Econ 802

Lecture Notes on Chapter 17

Greg Dow December 3, 2020

In a partial equilibrium model, we focus on one or a few markets, while holding prices (and quantities) in all other markets fixed. Thus the prices in some markets are treated as exogenous.

In a general equilibrium model, we consider all markets simultaneously, treat all prices as endogenous, and require that all markets clear. The exogenous variables are preferences, technology, and endowments.

Why is GE important? Several reasons.

1. We often want to know whether a model is logically consistent – would it be theoretically possible for all markets to clear simultaneously?
2. By treating more prices as endogenous, we get a more comprehensive picture of the interactions across markets, so possibly more accurate predictions
3. Many normative questions arising in welfare economics are best handled in a GE framework (for example, is a competitive equilibrium Pareto efficient?)

Applied GE models come up in trade (effects of tariffs) public economics (effects of taxes), etc.

Some basic questions about GE:

1. does an equilibrium exist? (This is not an empirical question, it is a question about the internal logical consistency of a model)
2. is equilibrium unique? (often important when thinking about comparative static issues)
3. is equilibrium stable? (only stable equilibria are likely to be observed, so only these are empirically relevant)

I will not address uniqueness or stability here, but I do address the existence issue.

<u>Pure exchange versus Production</u>

1. In a pure exchange economy (Ch. 17 of Varian), each consumer has an initial endowment of goods. These are traded at specific prices according to utility maximization. We usually want to know the equilibrium prices, quantities traded, and final consumption bundles. In a model of this kind, there are no firms and no new goods are produced.

2. GE with production (Ch. 18 of Varian). Here we introduce some firms as well as consumers. The consumers get income by selling inputs to firms and receiving profits from firms. We must have supply = demand both in product markets and input markets.

In these notes, I deal with the pure exchange case.

The pure exchange model: agents and goods

Consumers (agents) are indexed by $i = 1 \ldots n$

Goods are indexed by $j = 1 \ldots k$.

A consumption bundle for agent i is $x_i = (x_{i1} \ldots x_{ik})$

So x_{ij} is the quantity of good j consumed by agent i.

An allocation is a list of n consumption bundles, one for each agent, where each specifies a list of goods: $x = (x_1 \ldots x_n)$

Note that the dimensionality of the allocation x is nk; for consumer $i = 1$ we have the consumption bundle $(x_{11} \ldots x_{1k})$ with k elements, and likewise for $x_2 \ldots x_n$.

Each agent has a utility function $u_i(x_i)$ and an endowment $w_i = (w_{i1} \ldots w_{ik})$ stating how much of each good agent i initially possesses (before trading begins)

Feasiblity: $\sum_{i=1}^{n} w_{ij}$ is the total amount of good j available in the economy (you can also call it total supply).

This is the sum of each person's endowment of j.

In vector notation,

$$\sum_{i=1}^{n} w_i = \left(\sum_i w_{i1} \ldots \sum_i w_{ik}\right)$$

An allocation is feasible if $\sum_{i=1}^{n} x_i \leq \sum_{i=1}^{n} w_i$

or for each good,

$$\sum_i x_{ij} \leq \sum_i w_{ij} \quad \text{for all } j = 1 \ldots k.$$

The interpretation of $\sum_{i=1}^{n} x_i$ is that it is the vector of goods available after trade, while $\sum_{i=1}^{n} w_i$ is the vector of goods available before trade.
Note that we are allowing free disposal of excess endowments (any quantity that is not consumed).

Prices. There is a price vector $p = (p_1 .. p_k)$, one price for each good. The agents are price takers and at given prices, they choose a consumption bundle to max utility subject to a budget constraint:

$$\max_{x_i} u_i(x_i) \text{ subject to } p x_i \leq p w_i$$

The constraint says that the value of the consumption bundle x_i cannot exceed the value of the consumer's endowment vector w_i.

So in this model, $p w_i$ plays the role of income (m_i). We can find the Marshallian demands $x_{ij}(p, m_i)$ in the usual way, and then substitute $m_i = p w_i$ to get $x_{ij}(p, p w_i)$ which is i's demand for good j.

Total market demand for good j is $X_j(p, w_1 .. w_n)$

$$= \sum_{i=1}^{n} x_{ij}(p, p w_i)$$

Equilibrium in market j requires

$$\sum_{i=1}^{n} x_{ij}(p, p w_i) \leq \sum_{i=1}^{n} w_{ij}$$

(we allow excess supply but not excess demand)

We say that p^* is a <u>Walrasian equilibrium</u> price vector if

$$\underbrace{\sum_{i=1}^{n} x_i(p^*, p^* w_i)}_{\substack{\text{The vector of demands} \\ \text{for goods } j = 1..k}} \leq \underbrace{\sum_{i=1}^{n} w_i}_{\substack{\text{The vector of supplies} \\ \text{for goods } j = 1..k.}}$$

At the equilibrium prices, each individual is maximizing utility, and all consumer demands can be satisfied simultaneously.

Edgeworth box

To get some intuition about Walrasian equilibrium, we consider a simple model with two consumers (A and B) and two goods (1 and 2). This leads to the following graph, known as an Edgeworth box.

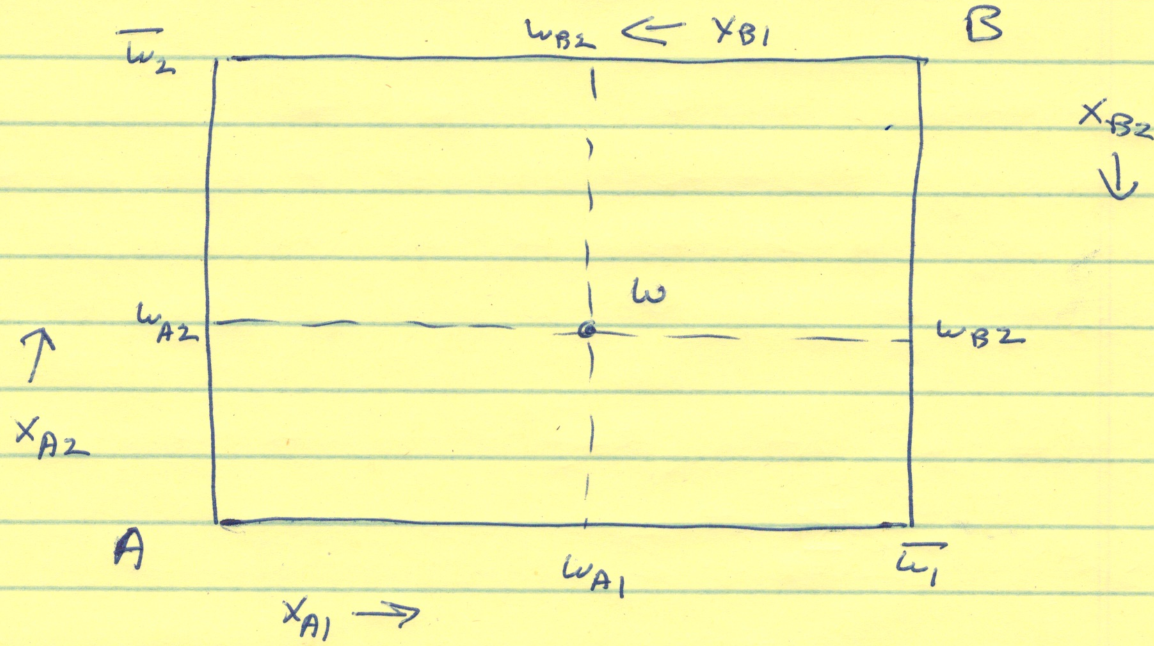

Things to observe about this:

1. The width of the box $\bar{\omega}_1$ is the total supply of good 1; the height of the box $\bar{\omega}_2$ is the total supply of good 2.
2. Person A's origin is the lower left corner, and we measure A's consumption of each good $x_A = (x_{A1}, x_{A2})$ from this origin.
3. Person B's origin is the upper right corner and we measure B's consumption $x_B = (x_{B1}, x_{B2})$ from this origin.
4. Each point x in the box is a feasible allocation: it gives a bundle x_A for A and (simultaneously) a bundle x_B for B. By construction we have $x_{A1} + x_{B1} = \bar{\omega}_1$ and $x_{A2} + x_{B2} = \bar{\omega}_2$.
5. The endowment point ω tells us the starting point for trading. A initially has the bundle $(\omega_{A1}, \omega_{A2})$ and B initially has the bundle $(\omega_{B1}, \omega_{B2})$.

In order to know what trades each person is willing to make, we have to include indifference curves for each person. We normally assume these involve strictly convex preferences for each person. A is better off as we move to the northeast, and B is better off as we move to the southwest.

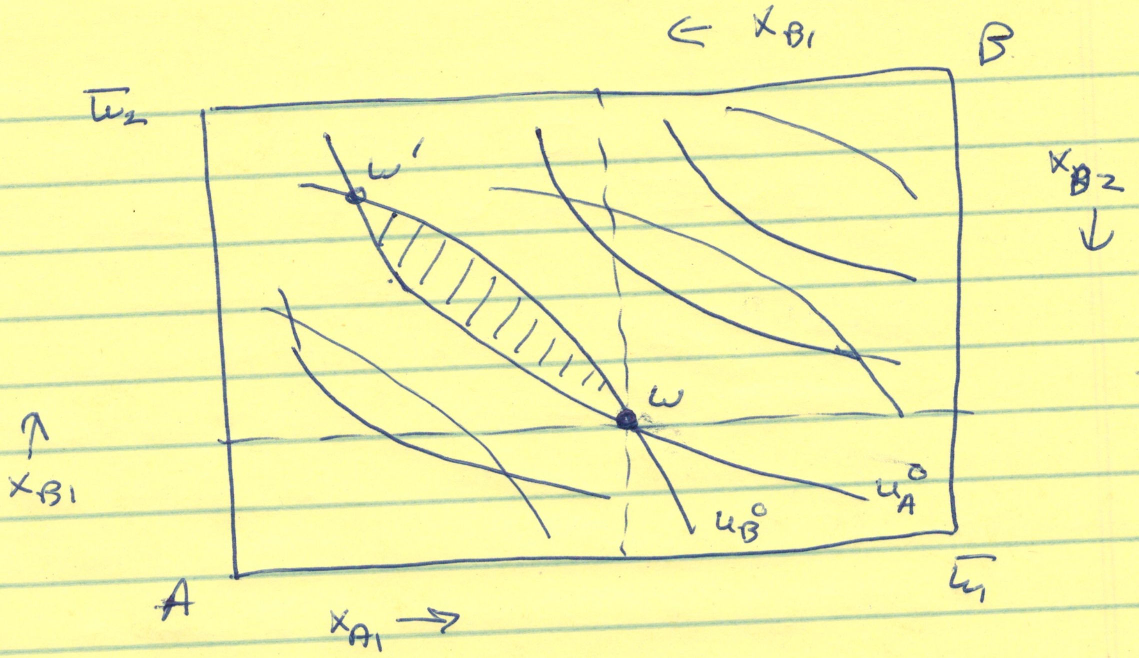

Starting from the endowment point w, person A is on indifference curve u_A^0 and B is on u_B^0. Assuming the slope of u_A^0 is flatter than the slope of u_B^0, there is a region (shaded) where both people are better off than at w. These allocations are Pareto improvements relative to w.

The same would be true if we started from an endowment point like w', where the slopes differ in the opposite way.

Now define an allocation to be Pareto efficient if it is impossible to find a Pareto improvement.

By process of elimination, a Pareto efficient allocation will need to have equal slopes for the indifference curves of A and B. (There will have to be a tangency point).

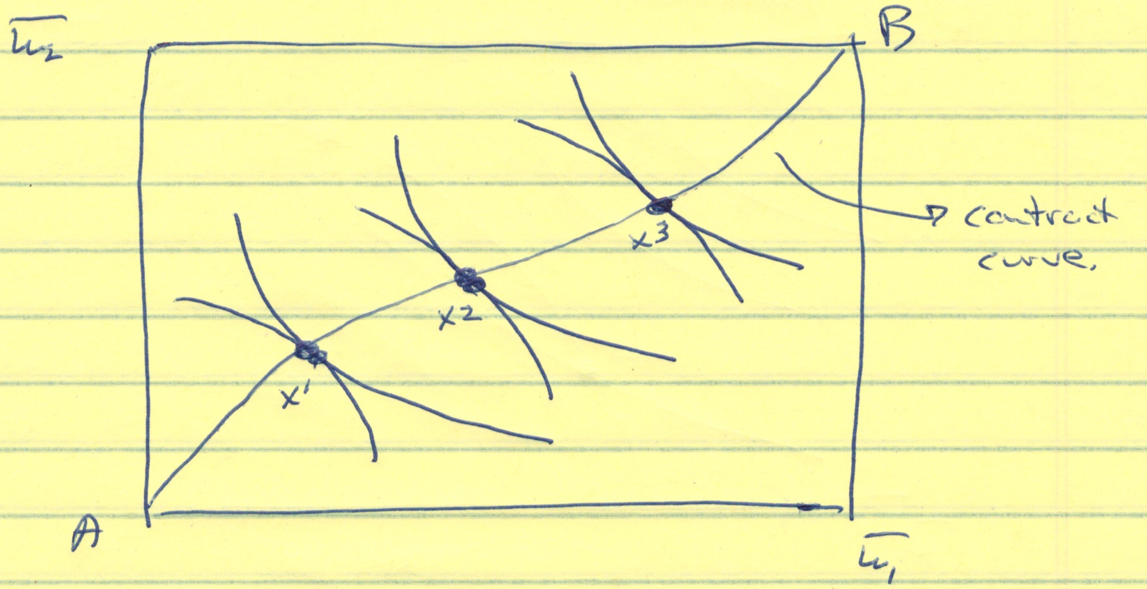

In general, there are many Pareto efficient allocations. The set of all Pareto efficient points is called the contract curve. You should be able to convince yourself that if you are starting from any point like x^1, x^2, or x^3, it is impossible to move to some other point that makes both people better off simultaneously (so it is impossible to have a Pareto improvement).

If the endowment point is not on the contract curve, it is reasonable to think that A and B will keep trading until they reach some point on the CC. Once they get there, trading stops.

Starting from ω, trade could lead to any point on CC between M and N (other points on CC would make someone worse off than at ω, so would not be acceptable to that person).

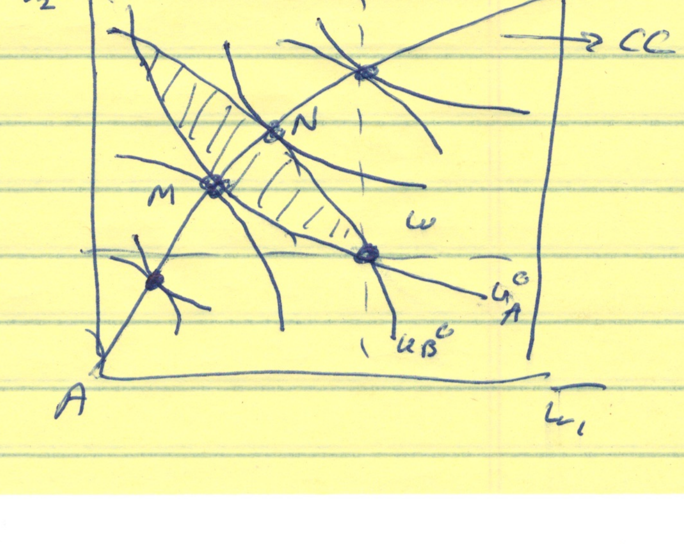

In general, we don't know where on the contract curve between M and N we will end up, starting from ω. One way to approach this problem is to construct a model based on bargaining theory. However here we are interested in what happens when the consumers have to trade at given prices

Suppose we have prices p_1 and p_2. Person A is limited to bundles x_A that satisfy

$$\underbrace{p_1 x_{A1} + p_2 x_{A2}}_{\substack{\text{value of} \\ \text{consumption bundle}}} = \underbrace{p_1 \omega_{A1} + p_2 \omega_{A2}}_{\text{income = value of endowment}}$$

Note: I will assume local non-satiation so all income is spent.

This gives a budget line for A that passes through ω and has the slope $-\frac{p_1}{p_2}$:

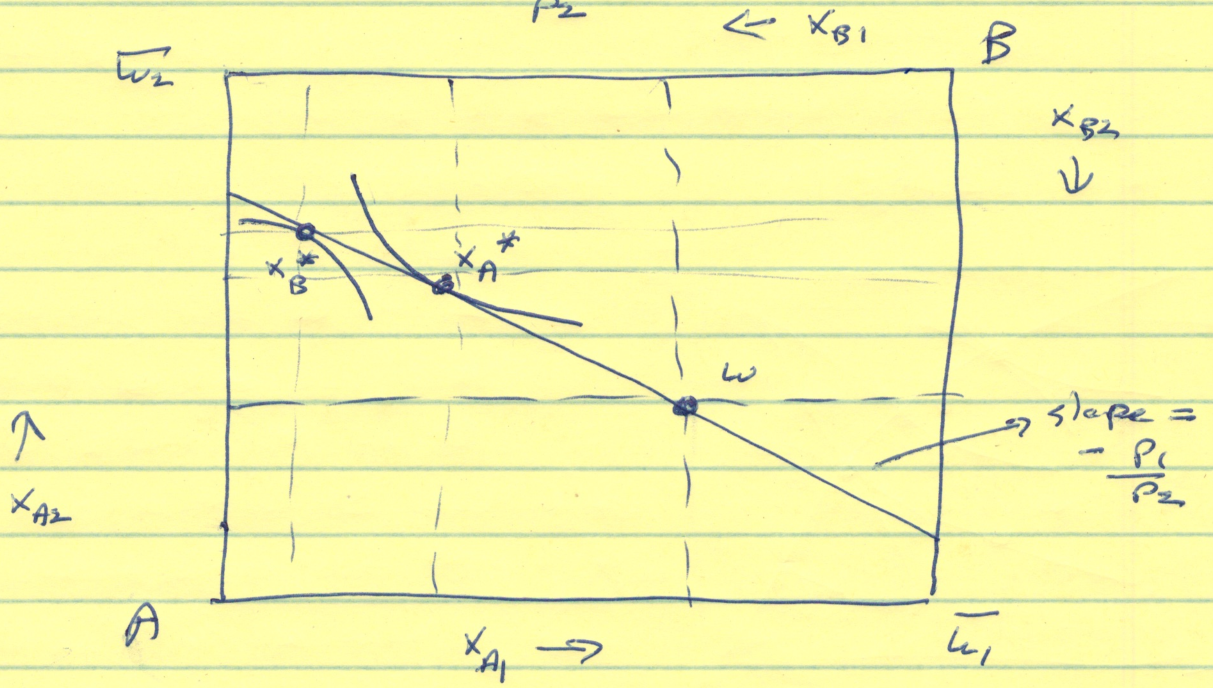

It should be clear that A can always afford to keep the endowment bundle ω_A (A can choose $x_A = \omega_A$) so the budget line must go through the endowment point. When A maxes utility in the usual way, this yields some optimal bundle like x_A^*.

Now think about the situation from B's perspective. B's budget constraint is $p_1 x_{B1} + p_2 x_{B2} = p_1 \omega_{B1} + p_2 \omega_{B2}$. This also has the slope $-\frac{p_1}{p_2}$ and it also passes through the endowment point (think of everything from the point of view of B's origin in the upper right). So B's budget line is identical to A's budget line. In general B's optimal bundle could be something like x_B^*.

If x_A^* and x_B^* look the way they do in the preceding graph, we do not have a Walrasian equilibrium. The reason is that when we add up $x_{A1}^* + x_{B1}^*$, we get a total that exceeds $\overline{\omega}_1$, so there is excess demand for good 1. At the same time, we have $x_{A2}^* + x_{B2}^* < \overline{\omega}_2$ so there is excess supply for good 2.

In general, there is always excess demand for one of the goods if x_A^* and x_B^* are different points on the budget line. However, consider what happens if x_A^* and x_B^* are the same point:

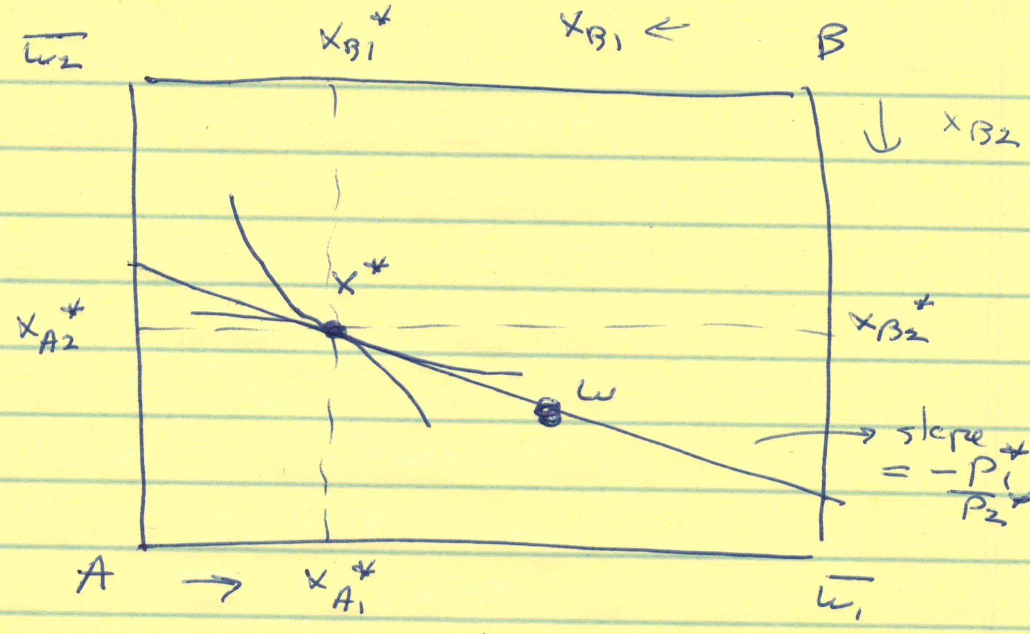

In this situation, $x_{A1}^* + x_{B1}^* = \overline{\omega_1}$

and $x_{A2}^* + x_{B2}^* = \overline{\omega_2}$

So supply = demand for both goods simultaneously. If we can find a price vector $p^* = (p_1^*, p_2^*)$ such that this is true, then we have a <u>Walrasian</u> <u>equilibrium</u>.

Important note: This only occurs if x^* is a point on the contract curve. So immediately, we have a glimpse of the first theorem of welfare economics: a Walrasian equilibrium must be Pareto efficient. We'll generalize this idea later.

The big question is: how do we know there <u>is</u> such a price vector? Can we be sure that a Walrasian equilibrium exists?

Existence of Walrasian Equilibrium

Now go back to the previous model with consumers $i = 1 \ldots n$ and goods $j = 1 \ldots k$. The main tool for showing the existence of WE is the aggregate excess demand function, defined as

$$z(p) = \sum_{i=1}^{n} x_i(p, p\omega_i) - \sum_{i=1}^{n} \omega_i$$

Note that $z = (z_1 \ldots z_k)$ is a vector of goods. We define p^* to be a Walrasian equilibrium if $z(p^*) \leq 0$. This says that at the prices p^*, there is no excess demand for any good, so all consumer demands can be satisfied at those prices.

To rule out excess supply, we would need $z(p^*) = 0$. I will come back to this issue later.

In order to know whether such a p^* exists, we have to study the properties of the function $z(p)$. Here are some things we know.

① Continuity. If each consumer i has strictly convex preferences, then s/he has continuous Marshallian demand functions, which implies that $z(p)$ is also continuous.

(2) <u>Homogeneity of degree zero in prices.</u>
Recall that Marshallian demands $x_i(p, m_i)$ are homogeneous of degree zero in (p, m_i). Since $m_i = p w_i$, it follows that $x_i(p, p w_i)$ is homog. of degree zero in p. Therefore $z(p)$ is homog of degree zero in p.
This implies that if p^* is a WE, so is $t p^*$ for any $t > 0$. For example, if we only have two goods, it will only be possible to solve for the <u>ratio</u> p_1/p_2, not the absolute levels p_1 and p_2.
More generally, only relative prices matter.
People normally handle this by fixing one price arbitrarily (setting some $p_j \equiv 1$). This good is called the <u>numeraire</u> and the prices of all other goods are measured relative to this good.
Another way to handle it is to normalize the price level in some other way, for example by imposing $\sum_j p_j = 1$.

(3) <u>Walras's Law.</u> Given local non-satiation, we can prove that $p \cdot z(p) \equiv 0$ (the value of the aggregate excess demand vector is always zero)
Note: this is true for <u>all</u> price vectors, not just an equilibrium price vector.
The proof involves manipulation of the budget constraints:

$$p \cdot z(p) = p \left\{ \sum_{i=1}^{n} x_i(p, pw_i) - \sum_{i=1}^{n} w_i \right\}$$

$$= \sum_i \{ \underbrace{p x_i(p, pw_i) - pw_i}_{\substack{\text{This is zero for all } i \text{ due to} \\ \text{the budget constraints}}} \} = 0$$

This has various implications.

(a) Suppose we have some $p^* > 0$. If $k-1$ markets clear ($z_j(p^*) = 0$ for $j = 1 \ldots k-1$) then the last market must also clear ($z_k(p^*) = 0$). This allows us to delete one market (drop one equation). For example, in the Edgeworth box model, we only need supply = demand for one good; this automatically implies $S = D$ for the other good.

Due to homogeneity, we can drop one variable and work with $k-1$ prices (or price ratios);
due to Walrasian's Law, we can drop one equation and work with $k-1$ equilibrium conditions.
This is often a very useful simplication.

(b) Walras's Law implies that in equilibrium, if we have $z_j(p^*) < 0$ for some good j, then its price is $p_j^* = 0$. Thus goods in excess supply must be free. This follows from $\sum_j p_j^* z_j(p^*) = 0$ and $z_j(p^*) \leq 0$ for all j; if we had both $z_j(p^*) < 0$ and $p_j^* > 0$, the sum would be negative and we would contradict Walras's Law.

(i) Another implication: suppose it is true that whenever $p_j = 0$, we have $z_j(p) > 0$. This is often a reasonable assumption about consumer preferences: if something is free, the aggregate demand for it will exceed supply. This assumption is called "desirability". If all goods are desirable in this sense, then any WE price vector p^* must give $z(p^*) = 0$ so all markets clear exactly (no excess supply).

The reason is this. Suppose we had $z_j(p^*) < 0$ for some j. We already know from (h) above that this implies $p_j^* = 0$. But desirability for j says that whenever $p_j = 0$, we have $z_j(p) > 0$. Therefore $z_j(p^*) > 0$, which contradicts $z_j(p^*) < 0$. Thus $z_j(p^*) < 0$ is impossible, and we must have $z_j(p^*) = 0$ so there is no excess supply for j in equilibrium.

<u>Proof of Existence for Walrasian Equilibrium</u>

If we have ① continuity ② homogeneity and ③ Walras's Law, then it can be shown mathematically that there is some p^* such that $z(p^*) \leq 0$. The proof involves use of a fixed point theorem (see Varian for details). Also, if $z_j(p^*) < 0$ then $p_j^* = 0$. If all goods are desirable, then there exists a price vector p^* such that $z(p^*) = 0$.

The First and Second Welfare Theorems

We already saw from the Edgeworth box model that there is a close connection between Walrasian equilibrium (WE) and Pareto efficiency (PE). Here I will generalize these ideas.

There are various ways to define Pareto efficiency. It will be convenient to define it as follows:

The allocation x^* is <u>Pareto efficient</u> if there is no feasible allocation x' such that $x_i' \succ_i x_i^*$ for all $i = 1..n$

Note: the definition of PE has nothing to do with prices - it only involves preferences and aggregate resource endowments.

This says that x^* is PE if it is impossible to make everyone strictly better off simultaneously. Note that $\succ_i$ is the preference ordering of consumer i.

<u>The First Theorem of Welfare Economics:</u>
If (p^*, x^*) is a WE then x^* is PE.

<u>Proof:</u> suppose x^* is not PE. Then there is a feasible x' such that $x_i' \succ_i x_i^*$ for all i. Because (p^*, x^*) is WE, it must be true that $p^* x_i' > p^* x_i^*$ for all $i = 1..n$. This follows from the fact that if $p^* x_i' \leq p^* x_i^*$ for some i, that consumer could have chosen the strictly preferred bundle x_i' but did not, which contradicts utility max in WE. Therefore we must have

$$\sum_i p^* x_i' > \underbrace{\sum_i p^* x_i^* = \sum_i p^* \omega_i}_{\text{using the budget constraints in WE}}$$

But we said x' is feasible, which implies

$$\sum_i x_i' \leq \sum_i \omega_i$$

and therefore

$$\sum_i p^* x_i' \leq \sum_i p^* \omega_i$$ (This requires a little bit of algebra)

Now we have a contradiction with the result at the top of the page. This shows that the original supposition that x^* is not PE must be false. Therefore x^* is PE.

QED

What about the converse? If we are given some Pareto efficient allocation x^*, can we find prices p^* such that (p^*, x^*) is a Walrasian equilibrium? In general the answer is yes, but we need to be able to choose the individual endowments ω_i (of course, we can't change the sum $\omega = \sum \omega_i$ since this is the aggregate supply of resources).

The proof I will give involves a short cut where I use the existence of a Walrasian equilibrium. Here is how it goes.

Second Theorem of Welfare Economics

Given an aggregate endowment vector ω if x^* is PE then there are prices p^* and endowments ω_i for $i = 1 \ldots n$ with $\sum_i \omega_i = \omega$ such that (p^*, x^*) is a WE.

Proof: Set $\omega_i = x_i^*$ for all $i = 1 \ldots n$, so the individual endowments are equal to the consumption bundles we are trying to achieve. This ensures $\sum_i \omega_i = \sum_i x_i^* = \omega$ because x^* is PE and thus it must be feasible.

We know a WE exists for these endowments (given continuity, homogeneity, and Walras's Law).

Call this WE (p', x').

We also know that $\underbrace{p' x_i' = p' x_i^*}_{\text{from budget constraints in WE}} \underbrace{= p' \omega_i}_{\text{from how we choose the individual endowments}}$ for all $i = 1 \ldots n$

Therefore $x_i' \succsim_i x_i^*$ for all $i = 1 \ldots n$. This follows from the fact that i can afford to buy both x_i' and x_i^* at the prices p' but actually choose x_i' in WE, so x_i' must be at least as good as x_i^* according to i's preferences.

We know x^* is PE by assumption and we know x' is feasible (because it is WE and markets clear)

If we had $x_i' \succ_i x_i^*$ for any i, this would imply that a Pareto improvement is possible, and would contradict the fact that x^* is PE.

Note: I have changed the definition of PE slightly, so that x^* is PE if it is impossible to make someone better off without making anyone worse off. This is a technical detail and you don't need to worry about it.

This line of argument shows that $x_i' \sim_i x_i^*$ for all $i = 1 \ldots n$. But we know (p', x') is WE. So x_i' maxes utility subject to $p'x_i \leq p'w_i$ for all i and therefore x_i^* also maxes utility subject to $p'x_i \leq p'w_i$ for all i.

Hence (p', x^*) is <u>also</u> a WE (everyone is maxing utility at the prices p' and markets clear by construction:

$$\sum_i w_i = \sum_i x_i^* = w \;)$$

This shows that if we are given any allocation x^* that is PE, we can construct prices and individual endowments such that x^* occurs in WE. (Sometimes economists say "x^* can be supported by a WE").

The main drawback of the preceding proof is that we had to assume existence of WE. There are other proofs involving separating hyperplanes that avoid this short cut (see Varian for details).

Calculus version of Walrasian Equilibrium and Pareto Efficiency (Varian sections 17.8–17.9)

Suppose we have differentiable utility functions. Let (p, x^*) be a WE. Then for each consumer i,

x_i^* solves $\max_{x_i} u_i(x_i)$ subject to $px_i = p\omega_i$
(assuming non-satiation)

FOC: $\dfrac{\partial u_i(x_i^*)}{\partial x_{ij}} = \lambda_i p_j$ all $j = 1 \ldots k$
where λ_i is the multiplier in consumer i's problem.

I will assume u_i is strictly concave. Therefore it is strictly quasi-concave, the problem has a unique solution, and the FOC are sufficient for a max.

Now consider the problem of maximizing "social welfare." A benevolent social planner might want to allocate resources to solve

$$\max_{x_1 \ldots x_n} \sum_{i=1}^{n} a_i u_i(x_i) \quad \text{subject to} \quad \underbrace{\sum_{i=1}^{n} x_{ij} = \sum_{i=1}^{n} \omega_{ij} \quad \text{for } j = 1 \ldots k}_{\text{physical feasibility constraints}}$$

where $a_i > 0$ for all i
(These are weights that indicate how much the planner cares about each consumer $i = 1 \ldots n$)

Notice that the planner's problem only involves physical quantities — there are no prices or budgets.

The Lagrangean for the planner is

$$L = \sum_{i=1}^{n} a_i u_i(x_i) - \sum_{j=1}^{k} q_j \left[\sum_{i=1}^{n} x_{ij} - \sum_{i=1}^{n} \omega_{ij} \right]$$

where the q_j are multipliers for the feasibility constraints

FOC: $\quad a_i \dfrac{\partial u_i(x_i^*)}{\partial x_{ij}} = q_j \quad$ all $i = 1..n$, all $j = 1..k$

It is easy to see that any allocation $x^* = (x_1^* .. x_n^*)$ that solves this problem must be Pareto efficient.

Proof: Suppose x^* solves the problem but is <u>not</u> PE. Then there is some other feasible allocation x' that is preferred by everyone, so it is possible to increase the objective function. This contradicts the assertion that x^* is a solution.

Another question is whether <u>every</u> PE allocation can be represented as the solution of such a problem for some choice of the weights $a_i > 0$ for $i = 1..n$? The answer is yes, but I won't prove it here.

<u>Note</u>: because the u_i are assumed to be strictly concave, the planner's objective is strictly concave, which ensures that the FOC are sufficient for a (unique) solution.

Now compare the FOCs for WE and social welfare:

WE: $\frac{\partial u_i(x_i^*)}{\partial x_{ij}} = \lambda_i p_j$ all i, j

SW: $\frac{\partial u_i(x_i^*)}{\partial x_{ij}} = \left(\frac{1}{a_i}\right) q_j$ all i, j

Both allocations satisfy the feasibility constraints

$$\sum_i x_{ij}^* = \sum_i \omega_{ij} \quad \text{for all } j = 1..k$$

(in the case of WE, this comes from market clearing)

Clearly if we set $\lambda_i = \left(\frac{1}{a_i}\right)$ for all i

and $p_j = q_j$ for all j

These conditions are identical and the same x^* satisfies both. Due to strict concavity, the FOCs are sufficient. Therefore:

A WE allocation x^* always maximizes social welfare for some choice of the weights $(a_1 \ldots a_n) > 0$, and therefore x^* is Pareto efficient.

This gives an alternative proof of the first welfare theorem.

A few comments:

① In equilibrium, the market behaves "as if" it is maximizing a weighted sum of utilities.

(2) The implicit weights a_i are $\frac{1}{\lambda_i}$, reciprocals of the marginal utility of income in the individual utility max problem [so if you have a big endowment and consume a lot, your MU of income is low, so λ_i is low and a_i is big; it is like having a social planner who cares about you a lot]

(3) The prices of goods p_j can be interpreted as Lagrange multipliers on the feasibility constraints. To put it another way, they are "shadow prices" that indicate the derivative of social welfare with respect to the total supply of a good.

(4) The same methods can be used to prove the second welfare theorem. Start with any PE allocation x^*. It can be shown that there are some weights $(a_1, \ldots a_n) > 0$ such that x^* maximizes social welfare. Choose the endowments so that $w_i = x_i^*$ for all i. This will ensure that each i can afford to buy x_i^*. Finally, set the prices so that $p_j = q_j$ for all j. This will ensure that the FOCs for WE are satisfied, and by feasibility all markets will clear. Given strict concavity, the FOCs are sufficient, and we get a Walrasian equilibrium.

(5) A key element of both WE and PE is that the marginal rate of substitution between any pair of goods must be identical for all consumers.

Note: all of This is much more general Than what we did in Chapter 13, where we used quasi-linear utility functions, two goods, and maximized The sum of The utilities. Here we allow any strictly concave utility functions, any number of goods, and any set of positive weights in The social welfare function.

Closing Thoughts:

① PE does not imply fairness. There could be a very unequal distribution of goods or utilities. For two people, we get something like This:

The set of utility pairs along The UPF corresponds to The set of Pareto efficient allocations. So PE could give a lot to person 1 and not much to person 2 or vice versa.

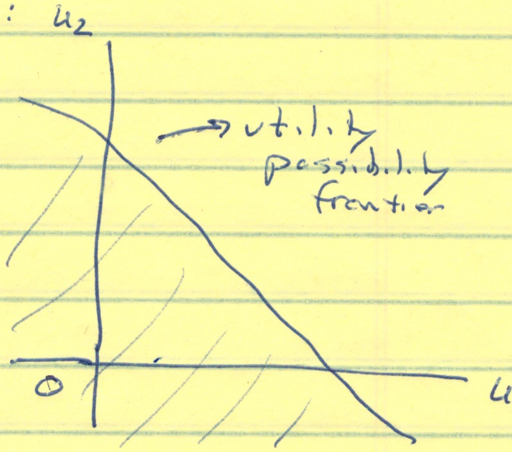

[Think about The differences in utilities as you move along The contract curve in an Edgeworth box]

② in WE, endowments determine which PE allocation we reach; a big endowment means you will get a lot of utility and vice versa.

③ if a different PE allocation is desired, in principle we could achieve This by ~~the~~ redistributing endowments

In practice, there are limits to direct redistribution of endowments. This is especially true for endowments of time, human capital, and so on. Real governments try to deal with this by equalizing access to human capital through education, redistributing money rather than physical goods, etc. However, taxes and transfers are hard to do in a "lump sum" way that doesn't affect prices.

There are also information problems: how do we compute PE allocations or equilibrium prices, given that people may lie about their true preferences (and the computational problems would be huge even if everyone told the truth). In practice, we typically just move endowments in a given direction rather than aiming at a specific WE.

Of course, there are political issues with redistribution; some countries do a lot more of it than others.

Finally, keep in mind that everything we have done here assumes price-taking behavior. The welfare theorems do not apply in situations with monopoly power, public goods, externalities, or informational asymmetries. The theorems are interesting, but the real world is much more complicated!