Microeconomic Theory
Lecture 5: General Equilibrium and Welfare Theorems

Marcel Fafchamps
Oxford University

1. General Equilibrium

Readings: Varian, Chapters 17 and 18.

One of the most fundamental contribution of economics to the world is an understanding of general equilibrium effects. This goes back all the way to Adam Smith and his insightful observation that even though the shoemaker and farmer all believe they seek their self-interest, by trading with each other they all de facto work for each other and pursue the common good: the shoemaker does the shoemaking, which he is good at, and the farmer does the farming, which he is good at, so that tasks that are required for the common good are taken up by the best qualified people. In the end the farmer consumes the shoemaker's shoes and the shoemaker consumes the farmer's food. This kind of issue is what attracted the attention of the classical economists, at a time where the reach of the market was expanding thanks to new infrastructure (canals and shipping, mainly). They needed a way to think about the aggregate outcome of many decentralized individual actions, and what they came up with was economics.

Another important early insight is that there is no way for a group of people to consume more goods than what they produce themselves or purchase from the rest of the world. This apparently obvious observation is nevertheless an endless source of confusion. This is because what is true for individuals – one needs money for purchase goods – is not true in aggregate: money does not 'buy' anything in the aggregate because agents can only consume what they produce or exchange from the rest of the world. An immediate consequence of this observation is that if the demand for a good exceeds its supply (the quantity available), then price must go up to equilibrate the market. Similarly, if the demand for all goods exceeds aggregate supply for all goods, then all prices must go up to equilibrate the market.

It is also possible to equilibrate the market through other means, such as rationing and coercion. There are various forms of rationing (e.g., queuing, random allocation, allocation based on self-proclaimed needs, fighting). All rationing allocations are fragile because side-deals can nearly always be found that increase individual – and aggregate – efficiency. For instance, allocation based on self-declared needs can be manipulated by overstating one's needs. If everyone overstates and allocation is pro-rated, the only incentive-compatible statement of needs is infinity; hence the system collapses.

To take another kind of rationing, suppose there is a queuing system in place for, say, gasoline. In this system, it is in the interest of the gasoline owner to take a bribe to let someone ahead of the queue.[1] It is also in the interest of hurried consumers to offer such a bribe. Suppose

[1] It is also in the interest of an *employee* of the gasoline distribution firm to take a bribe to let someone ahead of the queue. In this case, the bribe goes to the employee, not to the owner of the gasoline. But this problem is not specific to rationing: it is always in the interest of employee of a firm to divert firm resources to their pocket. Disciplining workers is a problem common to rationing and market equilibria.

the person most in a rush pays a bribe first, then the person second highest in a rush, etc. Bribing in this case operates a little bit like an auction system: gasoline goes to those who are willing to pay the most for it. Since this is precisely what would happen in a market anyway, bribing generates an allocation that is not very different from the market, except that more rent is captured by the gasoline owner. We revisit this issue later in the course. In contrast, market exchange is not vulnerable to side-deals because it is voluntary for both sides and competition ensures that if a trade takes place, there does not exist another buyer who would pay more for the same good or a seller who would sell for less.

General equilibrium theory is an effort to formalize all these ideas and to derive conditions under which markets achieve efficient outcomes, that is, outcomes that cannot be improved upon by side-deals or by government intervention. Imperfect information and enforcement issues are ignored throughout this section.

2. A Pure Exchange Economy

Throughout this lecture, we ignore the possibility of trade with the rest of the world. Trade can easily be added along the lines of the household with missing markets studied in lecture 2. But for the simplicity of exposition, we limit our presentation to closed economies.

We begin with what is known as a pure exchange economy. In this economy, there is no production. Individuals are endowed with certain consumption goods and with preferences over these goods. Their joint welfare can be improved upon by exchanging these goods among themselves. To follow Varian's notation, let the initial endowment vector of individual i be denoted ω_i (instead of T_i as in lecture 3) and let the consumption vector of the same individual be denoted x_i. The collection of all consumption bundles is called an allocation. The theory of general equilibrium is concerned with the process of going from endowments ω_i to allocations x_i.

A feasible allocation is one that is physically possible:

$$\sum_{i=1}^{n} x_i \leq \sum_{i=1}^{n} \omega_i \tag{2.1}$$

The inequality is there in case some goods are in excess supply in equilibrium – e.g., garbage, air. If all goods are consumed, then equation 2.1 is an equality. [Note that equation 2.1 is written in vector notation. It implies that the aggregate consumption of each individual good, say milk, is equal to the total availability of that good, e.g., milk, in the economy. This is true for each good.] With two goods and two agents, the endowment-allocation mapping can conveniently be represented by an Edgeworth box (Figure 17.1 in Varian). (Discuss how the box is constructed and what it means.)

Assume that agents are self-interested (no altruistic motives). If agents cannot be forced to exchange with others, i.e., if they must trade voluntarily, then they will probably block all allocations that give them less than their autarchic payoff, that is, less than the utility they would get from consuming their initial endowment. This means that, in a voluntary exchange economy, initial endowments determine the level of utility below which agents cannot be forced. Allocations that satisfy this requirement are said to belong to the core. If endowments are unevenly distributed, this may imply that all equitable distributions can be blocked by well endowed individuals. By this I mean that individuals with high initial endowments can refuse exchanges that do not give them a high utility, thereby reducing the amount of stuff that can be redistributed to other, less fortunate agents. There is therefore a limit to the amount of redistribution that can be achieved through market forces.

2.1. Walrasian Equilibrium

Following Walras, we now examine the allocation that would arise if all agents pursue their self-interest. Formally, we assume that each agent solves the following problem:

$$\max_{x_i} \; u_i(x_i) \quad \text{subject to} \quad px_i = p\omega_i \tag{2.2}$$

The answer to this maximization problem is a set of demand functions $x_i(p, p\omega_i)$. This is similar to the consumer demand function studied in lecture 1, except that here the dependence of cash-in-hand m_i on the price of endowments is explicitly recognized (think of the price of endowments in lecture 3 pT_i). The aggregate demand for all goods, expressed in vector notation, is $\sum_i x_i(p, p\omega_i)$. The aggregate supply is $\sum_i \omega_i$. A Walrasian equilibrium is one in which demand does not exceed supply. Formally:

$$\sum_i x_i(p, p\omega_i) \leq \sum_i \omega_i \tag{2.3}$$

(Show graphically using an Edgeworth box.) The Walrasian equilibrium describes the equilibrium that would be attained if all agents take prices as given, know the quality and type of what they buy, and can costlessly enforce agreements among themselves. In this sense, it represents the equilibrium that is achieved by an idealized economy, often called the perfect competitive economy. By extension, it is also called a market equilibrium because the allocation is achieved in an entirely decentralized manner via voluntary trade among individual agents. Note, however, that actual markets fail to satisfy many of the conditions for a Walrasian equilibrium.

Does a Walrasian equilibrium exist? The answer is in general yes, provided preferences respect some intuitive and technical conditions. Let's define the aggregate excess demand function $z(p)$ as the difference between demand and supply (endowments).

$$z(p) \equiv \sum_i [x_i(p, p\omega_i) - \omega_i]$$

We begin by noting that, since each demand function is homogenous of degree 0 in all prices, so is $z(p)$. Furthermore, the aggregate excess demand function must satisfy Walras Law, namely, for any price vector p, we have $pz(p) \equiv 0$, the value of excess demand is identically zero. The reason is that each consumer satisfies his individual budget constraint: Walras Law is instantly obtained by summing up all budget constraints of all agents. In other words, Walras Law is nothing else than an accounting identity. Walras Law is not an equilibrium condition: it is satisfied simply because individual agents can only spend what they earn/the value of their endowments. In this sense, it is 'always true' even if the equilibrium is not Walrasian, provided each agent is on its budget constraint (does not spend more than what it has).

If we combine Walras Law with the definition of a Walrasian equilibrium, we have a market clearing condition which says that, if demand equals supply in $k-1$ goods, and $p_k > 0$, then demand must equal supply in the kth market as well.

Another intermediate result of importance is that, if all goods are desirable in the sense that if $p_k = 0$, then $z_k(p) > 0$ for all goods, then the aggregate excess demand function is equal to 0 in equilibrium, i.e., $z(p^*) = 0$ for p^* an equilibrium. This result is fairly intuitive: if all goods are desirable, there should not be unconsumed endowments in equilibrium. Another corollary is that in equilibrium undesirable goods k have a zero price and excess supply in equilibrium, i.e., $p_k = 0$ and $z_k(p) < 0$.

2.1.1. Existence of an equilibrium

All existence proofs rely on the same kind of argument, called a fixed-point theorem. The specific kind of fixed point theorem may vary, however, – and the complexity of the mathematical argument occasionally gets a bit out of hand. The basic argument, however, is extremely simple. It goes like this. Since only relative prices matter, define normalized prices:

$$p_k = \frac{\hat{p}_k}{\sum_{m=1}^{M} \hat{p}_m} \tag{2.4}$$

where $\hat{p}_k$ now stands for absolute (i.e., nominal) price. By construction, normalized prices sum to 1. Finding an equilibrium thus boils down to finding a set of numbers that sum to one and satisfy the Walrasian equilibrium conditions. Brouwer's fixed point theorem states that any continuous function $f(.)$ that maps a set of numbers onto itself (i.e., lives in a finite space) has at least one fixed point at which $f(x) = x$. This is actually a very simple and intuitive result, as can easily be shown graphically.

This theorem is used to demonstrate the existence of a Walrasian equilibrium by constructing a function $f(p)$ that returns a higher price for good i if there is excess demand for that good, and a lower price if there is excess supply (negative excess demand). One easy example is simply $f(p) \equiv z(p) + p$. It is of course possible to add more complications and worry about pathological cases. But the argument remains the same: construct a continuous function that returns a higher price if there is excess demand, a lower price if there is excess supply, and the same price if excess demand is zero. Any such continuous function will do. There is a price vector so that the function returns itself, and this is the Walrasian equilibrium.[2]

That an equilibrium exists does not imply that this equilibrium is unique. (Discuss with Edgeworth box.) Unicity of equilibrium usually requires stringent conditions, especially when production is added to the model. We revisit this issue later.

2.2. Pareto Efficiency

Following Pareto, economists define efficiency as follows. A feasible allocation x is a weakly Pareto efficient allocation if there is no feasible allocation x' such that all agents strictly prefer x' to x. A feasible allocation is a strongly Pareto efficient allocation if there is no feasible allocation x' such that all agents weakly prefer x' to x, and some agent(s) strictly prefers x' to x. In practice, we usually work with models in which the two efficiency concepts coincide, so that we will speak of Pareto efficiency (see Varian for details).

Pareto efficient allocations can be found by solving problems of the following form (assuming two agents only):

$$\max_{x_1, x_2} u_1(x_1) \text{ subject to } u_2(x_2) \geq \bar{u}_2 \text{ and } x_1 + x_2 = \omega_1 + \omega_2 \tag{2.5}$$

The above maximization problem treats the utility of agent 2 $\bar{u}_2$ as given and seeks to maximize the utility of agent 1, subject to the feasibility condition. Since the choice of agent is arbitrary, Pareto efficient allocations can also be found by maximizing 2's utility subject to 1's being held constant. All Pareto efficient allocations can be found by changing $\bar{u}_2$. (Illustrate with Edgeworth box.)

[2] This does NOT imply that iterating on any arbitrary such function will necessarily lead to the equilibrium. In other words, the existence theorem does not imply that an equilibrium search algorithm based on such function would necessarily converge to an equilibrium. Convergence typically requires that the function be chosen very carefully, but this is not necessary to demonstrate existence.

2.2.1. First Welfare Theorem:

A Walrasian equilibrium is Pareto efficient. The proof follows from the fact that, given that everybody optimizes, if x_i' is preferred by agent i to x_i, then it must not be affordable, $px_i' > p\omega_i$. This theorem is truly a remarkable result. Not surprisingly, it has been extensively used as justification for market liberalization the world over. Note, however, that the theorem does not say anything about which Pareto efficient allocation is picked. In practice, the resulting allocation may be quite inequitable.

2.2.2. Second Welfare Theorem:

The second welfare theorem is arguably the theoretical result that has had the most dramatic effect on economic thinking. It goes as follows:

(First form) *Every Pareto efficient allocation can be supported as a Walrasian equilibrium with the right kind of initial endowments.* More formally, suppose that x^* is a Pareto efficient allocation in which the agent holds a positive amount of each good. Suppose that preferences are convex, continuous, and monotonic. Then x^* is a Walrasian equilibrium for the initial endowments $\omega_i = x_i^*$ for $i = 1, ..., n$. The theorem may seem a bit vacuous since it requires that all agents receive the Pareto efficient allocation as endowment. We will see later that this requirement can be somewhat reduced. But the interest is elsewhere: ANY Pareto efficient allocation can be supported as a competitive equilibrium. This means that any efficient allocation can be decentralized, provided endowments are right. Another implication is that equity concerns should be addressed by affecting endowments directly, and letting agents trade freely afterwards.

This version of the second welfare theorem has been used to justify land reform as an efficient way of dealing with poverty. Free schooling can also be seen in this light as a way of redistributing human capital to the poor – relative to the level of education they would receive if they had to pay the market price for it.

The attractiveness of the two welfare theorems comes from the fact that the market mechanism is a completely decentralized mechanism in which every agent is convinced he or she is pursuing its own interest. The resulting allocation is thus robust to attempts to manipulate the outcome, i.e., it is not subject to corruption or bribery. Of course, the redistribution of endowments by the state is subject to bribery.

There is another version of the second welfare theorem in which the redistribution of physical endowments is replaced with monetary transfers and taxes. The key to this alternative version of the theorem is that transfers and taxes should not distort the economy – i.e., they should be lumpsum transfers and taxes.

(Second form): *Any Pareto efficient allocation x^* can be supported as a Walrasian equilibrium by giving to each agent the cash-in-hand they need to purchase x^* at market equilibrium prices.*

Here I have somewhat generalized the second welfare theorem as follows. Suppose we are given a Pareto efficient allocation and we are told to support it as a market equilibrium. From the first form of the second welfare theorem, we know that if we distribute the allocation itself to agents, then they can certainly afford it. In practice, this may be a bit difficult though. The second form says that we can achieve the same result by redistributing endowments such that each agent has cash-in-hand equal to px^*, the value of the allocation at market equilibrium prices. In other words, we do not need to redistribute all the goods themselves, we only have to ensure that each agent receives a lumpsum transfer equal to $\sum_g p_g[x^{g*} - \omega^g]$ where g is the index of good g.

In this modified form, the second welfare theorem has been used extensively to argue that issues of efficiency can be dealt with separately from issues of equity. It has also been used to argue that taxes and transfers should be as neutral as possible in terms of incentives. The basic idea is that if an equity issue arises and becomes politically sensitive (e.g., farmers and a livestock epidemic), it should be dealt with via transfers not via distortionary measures such as price support, tariffs, quota, and the like. These views have had a major influence on the reorganization of the CAP and of WTO negotiation rounds, for instance.

2.2.3. Welfare Weights and Social Welfare Function

Another way of finding Pareto efficient allocation is to solve the following maximization problem:

$$\max_{x_i} \sum_{i=1}^{n} a_i u_i(x_i) \text{ subject to } \sum_{i=1}^{n} x_i^g \leq \sum_{i=1}^{n} \omega_i^g \text{ for } g = 1, ..., k \tag{2.6}$$

where the superscript g denote an individual good. This problem differs from the previous one in that what is maximized is the weighted sum of all individual utilities. Welfare weights are a_i. This objective function is sometimes referred to as a social welfare function. A theorem states that any allocation that solves the above maximization problem is Pareto efficient. Another theorem states that any Pareto efficient allocation is the solution of the maximization problem above for a suitably chosen set of welfare weights. Finally, a theorem states that these suitable welfare weights are the inverse of the marginal utility of income at the Walrasian equilibrium that corresponds to this efficient allocation (Do we know that such an equilibrium exists?)

3. An Economy with Production

The analysis of general equilibrium can be extended to the case where agents can transform some of their endowments to produce other goods. We now consider an economy with producers and consumers. The state is ignored. So, strictly speaking, the model we write below is more appropriate for a village or a small group of people than for an entire country. But the state can also be added, as is typically done in Computable General Equilibrium models.

3.1. Producers

Producers are assumed to behave in a competitive manner, that is, to regard prices as given. To keep notation as compact as possible, we define a net output vector y_j for each firm j. As in lecture 2, this vector has negative entries for inputs and positive entries for outputs. The profit definition is simply py_j for firm j. The price vector p includes all input and output prices. We ignore fixed factors of production in our notation. Profit maximization yields an optimal production plan given the price vector p as well as a net supply function (output supply and input demand) $y_j(p)$ for each firm j.

We define the aggregate net supply function as

$$y(p) \equiv \sum_j y_j(p)$$

Our first result of interest is that aggregate profits are maximized if each firm individually maximizes its own profit – keeping all prices constant. More formally, an aggregate production plan y maximizes aggregate profit py if and only if each firm's production plan y_j maximizes its individual profit py_j. The logic of this proposition is clear: suppose that a firm is not maximizing its profit. Keeping all prices unchanged, aggregate profit can clearly be increased by forcing this firm to maximize its profit.

It is convenient to assume that the aggregate net supply function is well-behaved, but a number of conditions are required for this to be the case. Without going into details, we can say that in general problems are most likely to arise when returns to scale are increasing, or when there are economies of scope (that is, some kind of joint production externality between different goods, such as when producing honey helps the production of apples). In these cases, the production possibility set of the economy as a whole is not convex and the aggregate net supply function need not be a function at all, but rather a correspondence – multiple net supply quantities correspond to a single price vector. We will not cover these difficulties here. We focus our attention exclusively on economies where $y(p)$ is a function. This typically requires decreasing or constant returns to scale at the firm level, and no joint production externalities.

3.2. Consumers

We assume that each consumer i is a price taker and solves the following optimization problem:

$$\max_{c_i,l_i} u_i(c_i, l_i) \text{ subject to}$$

$$pc_i + pl_i = p\ \bar{c}_i + w\bar{L}_i$$

where l denotes leisure, $\bar{L}$ denotes total time, and $\bar{c}$ denotes an initial endowment of good c. You will notice that so far the notation is quite similar to the household model with missing markets that we studied in lecture 2. Consumers need not have the same preferences. The above consumer's optimization problem can be rewritten more compactly as:

$$\max_{x_i} u_i(x_i) \text{ subject to } px_i = p\omega_i$$

where ω_i stands for a vector of endowment $(\bar{c}_i, \bar{L}_i)$ and x_i stands for (c_i, l_i).

Something is missing from the above formulation: profits. We assume that profits go to consumers. Let T_{ij} denote the share of profits of firm j going to consumer i. The revised budget constraint now becomes:

$$px_i = p\omega_i + \sum_j T_{ij} p y_j(p) (\equiv m_i)$$

The consumer is assumed to maximize his utility $u(x)$ subject to the above budget constraint, taking as given his total income m_i. This yields a vector of consumer demand $f_i(p, m_i)$. Since m_i is itself a function of prices, the consumer demand function can be written more compactly as

$$x_i(p) \equiv f_i(p, p\omega_i + \sum_j T_{ij} p y_j(p))$$

3.3. Existence of an equilibrium

Let us define aggregate consumer demand and aggregate endowments as:

$$x(p) \equiv \sum_{i=1}^{n} x_i(p)$$

$$\omega = \sum_{i=n}^{n} \omega_i$$

where n is the number of consumers in the economy. The aggregate excess demand function is defined as:

$$z(p) \equiv x(p) - y(p) - \omega$$

Note that $z(p)$ is positive if demand exceeds production plus endowments, and vice versa.

It is easily verified that Walras Law holds in the sense that $pz(p) = 0$ for all p. Again, this is an accounting identity. It arises because all profits are distributed to consumers and all consumers spend all their income and no more.

The existence of a Walrasian equilibrium can be demonstrated using the same fixed-point argument. A series of conditions are nevertheless required to ensure that consumption demand and net output supply functions exist and are continuous (continuity is indeed a necessary condition to be able to apply a fixed-point theorem). Varian lists all the technical conditions and provides some intuition on pages 344 and 345. For our purpose it is sufficient to remember that if consumption and production are well-behaved, an equilibrium exists. There may be multiple equilibria, though.

3.4. Welfare theorems

The two welfare theorems hold with production as well, provided a series of stringent conditions are fulfilled. In practice, these conditions are never completely satisfied. There are two schools of thought regarding how to conduct welfare analysis – and economic policy – when the conditions for the welfare theorems are violated.

The first approach is 'second best' reasoning (first best is when all the conditions are satisfied). Second best reasoning imposes constraints on what can reasonably achieved and seeks to maximize welfare conditional on these additional constraints. Depending on what constraints are imposed, second best reasoning can be used to justify pretty much any policy intervention, including the most distortionary forms of intervention. The Latin American structuralist school is a good is a good illustration of this approach.

Another approach is to seek to redress the constraints themselves. It is a little bit like second welfare theorem reasoning where, if you do not like the distributional aspects of a competitive equilibrium, you don't mess with the market, but you simply redistribute endowments directly. This kind of "let's get to first best" reasoning has been extremely influential in recent years. Here are some illustrations.

3.4.1. Increasing returns and public utilities

For the two welfare theorems to hold the economy must be 'convex', that is, cannot have increasing returns to scale in production or non-convex preferences.[3] If the economy is not convex, a competitive equilibrium may not be efficient. One good example is one in which all production is in the hands of a single producer with constant marginal cost but a large fixed cost. Setting the price equal to marginal cost generate negative profits for the producer, and hence bankruptcy.

Public utilities are the typical application of this problem. The "let's get to first best" solution to this problem is not to distort the price at which the good is sold. Rather, it is to make a transfer to the producer equal to the fixed cost, combined with a regulated price equal

[3] Non-convex preferences arise when people prefer either A or B, but not some combination of the two – e.g., they want coffee or tea but not the two mixed together. In practice, non-convex preferences are not a serious concern because, in the presence of many consumers, these local non-convexities 'convexify' by averaging many individual preferences through the market. Arrow worked on this, among many other things.

to marginal cost. Events such as Railtrack troubles and the debacle of electricity provision in California a few years back suggest that effectively regulating public utilities is not an easy task.

3.4.2. Externalities

The economy also must not have 'externalities', that is, positive or negative effects that do not manifest themselves through markets. If the economy is convex but externalities are present, a competitive equilibrium typically exist but it is not Pareto efficient. One "let's get to first best" solution to externalities is to modify 'property rights' so that, say, polluters face the full cost of the damage they impose onto others.

Another approach is to institute taxes and subsidies (named Pigovian taxes after the French economist Pigou) so that producers of beneficial side-effects are compensated for the positive externalities they create and beneficiaries are charged for the benefit they receive – thereby mimicking the behavior of the market. This idea has influenced the debate surrounding the Kyoto agreement – i.e., the issue of tradable pollution rights. Varian discusses these issues in his Chapter 24.

3.4.3. Information asymmetries

For the welfare theorems to hold, there must not be information asymmetries and transactions costs. If consumers and producers have insufficient information or if there are transactions costs, efficiency will typically not be achieved. Coase and Hayek were the first to have brought these issues to light.

Here again the preferred solution is to seek to address the information asymmetry directly, e.g., by creating institutions that circulate information (e.g., employment office, stock market, credit reference agency, grading agency, ISO certification, Moody rating) and by punishing those who defraud others (e.g., penalties in opportunistic breach of contract).

We will revisit information asymmetries later in the course.

4. Empirical applications

There have been numerous applications of general equilibrium theory to development economics issues. On the methodology side, a whole tradition of work has emerged around the use of Computable General Equilibrium models as policy analysis tools. These models are generalized versions of the above general equilibrium model in that they include trade with the rest of the world, and a government. With trade, the general equilibrium model resembles closely the household model with missing markets that we discussed in Lecture 2. Open economy general equilibrium models are generalizations of the standard trade models.

CGE practitioners essentially construct general equilibrium sectorial models of real economies and simulate them to see what happens. It is great comfort to them to know that, if they make the right assumptions, an equilibrium exists. CGE models are used to investigate various trade and taxation issues. It is also possible to construct CGE's smaller entities such as regions, villages, or households – in which case the model resembles closely the one we developed in lecture 2.

As we have illustrated, the two welfare theorems have had a massive effect on economic thinking. What has been called the 'Washington consensus' (a convergence of views between the Reagan administration, the IMF, and the World Bank that emerged in the 1980's) reverts largely around the kind of arguments we discussed above. While in the 1960's the second welfare theorem was used to justify all kinds of 'second best' distortions, starting in the mid-1980's it has been used primarily to justify decoupling the efficiency of policy interventions from their social repercussions. It has also been used as a way of organizing corrective policy interventions (e.g., poverty is dealt with via transfers, natural monopolies via regulation, externalities via tradable rights or proportional taxes, information asymmetries via better institutions).