

Lecture Notes: Chapter 6: Building Blocks of the Flexible-Price Model

J. Bradford DeLong
<http://econ161.berkeley.edu/>
delong@econ.berkeley.edu

In the past two chapters we have looked at long-run growth—at how the economy develops and evolves over periods as long as generations. In this chapter we shift our point of view and take instead a “snapshot” view of the economy, looking at it over such a short period that its productive resources will be fixed. The key questions we will seek to answer are:

1. What determines the equilibrium level of real GDP (Y)?
2. What economic forces keep real GDP (Y) at its equilibrium level?
3. What determines the composition of real GDP—that is, the division of production and spending between consumption goods (C), investment goods (I), government purchases (G), and net exports (NX)?

To answer the first two questions we will assume that wages and prices are sufficiently flexible that markets clear—that every buyer finds a willing seller and every seller finds a willing buyer. This flexible-price assumption means most importantly that supply equals demand in the labor market: no firms wishing to hire workers are left unsatisfied, and no workers who are willing to work are left unemployed. This type of analysis is a full-employment analysis.

In answering the third question, what determines the composition of spending, we will be assembling the building blocks for the analysis of the following chapter. In Chapter 7 we will put these building blocks together and show in detail how the answers to the three key questions are consistent, and what are the economic forces that ensure that a flexible-price economy reaches and stays at its equilibrium.

Potential Output and Real Wages

In the flexible-price model of the macroeconomy to be developed in this section, two sets of factors determine the levels of potential (and actual) output and of real wages; the production function and the balance of supply and demand in the labor market.

The Production Function

Chapter 4 introduced the *production function*, the rule that tells us how much the economy can produce given its available productive resources. In the Cobb-Douglas form of the production function, we learned, potential output (Y^*) is determined by (1) the size of the labor force (L), (2) the economy's capital stock (K), (3) the efficiency of labor (E), and (4) a parameter α that tells us how fast returns to investment diminish. The production function tells us that potential output is:

$$Y^* = (K)^\alpha (L E)^{1-\alpha}$$

The assumption in this chapter that wages and prices are flexible was commonly made by the so-called “classical” economists who wrote back before World War II. Thus this assumption is also called the *classical assumption*. The classical assumption guarantees that markets work—that prices adjust rapidly to eliminate gaps between the quantities demanded and the quantities supplied. Thus no businesses find their inventories of unsold goods piling up. Thus there is full employment: everyone who wants a job (at the market-clearing level of wages) can get a job, and every business that wants to hire a worker (at the market-clearing level of wages) can hire a worker. And because there is full employment, actual output is equal to potential output: there is no gap between the economy’s productive potential and the level of output the economy does produce.

The classical assumption made in this section means that this section is devoted to full-employment flexible-price macroeconomics.

The flexible-price assumption it is not always a good one. Experience has shown that a market economy does not always work well and does not always produce full employment. So while the flexible-price assumption is key in this section, starting in Section III we will drop it and make instead the “Keynesian” assumption that wages and prices are sticky (see Table 6.1).

If the classical flexible-price assumption is not always a good one to make, why make it? It *is* a good assumption if wages and prices are relatively flexible and have enough time to adjust in order to balance supply and demand. The classical assumption simplifies the analysis of several issues, making how the macroeconomy works easier to grasp. In general it is better to start with the simpler cases before looking at more complicated ones. Moreover, the way an economy if the flexible-price assumption held provides a useful baseline against which to assess economic performance. Nevertheless, we must remember that this section presents only one model of the economy: the classical model. The Keynesian sticky-price model behaves very differently in a number of ways.

The Labor Market

Economists try to suppress every detail and difference that does not matter to the overall result in order to simplify the analysis and focus it on the key, important factors that count. Because differences between businesses will not matter, let's think about an economy with K typical--identical--competitive firms, each of which owns one unit of the economy's capital stock. Each of these typical competitive firms hires L workers and each firm pays each worker the same wage W . Each firm sells Y units of its product at a per-unit price P .

The typical firm does not control either the wages it must pay or the prices it receives; those are determined by the market. The firm tries to make as much money as it can. The typical firm's profits are simply its revenues minus its costs, and its only costs are the wages it pays to workers. Therefore:

$$\text{Profits} = \text{Revenues} - \text{Costs}$$

$$\text{Profits} = P \times Y - W \times L$$

To figure out how many workers to hire, the firm follows two simple rules:

1. Hire workers to boost output.
2. Stop hiring when the extra revenue from the output produced by the last worker hired just equals his or her wage.

The value of the output produced by the last worker hired is the product price P times the marginal product of labor [MPL]. The cost of hiring the last worker is his or her wage W . The firm will keep hiring until:

$$P \times \text{MPL} - W = 0$$

The marginal product of labor is the difference between what the firm can produce with its current labor force L_{firm} , and what it would produce if it hired one more worker (see Figure 6.2):

$$\text{MPL} = F(1, L_{\text{firm}} + 1) - F(1, L_{\text{firm}})$$

The MPL for the Cobb-Douglas production function is:

$$\text{MPL} = (K_{\text{firm}})^{\alpha} \times E^{1-\alpha} (L_{\text{firm}} + 1)^{1-\alpha} - (K_{\text{firm}})^{\alpha} \times E^{1-\alpha} (L_{\text{firm}})^{1-\alpha}$$

Again, we take how much the firm would produce if it hired one more worker and subtract how it produces now, with its current labor force. Since the firm has only one unit of capital, we can rewrite this equation as:

$$\text{MPL} = (1)^{\alpha} \times E^{1-\alpha} (L_{\text{firm}} + 1)^{1-\alpha} - (1)^{\alpha} \times E^{1-\alpha} (L_{\text{firm}})^{1-\alpha}$$

$$\text{MPL} = E^{1-\alpha} \left[(L_{\text{firm}} + 1)^{1-\alpha} - (L_{\text{firm}})^{1-\alpha} \right]$$

The term inside the brackets looks like the rate-of-growth of a variable growing by one unit (the firm's labor force L_{firm}) and then raised to a power (the term $1-\alpha$). We have a standard rule-of-thumb (look back in Chapter 2 at Box 2.3: Useful Mathematical Tools) for dealing with such a situation. The proportional growth rate of a variable raised to a power is the proportional growth rate of the variable multiplied by the power to which the variable is raised. This tells us that the term inside the brackets is:

$$\left[(L_{\text{firm}} + 1)^{1-\alpha} - (L_{\text{firm}})^{1-\alpha} \right] = (1-\alpha) \times \frac{1}{(L_{\text{firm}})^{\alpha}}$$

So the MPL is:

$$\text{MPL} = \frac{(1-\alpha)E^{1-\alpha}}{(L_{\text{firm}})^{\alpha}}$$

There is nothing deep in this math. Indeed, the Cobb-Douglas function was carefully tweaked so that it would yield such simple forms for quantities like the MPL. That is why economists use it so often. If the Cobb-Douglas production function produced more complicated expressions, we would not use it.

The firm hires workers up to the point where the product price times the marginal product of labor equals to the wage:

$$P \times MPL - W = 0$$

Substituting for the MPL in this equation, we get

$$P \times \frac{(1 - \alpha)E^{1-\alpha}}{(L_{firm})^\alpha} = W$$

Next, we rearrange this equation to see that the typical firm's demand for workers is:

$$L_{firm} = \left(\frac{(1 - \alpha)E^{1-\alpha}}{(W/P)} \right)^{1/\alpha}$$

Because there are K firms in the whole economy, total economy-wide employment is equal to K times the typical firm's demand for labor:

$$L^d = K \left(\frac{(1 - \alpha)E^{1-\alpha}}{(W/P)} \right)^{1/\alpha}$$

What is the labor supply? The answer is simple: it is the number of workers who want to work. The labor market will be in equilibrium when firms' total demand for workers is equal to the labor force.

Can the labor market not be in equilibrium if wages and prices are flexible? Think about what would happen if supply were not equal to demand. Suppose there are more workers than firms wish to hire at current wages and prices. Then some of the unemployed will underbid their employed fellow workers: offer to take their jobs and work for less. Those workers who are employed will respond by offering to accept lower wages to keep their

jobs. The wage W will decline relative to the price level P , and the real wage W/P will fall. As the real wage falls, firms will hire more workers.

Suppose firms want to hire more workers than there are people in the labor force. Some firms will try to bid workers away from other firms by offering higher wages. The real wage W/P will rise. As the real wage rises, employers will reduce the quantity of labor they demand.

Thus in equilibrium, labor demand L^d will equal the labor force L (see Figure 6.4):

$$L = L^d = K \left(\frac{(1-\alpha)E^{1-\alpha}}{(W/P)} \right)^{1/\alpha}$$

Labor demand is equal to the labor force when the real wage W/P is:

$$\frac{W}{P} = ((1-\alpha)E^{1-\alpha}) \left(\frac{K}{L} \right)^\alpha = (1-\alpha) \left(\frac{Y}{L} \right)$$

and each of the K firms in the economy employs L/K workers. As long as wages and prices are flexible enough for this adjustment process to work, the economy will remain at full employment.

Note that a full employment economy is not necessarily the best or even a good economy. The real incomes of those who don't own chunks of the capital stock are their real wages: $W/P = (1-\alpha) \times (Y/L)$. If α is large, their real incomes will be small, and social welfare may be low.

When the labor market is in equilibrium the typical firm produces a level of output equal to:

$$Y_{firm} = (1)^\alpha (E)^{1-\alpha} (L/K)^{1-\alpha}$$

Because there are K firms, total output Y is simply K times the typical firm's output: $Y = K \times Y_{firm}$. This is the same *potential output*, Y^* in the Cobb-Douglas form of the production function analysis

$$\begin{aligned} Y &= K \times Y_{firm} = K \times (1)^\alpha (E)^{1-\alpha} (L/K)^{1-\alpha} = (K)^\alpha (E)^{1-\alpha} (L)^{1-\alpha} \\ &= (K)^\alpha (LE)^{1-\alpha} = Y^* \end{aligned}$$

Domestic Spending

Real GDP has four components:

1. Consumption spending (C)
2. Investment spending (I)
3. Government purchases (G)
4. Net exports, the balancing item (NX).

These four components add up to national income, which according to the circular-flow principles is the same as real GDP, Y:

$$C + I + G + NX = Y$$

Consumption Spending

Individual households make the spending and saving decisions that ultimately determine the flow of consumption spending. In this section we will first examine how households divide their income up among taxes, saving, and consumption spending. Then we will see how consumption spending varies in response to changes in income and in other aspects of the economic environment.

The wages of workers plus the profits--rent, interest, dividends, and retained earnings--of property owners add up to national income. Because at this level of analysis we are uninterested in any accounting distinctions between national income and real GDP, we use the letter "Y" to represent both. (Recall that the circular flow principle guarantees that whatever businesses produce and sell must show up as income for households.

Households pay some of their income to the government in net taxes--taxes less transfer payments from the government--which we will write as T. To keep the analysis simple, throughout this book we will assume that net taxes are equal to the constant average tax rate t multiplied by national income:

$$T = t \times Y$$

In the real world taxes are not proportional to income. Our tax system is somewhat progressive, which means that richer taxpayers on average pay more of their income in taxes than do the relatively poor. Once again, however, the complications induced by the fact that our tax system is not proportional to income are not central to the analysis. And so we follow economists' standard practice of simplifying wherever possible.

What is left after households pay their taxes is their *disposable income*, written Y^D :

$$Y^D = Y - T = (1-t)Y$$

Households also save some of their income to boost their wealth and future spending. We will represent these private household savings by S^H — S for “savings” and H for “household.” (Note that these household savings include the retained earnings of corporations: the NIPA treats corporate earnings not distributed but retained by the corporation as if they were distributed to the shareholding households and then immediately reinvested back into the corporation.) Households spend the rest of their income--everything that is not saved or paid to the government in taxes--buying consumption goods:

$$C = Y^D - S^H = Y - T - S^H$$

In the United States today, consumption spending C --purchases by households for their own use, from pine nuts and flour to washing machines and automobiles--adds up to roughly two-thirds of GDP.

We will break consumption spending down into a baseline level of consumption (defined to be the value of a parameter C_0) plus a fraction (a parameter C_y) of disposable income Y^D , or a fraction $C_y \times (1-t)$ of total income Y .

$$C = C_0 + C_y \times Y^D = C_0 + C_y \times (1-t)Y$$

Thus we assume that consumption spending C is a linear function of real GDP Y .

Notice that in writing this particular consumption function, we have once again followed economists' principle (or vice) of ruthless simplification. In this complicated world, consumption spending does not depend on disposable income alone. Other factors affecting it include changes in the real interest rate, in household total stock market and

real estate wealth, in the demographic structure of the population, in income distribution, in consumers' relative optimism, in expected future income growth, in tolerance for risk, and in whether consumers see changes in disposable income as transitory or permanent. (If consumers expect an income increase to be transitory, they will save most of it and spend only a little; if they expect an income increase to be permanent, they will spend most of it.) But here and throughout the book we will sweep these complications under the rug. We will think only about baseline consumption C_0 , the marginal propensity to consume C_y , and disposable income Y^D as the determinants of consumption spending (although we will occasionally sneak in other factors by saying that they change baseline consumption C_0).

The baseline level of consumption, the parameter C_0 , is the amount households would spend on consumption goods if they had no income at all. That is, it is the amount by which they would draw down their wealth in the absence of income, in order to keep body and soul together.

The marginal propensity to consume, also called the MPC, the same parameter C_y in the consumption function, is the amount by which consumption spending rises in response to a \$1 increase in disposable income (see Figure 6.9). We are sure that C_y is greater than zero: if incomes rise, households will use some of their extra income to boost their consumption spending. We are also sure that C_y is less than one: as incomes rise, households will increase their savings as well; they will not spend all their extra income on consumption goods.

Investment Spending

In the United States today investment spending averages roughly 17% of GDP. But investment spending is the most volatile and variable component of GDP. (Note also that economists' definition of "investment spending" is probably not what you think it is; Boxes 6.2 and 6.3 explain how economists define and calculate investment spending.)

Fluctuations in economy-wide investment spending have two sources. First is the interest rate: the higher the real interest rate, the lower is investment spending. A higher real interest rate makes investment projects more expensive for firms to undertake, so they undertake fewer of them. Second is business managers' and investors' confidence--what John Maynard Keynes called their "animal spirits."

To model the inverse relationship between the level of investment spending and the long-term real risky interest rate, set investment spending I equal to the baseline level of investment (the value of the parameter I_0) minus the real interest rate r times the slope-of-the-investment-function parameter I_r (see Figure 6.10):

$$I = I_0 - I_r \times r$$

Details: The Stock Market

An alternative way of looking at investment—one that would complicate our models too much for us to use it here—sees the level of investment spending as a function of the level of the stock market. Think about what determines stock market values. Most investors in the stock market face a choice between holding stocks--shares of ownership of a corporation that also give you ownership of that corporation's profits or earnings--or holding bonds: a piece of paper that represents a that pays interest. If you invest your money in bonds, you earn the real interest rate r . If you invest in shares of stock, your return is equal to your share of the profits of the companies in which you have invested.

When expected future profits are high, investors will find stocks more attractive than bonds and will bid up stock prices. The stock market will rise. When the real interest rate falls, investors will find stocks more attractive than bonds and will bid up stock prices as well. In either case, the stock market will rise.

However, when expected future profits are high, businesses will invest more. When the real interest rate falls, businesses will find investment projects cheaper and will invest more. The same things that determine the value of the stock market also determine the level of investment spending. The stock market and investment move together: what raises or lowers one raises or lowers the other.

The only significant difference is that causes of fluctuations in investment affect the stock market first and investment spending second. The stock market is thus a very useful leading indicator of investment spending. Thus keep a close watch on the stock market if you want to forecast the level of investment spending.

Notice the pattern used for parameters so far: C_0 , C_y , I_0 , I_r . This should make the symbols used in algebraic equations clearer and easier to remember. The capital letter in the name of each parameter tells you what variable is on the left-hand-side of the equation in which the parameter appears. A “C” means that this parameter is part of an equation determining the level of consumption spending C ; an “I” means that this parameter is part of an equation determining the level of investment spending; and so forth. The subscript tells you by which variable the parameter is multiplied in that equation. For example, I_r is the amount by which investment spending I changes in response to a change in the real interest rate r .

Like the consumption function, the investment function is an enormous simplification of real-world investment patterns.

Government Purchases

The federal government buys the labor of government employees—judges, air traffic controllers, customs inspectors, FBI agents, National Oceanic and Atmospheric Administration researchers, and others—as well as military hardware, sections of the interstate highway system, and other goods and services. All these expenditures make up the *government purchases* component of GDP. Such government purchases of goods and services add up to about 25 percent of GDP, counting together the purchases of local, state, and the federal government.

Note that government *spending* is larger than government *purchases*. The government also spends by transferring money to citizens through Social Security and other payments, disability benefits, food stamps, and other *transfer payments*. Because these transfer payments are not themselves demand for final goods and services, they do not show up directly as a piece of GDP in the government purchases total G . Rather, transfer payments show up in the NIPA as negative taxes. The variable T --taxes--represents *net* taxes, taxes less transfer payments. It is the net amount by which the government’s tax and transfer system reduces disposable income.

Recap: Domestic Components of Aggregate Demand

Consumption spending depends on four factors: (a) the baseline level of consumption C_0 , (b) the marginal propensity to consume [MPC] C_y ; (c) the tax rate t ; and (d) the level of real GDP Y .

$$C = C_0 + C_y \times (1-t)Y$$

Investment spending depends on three things: (a) the baseline level of investment I_0 ; (b) the interest sensitivity of investment I_r ; and (c) the real interest rate r .

$$I = I_0 - I_r \times r$$

We leave the determinants of government spending to the political scientists.

International Trade

The final component of GDP is net exports--the difference between gross exports and imports. Gross exports are made up of goods and services that are produced in the home country and then sold abroad. GDP is a measure of production, and since gross exports are part of production, they need to be counted in GDP. But first imports need to be subtracted from GDP. Not all the goods and services that make up consumption, investment, and government purchases are produced domestically. Consumption, for instance, includes spending on Chinese toys, Irish computers, Brazilian coffee, and Scottish tweeds as well as on domestically-made consumption goods. So adding up C , I , and G overestimates domestic demand for U.S.-made products. By adding *net* rather than *gross* exports to $C+I+G$, economists (a) take account of goods made here that are sold to foreigners and don't show up in $C+I+G$, and (b) correct $C+I+G$ for the amount of foreign-made goods it counts.

The volume of gross exports from the United States depends on two variables. The first is the real GDP of our trading partners—call it Y^f , for foreign GDP. The second is the real exchange rate—call it ϵ . The higher the value of the real exchange rate—the more expensive foreign currency—the cheaper U.S.-made goods are to foreigners, and the more of them they buy.

$$GX = (X_{yf} \times Y^f) + (X_\epsilon \times \epsilon)$$

We model gross imports IM as a constant share—a share that is the propensity to import IM_y --of real GDP Y .

$$IM = IM_y \times Y$$

Net exports NX are the difference between gross exports and imports:


$$NX = GX - IM = (X_{yf} \times Y^f) + (X_{\epsilon} \times \epsilon) - (IM_y \times Y)$$

What determines the exchange rate?

Consider foreign exchange speculators whose job it is to trade currencies and make money. They spend their days glued to computer terminals, watching the prices of bonds denominated in different currencies flash across the screen. They buy and sell bonds and stocks of different countries and governments denominated in different currencies--dollars, euros, pounds, yen, pesos, ringgit, and more than one hundred others.

Their lives are ruled by greed and fear.. *Greed*: Suppose a trader sees higher interest rates paid on the bonds of U.S. companies denominated in dollars \$ than of German companies dominated in euros €.; then there is money to be made by selling—going short--German companies' bonds, buying—going long--American companies' bonds, and pocketing the extra interest. *Fear*: Suppose that the trader is long dollar-denominated bonds and short euro-denominated bonds, and the U.S. exchange rate rises. At a higher real exchange rate, each dollar is worth fewer euros, and whatever profits were expected from the interest-rate spread have been wiped out by the capital loss caused by the exchange rate movement. If today's value of the exchange rate is different from long-run historical trends, the fear that exchange rates will return to their normal relationships and impose large foreign exchange losses will be immense.

Greed and Fear in Foreign Exchange Markets



The greater the difference in interest rates in favor of dollar-denominated securities, the higher the greed factor. And the higher the greed factor, the lower must the value of the exchange rate be in order for fear to offset greed. The enormously liquid, enormously high-volume, enormously volatile foreign exchange markets settle at the point where greed and fear balance. Thus the real exchange rate ε is equal to the average foreign exchange trader's opinion ε_0 of what the exchange rate should be if there were no interest rate differentials, minus a parameter ε_r times the interest rate differential between domestic real interest rates r and foreign real interest rates r^f :

$$\varepsilon = \varepsilon_0 - \varepsilon_r \times (r - r^f)$$

The longer interest rate differentials are expected to continue, and the more slowly real exchange rates are expected to revert to trend, the higher ε_r will be, and the larger will be the effect of a given interest rate differential on the exchange rate.

Remember: the exchange rate is the value of foreign currency: if foreign currency becomes more valuable, the exchange rate rises; if domestic currency becomes more valuable, the exchange rate falls. Often you will hear people talk of an appreciation or revaluation of the dollar, or of a depreciation or devaluation of the dollar. An appreciation or revaluation of the dollar is a *reduction* in the value of the exchange rate. A depreciation or devaluation of the dollar is an *increase* in the value of the exchange rate.

If we take the equation for net exports:

$$NX = GX - IM = (X_{yf} \times Y^f) + (X_\varepsilon \times \varepsilon) - (IM_y \times Y)$$

and substitute into it the equation for what the value of the exchange rate, the result is the relatively unappetizing:

$$NX = GX - IM = (X_{yf} \times Y^f) + (X_\varepsilon \times \varepsilon_0) - (X_\varepsilon \times \varepsilon_r \times r) + (X_\varepsilon \times \varepsilon_r \times r^f) - (IM_y \times Y)$$

This equation is unappetizing because it is complex. It is nevertheless valuable because it contains a lot of information. It tells us directly how domestic and foreign interest rates affect net exports, without requiring us to go through the intermediate step of calculating the real exchange rate. This direct equation is sometimes valuable because removing the exchange rate means that (for the moment at least) we have one fewer thing to keep track of. Tweaking the model to rid it of complicating variables is a standard technique of economists.