual m and the prices of the two goods, p_1 and p_2 . If we now substitute these demand functions back in the direct utility function we get a new utility function, called the *indirect utility function*, as follows:

$$V(m, p_1, p_2) = U(q_1, q_2) = U(f_1(m, p_1, p_2), f_2(m, p_1, p_2))$$

What does this tell us? It tells us the utility of the individual when faced with prices p_1 and p_2 for the two goods and having income m , and when the individual chooses the optimal demands. In other words it tells us the *maximum* utility of the individual when faced with prices p_1 and p_2 for the two goods and having income m . It is called the *indirect* utility function of the individual, because the individual does not get utility directly from money and prices – but indirectly from the goods that they buy. Notice the arguments of $V(\cdot)$ – the first is p_1 , the second p_2 and the third m . The function is obviously increasing in m and decreasing in p_1 and p_2 .

We can use this indirect utility function to tell us the utility that the individual gets at any given income and prices. We can thus use this to calculate the compensating and equivalent variations. Suppose the individual starts with income m facing prices p_1 , and p_2 . The individual gets utility $V(m, p_1, p_2)$. Suppose now the price of good 1 rises to P_1 . The individual's utility is now $V(m, P_1, p_2)$. As we have assumed that P_1 is greater than p_1 , it follows that $V(m, P_1, p_2)$ is less than $V(m, p_1, p_2)$. To calculate the *compensating variation* we need to find the increase in income necessary to restore the utility to its original level. If we denote the compensating variation by cv it is defined by:

$$V(m+cv, P_1, p_2) = V(m, p_1, p_2)$$

So with the compensation the individual has the same level of utility at the new price as he or she originally had at the old price. Similarly we can define the *equivalent variation* as follows:

$$V(m, P_1, p_2) = V(m-ev, p_1, p_2)$$

Reducing the income by ev at the old prices reduces the individual's utility to that implied by the new prices.

Now let us give a particular example. This is in fact the same example as we used in the sections above. There the preferences were assumed to be symmetric Stone-Geary with subsistence levels of the good equal to 5 and the subsistence level of money equal to 10. From Chapter 5 we have the direct utility function

$$U(q_1, q_2) = (q_1 - 5)^{0.5} (q_2 - 10)^{0.5}$$

and from equation (6.7) of Chapter 6 we have the optimal demands, where we denote the price of the good by p and the price of money equal to 1:

$$q_1 = 5 + (m - 5p - 10)/2p \text{ and } q_2 = 10 + (m - 5p - 10)/2$$

If we substitute these demands back into the direct utility function we get the indirect utility function:

$$V(m, p) = (m - 5p - 10)/2p^{0.5} \quad (19.1)$$

where we have suppressed the p_2 argument of the function because it is assumed to be constant.

Equation (19.1) is the indirect utility function of the individual. It tells us the *maximum* utility of the individual for any price p and income m . We can use this to calculate the compensating and equivalent variations of any price change.

We use the example considered in the text. We start with an income $m = 70$ and with price $p = 0.8$. Then the individual's original utility level is $V(70, 0.8)$ which we can calculate using (19.1) to be equal to 31.30495.

Now suppose, as in the sections above, that there is a change in the price of the good – from 0.8 to 1.25. So p rises from 0.8 to 1.25. This causes a fall in the level of utility from $V(70, 0.8)$ to $V(70, 1.25)$ which we can calculate using (19.1) to be equal to 24.0377. As a consequence of the price rise, utility falls from 31.30495 to 24.0377. We note that these numbers are meaningless as they depend upon the particular utility representation we have chosen – if we change the representation we change the numbers. However we can get some numbers that mean something if we calculate the compensating and equivalent variations.

Let us start with the first. Let us denote the compensating variation by cv . We know that this is the amount of money that we should give to the individual to compensate him or her for the price rise. So we know that the utility of the individual at the new price but with his or her income increased by cv should give exactly the same utility as the individual had originally. Now we know that this latter is 31.30495. So cv satisfies the equation

$$V(70+cv, 1.25) = 31.30495$$

That is, from (19.1)

$$(70+cv - 5 \times 1.25 - 10)/2(1.25)^{0.5} = 31.30495$$

If we solve this equation for cv we find that $cv = 16.25$. The compensation variation is 16.25 – if, at the new price, we increase the individual's income by 16.25 then we return the individual back to his or her original level of utility. Note that this is almost equal to the approximate value (a little over 16) we found from the graphical analysis of section 19.3.

Let us denote the equivalent variation by ev . The equivalent variation is the amount of money that we should take away from the individual at the original price to have the same effect on his or her utility as the price rise. So we know that the utility of the individual at the original price but with his or her income decreased by ev should give exactly the same utility as the individual has in the new situation. Now we know that this latter is 24.0377. So ev satisfies the equation

$$V(70 - ev, 0.8) = 24.0377$$

That is

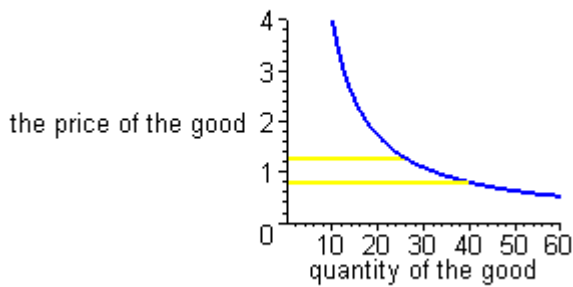
$$(70-ev - 5 \times 0.8 - 10)/2(0.8)^{0.5} = 24.0377$$

If we solve this equation for ev we find that $ev = 13$. The equivalent variation is 13 – if, at the original price, we decrease the individual's income by 13 then there is the same effect on his welfare as the price rise. This is exactly the same figure as the one we got graphically in section 19.4.

So far so good – we have calculated the compensating and equivalent variations and have shown that the former is larger than the latter. There is actually quite a big difference between them – but that is because the price rise that we have considered is quite large.

What about the change in the surplus? Well we know that the original surplus is the area between the price of 0.8 and the demand curve for the good while the new surplus is the area between the price of 1.25 and the demand curve for the good. So the change in the surplus is the area between prices 0.8 and 1.25 and the demand curve. See figure 19.3, which draws the demand curve for the good, which we know is given by $q_I = 5 + (70 - 5p - 10)/2p$. (Note that the demand at a price of 0.8 is 40 and the demand at a price of 1.25 is 27.5.)

19.14: the loss in consumer surplus



The loss of surplus is the area between the two horizontal lines (at prices 0.8 and 1.25) and the demand curve. Now the demand at a price of 0.8 is 40 and the demand at a price of 1.25 is 27.5 – so the area we need to calculate is somewhat less than a trapezium of height 0.45, a base of 40 and a top of 27.5. This trapezium has area 15.1875. So the change in surplus is somewhat less than this. Of course we can always calculate this area precisely – either graphically or using calculus. With the latter we see that the area is equal to the integral, with respect to p , of $2.5 + 30/p$ (the demand) between 0.8 and 1.25. This is equal to $2.5p + 30 \ln(p)$ evaluated between 0.8 and 1.25. This is equal to 14.52. The loss in the surplus is 14.52.

To summarise we have for the effect of this price rise:

the compensating variation = 16.25

the loss in the surplus = 14.52

the equivalent variation = 13.00

We see that the change in the surplus is *between* the compensating and equivalent variations. This is a general result⁷ that suggests that the change in the surplus might be a good compromise measure of the welfare effect of the price change. This is one reason that we have been advocating it so strongly. A second reason is that it is easy to calculate if we know the demand curve. In practice we often have a good estimate – and we usually know more about demand than we do about preferences (if only for the reason that we can observe demand but not preferences).

19.8: Another Way of Doing the Decompositions

We should note before finishing this chapter that you will find that other authors do the two decompositions we have done in a different way⁸. In essence, when finding the intermediate budget constraint, they consider it a rotation, either around the original *optimal point* or around the new *optimal point*, rather than around the original indifference curve or around the new indifference curve. I prefer not to do this decomposition. This is for two reasons: when we rotate the budget line around the original optimal point, rather than around the original indifference curve, it is not true to say that the individual remains just as well off as before. Similarly, when we rotate the budget line around the new optimal point, rather than around the new indifference curve, it is not true to say that the individual remains just as well off as in the new position. In fact a rotation around the

⁷ Which is too difficult for this book to prove.

⁸ The way we have done the decomposition is usually attributed to the economist Hicks and the new way we are describing in this section to the economist Slutsky.

optimal point is bound to make the individual better off⁹. I would call this an increase in his or her real income.

The second reason is that we recover two different monetary measures of the welfare effect of the price change from this decomposition. The appropriately equivalent compensating variation would tell us the amount of money that we need to give to the individual to enable him to purchase the original quantities of the goods. (Not that he or she would do given that amount of money – he or she would buy different amounts and be better off than originally.) Similarly the appropriately equivalent variation would tell us the amount of money that we should take away from the individual at the original price so that he or she could buy the new quantities of the goods. (Not that he or she would do - he or she would buy more and be better off than in the new situation.)

So there seem to be good intellectual reasons against this alternative decomposition. There may be a practical reason in its favour – it may be easier to use. But given that we are recommending that we use something even easier - the change in the surplus – this reason is not too overwhelming.

19.9: Summary

This chapter has done a lot. In particular we have clarified how we might get a monetary measure of a welfare effect of a price change with non-quasi-linear preferences. We also looked at the effect of the price change on behaviour.

The effect of a price change on behaviour can be decomposed into a part that is due to the change in relative prices of the goods and a part that is due to the change in the real income of the individual.

We saw that we can do this decomposition in two different ways. This lead us to realise that we can measure the welfare effect of a price change in two ways.

..the compensating variation which is the amount of money that the individual would need to be paid at the new price to compensate him or her for the adverse effects of the price change;

..the equivalent variation which is the amount of money that needs to be taken away at the original price to reduce the individual's welfare by the same amount as the price rise.

In general if the demand for the good increases with income the compensating variation is larger than the equivalent variation - and the change in consumer surplus is in between these two measures.

19.10: A tool for deciding whether policies should be implemented?

This chapter has enormous potential since it provides a way for deciding whether particular policies should be implemented or not. Some policies are straightforward – if they benefit a set of people and do not harm anyone else, then surely they should be implemented. But usually policies are not

⁹ Try it! Draw any optimal point on any original budget constraint. Then rotate the budget line around this optimal point and ask yourself which the individual prefers – before or after the rotation. But ask yourself the same question about a rotation about the original indifference curve.

like this – they make some people better off and others worse off. How can we decide whether they should go ahead or not?

This chapter has shown a way of measuring, in money, how much better or worse off is someone if a price changes. Some policies affect prices and only prices and we can therefore use the apparatus of this chapter to work out the effects on people. We will consider such a case here, though the method can be extended to other cases. This type of analysis is usually called *cost-benefit* analysis, since it counts the costs and benefits to the people affected by the policy change.

Consider then a proposed policy that will increase the price of some good. Those people who buy the good would be worse off if the policy were implemented while those who sell the good would be better off. We can measure and aggregate the costs (to the buyers) and the benefits (to the sellers) of the price change, using the methods of this chapter. Suppose there are I sellers, $I = 1, 2, \dots, I$, and J buyers, $j = 1, 2, \dots, J$. For each of the buyers we can calculate the *compensating variation* – that is, the amount of money that we should give to the buyer to compensate him or her for the price rise. Denote this by c_j for buyer j . This can be considered the cost imposed on him or her by the policy implementation. For each of the sellers, we can calculate also the *compensating variation*, the amount of money that we should take away from the seller to ‘compensate’ him or her for the price rise – that is, reduce his or her welfare back to what it was before he or she was better off as a consequence of the price rise. Denote this by b_i for seller i . This can be considered as the benefit he or she gets if the policy is implemented. An aggregate measure of the costs of the proposed policy is therefore $c = c_1 + c_2 + \dots + c_I$. An aggregate measure of the benefits of the proposed policy is therefore $b = b_1 + b_2 + \dots + b_J$. It could be argued that if $b > c$ then the policy should be implemented.

Why? Suppose that $c = ab$ where a is some fraction between 0 and 1. Then, if the policy were implemented the sellers would be sufficiently better off in that they could all give a fraction a of their ‘gain’ to the buyers, and this collectively would be sufficient to compensate the buyers for being worse off. Thus, with the policy implemented, the sellers would be better off and the buyers no worse off than before. We are back to the simple type of policy that we discussed in the first paragraph of this section. Alternatively the sellers could all be ‘charged’ their benefit and the total benefit could be distributed appropriately to the buyers to make them better off than they were originally. The buyers would be better off and the sellers no worse off than before. Again we are back to the simple type of policy. The reason of course is clear: in each case, that the benefits outweigh the costs. Clearly if this was not the case, that is, if $b < c$, we cannot argue in this way for the policy to be implemented. Indeed, one might think that it should not be implemented. But if the benefits outweigh the costs, that is, if $b > c$, we could argue that the policy should be implemented – we can impose various payments from those who benefit to those who suffer (as discussed above) and everyone would be better off. Indeed, there are some economists who would argue that the policy should be implemented even if these payments are not made, for the total surplus is obviously higher even if its distribution is changed. This latter can be rectified by some appropriate taxation policy¹⁰.

There is another way of looking at the problem. The above assumes that the *status quo* is the pre-policy implementation situation, and therefore that the onus should be on those who want the policy implemented to show that it should be. But we could take the alternative perspective and argue that the *status quo* is the implementation of the policy. We could then ask, starting from this position, what would be the implications of dismantling it – that is, going back to the pre-policy implementation situation? Continuing our example from above, it is clear that the buyers would be

¹⁰ You probably realise that this is, in essence, the argument we have used for the superiority of competition over, for example, monopoly – namely that the total surplus is higher.

better off if the policy were dismantled and the sellers worse off. We know, from the material of this chapter, how to measure how much better off each buyer would be – by his or her *equivalent variation* of the original price rise. Denote this by B_j for buyer j . This is the maximum that buyer would pay to reverse the price rise. We also know how to measure how much worse off would be each seller – by his or her *equivalent variation* of the original price rise. Denote this by C_i for seller i . This is the minimum compensation that he or she would accept to reverse the price rise. So the total benefit from reversing the price rise (from dismantling the policy) would be $B = B_1 + B_2 + \dots + B_J$ and the total cost from reversing the price rise would be $C = C_1 + C_2 + \dots + C_I$. We could argue, as before, that the policy should be dismantled if $B > C$. The story is the same. Suppose $B = AC$ where A is some fraction between 0 and 1. Then, if this condition is satisfied, either (1) each buyer pays a fraction AB_j of the benefits of reversing the policy and each seller receives his or her cost from so doing C_i ; or (2) all buyers pay B_j and the total is distributed over the sellers so that each receives at least their C_i ; or (3) something in between. Under (1) the buyers are all strictly better off if the policy is dismantled and all the sellers are no worse off; under (2) all the buyers are strictly better off and all the buyers are no worse off; under (3) all buyers and sellers could be better off. The policy should be dismantled. It could then be argued that the policy should not be implemented as there would be clear gains from its dismantlement.

So whether the policy should be implemented or not depends on the *status quo*, and upon whether either of the conditions $b > c$ or $B > C$ are satisfied. But you might realise that these conditions are themselves inter-related. To focus our minds, consider the key case of quasi-linear preferences. We know that in this case, the equivalent and compensating variations are equal. Thus we have that $c_j = B_j$ for all j (the two variations are equal for each buyer) and that $C_i = b_i$ for all i (the two variations are equal for each seller). It therefore follows that (1) $c > b$ if and only if $C < B$; and (2) $c < b$ if and only if $C > B$. What does this mean? If (1) is the case, then the argument for the implementation of the policy says that it should not be implemented (the costs outweigh the benefits), and the argument for the dismantlement of the policy says that it should be dismantled (as the benefits outweigh the costs). On either argument the policy should *not* be implemented. If, instead (2) is the case, then the argument for the implementation of the policy says that it should be implemented (the benefits outweigh the costs) and the argument for the dismantlement of the policy says that it should not be dismantled (as the costs outweigh the benefits). On either argument the policy *should* be implemented. With quasi-linear preferences the policy advice is clear: implement the policy if $b > c$ (or equivalently if $C > B$) and do not implement it otherwise. Of course, we could use instead a comparison of the surpluses gained and lost (as we know with quasi-linear preferences these changes in surpluses are equal to the equivalent or compensating variation). This makes policy implementation simple: we simply need to calculate the changes in surplus and to do this all we need to know are the relevant demand and supply schedules.

We can, of course, do the same with non-quasi-linear preferences – assuming the two arguments go the same way. If $b > c$ and $B < C$ then the policy should be implemented; if $b < c$ and $B > C$ then the policy should not be implemented. In the first of these two cases, the total increase in the surplus is positive; in the second of these two cases negative. We can measure these surpluses from the appropriate demand and supply curves.

However, it is crucial to note that, since the surplus is *between* the equivalent and compensating variations when they are not equal, we may have cases in which the net surplus is positive but the two conditions $b > c$ and $B < C$ are *not* both met. Contrarywise, we can have cases in which the net surplus is negative but the two conditions $b < c$ and $B > C$ are *not* both met. These are interesting cases and leave us in a policy quandary.

Consider, for example, a situation in which $b > c$ but $B > C$. What does this imply? The first condition leads to the argument that the policy should be implemented; the second condition to the argument that, once implemented, it should be dismantled! We have a paradox. A similar paradox emerges if we have $b < c$ but $B < C$. This would imply that the policy should not be implemented, but if it is, then it should not be dismantled. In these two paradoxical cases we have a real quandary as to what to – if we are to be guided solely by the arguments of this section. We leave it to the politicians to resolve this quandary. We should note, however, that such cases are likely to be rare. The more usual cases are: (1) where $b > c$ and $B < C$ and the net change in surplus is positive; and (2) where $b < c$ and $B > C$ and the net change in surplus is negative. In (1) the policy *should* be implemented; in (2) it *should not* be implemented.

