32. Exchange

What has been done so far (particularly in chapters 16, 25, 26, and 28) is **partial equilibrium analysis**. In each of those chapters, we have tried to analyze the equilibrium in a market, taking into account the behaviour of the profit–maximizing firm or firms which produce the product, and the utility–maximizing behaviour of the consumers who buy the product.

But the analysis in these chapters focuses on a single market. In particular, I have ignored any interaction among markets. In many cases, those interactions are very important, and that means that partial equilibrium analysis may be misleading.

For example, consider the impact of Trump's protectionist policies in the softwood lumber industry. If Trump puts a tariff on imports of softwood lumber to the United States, we good analyze those effects on the price, and quantity, for Canadian softwood lumber sold in the US, using the models of chapters 16 or 25 or 28 — depending on whether we think Canadian suppliers are perfect competitors, monopolists or oligopolists. But we expect that tariffs on imported Canadian softwood will also have a big impact on the market for domestically produced American softwood lumber, which is a different market (which would not be subject to the tariff). That impact is the reason Trump wants the tariff. We also expect that tariff to have a big effect on the housing industry in the US, because lumber is an important input to housing construction. That's one reason why Trump's idea may be a bad one for the US as a whole : tariffs on Canadian lumber drive up the cost of houses in the US. The tariff would also have a big effect on a lot of other markets in Canada, particularly in British Columbia. If the demand curve for Canadian softwood lumber shifts in due to the tariff, that will cause Canadian firms to lay off workers. So an input market — the market for Canadian forest products workers — will be affected. And I expect that effect will include a fall in income of Canadian lumber workers, which means that the markets for a lot of consumer goods in British Columbia will be affected : less income for the lumber workers means a shift left in the demand curve for consumer goods in areas of British Columbia which depend on the softwood industry.

By definition, these spillovers from one market to another are ignored in partial equilibrium analysis. Worse, these interactions keep going : auto workers buy groceries, and supermarket managers buy cars. So a tariff on cars effects the market for groceries, and the induced changes in the grocery market feed back into the car market.

If we do not want to ignore the fact that everything depends on everything else, we need to analyze the effects of a policy change, considering all markets. That's general equilibrium analysis, which is the topic for chapters 32 – 34. In general equilibrium we take into account the fact that consumers have budget constraints, so that policy changes in one market affect every market in which consumers buy. We also take into account that consumers are buyers of consumption, but also suppliers of labour — and of land and capital. That is we consider both the uses of consumer income, and its sources.

So the good news is that now, in these chapters on general equilibrium, we're going to consider the entire economy, taking into account the sources and uses of spending of all people, and of all firms.

The bad news : if we're going to consider the whole economy in detail, it's got to be a pretty small economy. How small? In chapter 32, there will be only 2 people, and only 2 goods which the people consume. Why 2? If there is only 1 good, then the consumer has no choice at all, so there is really no economic problem at all to consider. So we need at least 2 goods to ensure there is anything interesting to analyze. If there is only 1 person, then we won't have to worry about a policy making some people better off and some people worse off. We need at least 2 people to do any analysis at all of income distribution. Hence 2 goods and 2 people is the bare minimum for a problem in which economics matters. And we're going to stick with the bare minimum.

The production of goods will be totally ignored in chapter 32. In chapter 33, we introduce production, and start to consider that each consumer spends money buying goods from firms, but also earns money selling factors of production, such as labour, or capital, or land, or wheat, to firms. How many of these factors of production will there be in chapter 33? Again, 2. If there were only one input to production, then again the production side of the economy becomes trivial. 2 is the smallest number of factors of production which makes the production side of the economy non–trivial. And again we'll go with the bare minimum.

In chapter 32, there is no production. This sort of world is usually described as an "exchange economy". That name answers the question : "if there is no production on this planet, what economics is there?". So this is a planet in which gods (and services) are not produced by workers and machines. There is a given **endowment** of the goods, which just appears on the planet. So in the planet depicted in Figure 32.1 in Varian, 10 units of good #1, and 20 units of good #2 are available. The quantities are just there. And there's no way of getting more of the good. (How bizarre is that? Think about oil, or other minerals. On our planet, there is a given quantity of oil, and of gold, and of copper. They got here millions of years ago. And — unless you believe in alchemy — there is nothing we can do to change those quantities.)

With a fixed endowment of 10 units of good #1 and 20 units of good #2, what economic activity is there for the two people who live on the planet? The one economic question is : who gets the goods? More particularly : how much of each good does each person get?

So an **allocation** is simply a way of dividing up the goods. An example of an allocation for the economy depicted in figure 32.1 is $\mathbf{x}^A = (6, 11)$ and $\mathbf{x}^B = (4, 9)$. That means that person A gets 6 units of good #1, and 11 units of good #2 ; person B gets 4 units of good 1 and 9 units of good 2. Notice that the quantities chosen here add up to the available endowments of the two goods : $6+4 = 10$ and $11+9 = 20$. An allocation is said to be **feasible** if it is — feasible, that is if it does not promise people more goods than are available.

More generally, if we have X_1 units of good 1, and X_2 units of good 2, then an allocation $\mathbf{x}^A = (x_1^A, x_2^A)$, $\mathbf{x}^B = (x_1^B, x_2^B)$ is feasible if

$$x_1^A + x_1^B \leq X_1$$
$$x_2^A + x_2^B \leq X_2$$

That is : it's feasible if the quantities allocated of each good do not exceed the endowments of the goods. In the equations above, I wrote each equation as a "less than or equal" inequality. But the only useful feasible allocations are those where we have strict equality. That is, if we have endowments of the goods, and if people like the goods, it would be foolish to waste the endowments by allocating less of a good than we have available.

Who decides on an allocation? That depends on the political and economic organization of the planet. But the only physical and technological restrictions on allocations are : they have to be feasible.

The basic question of "welfare economics" is : "given some endowments, what's a good way of allocating the goods to the people?".

What do I mean by "good"? Surely person A likes allocations which give her a lot of each good, but person B likes allocations which give him a lot of both goods. So A might think $\mathbf{x}^A = (9, 17)$, $\mathbf{x}^B = (1, 3)$ is a pretty good allocation [but not as good as $\mathbf{x}^A = (9, 19)$, $\mathbf{x}^B = (1, 1)$.] Person B would disagree.

To get to my (and Varian's) concept of a good allocation, think first of what's a bad allocation. Here's a bad allocation : Good 1 is bread, good 2 is milk. Person A is gluten intolerant, person B is lactose intolerant. The allocation is $\mathbf{x}^A = (9, 3)$, $\mathbf{x}^B = (1, 17)$. Why give most of the bread to the person who can't eat it, and most of the milk to a person who can't drink it?

I'll call an allocation like the one in that little example "Pareto dominated", which is just another term for "bad". Specifically, an allocation $\mathbf{x}^A$, $\mathbf{x}^B$ is Pareto dominated if there is some other feasible allocation $\mathbf{z}^A$, $\mathbf{z}^B$ which both people like better : person A likes $\mathbf{z}^A$ better than $\mathbf{x}^A$ and person B likes $\mathbf{z}^B$ better than $\mathbf{x}^B$. In my gluten–lactose example, if $\mathbf{z}^A = (2, 18)$ and $\mathbf{z}^B = (8, 2)$, we'd expect each person would prefer the "z allocation" to the "x allocation", so that the allocation $\mathbf{x}^A$, $\mathbf{x}^B$ is Pareto dominated by $\mathbf{z}^A$, $\mathbf{z}^B$.

So the way I'm looking for "good" allocations is somewhat roundabout. I am first trying to rule out allocations which everyone agrees are "bad". Then I'm going to look at what's left, that is at allocations which are not (obviously) bad.

In particular, I'm going to concentrate on allocations $\mathbf{x}^A$, $\mathbf{x}^B$ which have the following property

property : there is **no** other feasible allocation $\mathbf{z}^A$, $\mathbf{z}^B$ that both people like better ; that is, there is no other feasible allocation $\mathbf{z}^A$, $\mathbf{z}^B$ which **Pareto dominates** my original $\mathbf{x}^A$, $\mathbf{x}^B$

If my original allocation $\mathbf{x}^A$, $\mathbf{x}^B$ has this property, that is if my original allocation $\mathbf{x}^A$, $\mathbf{x}^B$ is not Pareto dominated, then I'll say that this original $\mathbf{x}^A$, $\mathbf{x}^B$ is **Pareto efficient** (or "Pareto optimal" ; the terms are equivalent).

In other words

definition an allocation is Pareto efficient if there is no other feasible allocation which is preferred **unanimously** to it

Notice the philosophy underlying these definitions. This was the idea of Vilfredo Pareto, for whom all these terms are named. He was avoiding — totally — making comparisons between people. He's asking how to define good allocations, without trying to balance one person's gain against another person's loss. In effect, the concept of Pareto dominance is restricted to cases in which there is no conflict among people, in which everyone agrees that one allocation is better than another.

It might seem that the only allocations we're going to get rid of this way, the only "bad" allocations, are obviously stupid mismatches, such as my example of gluten and lactose intolerance.

It turns out that guess is wrong. The good news is that the concept of Pareto efficiency is not trivial. Among all the possible feasible allocations, only a small subset are Pareto efficient [if I define "small" appropriately]. The bad news is that my small set still has a lot of choices in it. There are going to be many Pareto efficient allocations, so that the concept of Pareto efficiency will not find us a single "best" allocation in an exchange economy.

Perhaps the best way to see which allocations are, or are not, Pareto efficient is using a diagram, in particular the **Edgeworth box diagram** Figure 32.1 in Varian is an example of an Edgeworth box diagram. This is a diagram in which we can put all the feasible allocations in one graph. Every point in the box in figure 32.1 depicts a feasible allocation, and every feasible allocation for this exchange economy can be represented by a point in the box in figure 32.1.

This trick will work only for exchange economies. And it will work only if there are exactly 2 people in the economy (called Ms. A and Mr. B in the example in figure 32.1). And if there are exactly two goods (called good 1 and good 2 in the example in figure 32.1).

Remember, in an exchange economy there are fixed endowments of the 2 goods. So suppose that the — fixed — total endowment of good 1, X_1, is 24 kilograms. Then the width of the Edgeworth box is 24. In other words : one of the goods is represented on the horizontal axis, and the width of the box is the total quantity available of that good. If there are 10 litres available of good B, then the height of the Edgeworth box is 10. The height of the Edgeworth box is the total quantity available of the good represented on the vertical axis.

So the two axes, horizontal and vertical, are used to measure quantities allocated of the 2 goods.

But there are two people here : any allocation gives some of those 24 kilograms of good 1 to Ms. A, and gives the rest of the endowment of good 1 to Mr B.

The convention here : we will measure Ms. A's consumption from the bottom left corner. In other words, if a allocation gives 8 units of good 1, and 6 units of good 2, to Ms. A, we will represent her consumption by the point $(8, 6)$ in the diagram, the point which is 8 kilograms to the right of the origin, and 6 litres above the origin.

This is just a convention, but it's the convention we always use in drawing Edgeworth boxes.

convention : In the Edgeworth box, Ms. A's consumption of the 2 goods is measured from the origin (at the bottom left of the Edgeworth box). Moving right means giving Ms. A a bigger allocation of good 1. Moving up means giving her a bigger allocation of good 2.

How about the other person, Mr. B? Since there are only 2 people in this exchange economy, what he gets is what is left over from the endowment, after Ms. A's allocation is taken out. If Ms. A gets 8 kilos of good 1, and the total endowment of good 1 is 24 kilos, that leaves $24 - 8 = 16$ kilos for Mr. B. Now the original allocation of $(8, 6)$ to Ms. A was represented by some point in the Edgeworth box, a point 8 kilos to the right, and 6 litres up, from the origin.

Now note that Mr. B's allotment of good 1 is the distance of this point from the right side of the box : 16 kilos. Similarly, his allotment of good 2 is what is left over after Ms. A gets her 6 litres, $10 - 6 = 4$ litres. So Mr. B gets $(16, 4)$ if Ms. A gets allocated $(8, 6)$. And that dot in the diagram measures Ms. A's allocation, measured from the bottom left. But it also measures Mr. B's allocation, if we measure from the top right.

And that's going to be true for any allocation. If Ms. A gets an allocation (x_A^1, x_A^2) of the two goods, and we draw a point at (x_A^1, x_A^2) in the diagram — measured in the usual way from the origin at the bottom left — then what is left for Mr. B, which is $X_1 - x_A^1, X_2 - x_A^2)$ is also represented at the point, if we measure not from the usual origin, but from the **top right** corner.

So the trick in the Edgeworth box diagram is that Mr. B's consumption is measured upside down and backwards, from the top right corner of the box.

That means that Ms. A likes allocations which are further right, and further up, in the box. But Mr. B likes allocations which are further left, and further down. [As usual, I am assuming that both consumers are selfish people, and care only about their own consumption.]

What about allocations that are further right and further down? That means Ms. A gets more of good 1, but less of good 2. And Mr. B gets less of good 1 and more of good 2. How either of them would feel about such a move depends on their preferences, how they feel about the value of one good relative to the other.

As usual, we represent their preferences by indifference curves. As usual, indifference curves will be drawn to represent preferences that have 2 properties : (1) each person prefers more of each good ; (2) each person has convex preferences, meaning that his or her willingness to pay for a good gets smaller, the more he or she has of that good.

For Ms. A, those 2 properties mean that her indifference curves are like the blue indifference curves in figure 32.1 : they slope down, and they get less steep as we move down and to the right.

(Notice, as usual, there's a whole family of blue indifference curves for Ms. A, one through any given allocation.)

But Mr. B's indifference curves, the black ones in figure 32.1, look a little different. They still slope down, but they get more steep as we move down and to the right. That's because Mr. B's indifference curves are being measured from the top right : turn the textbook backwards and upside down, and Mr. B's indifference curves look normal (and Ms. A's look weird).

So take some allocation, such as the one marked "W" in figure 32.1. (You can ignore the label "endowment".) That gives Ms. A a consumption bundle which Varian has labeled (ω_A^1, ω_A^2). And it gives Mr. B the rest, $(X_1 - \omega_A^1, X_2 - \omega_A^2)$.

Varian has also drawn indifference curves through the point W. He's drawn two of them, the blue one for Ms. A and the black one for Mr. B. And the indifference curves have the interpretation : Ms. A is indifferent between her consumption bundle (ω_A^1, ω_A^2), and any other consumption bundle on the blue indifference curve. Anything that's above that indifference curve (for example, the point labeled M in the diagram) is something she likes better than W.

Similarly, Mr. B is indifferent between W, and any other allocation on his black indifference curve through W. But he prefers (to W) anything which is **below** the black indifference curve : the direction of allocations which he likes better is the direction in which he gets more of the goods, down and to the left.

Notice that Mr. B, like Ms. A, prefers M to W. That is, in the diagram the allocation M Pareto dominates to the allocation W. In fact, any allocation in the dark blue lens–shaped area — between the two indifference curves — Pareto dominates W.

So W is **not** a Pareto efficient allocation, since there are other feasible allocations, such as M, which both people prefer to W.

Now here is the value of the Edgeworth box diagram, in testing whether an allocation is Pareto efficient or not. Why was the allocation W Pareto–dominated? Because there was a dark blue lens–shaped area, of feasible allocations which both people preferred to W.

Why is there such a lens–shaped area of feasible allocations which Pareto dominate? Because Mr. B's (black) indifference curve through W is more steep than Ms. A's (blue) indifference curve. For **any** feasible allocation inside the Edgeworth box, if Mr. B's indifference curve is steeper than Ms. A's, then there will be a lens–shaped area, above and to the left of the original allocation, which Pareto dominate the original allocation. We just need an allocation which is above Ms. A's indifference curve, and below Mr. B's. And as long as Mr. B's indifference curve (through the original allocation) is steeper than Ms. A's, there will be a few such allocations, just above and to the left of the original allocation.

So the example of allocation W in the diagram indicates that : any time that Mr. B's indifference curve through the allocation is steeper than Ms. A's, the allocation will be Pareto-dominated by some other allocation, which is above it and to the left in the Edgeworth box diagram.

On the other hand, suppose that Mr. B's indifference curve through an allocation is less steep than Ms. A's. (Consider, for example, the other intersection point of the two indifference curves through W in figure 32.1.) Any allocation above Ms. A's blue indifference curve is preferred by her to the original allocation ;

anything below and to the left of Mr. B's indifference curve is preferred by him. And if his indifference curve is less steep than hers, there must be some allocations, below and to the right of the original allocation, which Pareto dominate the original allocation.

So that's what I meant by saying that Pareto efficient allocations were only a small subset of all the feasible allocations. The only Pareto optimal allocations inside the Edgeworth box are allocations like M in figure 32.2. At M, the indifference curves of the two people have exactly the same slope. They are **tangent** to each other in the diagram. The allocations which Ms. A prefers to M are those above her indifference curve, those in the blue–shaded set in figure 32.2. The allocations Mr. B prefers are those below his indifference curve, those in the gray set in figure 32.2. Because the indifference curves of the two people, through M, are tangent to each other, there is **no** allocation which is in both the gray and the blue sets. That means that there is no allocation which Pareto dominates M, meaning that M is Pareto efficient.

Summarizing the conclusions from the Edgeworth box diagram :

summary An allocation inside the Edgeworth box is Pareto efficient if — and only if — the two people's indifference curves through the allocation are tangent to each other at the allocation.

How many such Pareto efficient allocations are there? Quite a few, actually, but I will return to that question shortly.

First, what does it mean when the indifference curve of the two people are tangent at some allocation? The absolute value of the slope of a person's indifference curve is her **marginal rate of substitution**. That's the rate at which she is willing to trade (small quantities of) one good for the other.

So if Ms A's MRS through some allocation is 3, that means that she is willing to trade 3 grams of good 1 in exchange for 1 millilitre of good 2. And if Mr B's MRS was different, say 2, that mean that he is willing to trade 2 grams of good 1 in exchange for 1 millilitre more of good 2 — or vice versa. When these MRS's are different, as in this example, then there is room for a deal. What if Mr B offers to give Ms A 1 millilitre of good 2, in exchange for 2.5 grams of good 1. He likes the deal, which is why he's proposing it. He is willing to give up the 1 millilitre of good 2 in exchange for 2 grams of good 1, but the deal that he is proposing gives him more than that, 2.5 gram. On the other hand, Ms. A is willing to give up a lot of good 1, 3 grams,to get one more millilitre of good 2. This deal requires her only to give up 2.5 grams.

So the fact that the two people had different MRS's at some allocation (such as W in figure 32.1) means that there is room for a deal. In this example, the deal involved Mr. B trading good 2 to Ms. A in exchange for good 1, because he is the one with the (relatively) lower value on good 2 [or the (relatively) higher value for good 1]. That's a move up and to the left — into the region of Pareto dominant allocations in figure 32.1.

Rational people will agree to a deal only if it makes them better off. So if we start at some allocation (such as W in figure 32.1), and if there is some deal possible which makes both people better off, then the original allocation can't be Pareto efficient. On the other hand, if the two people have exactly the same MRS (as at M in figure 32.2), then there are no deals possible that make both people better off.

That means that the following statements are all equivalent :

(*i*) an allocation W (in the interior of the Edgeworth box) is **not** Pareto efficient

(*ii*) the two people's indifference curves through W have different slopes

(*iii*) the two people's marginal rates of substitution between the two goods, at W, are different

(*iv*) a deal is possible, in which the two people trade small amounts of the goods, which makes them both better off compared to W

A feasible allocation is Pareto efficient if it does **not** have the 4 (equivalent) properties listed above. So an allocation is Pareto efficient only if there are no deals possible which make everyone better off. In a sense, Pareto efficient allocations are where we'd get if we used up all the mutually beneficial deals we could find.

As I mentioned, there are a lot of Pareto efficient allocations in a 2–person, 2–good exchange economy. The set of Pareto efficient allocations in an Edgeworth box is usually called the **contract curve**, as in figure 32.2 in Varian. The term is perhaps not the most appropriate, but it is an old one, due to Edgeworth himself.

In figure 32.2 the contract curve is a 1–dimensional curve, which goes from the bottom left corner of the Edgeworth box to the top right. That's usual. The exact shape of the contract curve depends on the preferences of the individuals. But it always includes the bottom left and top right corners. And if both goods are normal goods for both people, the contract curve must slope up.

So if there are infinitely many Pareto efficient allocations, how do we choose among them? Unless we are willing to make value judgments about the relative merit of Ms. A and Mr. B, we don't. That was Pareto's point. He wanted to see how far we could go in assessing economic policies, if we did not want to choose between different people's potentially conflicting interests.

As we move up the contract curve in figure 32.2, Ms. A gets better off and Mr. B worse off. So if we ask the two individuals which are the best of the many Pareto efficient allocations, their opinions would be completely opposite to each other. Naturally, each person wants as much as he or she can get of the goods.

Repeating the above argument, why is the top right corner of the Edgeworth box on the contract curve in figure 32.2? An allocation is on the contract curve if it is Pareto efficient — that is if there is no other feasible allocation which makes both people better off.

Suppose we started with an allocation at the top right corner of the Edgeworth box. That is, supposed the proposed allocation was simply : give everything, all of each good, to Ms. A, and nothing to Mr. B. Is this allocation Pareto dominated by some other feasible allocation? To Pareto dominate the allocation, we have to find a way of making both people better off. That's impossible : we can't make Ms. A better off with some other feasible allocation, since we have already given her everything. So there's no way of making her better off than she is at the top right hand corner of the box.

Another version of the definition was to look at possible trades. Is there a trade possible which makes both people better off, starting at the top right corner? Mr. B has nothing to trade, since he's been allocated nothing. Which means that there is no mutually beneficial trade possible, so that the allocation at the top right corner of the Edgeworth box must be Pareto efficient.

More generally, take any indifference curve for Mr. B, say the one near the bottom left corner of the Edgeworth box in figure 32.2. What do we know about Ms. A's (blue) indifference curves? They start out very steep near the vertical axis, and they get less and less steep. So start on Mr. B's black indifference curve, where it hits the left axis. Ms. A's indifference curve through this point [the point where Mr. B's black indifference curve hits the left vertical axis] is going to be very steep, steeper than Mr. B's indifference curve. [It is technically possible that Ms. A has a really strong taste for good 2, so that her indifference curves are never very steep, but that's not a very realistic case.] So as we start out along the black indifference curve of Mr. B, the allocation is not Pareto efficient, since Ms. A's indifference curve through the allocation is steeper than Mr. B's. Now start moving down Mr. B's indifference curve. As we move down and to the right, Ms. A's indifference curves [the blue ones] get less and less steep, until they are very shallow, on the bottom axis. So we start out on the top left of Mr. B's indifference curve, where his indifference curve is less steep than Ms. A's. And if we move far enough down along that indifference curve, then his (black) indifference curve will be more steep than Ms. A's. Since everything here is continuous, that means that there must be some point along the indifference curve that is similar to point M, where Mr. B's black indifference curve is just tangent to Ms. A's blue indifference curve through the point.

Conclusion : if I start out by picking some arbitrary indifference curve for one person [Mr. B here], then there is some point [exactly one point, in fact] on that indifference curve which is Pareto efficient, where his indifference curve is exactly tangent to the other person's.

In fact, what I have just described is the following mathematical problem. Pick some level $\bar{u}^B$ of utility for one person ; that's the same as picking an indifference curve. Then try and make the other person's utility as high as possible, subject to the allocation being feasible, and subject to getting Mr B the "guaranteed" level of utility $\bar{u}^B$. Since Mr B's consumption bundle is just the quantities of the two goods left over, after what Ms A gets, then

$$x_1^B = X_1 - x_1^A$$
$$x_2^B = X_2 - x_2^A$$

so that the mathematical problem is to maximize Ms. A's utility $U^A(x_1^A, x_2^A)$, subject to the constraint that Mr. B gets his specified level of utility $\bar{u}^B$, or

$$U^B(X_1 - x_1^A, X_2 - x_2^A) \geq \bar{u}^B$$

That is a standard calculus problem of constrained maximization

choose (x_1^A, x_2^A) to
maximize $U^A(x_1^A, x_2^A)$
subject to the constraint

$$U^B(X_1 - x_1^A, X_2 - x_2^A) \geq \bar{u}^B$$

The solution to this problem is a Pareto efficient allocation. If I solve this constrained maximization problem for different levels of the "guarantee" $\bar{u}^B$ given to person B, I will derive the contract curve for this economy.

The standard mathematical way of solving this kind of problem, a maximization which is subject to one or more constraints, is to set up a Lagrangean. The Lagrangean function for the maximization problem above is

$$\mathcal{L}(x_1^A, x_2^A, \lambda) = U^A(x_1^A, x_2^A) + \lambda[U^B(X_1 - x_1^A, X_2 - x_2^A) - \bar{u}^B]$$

and we solve the maximization in the usual way, by taking the derivatives of $\mathcal{L}(x_1^A, x_2^A, \lambda)$ with respect to the variables x_1^A, x_2^A and λ, and setting the derivatives equal to zero. Doing that,

$$\frac{\partial U^A}{\partial x_1^A} - \lambda[\frac{\partial U^B}{\partial x_1^B}] = 0$$

$$\frac{\partial U^A}{\partial x_2^A} - \lambda[\frac{\partial U^B}{\partial x_2^B}] = 0$$

$$U^B(X_1 - x_1^A, X_2 - x_2^A) - \bar{u}^B = 0$$

where I have used the facts that

$$x_1^B = X_1 - x_1^A$$

$$x_2^B = X_2 - x_2^A$$

Those first two first–order conditions imply that

$$\frac{\partial U^A / \partial x_2^A}{\partial U^A / \partial x_1^A} = \frac{\partial U^B / \partial x_2^B}{\partial U^B / \partial x_1^B} \qquad (*)$$

The left side of equation $(*)$, the ratio of Ms A's marginal utility of consumption of the two goods is — precisely — her *marginal rate of substitution* between the goods, as defined in chapter 4. [Well, Varian insists on sticking a minus sign in front, but otherwise it's the same.]

That's the slope of an indifference curve. So all equation $(*)$ is is a condition derived graphically earlier : for an allocation to be Pareto efficient, each person's indifference curve, through her own allocation, must have the same slope as any other person's.

The various characterizations I have of Pareto efficiency apply only in the world of the Edgeworth box, an exchange economy with two consumption goods and two consumers. But the good news is that these characterizations are completely general. They hold in a (more realistic) world with many different commodities, and many different consumers, and they hold as well in the more realistic world of chapter 33, in which goods and services have to be produced, using labour and land and raw materials.

In a many–person, many–good world, an allocation is Pareto efficient if it is not Pareto–dominated : if there is **no** other feasible allocation which **every** person likes better. Alternatively, the allocation is Pareto efficient only if there is **no** other feasible allocation which everyone prefers.

And an interior allocation will be Pareto efficient only if everyone's MRS between goods i and j is the same. So suppose we have 4000 goods and 100 million people. Take any 2 of those 4000 goods, say good 376 and good 2822. An allocation is Pareto optimal only if the MRS between good 376 and good 2822 is the same for each of the 10 million consumers.

I have talked a long time, and in a lot of different ways about Pareto efficiency. I have talked about making people better off, and about allocating goods and services, and about potential trades among people.

But there are a few economic terms I have not mentioned at all in this chapter : prices, income, demand, supply.

There's a reason I have not mentioned these concepts, which are so important in nearly every other section of the course. Those concepts are based on a particular type of economy, a market economy in which goods are bought and sold at market prices. That is a way of organizing the economy. It's not the only way. We could have a dictatorship, in which the dictator decided what was allocated to whom. We could have a type of democratic socialism, in which no person owned any property, and in which the democratically elected assembly decided how to allocate goods and services among people.

The idea of a "good" way of allocating the endowment of goods and services among people should not be specific to one particular type of economic organization. In fact, one criterion for choosing how we might want to organize some economy might be to look at how the different types of organization allocated those goods and services. If dictatorship resulted in an allocation which Pareto dominated the outcome under democratic socialism, that might be a strong argument in favour of dictatorship.

So the biggest question address in chapter 32 is "what types of organization lead to Pareto efficient allocations?".

I'm going to look at one particular type of organization, because it's one we pay a lot of time with in microeconomics. That's a market economy, in which there are many consumers, in which goods and services are bought and sold at market prices [at which quantity demanded equals quantity supplied], and in which there is private property. In particular, in a market economy, the endowments of goods and services are privately owned.

So each of the X_1 units of good 1 are owned by someone, either Ms. A or Mr. B in my 2–person economy. As in Varian, I will use the greek letter ω to represent endowment ownership : ω_2^A is Ms. A's initial ownership of good 2. In a two–person market economy, it must be true that

$$\omega_1^A + \omega_1^B = X_1$$

$$\omega_2^A + \omega_2^B = X_2$$

since each unit of each good is owned by one of the two people.

But, in a market economy, you do not have to consume what you own. You can sell some of your endowment of a good to someone else, and you can buy additional units (above your endowment) of a good from someone else. And you do that buying and selling at market prices, which you take as given.

Where do the market prices come from? As in partial equilibrium, market prices are determined by the interaction of people's demand and supplies. These market prices clear the markets : they will be equilibrium prices in a market economy if quantity demanded of a good equals quantity supplied. But this is a general equilibrium economy, in which we look at all markets together. If we had I goods, there are I equilibrium prices, one for each good, and these I prices will be a vector of equilibrium prices if quantity demanded of **any** good i equals quantity supplied of good i, where i is any one of the I goods.

So what happens in market economy, as I have defined it. Each person has a **given** endowment of each good : (ω_1^A, ω_2^A) for person A, (ω_1^B, ω_2^B) for person B. I don't know how this ownership pattern got arranged : a random distribution, history, some fight or contest. The main things are that these ownership patterns are clear and established, and everyone knows what she owns, and anyone can sell off pieces of what she owns.

Then there are **prices** for each good. These prices are not given : they are determined by how much people want to supply or demand. But to any single individual, they might as well be given. There are enough people in this economy that each person knows she has virtually no influence on any prices, and so takes them as being out of her control, just as her ownership of endowments is out of her control.

What is under a person's control? How much she actually consumes of each good. For example, she could sell some of her endowment of good 1 to buy extra units of good 2. She'll choose how much to buy and sell, based on her tastes, based on the prices she sees, and based on how much she owns.

Now there is no outside source of income in this world. That's the point of general equilibrium : the entire working of the economy is going to be considered, so we can't have any extraneous outside source of income.

So the only way that I can get any money to buy good 2, is to sell off some of my endowment of good 1. As I did earlier in the chapter, I'll let x_1^A denote Ms. A's **consumption** of good 1, which may be very different from her endowment of the good.

Suppose that she has endowment of ω_2^A units of good 2, but she wants to consume $x_2^A > \omega_2^A$ units of good 2. That means that she has to buy $x_2^A - \omega_2^A$ units of good 2 on the market. This amount she has to buy, $x_2^A - \omega_2^A$ is called her *net demand*, or her *excess demand* for good 2.

To buy $x_2^A - \omega_2^A$ litres of good 2 on the market, Ms. A is going to need to pay for her net demand. In particular, she needs $p_2[x_2^A - \omega_2^A]$ dollars to pay for these market purchases of good 2, where p_2 is the market price per litre of good 2. In this calculation, x_2^A is the quantity she wants to consume of good 2 ; that's something she herself is choosing. Her original endowment ω_2^A is not something she chooses ; it's just given. The market price p_2 of good 2 is also something out of her control, being determined by the overall market interaction of all the people's supply and demand decisions.

How can she get $p_2[x_2^A - \omega_2^A]$ dollars to pay for her market demand of good 2? The only way she can get that money is to sell off some of her endowment of the other good, good 1.

in a market economy, the only source of money to pay for (positive) excess demands for a good is to sell of some of one's endowment of some other good

So, for this person, selling some of her endowment of good 1 to buy some of good 2, she's going to collect $p_1(\omega_1^A - x_1^A)$ dollars from selling off some of her endowment, and she's going to use that money to buy some more of good 2. That quantity $\omega_1^A - x_1^A$ could be called her net supply, or her excess supply, of good 1. She's not selling all her endowment of good 1, just some of it, and her actual consumption x_1^A of good 1 is what's left over after she's sold off some of the good. Again ω_1^A is exogenous here (as is p_1), and her own consumption x_1^A is what she is choosing.

So person 1's budget constraint is the equation

$$p_2(x_2^A - \omega_2^A) = p_1(\omega_1^A - x_1^A) \tag{bc}$$

She picks her actual consumption (x_1^A, x_2^A), but this consumption bundle must obey the budget constraint (bc).

But I had assumed that she wants to sell off some of her endowment of good 1, in order to buy more of good 2. What if she decides to do the opposite, to sell off some of her endowment of good 2, and buy more of good 1. Then her budget constraint would be

$$p_1(x_1^A - \omega_1^A) = p_2(\omega_2^A - x_2^A) \tag{BC}$$

But (bc) and (BC) are the **exact same** equation. Each of them could be written

$$p_1 x_1^A + p_2 x_2^A = p_1\omega_1^A + p_2\omega_2^A \tag{**}$$

That equation says that the market value of what I choose to consume (the left side of equation $(**)$) cannot exceed the market value of what I own (the right side of equation $(**)$).

This equation isn't written in chapter 32 of Varian. But it is written — in exactly this form — in section 9.2 of Varian. The right hand side of equation $(**)$ is the income the person has available to spend on consumption, which is the total market value of her endowment. In fact, another way of describing this budget constraint is to have two stages : first, she sells all her endowment to get money ; the amount of money she gets is the right side of equation $(**)$; then she goes and spends that money on consumption, picking the bundle she likes best from her budget set.

So, with one difference, she is behaving exactly like the consumer in chapters 2 – 8 of Varian. The one difference is that the income she has available to spend on consumption is not some given amount m ; now her income, $p_1 x_1^A + p_2 x_2^A$ depends on the prices of the goods she owns. But it is still true that the right side of equation $(**)$ is exogenous to her ; it is still true that equation $(**)$ defines a budget line of (x_1^A, x_2^A) bundles which she can afford ; it is still true that the slope of her budget line is $-p_1/p_2$; and it is still true that the best bundle for her, from all the bundles she can afford, is the bundle (x_1^A, x_2^A) at which her indifference curve is tangent to the budget line.

Like consumers in chapters 2 – 9, each person in this economy has a demand function for each good, derived from the consumer finding the consumption bundle which she likes best among all the bundles on her budget line. Equation (∗∗) defines the person's budget line. The slope of the budget line (with x_1^A on the horizontal, x_2^A on the vertical) is $-p_1/p_2$. Changes in the relative prices of the 2 goods will change the slope of the budget line. To see how price changes affect the location of the budget line, note that the consumption bundle

$$x_1^A = \omega_1^A, x_2^A = \omega_2^A$$

is always right on the person's budget line, no matter what are the prices of the 2 goods. That consumption bundle is what she gets if she doesn't buy or sell anything, just consumes her endowment. That's always feasible for her, no matter what the prices, because she does not have to buy or sell anything to consume this bundle. It's also right on the budget line, since she can't get any more of either good unless she chooses to sell some of the other good.

So here are 2 properties of the budget line defined by equation (∗∗) :

(i) only the *relative* prices p_1/p_2 matter

(ii) changes in this relative price p_1/p_2 of good 1 pivot the line around the endowment point (ω_1^A, ω_2^A) [increases in p_1/p_2 make the budget line more steep, decreases in p_1/p_2 make it less steep]

Each person's quantity demanded of any good depends on the prices of all the goods (and on the person's endowments, and on her tastes). The total quantity available of a good is fixed, at X_i, in an exchange economy. So the market for good 1 (for example) will clear only if the total quantity demanded of the good (by all people) equal the total endowment of the good. If the total quantity demanded were less than the total endowment, there would be excess supply in the market, as total quantity demanded would be less than the available endowment. If total quantity demanded is more than the endowment, there is excess demand for the good. So if I take any good, say good i, then the market for good i will be in equilibrium if and only if

$$x_i^A(p_1, p_2, \ldots, p_I) + x_i^B(p_1, p_2, \ldots, p_I) + \cdots + x^Z(p_1, p_2, \ldots, p_I) = X_i = \omega_i^A + \omega_i^B + \cdots + \omega_i^Z \qquad (***)$$

if there were 26 people, named A,B,. . .,Z.

Varian writes my equation using the *excess demands* for each good. The excess demand of a person (say person A) for good i was defined as $x_i^A(p_1, p_2, \ldots, p_I) - \omega_i^A$. So another way of writing equation $(***)$ is to say that the sum of all people's excess demands for good i — Varian writes this as $z_i(p_1, p_2, \ldots, p_I)$ — must equal zero for the market to be in equilibrium.

In general equilibrium, since everything depends on everything else, the concept of an equilibrium refers to every market being in equilibrium. Adjusting p_2 changes not only the excess demand for good 2, but the excess demand for every other good.

A vector of prices $(p_1, p_2, \ldots, p_I)$ is an equilibrium price vector if (and only if) excess demand for each of the I goods is zero, when the prices of the goods are $(p_1, p_2, \ldots, p_I)$.

I will put off for a little while the question of whether there must be some vector of prices that clears all the markets, and the question of how we might find such a price vector if there is one, and the question of how markets might adjust to equilibrium if an equilibrium is possible.

Instead, I look first at what people's consumption of the different goods and services must be, if we are in equilibrium.

Each person chooses a consumption bundle which is best for her, given her tastes and endowments, and given the prices. That bundle is the point on her budget line where her indifference curve is tangent to the budget line. [This property really does not depend on the prices being equilibrium prices ; it's just what's best for the person, if she thinks that she can buy and sell at the given prices.]

So

property Person A always chooses a consumption bundle(x_1^A, x_2^A) such that her marginal rate of substitution between the two goods equals the slope of her budget line :

$$|MRS^A| = \frac{p_1}{p_2}$$

This property holds for person B as well, and for any other people in the economy. So

$$|MRS^A| = \frac{p_1}{p_2} = |MRS^B| \qquad (****)$$

What equation $(****)$ says is : if people are price takers, and if they all face the same prices, then they will choose consumption bundles such that each person's MRS equals the price ratio.

That means that in any equilibrium to a market economy, each person's MRS will be the same, and will equal the price ratio. But that equality of people's MRS's was exactly the condition for an allocation to be Pareto efficient. Therefore

First Theorem of Welfare Economics : Any equilibrium in a market economy must be Pareto efficient.

This theorem is one reason why economists pay some much attention to perfectly competitive markets. The theorem says that, if you got to set up the economic organization of a country, that having private ownership of goods, and organizing competitive markets to allocate those goods, is a pretty efficient way of doing things. In fact it says that there is no other way of organizing an economy that will do better than this market economy.

In particular, it says that **any** policy of government intervention — such as subsidizing some goods, any taxes which redistribute income, any controls on prices, any direct government provision of some goods or services — must make some people better off. There is no such thing as a policy intervention which makes everybody better off.

Two very important qualifications here, before I say that competitive markets, private property, and no government intervention are the only good policy.

Qualification number 1 is that there are some assumptions built into my model of consumer behaviour. I assumed that consumers know their own tastes, that one consumer's behaviour does not harm directly anyone else, that markets exist for everything we'd want to consume, and that consumers are perfectly informed. [More importantly, that no person in the economy is better–informed than any other.] Those are strong, unrealistic assumptions, and if they don't hold — and they don't — the first fundamental theorem does not hold. Chapters 35, 37 and 38 provide a brief indication of how things change when we try and get rid of some of those strong assumptions.

Qualification number 2, against the first fundamental theorem of welfare economics, is that we may not feel that it is such a tragedy to make some person worse off. The First Theorem of Welfare Economics does not say where the competitive equilibrium allocations are on the contract curve, just that they are on the contract curve. A competitive equilibrium might be in the lower left corner area of the Edgeworth box, meaning that Mr B is doing very well and Ms B not so well. In fact, this must be the case if Mr B happens to own most of the endowments of the goods. I can put the people's original endowment as a point in the Edgeworth box, as Varian does in figures 32.1 and 32.2.

Fact : in any competitive equilibrium, a person must be at least as well off as she would be consuming her own endowment. Why is that true? A person is always getting to choose how much to buy and sell. No matter what the equilibrium prices are, a person can always choose not to buy, or sell, any of the goods. That is, the person's endowment point is always on her budget line, no matter what are the prices. The person always chooses the point she likes best on her budget line. So she always gets to choose a consumption bundle which is on at least as high an indifference curve as her initial endowment.

So if the initial endowment were on the best of Mr B's indifference curves, the furthest left of the black indifference curves in figure 32.1, then the competitive equilibrium allocation or allocations must be below and to the left of that black indifference curve. None of those allocations are very attractive to Ms. A.

So if I am an outside observer of this economy, and I see the result of perfect competition is some point on the contract curve, near the bottom left corner, I may not think that result is very attractive. I might want to make Ms. A better off. What does the First theorem say about that? It says that there is no way that we can make Ms. A better off — unless we make Mr. B worse off. And that might be something we might want to do. Even if I like Mr. B, if I think he is doing way better than Ms. A, then I might be willing to make him a little worse off, if that's the only way of making Ms. A better off.

So the first theorem does not really say that any competitive equilibrium is better than any other feasible allocation, or the result of any other economic system. It just says that you can't improve on the equilibrium outcome for one person without making someone else worse off. Remember that Pareto, for whom the term was named, was looking at how to rank economic outcomes without making any interpersonal comparisons of well–being. He did not want to try and make such statements as "person C is twice as well off as person D", or "if I take some food from person E and give it to person F, the gain in person F's happiness is twice as large as the loss in person E's happiness".

Ultimately we (or someone) may have to judge different people's well–being. If and when we do that, we'll still want a Pareto efficient allocation — unless we don't trust people's own tastes, or unless we actually hate some people and want to see them suffer. But we may not the particular Pareto efficient allocation resulting from perfect competition, for some given set of initial endowments.

If that's the case, then the Second Theorem of Welfare Economics is a very powerful result, and a much more convincing argument in favour of perfect competition and private ownership. That's the theorem Varian mentions in section 32.11.

It's sort of the reverse of the first theorem. The first theorem started with some competitive equilibrium (resulting from some initial ownership pattern for endowments), and then showed it must be Pareto efficient. The second theorem starts with a Pareto efficient allocation, any Pareto efficient allocation, and goes in the other direction. It says that any Pareto efficient allocation can be achieved as a competitive equilibrium, for some particular initial endowment pattern.

That's a very strong argument for competitive markets. It says to a prospective political planner : "Pick the allocation which you think is best, based on your notions of fairness and justice. Unless you're pretty twisted, that allocation is going to be Pareto efficient. And one way you can achieve this best of all feasible allocations is to redistribute endowments — and then go away, do nothing, and let markets work."

To prove this second theorem, take any Pareto efficient allocation, such as point M in figure 32.2. A Pareto efficient allocation is an allocation at which both people's indifference curves have the same slope — that's the slope of the tangent indifference curves in the Edgeworth box diagram.

So any Pareto efficient allocation has a slope associated with it, the slope of the indifference curves which are tangent at the allocation. Let that slope be the price ratio p_1/p_2 (as illustrated in figure 32.7).

So take the allocation you think is best. Ms. A will get a consumption bundle (x_1^A, x_2^A) in that allocation. Draw a budget line for her, which goes through this point, and which has the slope p_1/p_2 — where that's the slope of the tangent indifference curves. We know that person A will choose the consumption bundle (x_1^A, x_2^A), is the price ratio is p_1/P_2, and if the bundle (x_1^A, x_2^A) is on her budget line. Why? It's on her budget line, **and** her indifference curve is tangent to it at (x_1^A, x_2^A), since p_1/p_2 was constructed to be the slope of her (and Mr B's) indifference curve through (x_1^A, x_2^A).

That means : if her initial endowment is on that budget line, and if p_1/p_2 is the price line, then she will choose (x_1^A, x_2^A) as her consumption bundle. That is, if the initial endowment for her is anywhere on that budget line, then there is a competitive equilibrium, with relative prices p_1/p_2, where she consumes (x_1^A, x_2^A).

What about person B? Everything is the same for him, but upside down (if we are using the Edgeworth box diagram). If his endowment is on the budget line constructed through the chosen Pareto efficient allocation, then he will choose that consumption bundle (x_1^B, x_2^B) if the price ratio is p_1/p_2.

In other words, if you think that the point X in figure 32.7 is the best allocation, here's how to achieve it. Distribute the endowments so that they are on the straight line through X (with slope tangent to the indifference curves) in figure 32.7. Any point on the line will do. Then go away, and let people buy and sell. The resulting market equilibrium will be the point X.

Three important issues will not be covered here. Varian discusses one of them, although fairly briefly.

One is my use of **relative prices** in discussing competitive equilibrium. I started by talking about a vector $(p_1, p_2, \ldots, p_I)$ of prices, which would be equilibrium prices if excess demand was zero for each of the I goods, when the I prices were $(p_1, p_2, \ldots, p_I)$. But then, when I went back to my diagrams, with only 2 goods, I mentioned only the price ratio p_1/p_2.

Is the price ratio all that matters? Does it matter whether the prices of the 2 goods are $(2, 1)$, $(20, 10)$, or $(0.6667, 0.3333)$? No, it doesn't matter. And the reason is each person's budget constraint

$$p_1 x_1^A + p_2 x_2^A = p_1 \omega_1^A + p_2 \omega_2^A \qquad (**)$$

If I multiply all prices by some constant — 10, or 1/3, or anything else — I don't change the budget equation or the budget set. And that's true for all people. And that's true whether there are only 2 goods, as in my pictures, or whether there are many more goods. That means that if $(p_1, p_2, \ldots, p_I)$ is an equilibrium price vector, then so is $(tp_1, tp_2, \ldots, tp_I)$, for any scalar t. In particular, one thing I could do is to set $t = 1/p_1$, so that my price vector is now $(1, p_2/p_1, \ldots, p_I/p_1)$. Doing that trick is called making good 1 the *numéraire* : measuring everything in terms of good 1. That is not the only way of choosing a price vector. But making good 1 [or any other good] the numéraire makes clear that everything is relative. It also makes clear that there are really only $I - 1$ prices to find, in searching for an equilibrium, since I can always set one of the prices equal to 1.

The second issue is the existence of equilibrium prices. I defined an equilibrium price vector as a price vector which cleared the markets for all I goods. For given endowments, and given the tastes of the people, how do I know that there is such a price vector? How do I find it?

Now Mr Walras, who first considered this problem, simply counted equations and unknowns. There are I goods, which means I equilibrium conditions, that the quantity demanded of each good equal the quantity supplied. And there are I prices. So I equations in I unknowns. There must be a solution.

Mr Walras was wrong in several ways. First of all, it's not generally true that there must be a solution to I equations in I unknowns. Second, we need the solution to make economic sense : we can't have the prices negative, for example. Third, my previous point, about relative prices, said that there were really only $I - 1$ unknowns.

Mr. Walras worked his way around that last objection. There are really only $I - 1$ prices, but there are really only $I - 1$ unknowns. That's the result of what we now call Walras's Law. Walras's Law implies the following : (i) take some price vector ; (ii) suppose that, with this price vector, the market clears [excess demands are zero] for goods $1, 2, \ldots, I - 1$; (iii) then it **must** be true that there is zero excess demand for the remaining good, good I.

Walras's Law states that the **value** of the some of all people's excess demands must equal zero, no matter what the prices are. The prices don't have to be equilibrium prices. Just any prices. I have to stay on my budget line, so that the value of what I buy on the market (goods for which I have positive excess demand) must equal the value of what I sell on the market (goods for which I have negative excess demand). Add that up over all people, and we have this law, which is just a piece of arithmetic : the sum of the value of all people's excess demands for all goods must equal zero, no matter what the prices are.

Because of Walras's Law, if excess demands equal zero for each of $I - 1$ goods, it is guaranteed that the excess demand is zero for the remaining good (as spelled out in 32.6 of Varian). So we need only check excess demands for $I - 1$ goods, in trying to find an equilibrium.

That still does not establish that there is at least one price vector that clears all I markets. Varian babbles a bit about the mathematics of this problem in section 32.8. But it is a difficult mathematical problem, and Gérard Debreu got the Nobel Prize in 1983 just for solving it.

Finally, the third of my problems is one which Varian doesn't discuss, which is good because it too is a messy problem. I have referred a few times to "an equilibrium allocation", for given endowments. Problem 2 : starting with some allocation of endowments, is there a set of prices which clears all I markets? Debreu's answer was "yes". So that set of prices leads to an equilibrium allocation, when people buy and sell however much they want, and reach their preferred consumption points. Problem 1 showed that there were a lot of price vectors which worked. But those price vectors were all proportional to each other, and all led to exactly the same allocation. Problem 3 is : with 2 goods, could there be two **different** price ratios, say $p_1/p_2 = 0.7$ and $p_1/p_2 = 0.3$, which were **both** equilibrium price ratios. They would lead to different allocations, since people will choose different consumption bundles when the slope of their budget line changes. But could two different sets of relative prices each lead to markets clearing?

And the general answer is "yes". That is, there has to be at least one equilibrium allocation for a given allocation of endowments. (That was problem 2.) But, in some cases, there might be more than one. In that case **each** of the equilibrium allocations will be Pareto efficient : the first theorem says that any competitive equilibrium allocation is Pareto efficient. It's messy, annoying and potentially confusing, but it is mathematically possible that there might be multiple equilibria for the same initial allocation of endowments.