

Chapter 5: Market Clearing Prices

I. Introduction: Equilibrium Prices

The previous three chapters developed the core neoclassical models of consumers and firms. In this chapter, implications from those models are used to develop the neoclassical theory of equilibrium prices in competitive markets. In the century leading up to World War II, a variety of mathematical models of consumer and firm behavior were worked out. A subset was compatible with one another and could be integrated into models of price determination in both output and input markets.

Although mathematical modeling was not the most commonplace mode of economic analysis during that period (closely reasoned verbal accounts were far more commonplace), towards the end of that period, economists who used mathematical models demonstrated many of the models could be combined to provide an internally consistent and rigorous characterization of equilibria in final goods, input markets, and even entire commercial networks. By the middle of the twentieth century, those more complete models came to be regarded as major improvements in our understanding of how markets operated and largely displaced verbal accounts. The same models also undercut, at least to some extent, the various critiques of neoclassical economics that had been developed in the first half of the century.

Although the core neoclassical models provided a variety of new insights, it bears keeping in mind that the mathematical models make a variety of implicit assumptions in addition to their explicit mathematical assumptions. For example, they implicitly focus on routine choices—that is to say, choices concerning well-understood goods and services that are produced by methods that are also well-understood. Such choices are, in a sense, perfectly informed because those making relevant choices have completely accurate expectations about the nature of the goods and services at issue, their normal prices, and their manner and cost of production. They are choices for which a consumer's transaction costs tend to

be minor because of past experience with the products and the merchants selling them. In such cases, transaction costs can be neglected without significant reductions in generality.

Not all choices are routine in this sense, but many of the implications of such models carry over to settings in which decisions are “almost” routine. For example, the models worked out in the previous chapters clearly demonstrate that fundamental relationships exist between consumer preferences, production methods (technology), input prices, and the prices of final goods and services. Conditions at the margin—marginal net benefits—provide a very general starting point for understanding both consumer and firm decisions in complex circumstances as well as simple ones.

During the second half of the twentieth century, generalizations of the partial equilibrium models were developed that proved that equilibrium prices can exist for entire market networks—which is not intuitively obvious. In practice, prices are changing all the time and the idea that a single price vector could clear all markets simultaneously was a major insight that suggested that the price changes observed were adjustments that tend to converge toward such equilibrium price vectors. A calculus-based proof of the existence of such general equilibria is provided in the appendix of this chapter. For more general characterizations, Debreu’s (1959) classic monograph, *Theory of Value*, provides one of the first very general demonstrations that such price vectors exist.

Unfortunately, extensions of those models also demonstrated that equilibrium prices do not necessarily emerge in all market settings. Arrow and Hurwicz (1960), for example, demonstrate that the Walrasian tâtonnement processes often said to generate equilibrium prices in competitive markets do not always converge to such prices. Moreover, because the basic models are grounded in price-taking behavior, the price theory of perfectly competitive markets includes no actors that can directly alter prices. Thus, this theory is less a complete theory of price determination than a theory of the properties of market-clearing prices.

Both general and partial equilibrium models demonstrate that prices exist that can coordinate the decisions of firms and consumers so that the quantities of goods and services

brought to market by firms exactly equal those demanded by consumers. That equilibrium can occur even though the thousands of decisions by the persons and organizations making and purchasing the goods brought to market are all independently made and advance a variety of ends. Neither firms nor consumers have to act in unison because of shared interests for such equilibria to exist.

That competitive markets tend to converge to such prices is evident in that we rarely observe markets in which there are large unrealized demands (excess demand) or great surpluses of the goods produced for sale (excess supply). This would not be the case if prices did not usually converge toward equilibrium prices. Thus, many of the prices that we observe have the properties that competitive theory implies.

In markets populated by price-making firms, the origins of price and equivalence between quantities demanded and supplied are less mysterious. They are simply the result of profit-maximizing choices by firms, as demonstrated in the models worked out toward the end of Chapter 3. Prices in markets populated by price-making firms have somewhat different properties than those that emerge in perfectly competitive markets. For example, prices still tend to equal marginal benefits for consumers, but they no longer equal marginal revenues or marginal production costs for firms.

Before providing an overview of the core results of neoclassical price theory, it should be acknowledged that the neoclassical characterization of prices predicts a narrower range of prices than actually observed in practice—the so-called **law of one price**. This is, of course, in large part because the models abstract from informational problems, transaction costs, and subtle differences among goods that might account for the variation in prices. The core competitive models do not attempt to explain every possible event in an economy, only the typical ones that tend to emerge in settings where price competition is relatively intense, and there are large numbers of well-informed consumers and firms.

As models, their construction necessarily abstracts from many details of economic life in order to illuminate what are believed to be the most common relationships and

determinations of choices by firms and consumers. These common features allow a good deal to be abstracted from without causing major mistakes or biases in predictions about the tendencies of consumer choices, firm decisions, and equilibrium prices. Nonetheless, as will be seen in parts II and III, bringing additional factors into the models can often improve our understanding of how markets operate.

It should also be acknowledged that a lack of realism—in the sense of departures from obvious features of markets—is sometimes generated by assumptions that make the mathematical models more tractable. For example, the assumptions of strict concavity and differentiability allowed us to use calculus to characterize unique optimal choices by firms and consumers. The assumption that all goods are available in any quantity that one might want is clearly incorrect. Yet, the discreteness of the products available does not greatly alter the decision-making that takes place. If one maximizes net benefits, then conditions at the margin are still ones where the last unit purchased or produced has marginal benefits approximately equal to marginal cost.

Nonetheless, the fact that mathematical models are not all-inclusive probably accounts for their limited acceptance during the period in which those models were first developed—and why such models remain controversial among a subset of behavioral and Austrian economists today.

Nonetheless, by rigorously illustrating that equilibrium prices exist, how they simultaneously affect the decisions of millions of consumers and producers scattered around the world, and how they tend to change when circumstances change, the core models help us to understand many fundamental properties of markets. They provide clear answers to questions such as why market networks exist, why markets are not chaotic, why markets tend to improve life for their participants, and why prices are themselves important phenomena.

Verbal and geometric illustrations had been undertaken in Smith's *Wealth of Nations* (1776) and in Marshall's (1920) microeconomics textbook well before the mathematization and generalizations of those models took place. What the mathematical models of the twentieth

century added was rigorous support for earlier intuitive and geometric claims about market equilibria. It turned out that many of these conclusions were correct—and that at least some of the verbal discussions of how markets operate were also more internally consistent and coherent than recognized by their critics.

The analytics of this chapter follows the pattern of the previous three chapters. It begins with geometric illustrations, proceeds to demonstrations using concrete functional forms, and concludes with more abstract characterizations of market prices. The last part of the chapter provides a model of prices that tend to emerge in intermediate cases between perfect competition and non-competitive forms of monopoly. The appendix provides a short overview of general equilibrium price theory.

II. The Geometry of Supply and Demand

The mathematical models developed in Chapters 2, 3, and 4 deepened our understanding of the geometric models. For example, they showed why variables other than the price of the good or service of interest affect consumer and firm choices. Consumer income, the prices of other goods, and what are usually referred to as “tastes” (e.g., the shapes and arguments of individual utility and benefit functions) affect the location and shape of individual demand curves and, thus, market demand curves. Similarly, the prices of inputs and readily available production technologies affect the location and shape of a firm’s supply curve through effects on decisions about how, what, and how much to produce for market. The same factors are relevant for both short and long-run decisions. When demand and supply curves are drawn based on intuition, it is not clear what ultimately determines their location or shape.

When demand curves are based on strictly concave utility functions (with positive cross partials), they slope downward. This reflects the finite resources of consumers (their income or wealth) and both the prices and nature of products for sale. Similarly, when supply curves are grounded in production functions that exhibit diminishing marginal returns, they slope upwards and reflect the marginal cost of producing the good or service of interest. The shape of a firm’s supply function is a joint product of its input costs and production

technology. Profit maximization assures that production costs are minimized, given input prices and the available production methods.


The refinements generated by the mathematical models include (i) explicitly characterizing the demand and supply curves as functions with other variables being held constant (e.g., having particular values, as denoted with the “0” superscripts); (ii) making their trace in the $P \times Q$ domain curves rather than straight lines, which reflects the more general shapes implied by general mathematical models of consumer and firm choices; and (iii) having the market supply curve slope upwards at a somewhat increasing rate, reflecting diminishing marginal returns (at least in Ricardian markets). Although supply and demand diagrams look simple, a good deal of analysis lies behind them.

Equilibrium prices, in turn, bring market demand into equivalence with market supply.

Figure 5.1 illustrates a market equilibrium for a final good that takes account of the insights generated by the mathematical models of the previous three chapters.

Figure 5.1 also illustrates one of the more powerful implications of the supply and demand model of price determination—its ability to explain the effects of changes in the variables being held constant in a demand and supply diagram. It illustrates the case where average consumer income in the market of interest falls ($Y^1 < Y^0$). The usual result is that the market demand curve shifts back to the left (rather than “down”), because the quantities purchased by the typical consumer will fall at every price if the good in question is a normal good (and most goods are). The diagram demonstrates that the predicted effect of a reduction in average consumer income is that both the prices and unit sales tend to fall, other things being equal. (The “other things” include all the other variables that affect the shape and location of the market demand and supply curves—many of which are left out of the core models, as developed in parts II and III of this book.)

Figure 5.1 Equilibrium Prices In a Competitive Market for a Final Good or Service
(Effect of reduced average consumer income)




To show why market prices affect the decisions of individual price-taking firms and consumers, one simply takes the equilibrium price from the market diagram and incorporates it into the individual decision diagrams of Chapters 2 and 3 that underpin the market demand and supply curves. Those diagrams demonstrated why individuals and firms change their purchasing and production decisions when prices change.

The models of Chapters 2 and 3 demonstrate why changes in the choice-influencing factors being held constant to draw a particular pair of supply and demand curves alter a market equilibrium. Such “shocks” do so by altering the circumstances of consumers and/or firms, in a manner that alters their net-benefit maximizing decisions. Consequently, if one or more of those other factors change, individual consumer and firm behavior changes, and if enough decisions are changed, then the overall market demand and supply curves change, and prices and unit sales change systematically.

The (P_1, Q_1) outcome of Figure 5.1 reflects the original purchasing and production decisions of consumers and firms that are active participants in the market of interest. Their typical choices are illustrated in Figure 5.2, which shows the typical choice of a consumer and firm facing price P_1 . Each consumer and each firm selects the quantity that best advances their self-interest (as it is understood by the consumer and firm or firm owner), given the prevailing market price (P_1) and their own individual circumstances as characterized by wealth, prices, technology, and input costs. Consumers purchase the quantities of goods that maximize their net benefits. Firms maximize their profits (here long run profits) given the prevailing price—and other aspects of their choice setting including technology and the prices of inputs.

Figure 5.2 Individual Consumer and Firm Adjustments To changing Market Conditions



If the typical consumer's income declines, the typical individual's demand curve shifts back to the left. Being the sum of individual demand curves, the market demand does so as well. Because of price and supply responses to diminished demand—prices fall, although not quite as much as would have been induced by the reduction in income alone. Suppliers face lower prices for their outputs and therefore produce fewer units of output for sale and profit less from the amounts produced and sold.

The pattern of adjustments by firms and consumers is exactly those necessary (in the simplified choice settings modeled) to set market demand equal to market supply in the new

circumstances with lower average consumer income. This is Adam Smith's (1776) invisible hand and Hayek's (1945) coordinating price signal operating in a partial equilibrium setting. The assumption that the market supply curve does not shift when consumer income falls implies that the input prices for the relevant firms are not affected by the economic shock that caused the decline in average consumer income in this market. This would be the case, for example, if the products sold were produced in other regions of the world that were not affected by the income-reducing shock in the market of interest (e.g., for the market being analyzed in Figure 5.1).

All this should be familiar to most readers of this textbook. Most readers will also be familiar with the geometry of equilibrium prices in markets and understand that the slopes of the demand and supply curves may differ from those in Figures 5.1 and 5.2 without changing the basic results. Similarly, they will be aware that a wide variety of comparative static results can be modeled. Such changes may affect one or both sides of the market simultaneously, as when a local employment shock reduces demand for local labor, reducing wages in that locality, which causes both the demand and supply of new local real estate to fall, reducing average consumer wealth, and so on in the community affected by the shock.

In general, the qualitative effects of changes that mainly affect one side of the market (supply or demand) tend to be clear. However, that is not always true of changes that affect both sides of the market. Nonetheless, the prices that clear markets can be characterized geometrically as long as the directions that demand or supply curves move are clear—as they tend to be in the models developed in Chapters two through four.

A wide variety of market types and relationships can be illustrated with the geometry of demand and supply—as long as firms and consumers have no or very little control over prices in the markets of interest. When firms have significant control over prices, then the models developed at the end of Chapter 3 provide better characterizations of price determination, profits, and consumer surplus than competitive models.

III. Algebraically Solving for Market Equilibrium Prices

Given specific functional forms for a market's demand and supply curves, one can often characterize the equilibrium in ways that provide equations that can be estimated using conventional econometric methods. These include reduced-form types of results that avoid simultaneity bias and structural equations of the demand and supply variety. In many cases, this is a very straightforward exercise in algebra.

For example, one may assume that demand and supply curves are both linear. In such cases, the supply and demand curves of Figure 5.1 take specific functional forms, such as

$$D = -aP + bY - cP^0 \quad (5.1)$$

And

$$S = eP - fw - gr + hT \quad (5.2)$$

where a, b, c, e, f, g and h are the parameter estimates of one's statistical models of demand and supply. The variables of interest for estimating demand functions include the prices of final goods (P), income (Y), and the price of other goods that affect consumer choices, which is exogenously determined (P^0) in other markets. Production technology (T) is similarly exogenously determined by scientific and engineering progress at the time that firms make their production decisions. The variables of interest for estimating supply functions include prices (P), wage rates (w), and interest rates (r)—the latter being determined by other markets and for most purposes can also be regarded as exogenous variables in cases in which this market employs relatively small fractions of overall labor and capital. In the case where unbiased estimates of these curves have been undertaken, the coefficients would be particular numbers instead of letters.

The algebra necessary to characterize equilibrium prices is essentially the same in either case. To begin with, the market clearing price, P^* , has the property (by definition) of setting demand equal to supply. This implies that

$$-aP^* + bY - cP^0 = eP^* - fw - gr + hT \quad (5.3)$$

To solve this equation for P^* , shift the P^* terms to the left-hand side of the equation and the other terms to the right-hand side.

$$-aP^* - eP^* = -bY + cP^0 - fw - gr + hT$$

Factor the terms on the left and divide both sides by $-(a + e)$ to characterize the equilibrium price:

$$P^* = (-bY + cP^0 - fw - gr + hT)/-(a + e) \quad (5.4a)$$

$$\text{or equivalently as: } P^* = (bY - cP^0 + fw + gr - hT)/(a + e) \quad (5.4b)$$

Notice that this equation can be interpreted as a reduced-form equation for market price. It is linear in average consumer income, input prices, and the prices of other relevant final goods. Comparative statics can be analyzed by differentiating this equation with respect to the variables of interest. For example, the effect of an increase in average consumer income is:

$$\frac{dP^*}{dY} = \left(-\frac{b}{-(a+e)}\right) = b/(a + e) > 0 \quad (5.5)$$

This is the rate at which prices increase as average consumer income increases.

A disadvantage of the characterization of market price is that it implicitly assumes that there is an “average” consumer or that the distribution of consumer demands simply shifts as average income changes. Preferences, income, and the effects of prices of other goods may vary widely among consumers, and it is possible that a change in average income may affect the slope of the demand curve or its shape, as well as induce a shift in the demand curve. This is one explanation for the standard errors of the estimated coefficients of a typical demand function.

Fortunately, the derivation of equilibrium prices is not limited to cases in which consumers and firms are essentially homogeneous. To see this, suppose that every consumer has a linear demand curve similar to that used in the previous example but that each person i 's income

differs, and each may have demand functions with different intercepts and coefficients. In this case, the i 'th individual's demand function can be written as:

$$D_i = a_i P + b_i Y_i + c_i P^0 \quad \text{with } a_i < 0 \quad (5.6)$$

The market demand curve is the sum of the N individual demand curves of each individual in this market. No two individuals need be the same.

$$D = \sum_{i=1}^N (a_i P + b_i Y_i + c_i P^0) \quad (5.7)$$

Similarly, suppose that each supplier j in this market has a unique supply function because of variations in access to inputs, location, managerial talent, or technology.

$$S_j = e_j P - f_j w - g_j r + h_j T \quad (5.8)$$

Market supply is the sum of all these individual supply functions.

$$S = \sum_{j=1}^M (e_j P - f_j w - g_j r + h_j T) \quad (5.9)$$

The market clearing price, P^* , sets D equal to S in this market or:

$$\sum_{i=1}^N (a_i P^* + b_i Y_i + c_i P^0) = \sum_{j=1}^M (e_j P^* - f_j w - g_j r + h_j T) \quad (5.10)$$

To characterize this price, we have to solve for P^* , which requires isolating the terms with P^* in them from the others.

$$\sum_{i=1}^N (a_i P^*) + \sum_{i=1}^N (b_i Y_i + c_i P^0) = \sum_{j=1}^M (e_j P^*) + \sum_{j=1}^M (-f_j w - g_j r + h_j T)$$

Shifting the P^* terms to the left-hand side and the others to the right-hand side yields

$$P^* [\sum_{i=1}^N (a_i) - \sum_{j=1}^M (e_j)] = \sum_{j=1}^M (-f_j w - g_j r + h_j T) - \sum_{i=1}^N (b_i Y_i + c_i P^0)$$

Dividing to solve for P^* yields:

$$P^* = \{\sum_{j=1}^M (-f_j w - g_j r + h_j T) - \sum_{i=1}^N (b_i Y_i + c_i P^0)\} / [\sum_{i=1}^N (a_i) - \sum_{j=1}^M (e_j)] \quad (5.11)$$

Equation 5.11 characterizes the market clearing price for the market in which consumers and firms differ—indeed, they may live in different countries or continents. N may be a number in the millions, and M may be in the thousands.

Note that this single price, P^* , affects each of their choices, although price is not the only factor that matters to them. As long as preferences and other prices are reasonably stable, changes in the prevailing price of the final good have predictable effects on each person's and each firm's behavior.

Note also that changes in average income may affect different individuals in different ways. The effect of a change in income on market demand is not generated by average income but by the sum of the marginal effects of a broad income shock, Δ_i , on each individual, which equation 5.11 implies is

$$\frac{dP^*}{dY} = - \sum_{i=1}^N (b_i \Delta_i) / [\sum_{i=1}^N (a_i) - \sum_{j=1}^M (e_j)] \quad (5.12)$$

Note that an income shock may increase some consumer's income, have no effect on others, while decreasing it on the majority. Its overall effect on market prices and sales depends on the individual coefficients (b_i) and the extent and distribution of the changes in income (Δ_i).

Changes in equilibrium market outcomes occur not because buyers and sellers care about the other persons in the market (although they may) but because it is in their interest to adjust their purchases in response to the new prevailing price. Altering their purchases or production increases their utility or profits relative to not doing so. At the new market clearing price, all these individuals and organizations have adjusted their patterns of behavior to the new price in a manner that maximized their net benefits (whether measured in utils or some currency). Prices, in turn, adjust to make supply exactly equal to demand.

Of course, in the real world, the factors explicitly modeled are not the only ones that affect consumption and output decisions—for example, there may be government policies that inhibit or encourage purchases, sales, or production. There may be taxes and tariffs that affect sales prices and production costs. Weather may affect both the demand and supply of

some products, as rain increases the demand for umbrellas and droughts diminish supplies of many agricultural products.

Nevertheless, even with all these factors, there is a tendency for market prices to clear the market of interest—although the equilibrium prices may have somewhat different properties due to the effects of less than perfect information, internalized norms, or various government policies as developed in Parts II and III of this text. The price of gasoline might, for example, vary among countries or states because of differences in tax policies and environmental regulations, as well as locally relevant differences in production, transport, or selling costs.

Although the core models tend to focus narrowly on purely economic factors, most other factors can easily be incorporated into the analysis as demonstrated in the next two parts of the text.

IV. Abstract Characterizations of Market Clearing Prices and Comparative Statics

More general models of market equilibria and comparative statics can be developed using abstract families of demand and supply functions rather than linear ones. Given such functions, the implicit function theorem can be used to characterize equilibrium prices as functions of all the variables that influence the choices of firms and consumers.

To illustrate how this can be done, suppose that the market demand function is $D = d(P, Y, P^0)$, and the market supply function is $S = s(P, w, r, t)$, as in Figure 5.1. At the market clearing price, we know that $D = S$, thus we know that in equilibrium:

$$d(P^*, Y, P^0) - s(P^*, w, r, t) = 0 \equiv H \text{ at } P^*, \quad (5.13)$$

the market clearing price. Given a “zero equation” such as equation 5.13, and assuming that it is differentiable, we can use the implicit function theorem to characterize any of the variables in the zero equation as a function of the others. The implicit function theorem implies that we can write P^* as a function of the other variables in the demand and supply functions.

$$P^* = f(Y, P^0, w, r, t) \quad (5.14)$$

In econometric terms, equation 5.14 is the reduced form equation for the equilibrium price in the market being analyzed (that of Figure 5.1).

We can also use derivatives of the H function to analyze the comparative statics of equilibrium prices in that market by applying the implicit function differentiation rule. For example, the effect of an increase in average consumer income is:

$$\frac{dP^*}{dY} = \frac{\frac{dH}{dY}}{-\left(\frac{dH}{dP^*}\right)} = \left(\frac{dD}{dY}\right) / -\left[\left(\frac{dD}{dP^*}\right) - \left(\frac{dS}{dP^*}\right)\right] \quad (5.15)$$

We have assumed that the good in question is a normal good, so we know (by assumption) that the numerator is positive ($\frac{dD}{dY} > 0$). The denominator is also uncomplicated. The demand curve in Figure 5.1 is downward sloping, so we know that $\frac{dD}{dP^*} < 0$. The supply curve is upward sloping, so $\frac{dS}{dP^*} > 0$. Notice that this implies that the term inside the brackets is less than zero. The negative sign in front of that term implies that it has the opposite sign.

Thus, the sign of the numerator is positive, and the sign of the denominator is also positive, so the predicted effect of an increase in average consumer income on market prices for normal goods is also positive. Whenever demand and supply functions have the usual characteristics (e.g., the ones assumed in chapters 2, 3, and 4), an increase in average consumer income will cause higher prices in markets for final goods.¹

¹ Note that whether the usual assumptions are realistic for a given market or not can be subjected to statistical tests by estimating a concrete functional form of equation 5.14.

If the signs of the coefficients are not as expected, and the market is otherwise well behaved, one can conclude that one or more of the usual assumptions about the shapes of the utility function or production functions do not hold in the market of interest. For example, the models developed have all assumed diminishing marginal returns in production (at the margin) which implies upward sloping market supply curves in both the short and long run. Marshall's characterization of long run supply as generated by the entry of and exit

Similar steps can be undertaken to determine the predicted results for changes in the prices of substitutes (P^O) or in the input prices (w and r) or improvements in technology (t).

It turns out that the comparative statics of diagrams grounded in the mathematics of the last three chapters are quite general—but this is only known because of the mathematics undertaken in the chapters of Part I.

V. The Spectrum of Market Types and their Equilibrium Prices

The perfectly competitive model characterizes circumstances in which both firms and consumers tend to be price takers. This is not the only type of market, nor the only possible source of equilibrium prices. The rest of this chapter provides an overview of other core models of markets in which competition is less than perfect—which is to say, markets in which firms have some degree of price-setting ability. The price-making model of supply developed at the end of Chapter 3 is one such model.²

The price-making model can be used to characterize a spectrum of market settings in several ways. For example, different degrees of competition (substitutability) can be characterized by assuming different slopes of the market demand curves faced by individual price-making firms. The flatter (less downward sloping) a market demand curve is, the easier it is for consumers to substitute other goods for the product of interest. The flatter (more elastic) a firm's demand curve is, the smaller the difference between the selling price and the marginal cost tends to be at the margin. Conversely, the more steeply downward sloping the demand curve is, the greater is the difference between the profit-maximizing price and quantity and

of identical efficient sized firms would imply a horizontal long run supply function, unless variations in supply affected input prices. The Ricardian assumptions tend to imply upward sloping long-run supply functions. It is also possible that long run supply curves are downward sloping because of economies of scale or technological externalities (cost reducing effects of broader markets) between the firms of interest and the production methods of their input suppliers.

² As is true of many other parts of this text, the aim is to provide an overview of a few foundational models rather than to survey the field of industrial organization. The field of industrial organization is a large one, and much of it is beyond the scope of this text.

the marginal cost of producing that quantity and the greater is its monopoly power.

Variations in the slopes (or elasticity) of the demand functions confronted by price-making firms, thus, trace out a spectrum of monopoly power.

For firms selling similar products, the slopes of their individual demand curves tend to be quite flat, because of the ease of substitution among similar goods and services. Such market types are sometimes referred to as monopolistically competitive markets.

Another spectrum of market types can be created by focusing on the number of firms selling a particular homogeneous (identical) product. This way of thinking about the spectrum of market types focuses mainly on the number of suppliers (although the number of buyers may also be important in some markets). The greater the number of sellers, the more competitive a market is said to be. Contrariwise, the fewer sellers, the more monopolistic a market is said to be from this perspective.

Intuitively, one expects smaller average markups (differences between marginal cost and price) in markets with lots of sellers, because competition for customers will tend to be more intense and cooperation among suppliers less likely. Conversely, markups in markets where there are few sellers of the product of interest tend to be larger because competition among firms for customers tends to be less intense and cooperation somewhat more likely. Such effects can be illustrated with Cournot models developed below.

Duopolies

Product markets in which there are just two suppliers are called duopolies. There are three basic models of pricing and output for duopolies.

The **Bertrand Model** is the most competitive of the three. It occurs when each seller attempts to gain market share by undercutting the other's price. This tends to generate a sequence of declining prices that converge toward the lowest price at which production will take place. In equilibrium, both sellers set price equal to marginal cost, and thus results similar to the perfect competition model emerge even when there are just two firms.

The **Stackelberg Model** is a model in which there is a sequence of entry into a market. The first firm chooses a price that will maximize its profits, given what it expects the second firm to produce. After the second firm enters, the result is an

equilibrium price based on the output choices of both firms. Neither firm has a reason to alter its production decision.

The **Cournot Model** assumes two firms making profit-maximizing output choices independently of each other, with no anticipation of the other firm's choices. In equilibrium, both firms simultaneously are maximizing their profits given what the other firm has chosen. The result is an early precursor of what a century later would be called the Nash equilibrium of a noncooperative game. (The nature and usefulness of the Nash equilibrium concept are developed in Chapter 13. It can be argued that Cournot actually invented the concept of a Nash equilibrium approximately a century before Nash worked out his somewhat more general characterization, 1838 vs 1951.)³

For the purposes of this chapter, the Cournot model is of greatest interest because it can easily be extended to characterize a spectrum of market outcomes between monopolistic and perfectly competitive markets.

Cournot Duopolies

In the **Cournot** duopoly model, two firms produce identical goods and make their output decisions independently of one another. Each firm takes the other's output as given and selects its own best output given the other's output and the downward sloping market demand curve for the product of interest. Total output and market price are represented as equilibria to a “noncooperative” production game between the two firms.

Best-reply functions characterize each firm's ideal output levels for the various possible outputs of their rival. The equilibrium occurs when both (or all) firms are simultaneously on their best-reply functions—which can be geometrically represented as a point in the strategy space where the best-reply functions intersect one another.

As an illustrative example of a Cournot duopoly setting, suppose that the market in question has a demand curve:

³ This section uses some of the vocabulary of game theory, which many students will already be familiar with. For those who are not and for a broader discussion of game theory see chapter 13.

$$Q = 1000 - 10P. \quad (5.16)$$

The linear case makes the pricing of the Cournot model relatively easy to work out and also relatively easy to generalize to N firms, as we'll see in the next section of this chapter.

Assume that both firms are profit maximizers, sell identical products, and that profit is simply revenues from sales less the cost of producing the goods sold. For example, firm "A's" profit can be written as

$$\Pi^A = PQ^A - C^A. \quad (5.17)$$

In order to know (or estimate) their profits, they will have to know the market price and the output of the rival firm. The market price depends on the total amount brought to the market by both firms, not simply that brought by firm "A." Demand curves "slope downward," which implies that the more output is "brought to the market" the lower prices tend to be.

Given the assumed demand curve, $Q = 1000 - 10P$, the effect of total market output on market price can be written as $P = 100 - 0.10Q$. (As mentioned before, this way of writing a demand curve is called an inverse demand curve because it goes from quantities into prices, rather than from prices into quantities.) If there are just two firms, A and B, then the total amount brought to market is $Q = Q^A + Q^B$. In that case, firm A's profits can be written as:

$$\Pi^A = [100 - 0.10 * (Q^A + Q^B)] Q^A - C^A \quad (5.18)$$

To simplify a bit more, let us also assume that the cost function is the same for each firm and exhibits constant returns to scale. An example of such a cost function is:

$$C = 5Q$$

which implies that the profit of firm A is simply

$$\Pi^A = [100 - 0.10 * (Q^A + Q^B)] Q^A - 5Q^A \quad (5.19)$$

Notice that firm A's profit (the payoff in this game) is affected by the other player's output decision, as is typical in a game setting.

To find firm A's profit-maximizing output, we need to find the “top” of the profit function, which can be found where the slope of the profit function is zero. Differentiating with respect to Q^A and setting the result equal to zero yields:

$$100 - 0.20 Q^A - 0.10Q^B - 5 = 0 \quad (5.20)$$

Note that the first three terms are Firm A's marginal revenue function and the last is its marginal cost function. Isolating the Q^A term allows us to characterize the output level that maximizes firm A's profit for any output level that B might choose:

$$0.20Q^A = 95 - 0.10Q^B$$

$$\text{or} \quad Q^{A*} = 475 - 0.50Q^B \quad (5.21)$$

Equation 5.21 represents the best-reply function for firm A. It indicates how much firm A should produce to maximize its profits, given the output level of firm B.

In this choice setting, there is no dominant pure strategy. The optimal response varies with the particular quantity brought to market by firm B. Firm B's best reply function can be found in a similar way.

$$Q^{B*} = 475 - 0.50Q^A \quad (5.22)$$

At the Nash equilibrium, both firms are simultaneously on their best reply function—which is to say both firms are maximizing their profits, given the other firm's output. This implies that:

$$Q^{A*} = 475 - 0.50Q^{B*} \quad \text{and}$$

$$Q^{B*} = 475 - 0.50Q^{A*}$$

are simultaneously satisfied. Substituting for Q^B into Q^A allows us to find the Nash equilibrium levels of output for firm A.

$$Q^A = 475 - 0.50 (475 - 0.50Q^A) \rightarrow 0.75Q^A = 237.5$$

Gathering the Q^A terms and a bit of arithmetic yields $0.75 Q^A = 475/2$. Thus

$$Q^{A**} = (4/3)(237.5) = (950/3) = 316.667 \quad (5.23)$$

By symmetry, we also know that $Q^{B**} = 316.667$, which implies that the total output at the Cournot-Nash duopoly equilibrium is $Q^{**} = 633.333$. It is the sum of the outputs of the two firms where their best reply functions intersect.⁴

Substituting back into the profit function allows us to determine each firm's profits (payoff) at the equilibrium output and price. The profit of firm A is simply

$$\Pi^A = [100 - 0.10(Q^{A**} + Q^{B**})] Q^{A**} - 5Q^{A**}. \quad (5.24)$$

Or, substituting for the “Qs” we have $\Pi^A = [100 - 0.10 * (633.333)] 316.667 - 5(316.667)$. A bit of arithmetic yields:

$$\Pi^{A**} = (36.667)(316.667) - 5(316.667) = 10027.778$$

$$\Pi^{A**} + \Pi^{B**} = 2(10027.778) = 20055.556$$

In equilibrium, both firms are simultaneously maximizing profits, and because their choice settings are identical (e.g. it's a symmetric game), both will earn the same profits from their sales of identical quantities. Geometrically, the equilibrium occurs at the intersection of the two firms' best reply curves.

Note that total Cournot duopoly output is larger than that associated with an otherwise equivalent monopolist. Thus, market price is lower and profits are lower. Recall that a monopolist will also maximize profit. A monopolist's profit function in this case is:

$$\Pi = [100 - 0.10 * Q] Q - 5Q \quad (5.25)$$

Differentiating the profit function with respect to Q, setting the result equal to zero yields:

$$100 - .20Q - 5 = 0$$

⁴ Symmetry could also have been used to simplify the math a bit. We could have replaced both quantities with Q^{**} and solved the resulting equation for Q^{**} . The longer derivation is done above to illustrate why the symmetric equilibrium tends to emerge if the participants in a game or contest make choices in an identical context—as is true of most parlor games.

and solving for Q yields

$$.2Q = 100 - 5 \rightarrow Q^* = 475.$$

Substituting $Q=475$ into the profit function yields

$$\Pi^* = [100 - 0.10 * (475)](475) - 5(475) = 22562.5$$

Which is about 12.5% greater than the total profits realized by the equivalent Cournot duopoly, 20055.556.

The diminution of total profits at the Cournot equilibrium provides an incentive for the two firms to coordinate their production decisions rather than setting them independently—e.g., to form a cartel—but the implicit assumption here is that cartel agreements are difficult to consummate—partly because there are incentives for each firm to cheat on the agreement and because cartel contracts are not enforced by courts in many countries.

Cournot-Based Models of Industries with More than Two Firms

The Cournot model can be generalized in several ways. For example, the inverse demand curve and cost function can be made more abstract and general. One can also extend the Cournot approach to characterize markets in which more than two firms interact in the manner postulated for duopolists. This extension can be used to demonstrate that a continuum between monopoly and competitive markets may exist in a rather neat way. To do so, we'll use a slightly more abstract model to characterize the N -firm form of Cournot competition.

The mathematics of Cournot equilibria is easiest for the linear demand and constant cost cases. So, we'll assume that the inverse demand curve is linear: $P = a - bQ$ and that each firm's cost function is $C_i = cQ_i$.

We'll also assume that there are N identical firms participating in the market. Each firm makes its own decisions independently of one another. (That is to say, there is no cartel-like coordination of output strategies or efforts to read the minds of other firms).

We'll focus on firm 1 and regard the total output of the other N-1 firms to be Z to reduce the amount of notation we need to deal with. Firm 1's profit in this case is:

$$\Pi_1 = a - b(Q_1 + Z)]Q_1 - cQ_1 \quad (5.26)$$

Differentiating with respect to Q_1 yields:

$$a - 2bQ_1 - bZ - c = 0$$

This can be solved for Q_1 as a function of parameters of the demand function, the total output of the other firms, and firm 1's cost function. Shifting the Q_1 term to the right and a bit of algebraic arithmetic yields firm 1's best reply function:

$$Q_1^* = (a - bZ - c)/2b \quad (5.27)$$

Similar functions can be derived for all the other firms in the market of interest.

The easiest way to find the Nash equilibrium is to assume that there is a symmetric equilibrium. In that case, every firm's output is the same in equilibrium, as with $Q_1^{**} = Q_2^{**} = \dots Q_N^{**}$. Total output of the rival firms (Z) is thus simply $(N - 1)Q_1^*$. Substituting that into equation 5.27 for Z yields:

$$Q_1^{**} = (a - b(N - 1)Q_1^{**} - c)/2b$$

Adding $(N - 1)Q_1^{**}/2$ to each side and multiplying by 2 yields:

$$Q_1^{**} + (N - 1)Q_1^{**}/2 = (a - c)/2b \rightarrow (N + 1)Q_1^{**} = (a - c)/b$$

Solving for Q_1^{**} yields the symmetric Nash equilibrium output for each of the firms:

$$Q_1^{**} = (a - c)/b(N + 1) = [(a - c)/b] [1/(N + 1)] \quad (5.28)$$

Total market output for N firms is thus:

$$N Q_1^{**} = [(a - c)/b] [N/(N + 1)]$$

Substituting this into the inverse demand curve, gives us the market price:

$$P^{**} = a - bQ = a - b[(a - c)/b] [N/(N + 1)]$$

$$P^{**} = a - [(a - c)] \left[\frac{N}{N+1} \right] \quad (5.29)$$

Equation 5.29 characterizes equilibrium prices for markets with different numbers of firms. Note that $N/(N+1)$ converges to 1 as N becomes large, which implies that in the limit price converges on c , the marginal cost of production. Thus, Cournot competition tends to converge toward the pricing of perfectly competitive markets as the number of firms increases.

(To verify that we’ve made no algebraic mistakes, we can substitute the numbers used in the duopoly case in the previous section and see that the result is the same as found in that section.)

It can also be shown that firm profits converge to zero as N gets large. Perfect competition, thus, is a limiting case that can emerge from entry in a Cournot-Nash model of competition among homogeneous firms, as argued by Marshall and many others.

VI. Conclusions: The Power and Weaknesses of Neoclassical Price Theory

Market prices are a social phenomenon when there are large numbers of consumers and firms. They are generated by the independent decisions of millions or even billions of individuals making choices in various settings that affect the demand and supply of products sold in extensive commercial networks. What neoclassical economics shows is that the results of a huge number of independent decisions in market networks are not chaotic “in the large,” although they may appear so in “the small.”

If people are generally forward-looking and attempt to achieve the “best” that they can with the resources at their disposal—using internally consistent ideas about “best”, which may be different for each individual—the result is not chaos because the price system tends to coordinate decisions so that supply equals demand in every market—or at least it produces a tendency for prices to move towards such an equilibrium.

In prosperous societies, these predictions are largely validated by experience. Supermarkets and other retail outlets rarely run out of things that their customers want to purchase. Nor

do they routinely have huge piles of unwanted goods sitting in their warehouses. Similar balances between supply and demand are also commonplace in most input markets.

This is not to claim that markets are always in equilibrium or that shortages and surpluses never occur, but the models imply that such “disequilibrium” outcomes are temporary and rare, unless they are caused by other nonmarket factors. For example, government regulations may prevent prices from moving to market clearing levels. Or, knowledge about production costs and consumer demand may be less complete than assumed in the models, and thus mistakes may be made.

It is also possible that firms have strategies that do not simply maximize profits on a day-to-day basis but instead maximize their longer-term profits. The latter may, for example, induce firms to maintain reputations for “fair pricing” because their customers are more likely to return if they feel fairly treated. Also, transaction costs of various kinds often slow the adjustment of market-driven price adjustments sufficiently that unexpected changes in demand or supply may lead to temporary shortages—as is often the case after catastrophic weather events such as hurricanes and floods.

The models developed in Part I have used a variety of geometric and calculus-based models to illustrate the type of equilibrium prices that may emerge in markets. Those models imply that equilibrium prices share many properties. For example, in markets where consumers and firms are price takers, prices are marginal costs for consumers and marginal revenues for firms, and thus $MC(Q^*) = MB(Q^*) = P^*$ for consumers and $MC(Q^*) = MR(Q^*) = P^*$ for firms. At the margin, prices determine the behavior of both firms and consumers. Similar relationships hold for the sides of markets that are composed of price takers in markets where the other side of the market has price making abilities. In the latter case, $MC(Q^*) = MR(Q^*) < P^*$ if firms are price makers, and price is determined by the firm’s demand curve (at Q^*).

Nonetheless, as very lean models, what may be called the “neoclassical core” abstracts from many aspects of life in a commercial society. For example, there are no information

problems, and thus there are no opportunities for fraud, intra-firm shirking, or production mistakes. Also, the regulatory and tax environments have been neglected in the analysis to this point—except for the implicit assumption that civil law, with its contracts and property rights, exists and is well and fairly implemented. The latter assures that contracts are kept, fraud and theft are minimized, and no extortion or theft takes place.

The core models also fail to account for important features of markets such as why products change through time. There is, for example, no model of entrepreneurship, innovation, or research and development. The second and third parts of this book investigate how many of these neglected factors can be incorporated into the core models and how the extended models can be used to provide somewhat more complete explanations for the emergence of markets and how they develop through time.

The appendix of this chapter provides a short overview of general equilibrium theory, which demonstrates that price vectors that clear all markets simultaneously exist—at least within the environments modeled by GE theorists. It does not provide a complete overview, but simply provides two examples of how such models can be developed.

References and Suggested Readings

- Arrow, K. J., & Hurwicz, L. (1958). On the stability of the competitive equilibrium, I. *Econometrica*, 26(4), 522–552.
- Arrow, K. J., & Hurwicz, L. (1960). Some remarks on the equilibria of economic systems. *Econometrica*, 28(3), 640–646.
- Azar, J., Marinescu, I., & Steinbaum, M. (2022). Labor market concentration. *Journal of Human Resources*, 57(S), S167–S199.
- Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4), 841–890.
- Buchanan, J. M., & Yoon, Y. J. (Eds.). (1994). *The return to increasing returns*. University of Michigan Press.
- Cournot, A. A. (1897). *Researches into the mathematical principles of the theory of wealth* (N. T. Bacon, Trans.). Macmillan. (Original work published 1838)
- Debreu, G. (1959). *Theory of value: An axiomatic analysis of economic equilibrium* (Vol. 17). Yale University Press.
- Hayek, F. (1945). The use of knowledge in society. *American Economic Review*, 35(4), 519–530.
- Marshall, A. (1920). *Principles of economics* (8th ed.). Macmillan.
- Mukhin, D. (2022). An equilibrium model of the international price system. *American Economic Review*, 112(2), 650–688.
- Samuelson, P. A. (1947). *Foundations of economic analysis*. Harvard University Press.
- Smith, A. (1776). *An inquiry into the nature and causes of the wealth of nations*. W. Strahan and T. Cadell.
- Stigler, G. J. (1951). The division of labor is limited by the extent of the market. *Journal of Political Economy*, 59(3), 185–193.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3), 213–225.
- Varian, H. R. (2016). How to build an economic model in your spare time. *The American Economist*, 61(1), 81–90.
- Varian, H. R. (1992). *Microeconomic analysis* (3rd ed.). Norton.
- Walras, L. (1954). *Elements of pure economics: Or the theory of social wealth* (W. Jaffé, Trans.). Allen and Unwin. (Original work published 1874)

Appendix to Chapter 5: On the Existence and Nature of General Equilibria

Market Clearing Price Vectors

It is one thing to say that a single market may reach an equilibrium; it is quite another to say (or prove) that any finite number of markets can clear simultaneously, given relatively general assumptions about individual tastes and production functions. One of the most impressive contributions of neoclassical economics in the twentieth century is general equilibrium analysis.

General equilibrium models attempt to determine “sufficient conditions” for the existence of a vector of prices that simultaneously clears all markets by setting demand equal to supply in every market. Léon Walras (1874), one of the founders of neoclassical economics, developed the initial intuition and some of the first mathematical models of such an equilibrium in the late nineteenth century. By now, there is a broad range of general equilibrium models which vary in the extent to which they depart from the assumptions of models of perfect competition, and with respect to the restrictiveness of the mathematical assumptions relied upon. Many of the classic works were developed in the 1950s and 1960s, as part of the great neoclassical synthesis taking place in that period. Gerard Debreu won a Nobel Prize in economics (in 1983) for his model and proof of the existence of a market clearing price vector in that model, which was published in a short book called *Theory of Value* in 1959.

It is surprising that fairly modest assumptions about human behavior, production technology, and budget constraints are sufficient to assure the existence of a market clearing price vector. In this part of the chapter, we’ll provide a basic overview of some of the core ideas and mathematics that demonstrate the existence of a market-wide vector of market prices. Debreu’s book is recommended to those interested in a more general treatment.


The Edgeworth Box: a Simple Geometric GE Model

The Edgeworth box provides a two-person two-good illustration of a general equilibrium. In an Edgeworth box, two consumers are assumed to have well-defined transitive preferences over both goods within the domain of the “box” and an initial endowment of each of the

two goods. Prices are “called out” by a "Walrasian auctioneer" until a price is found at which the “excess demand” for both goods is zero. (The excess demand for good i is the sum of desired consumptions of good i less the sum of the original endowment of good i .) At the market clearing price, the quantity that each person wants to sell is exactly the amount that the other wants to buy.

The geometry of the Edgeworth box is generated by taking two indifference curve diagrams and, in effect, turning around one of them so that it appears in the upper right-hand corner. This geometry sheds very useful light on conditions for the existence of mutual gains from trade, market clearing prices, and Pareto superior moves.

Figure 5.3 The Edgeworth Box



The area that characterizes mutual gains from trade is “football” shaped and generated by the indifference curves (Alice’s and Bob’s) that pass through the point that characterizes the division of the two goods between them, their initial endowments. Note that the endowment point simultaneously characterizes both Alice’s and Bob’s endowments. Any move into the football-shaped area is Pareto-improving. Both players benefit and no one loses.

At the market clearing price, each person wants to sell exactly what the other wants to buy at the price called out by the imaginary Walrasian auctioneer. In the above Edgeworth box, Bob sells good X_2 - X_{2b}^* units of X_2 at price P_2 in order to purchase $X_{1b}^* - X_{1b}$ units of good X_1 . Alice does the reverse and consequently the P_1/P_2 price vector is an equilibrium price vector. An equilibrium price in the 2-person 2-good case occurs when both persons' indifference curves are tangent to the same point on a price line.

In a more informal process, any point inside the mutual gains from trade region where the indifference curves are tangent to one another can be a trading equilibrium, because at such points there are no further gains from trade. The collection of such points within the box is called the “contract curve.” However, in a general equilibrium model, it is the prices rather than one-on-one bargaining that determine the final outcome.

General Equilibrium Model of a Large Barter Society

To demonstrate how the logic of the Edgeworth box can be generalized, we'll develop a model of a large “barter economy” similar to that of the Edgeworth box, but with many traders and products. The model and mathematics are based on a proof developed in Varian (1992). Varian's proof is calculus-based and so can be constructed with only minor extensions of the models developed in Part I.

To simplify a bit, we'll follow Varian and assume that no production takes place. Including production is not much more difficult, but such models require the introduction of more mathematical notation and assumptions, and do not provide much additional insight about how existence proofs are developed. Typical assumptions of a calculus-based general equilibrium model (without production) include:

- (i) **Goods:** Goods are characterized by time, location, and state of the world. There are assumed to be a finite number of goods, k . Agent i 's consumption bundle is denoted x_i and is a k -dimensioned vector of the goods possessed by i . The amount of the j th good possessed by individual i is denoted x_{ij} .

- (ii) Endowments: An individual's initial endowment of goods is his "pre-trade" consumption bundle, w_i , $w_i = [w_{i1}, w_{i2}, \dots, w_{iK}]$

which is a $k \times 1$ vector. An individual's demand for goods at a particular vector of prices is also a $k \times 1$ vector describing his or her ideal consumption bundle, given prices and his or her endowment $x_i = [x_{i1}, x_{i2}, \dots, x_{iK}]$

A feasible allocation for the economy is one that is possible. In the pure exchange case of interest here, it is one where $\sum_{i=1}^N x_i = \sum_{i=1}^N w_i$. A feasible allocation is one in which the total demand for each good equals the economy wide initial endowment of that good.

- (iii) Agents: Each consumer i is described by a complete transitive preference ordering $>_i$ (which is used to derive a utility function U_i) and an initial endowment w_i .

Each consumer is a utility maximizing price taker. Thus, each consumer maximizes

$U_i(x_i)$ subject to $Px_i = Pw_i$. (P is a $1 \times k$ vector, and w_i and x_i are also $k \times 1$ vectors.)

Note that this describes a budget set or multi-dimensional budget constraint for the case where endowments are in goods rather than "wealth." (Matrix multiplication implies that $Px_i = Pw_i \rightarrow \sum_{j=1}^K P_j x_{ij} = \sum_{j=1}^K P_j w_{ij}$.)

(In a model with production, there will also be k production functions which describe how "inputs" can be transformed into "final consumption goods" and individual endowments include inputs as well as final goods.)

Individual i 's excess demand for good j , z_{ij} , is simply his ordinary demand for good j (his desired consumption) less his initial endowment of that good, $z_{ij}(P) = x_{ij}(P, Pw_i) - w_{ij}$.

His or her vector of excess demands, z_i , is a vector of his or her excess demand for all $j = 1, 2, \dots, k$ goods.

- (iv) **Equilibrium Prices.** A Walrasian equilibrium in a barter economy exists if a price vector P^* exists such that $\sum_{i=1}^N x_{ij}(P^*, P^* w_{ij}) - \sum_{i=1}^N w_{ij} = 0$, which is to say when the sum of the demands for each good exactly equals the total endowment in each good. In that case, the sum of the excess demands for each good is zero. (Suppliers have negative excess demands and Demanders positive excess demands.)

An equilibrium price vector sets aggregate demand equal to aggregate supply, or equivalently, aggregate excess demand equal to zero for each person in every market. The demand correspondence x_i is a vector representing the utility maximizing levels of all goods for individual i with initial endowment w_i facing price P^* .

(A correspondence is a mapping from one set into another, whereas a function is a mapping from one set into a single dimensioned set—usually, the real number line.)

Notice that a lot of the cleverness of a GE model is writing down a complete model of an economy in very few equations with relatively few behavioral assumptions in a way that will be mathematically tractable. Note also that very similar assumptions were used in our partial equilibrium models of consumer choice, but with a money endowment rather than a goods endowment and 2 rather than K goods.

Some Properties of the Model

- (i) The budget set is homogeneous of degree 0 in prices.

If all prices are multiplied by any constant C , there is no change in any individual's budget constraint, when the endowment is in goods rather than money. This implies that the demand correspondence x_i is also homogeneous of degree 0 in “all” prices. E.g. there is no money illusion.

- (ii) The excess demand function ($z_{ij}(P) = x_{ij}(P, Pw_i) - w_{ij}$) is also homogeneous of degree zero in “all prices” for the same reason. Moreover, since

the sum of homogeneous functions of degree k is also homogeneous of degree k , the aggregate excess demand function is also homogeneous of degree 0 in all prices.

- (iii) Because endowments are in goods, rather than money, the vector of demands is affected by relative prices but not the price level, C .

Some Foundational Results for the Existence Proof

- (i) Each individual i 's vector of desired consumption is determined in the usual way—by maximizing individual i 's utility subject to his budget constraint. The k -dimensional vector of aggregate excess demand is $z(P)$, composed of elements $j = 1 \dots k$ of $z_j(P) = \sum_{i=1}^N (x_{ij}(P, Pw_i) - w_{ij})$
- (ii) **Walras' Law.** (Varian's version) For any P in S_k (remember there are k goods) excess aggregate demand (in dollars) is always zero, $P \cdot z(P) = 0$. (S_k is the price space associated with a k -dimensioned commodity space.)

Proof: recall that the k -dimensioned vector of aggregate demands is $z(P) = \sum_{i=1}^N (x_i(P, Pw_i) - w_i)$, and also that each person's demand correspondence (vector x_i) is derived by maximizing utility given a budget constraint. The budget constraint implies that $Px_i = Pw_i$ for each individual for every price vector. An individual's ideal bundle is always on the budget line. This implies that the sum of all the Pw_i vectors has to equal the sum of all Px_i vectors, hence there is never excess demand (in dollars) although there will normally be for the goods themselves—otherwise every price would be an equilibrium price.

- (iii) If $K-1$ markets have cleared, the excess demand in the remaining market, $P_j z_j$, must be zero. The excess demand in all the markets that have cleared is zero (by definition). Walras' law implies that overall excess demand is always zero in money terms. Thus, either the remaining market has cleared ($z_j = 0$) or there is excess supply and its price is zero. (This is the usual version of Walras' Law.)

- (iv) Similarly, **if all goods are desirable at the margin**, their equilibrium prices are greater than zero. In order for the money measure of excess demand, $P_j z_j$, to be zero in every market, excess demand in quantities of goods, z_j , has to be zero in every market.

Summary: The aggregate money value of excess demand is always zero. At the Walrasian equilibrium, if there is an excess supply of a good (an undesirable good) its price has to be zero. **In all other cases, demand must equal supply for all goods at a Walrasian equilibrium.**

Proof of the Existence of a Walrasian Equilibrium Price Vector

- (i) The proof begins with **Brouwer's Fixed Point Theorem**. If $f: S_k \rightarrow S_k$ is a continuous function from the unit simplex to itself, there exists some x in S_k such that $x = f(x)$. Such a point is called a fixed point. In a one-dimensional case, the unit simplex is just the 0-1 closed interval. (In the two-dimensional case it is a 1x1 square, in the three-dimensional case it is a 1x1x1 cube, etc.)

A function that maps a one-dimensional unit simplex onto itself can be illustrated by creating a 1x1 square and drawing a line in that square that is continuous. Any line that you draw will cross a 45 degree line. Points where the function crosses that line are its fixed points.

The ingenious trick in most existence proofs is to construct a function based on the choice setting modeled that is a continuous function of the variables of interest onto themselves. (In the case of interest here that mapping will be from the k-dimensional unit simplex on to itself.)

- (ii) **One example of such a mapping is the following:** First, define the elements of k-dimensional vector g as $g_j(P_j) = [P_j + \max(0, z_j(P))] / [1 + \sum_{j=1}^k \max(0, z_j(P))]$ where the prices have been normalized as: $P_j =$

$P_j / \sum_{j=1}^k P_j$ (This normalization of prices will not affect aggregate demand as we have already established above.)

This mapping is continuous because both P_j and $\max(0, z_j(P))$ are continuous.

(Note also that at the Walrasian equilibrium, $\max(0, z_j(P)) = 0$.)

- (iii) This mapping lies in the k -dimensional unit simplex since both function g and the price vector are k -dimensional. Both normalizations assure that the mappings of each price and element of g are within the 0-1 interval for their dimension.
- (iv) By Brouwer's fixed point theorem there exists a k -dimensional price vector, P^* , such that $P_j^* = g_j(P^*)$ for all j .
- (v) At this fixed point,

$$P_j^* = [P_j^* + \max(0, z_j(P^*))] / [1 + \sum_{j=1}^K \max(0, z_j(P^*))]$$

- (vi) **P^* is a Walrasian equilibrium price vector.**

To see this, multiply both sides by the denominator of the right-hand side, which yields:

$$P_j^* [1 + \sum_{j=1}^K \max(0, z_j(P^*))] = [P_j^* + \max(0, z_j(P^*))]$$

Then multiply both sides by $z_j(P^*)$, which yields

$$z_j(P^*) P_j^* [1 + \sum_{j=1}^K \max(0, z_j(P^*))] = z_j(P^*) [P_j^* + \max(0, z_j(P^*))]$$

Adding these up across all goods yields:

$$\begin{aligned} \sum_{j=1}^k z_j(P^*) P_j^* [1 + \sum_{j=1}^k \max(0, z_j(P^*))] \\ = \sum_{j=1}^k z_j(P^*) [P_j^* + \max(0, z_j(P^*))] \end{aligned}$$

From Walras' law we know that the left-hand side equals zero, because the term in brackets can be factored out and placed in front of the new summation term. (Note that the term in brackets does not have a “j” counter, because the j in that expression has already been added up.)

We also know that the first term on the right has to be zero, from Walras' law again, $P \cdot z(P) = 0$.

This implies that

$$0 = \sum_{j=1}^k z_j(P^*) [\max(0, z_j(P^*))]$$

This can only be true if $z_j(P^*)$ is zero for every good j that is desirable at the margin. If there are undesirable goods, there will be excess supply (excess demand less than zero) in which case the max-term is 0.

(Such cases are allowed in Varian's characterization of Walrasian equilibrium, and accounts for the max term. Without it the logic would be more straightforward.)

This implies that **the fixed point must characterize a Walrasian price vector**, because it causes excess demand for every good to be zero. (Otherwise, $z_j(P^*) [\max(0, z_j(P^*))]$ would exceed zero.) Q. E. D.

The economic significance of this and other similar existence proofs is that a market clearing price vector exists for any pattern of demand and initial endowments (wealth). That is to say, given the usual assumptions about preferences (and in a more general model, production correspondences) a price vector exists that simultaneously clears all markets.

At this price vector, (a) the excess demand for all goods (all things with $P > 0$) is zero, and (b) all potentially tradable "things" with negative excess demand have zero prices.

In a model that includes production, the same types of results would hold for both output prices and input prices.