

How Economists Think

Economists don't have much of a reputation as delightful company. Victor Fuchs—a preeminent health care economist at Stanford University—likes to say, “Some people talk in their sleep, but economists talk in other people's sleep.”

With that kind of encouragement, why study economics? Economic issues are central to many elements of our lives, including not only jobs and income, but also health care, education, retirement prospects, and America's future position in the world economy. If you're going to participate in the great social conversation about economics, a conversation that is happening all around us, you need the ability to “talk the talk.” Maybe you've already discovered this the hard way: You're having a friendly argument over the minimum wage, or the budget deficit, or national health insurance, and at some point the other person sniffs and says, “Well, even the most basic economics shows that . . .” and then repeats whatever his or her argument is. Now, in my experience, people who make assertions about what “basic economics” shows have only a fifty-fifty chance of being correct. But if you don't know any economics, basic or otherwise, you can't dispute the claim. All you can do is nod or shrug. As British economist Joan Robinson (1978, p. 75) once wrote, the reason to study economics is “to avoid being deceived by economists.”

How much economics do you really need to know to participate in such conversations, whether socially or professionally? Brace yourself

for a shock: Herb Stein (1991, p. 6), who worked in various roles as a U.S. government economist for almost fifty years, noted that "Most of the economics that is usable for advising on public policy is about at the level of the introductory undergraduate course." All right, maybe that advice wasn't a huge shock: we do live in cynical times. But the point is that you don't have to be eligible for that tenured chair in economics at Harvard or Stanford to hold your own in most everyday economics debates. You just need to understand the economist's way of thinking.

Let's lay some groundwork, beginning with the three basic questions of economics:

- What should be produced by a society?
- How should it be produced?
- Who gets to consume what is produced?

These three questions are fundamental to every economic system and indeed every society: capitalist or socialist or communist, low-, middle-, or high-income. It's useful to think of the possible answers to these questions as falling along a spectrum. At one end of the spectrum is total government control: government institutions determine what is produced, how it's produced, and who gets it. At the other extreme, you can imagine a society in which individuals make all the decisions about what, how, and who. In the real world, of course, very few societies occupy either extreme.

Let's think about what it would mean to move along this spectrum. Putting aside pure anarchy, we could start at one end with a society in which the government provides only the basics for a market economy: prosecuting theft, enforcing contracts, and providing minimal infrastructure such as national defense. This is sometimes called the "night watchman state." Farther along the spectrum, you can imagine a society with a slightly broader scope for government, adding public services such as roads and education to the night watchman's responsibilities. The next step might be taking on what's sometimes

called a social safety net: a system of national pensions (such as Social Security) and nationalized health care. An even broader government might be responsible for supporting or even partial ownership of certain industries, such as steel or agriculture; it might control the distribution of food or basic consumer goods, such as housing. At the other extreme, you could imagine a government that hands out all the jobs, all the housing, all the food, one that determines what everyone makes and what everything costs.

In the great debate between government control and individual freedom, there's a long tradition of treating the people at the other end of the spectrum as if they were idiots—or monsters. But modern economics recognizes both that markets have their strengths and that, in some situations, markets don't work especially well and the government may have the ability to do something genuinely helpful. Modern economics also recognizes situations in which government intervention hasn't worked well and in which it would probably work better to let the market have a go. To think like an economist, you have to step beyond the ideological arguments about market versus government and get pragmatic. It's necessary to dig down into understanding how markets *really* work and what to do, in practical terms, when they don't work well.

With this idea of what economics is, it's useful to clear away some misconceptions about what economics is not. For one thing, economics is not about predicting the future. It's a common complaint that economists can't say when a recession is going to start or end, or when the stock market will rise or fall. But economists aren't fortune-tellers, able to predict every factor that might affect consumption or production in an economy.

Economics is also not about taking political sides. A lot of people ask me (often in a polite and coded way!) whether I'm a Republican or a Democrat or a Libertarian or a Green, but in teaching basic economics, political affiliation isn't relevant. Economics is not about supporting business or labor, or Democrats or Republicans. Economics is a framework for thinking about the questions.

As an entrée into how economists think, let's consider some statements that most economists would view as obvious but many non-economists wouldn't.

Trade-offs should be taken seriously. Well, everybody believes that, right? Actually, no, they don't. Think about the question of whether, if a government needs to raise additional revenue, it should raise the tax on individuals or on corporations. In public discourse, this tends to boil down to asking, "Which do you care about, corporations or people?" But an economist sees the bigger picture. If you raise taxes on corporations, where do the corporations get the money? They could raise the prices of their products sold to consumers; they could cut the bonuses of top executives; they could cut the dividends they pay to stockholders—all of which would mean less money in the pockets of some actual person. My point here is not that corporate taxes should or shouldn't be raised, but that any sensible discussion of corporate taxes should focus on which actual people are likely to end up paying the tax. Similarly, when the media report on economic issues, they tend to start their stories with a person. Perhaps it's Joe, who just got laid off from a job at a failing company, or Susan, who depends on a social program that's being cut. This is sometimes called "putting a face on the news," and it's effective journalism. But when I hear about Joe or Susan, I always wonder about all the people who aren't in the news story but who are affected in one way or another by the same issue. As economists sometimes say, the plural of "anecdote" is not "data." Many economic choices have the characteristic that they help some people and hurt others. Economists care about all the statistical people who are hurt or helped, not just the individual faces in a news report.

Self-interest can be an effective way of organizing a society. If you ask a number of people, "What would happen if everyone in a society behaved in a purely selfish way?" most of them would reply that it would lead to chaos. But many everyday market exchanges rely on self-interest: shopping around for the best deal, waiting for a good price before selling your house, and so forth. Adam Smith (1776, pp. 484–85), arguably the founder of the discipline of economics, wrote:

Every individual . . . generally, indeed, neither intends to promote the public interest, nor knows how much he is promoting it. . . . He intends only his own security . . . only his own gain, and he is in this, as in many other cases, led by an invisible hand to promote an end which was no part of his intention. Nor is it always the worse for the society that it was no part of it. By pursuing his own interest he frequently promotes that of society more effectually than when he really intends to promote it.

The idea of the "invisible hand" is that in pursuing your own self-interest, you may benefit others. By producing a better product, for example, you're improving the life of the person who uses it. Adam Smith knew perfectly well that the invisible hand is not a magical cure for all an economy or a society's woes. But economists view self-interest as a powerful force, which when appropriately directed, can provide broad social benefits.

For example, if you wanted to get people to conserve energy, how might you go about it? You might develop a huge public relations campaign and get the word out on TV and in the school curriculum. But an economist is more likely to say, "You want people to use less oil? Tax it; they'll use less. You want firms to develop more fuel-efficient cars? Subsidize that technology; the firms will then do the research and development that will make it happen. You want people to use more solar power in their homes? Give them a tax credit; they'll put the extra money toward the installation." If there's something you want less of, discourage it with a tax; if there's something you want more of, encourage it with a subsidy. Such choices in any individual case may be wise or unwise public policy for a variety of reasons (discussed in more detail later in this book), but at least they work with incentives, rather than ignoring them.

All costs are opportunity costs. When you make a choice, the thing you didn't choose is what economists call an "opportunity cost." For example, if you want to hire someone to clean your house, let's say it will cost you \$150 per cleaning, two cleanings per month. So you

could say it costs you \$3,600 a year to have your house cleaned, or you could say it costs you a week's vacation at a beach resort in Mexico to have your house cleaned. The true cost is not the money you've spent; it's the thing or things you give up. Thinking in terms of opportunity costs includes costs that aren't measured in terms of money. If you attend college full time, you are giving up time you could spend doing something else—including working for pay. That opportunity cost is part of the cost of attending college.

Prices are determined by the market, not by a producer. In everyday conversations, you've probably heard someone make a comment such as "my landlord raised my rent" or "those big oil companies raised fuel prices" or "the banks raised my interest rate." But when, say, gasoline prices drop, you probably don't hear anyone say, "Oh, those generous oil companies. So nice of them to give us hardworking folks a break!" Or when interest rates are low, people don't say, "Those generous banks—how sweet of them to give me more for my money." To an economist, the basic premise behind both the blame and the praise in these statements is faulty. Economists certainly agree that landlords and gas companies and bankers are greedy and are trying to make the most money they can, but they're greedy all the time. They raise rents and prices and rates not because they want to—they always want to—but because market conditions of supply and demand shift in a way that allows them to do so.

No person can have everything he or she wants. No society can have everything it wants. Trade-offs are unavoidable. In a modern economy in which people have a wide variety of skills and desires, the question is how to coordinate the decisions about what is produced, how it is produced, and for whom it is produced.

CHAPTER 2

Division of Labor

In our modern world, even seemingly simple consumer goods are often produced through a complex process that reaches across the world. Let's take the pencil, for example. In 1958, an economic educator named Leonard Read wrote an essay, "1, Pencil," that described the remarkable process of pencil production. The wood comes from Northern California, where it must be logged, then shipped and milled. The lead is a mixture of graphite produced in Ceylon and clay mined in Mississippi, which are combined in a process performed at yet another location. The yellow paint on the outside is made from castor beans—that's three more steps: growing, shipping, and paint making. The brass sleeve that holds the eraser is made from zinc, copper, and nickel, which must also be mined, shipped, and refined. The eraser is a mixture of vegetable oil from the Dutch East Indies, pumice from Italy, and various binding chemicals—imagine how many steps that is for the eraser alone. In the essay, Read claims that there is no single person in the world who could make a pencil from scratch, and he may well be right.

A pencil is disposable. If you drop one on the floor, you may well let it roll away without thinking twice. But what it takes to make a pencil, considered closely, is awe-inspiring. What's even more awe-inspiring is that almost everything in the economy is the result of this kind of near-miraculous economic coordination.

This division of labor creates substantial economic gains, both at

Supply and Demand

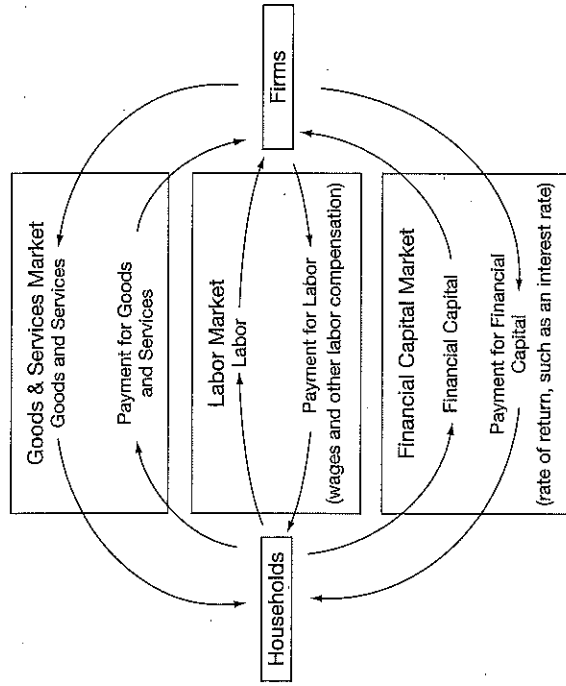
You should be starting to glimpse how economists see the world: the division of labor leads to exchange of goods and services; somehow a society has to coordinate all that production and consumption. All high-income societies of the world such as the United States, Canada, Japan, and the countries of Western Europe primarily coordinate their economies through market arrangements, influenced to a greater or lesser extent by government. Let's take a deeper look at how markets work together in the economy as a whole.

We're going to start with a circular flow diagram, which pictures the economy in terms of flows of goods, services, and payments between two groups, households and firms, through three markets: goods, labor, and financial capital.

The goods market includes all the items that households buy: food, clothes, furniture, haircuts, phone service, computers, and so on. In the goods market, goods flow from the firms, which produce the goods, to the households. Households make payments for the goods, which flow back to the firms. In the jargon of economists, firms are the suppliers of goods and households are the demanders of goods.

In the labor market, labor flows from households—that is, from people who work—to the firms who hire those workers. For example, the Target Corporation has about 350,000 employees. Payments flow

THE CIRCULAR FLOW DIAGRAM



from the company to the workers and their households in the form of wages and benefits. In this market, the roles of supply and demand are reversed from the goods market; firms demand labor, while households supply labor.

In the financial markets, households invest money—either directly by buying shares of stock or indirectly by putting their savings in a bank, which in turn invests in or lends to firms. The households are paid for those investments by the firms in the form of interest and dividends. Thus, households are suppliers of financial capital and firms are demanders. (It should be noted that firms may also supply capital, but they invest on behalf of their owners and shareholders—that is, households again.)

The circular flow diagram shows how all three circles run through households and firms, and thus, how all three markets are part of a larger, integrated macroeconomic whole. Later on we'll look at how government and other nations interact with these three loops, but for

starters, let's focus on these two sets of actors and three markets in which they interact.

In the first of three markets—the goods market—where do prices come from? When many noneconomists talk about prices, they talk about prices being “too high” or “too low,” which is best understood as a way of comparing the world as it is to the world as they think it ought to be. You’ll hear statements such as “nurses are paid too little” or “gasoline costs too much.” To an economist, this sort of judgment is like saying the weather today is “too cold” or “too hot.” It tells you something about the preferences of the person, but nothing about why things are as they are.

To noneconomists, prices are typically value-laden. Economists try to avoid those sorts of value judgments, which we call the diamond water paradox. This paradox goes back to the forefather of all economists, Adam Smith (1776 [1994], pp. 31–32), and *The Wealth of Nations*, in which he draws a distinction between “value in exchange” and “value in use.” Diamonds have great value in exchange. If you have one to trade, you can get a lot for it. But diamonds don’t have much value in use—you can’t eat them, they don’t trim your hedges for you, and they make lousy paperweights. They’re basically frivolous, a luxury. Water, on the other hand, is one of the basic necessities of life, not to mention water’s not-so-basic uses, such as transport and steam power. It has very high value in use. But water is also extremely cheap. In most places, it falls out of the sky for free, and in normal circumstances its value in exchange is correspondingly low.

Clearly, value in exchange and value in use don’t always line up. So when we look at the price of an item, which of these values are we talking about? When economists talk about price, we’re talking about the exchange value. The exchange value of a good is tied to its scarcity—how much of the good there is relative to how many people want it. Diamonds have a high price because, compared with how many diamonds there are, lots of people want them enough to pay a high price. Water has a low price because,


compared with how much is available, people aren’t willing to pay much for it. You can imagine a situation in which someone dying of thirst would be willing to trade diamonds for water, but that’s not the norm!

The dramatist Oscar Wilde (1891) once defined a cynic as “a man who knows the price of everything and the value of nothing.” That’s also a good description of economists, who focus on the price of everything and the intrinsic value in use of nothing. To think about price like an economist, you need to purge your mind of preconceptions about a good’s value in use. When you get used to it, separating prices from judgments about value is emotionally liberating. You don’t have to think about whether a price is “right” or whether it’s an accurate reflection of your personal values. Price is what happens out in the world from the interaction of supply and demand, from what’s available and what people want.

I’ve been tossing the terms “supply” and “demand” around rather loosely, but they actually have quite specific meanings. When economists talk about demand for a good, they’re referring to the relationship between the price of a good and the quantity of that good that’s demanded. For most goods, most of the time, as the price of the good goes up, the quantity demanded tends to drop.

This is easy to visualize with a graph. The quantity of the good is on the horizontal axis, and the price of the good is on the vertical axis. The curve representing demand slopes downward. That downward slope tells you that as the price drops further and further, the quantity demanded gets bigger and bigger.

At the gut level, this pattern makes sense, but what’s the actual cause? Economists have offered two specific reasons. One is the “substitution effect.” As the price of a good goes up and up, people tend to substitute other goods for it. For example, as the price of orange juice goes up, people substitute other drinks, or perhaps vitamin C pills. As the price of gasoline goes up, people might drive less, they might join carpools, or they might buy cars that have higher gas mileage.



The other reason is called an “income effect.” As the price of a good rises, your income has less buying power, so you can’t buy everything you did before—you buy either fewer things or less of the same things. For example, if you’re someone who likes to buy fancy coffee every morning on your way to work, and the price of your favorite beverage goes up to \$100 a cup, you probably can’t buy that cup of coffee every morning. The buying power of your income in terms of coffee is reduced. Even if the price rises only a little bit, the buying power of your income is diminished, and the income effect forces you to buy less of that good or other goods.

It’s important not to confuse the terms “demand” and “quantity demanded” as economists use them. “Quantity demanded” refers to the specific amount of a good that is desired at each given price. In 2009, about 120 million bags of coffee were sold at a price of \$1.15 per pound. “Demand” refers to the relationship between price and quantity demanded. It refers to how much is desired at any possible price or at every price. As the price of coffee rises, the quantity of

coffee demanded will decline. In terms of our graph, quantity demanded is a point, but demand is the curve.

So here’s a trick question: What makes demand rise or fall? The answer is not price; price affects the quantity demanded, but it doesn’t cause the demand relationship itself to move. When economists talk about “demand” shifting, we’re not talking about one point moving up or down; we’re talking about the whole demand curve shifting up or down by the same amount. We’re talking about a situation in which, at every given price on the axis, the same larger or smaller quantity is demanded. What could cause such a change?

- What if income rises for society as a whole? If everyone had more money, there would be greater quantities demanded for most goods across the board at pretty much every price.
- What if a society has a population boom? If you have more people demanding goods, that equates to higher quantities demanded at every price.
- What about tastes and fads? Certain things become more or less popular all across a society—such as people consuming more chicken and fish and less beef. In that example, the quantities of chicken and fish go up and the quantity of beef demanded goes down at any given price. The demand for chicken and fish rises, and the demand for beef declines.
- What about a change in the price of a substitute good? In the previous scenario, if most people think chicken is the best substitute for beef, its price goes up, and in response, people move away from chicken and demand more beef. Conversely, if the price of chicken goes way, way down, people buy more chicken, and the demand for beef goes down.

Now let’s turn to supply. Supply refers to the relationship between the quantity of a good supplied and the price of a good. As the price of the good goes up, the quantity supplied tends to rise, too, because as the good’s price rises, firms become more willing to supply

that good. So whereas our demand curve sloped downward, the supply curve slopes upward.

Once again, this probably makes some intuitive sense, but economists have tried to spell out specific reasons behind the pattern. First, existing firms want to produce more as the price rises because they can earn higher profits. Second, new firms decide to enter the market and start producing the good if the price rises enough.

Just as there's often confusion between "demand" and "quantity demanded," there's a parallel confusion between "supply" and "quantity supplied." Quantity supplied refers to the specific amount produced at a given price. Supply refers to how much is produced at every price. Quantity supplied is a point, and supply is a curve.

Here's that trick question again, this time from the supply perspective: What makes supply rise or fall? Again, it's not price. Price causes the quantity supplied to change, but it doesn't shift the entire supply curve. For supply to increase, the whole supply relationship has to move, so that at every given price, a larger quantity will be supplied. Conversely, for supply to fall, at every given price, a smaller quantity must be supplied. What are some examples of factors that could move supply in this way?

- What if there's a change in technology? A cheaper production method could mean that a greater quantity of a certain good could be supplied at every given price.
- What if production is affected by weather? This factor is particularly important in agriculture. Better weather means greater crop yields, which means more of a good supplied at every given price; bad weather means lower yields, which means less of a good supplied at every given price.
- What about a change in input prices? An input price is the cost that goes into making a good. If a firm uses a lot of oil or a lot of steel to make its product, and the price of oil or steel goes up, then the quantity supplied of their product at every given price is going to fall.

Now we're ready to look at how supply and demand interact. Let's start off by thinking about a basic good, such as pizza. Let's first consider the situation with a low price. At a low price, the quantity supplied is correspondingly low because nobody wants to produce the good, but the quantity demanded is likely to be relatively high because many people want to buy lots of pizza at that low price. As the price of pizza rises a little, the quantity supplied increases; restaurants want to produce more of the good. But when the price rises, the quantity demanded drops off; people are less willing to buy the pies. As quantity supplied rises and quantity demanded falls, at some point the quantity of pizza demanded is equal to the quantity supplied. That point is called the equilibrium.

What does equilibrium mean in practical terms? If a good's price is higher than the equilibrium, then the quantity supplied of that good will exceed the quantity demanded. Stuff starts piling up on the shelves. To get rid of it, the seller has to start cutting prices until people are willing to buy. The price starts to drop toward that equilibrium point, where the quantity supplied and the quantity demanded meet. Equilibrium is the point at which price and quantity are efficient in the specific economic sense that nothing is being wasted. Just as an efficient machine has no wasted motion, or no extra parts, an efficient market has no extra products sitting around, or no extra demand for products that are not there.

If the price of a good falls below the equilibrium, then quantity demanded will exceed quantity supplied. In that situation, people are practically lining up to buy that good before supplies at the local store run out. Suppliers notice this, and they begin to raise prices. As a result, quantity demanded begins to fall and quantity supplied starts to rise, until again, the two quantities are equal and the price reaches equilibrium.

Equilibrium is the point toward which a market economy tends—but that's not to say that markets are always at equilibrium. There are long-standing disputes over how long it takes markets to reach equilibrium, how close they usually are to equilibrium, and when or whether

market prices will overshoot equilibrium and need to bounce back. In the mid-2000s, housing prices in the United States clearly veered well above equilibrium for a period of several years before the pendulum began to swing the other way. But over time, markets typically *tend* in the direction of equilibrium.

Any change in demand or supply—remember, that's a change in the whole curve—will cause the equilibrium point to shift. Consider, for example, the market for beef. If income rises, then the demand for beef rises. The result would be a situation in which the new equilibrium has a higher price and quantity sold in the market. Now imagine the opposite situation. Let's say there's an outbreak of cattle disease, and the supply of beef drops as a result. The outcome would be a lower quantity and a higher price sold at equilibrium. A lot of an introductory economics class is devoted to thinking through the consequences of these demand and supply shifts. The specific goods and examples change, but the basic pattern is the same: Think about demand, think about supply. Start at equilibrium; think about what would happen if demand or supply shifted. Think about what new price and quantity would result at the new equilibrium. Supply and demand is a framework for discussing how prices and quantities are determined in markets and why those market prices and quantities are going to change. Understand this, and you're on your way to a solid grasp of the basics of economics.

In the real world, equilibrium means only that quantity demanded and quantity supplied are in balance; it doesn't mean that people feel content with the result. Some buyers will always say, "I think I'm paying too much." Some sellers will always say, "I can't believe it's selling for so little." Sometimes buyers or sellers will go to the government and lobby for a change in the price of a good, even when the market for that good is at or near equilibrium. In the next chapter, we'll talk about some consequences that arise when this happens.

One common complaint about the supply-and-demand model is that "real people don't think that way." At one level, this is obviously

true: most people don't use these terms or draw graphs in their heads. But as long as buyers search for what they prefer at the lowest possible price, taking their desires, their finances, and their possibilities for substitution into account, and as long as firms adjust their production in response to changes in price, the supply-and-demand approach will work reasonably well. The reality of supply and demand may not always be likable, or morally attractive, or desirable in any deep philosophical sense, but it is a useful tool—a powerful and accurate way of describing and understanding why prices are at the levels they are and why prices might be rising or falling. It works as a way of describing markets all over the world, at all different times in history, and for all sorts of products—from pencils to pizza.

Macroeconomics and Gross Domestic Product

Macroeconomics is the aggregated top-down view of the whole economy as one large organism: growth, unemployment, inflation, the trade balance, and more. A macroeconomic perspective is not just a scaled-up version of microeconomics, which has been the main focus of the earlier chapters. Microeconomics focuses on individual markets for goods, labor, and capital, and on issues that arise including monopoly, competition, pollution, technology, poverty, inequality, insurance, and governance. But microeconomics has no real language for talking about overall macroeconomic problems in the economy, such as growth and trade deficits. Microeconomics can explain why the price of one good might rise, or why one firm might hire or fire workers, but it cannot come to grips with macroeconomic issues such as inflation or unemployment for the entire economy. Moreover, microeconomics has no structure for talking about macroeconomic policies such as the federal budget and budget deficits or the actions of the Federal Reserve in affecting interest rates and credit.

In fact, behavior that is rational for individuals at the microeconomic level can lead to unexpected results when everyone in a group acts that way. Imagine you're in a stadium among a big crowd watching a concert. You want a better view of the antics onstage, so you stand up. Then some others stand up for a better view, and eventually everyone is standing up. Everyone acted rationally from a microeconomic,

individual point of view, but the end, macroeconomic result was that no one saw any better than they did before.

To provide a broad framework for thinking about macroeconomic policy, we will look at four broad goals of macroeconomic policy, then at an analytical framework for thinking about the relationships between these goals, and then at two sets of tools for accomplishing those goals. The four goals of macroeconomic policy are: (1) economic growth, (2) low unemployment, (3) low inflation, and (4) a sustainable balance of trade. The framework for discussing policy is called the aggregate supply/aggregate demand model. This framework helps to organize macroeconomic analysis and should allow us to analyze the trade-offs among growth, unemployment, inflation, and the balance of trade. The two main sets of tools for macroeconomic policy are fiscal policy and monetary policy. Fiscal policy is government tax and spending policy, including the federal budget and budget deficits. Monetary policy refers to the policies of the Federal Reserve, which affect interest rates, credit, and how much money is being loaned and borrowed in the economy.

Before embarking on the topic of macroeconomics, it's important to have a firm grasp of the concept of gross domestic product (GDP), which is the standard measure for the size of a nation's economy. GDP is defined as the total value of final goods and services produced in an economy in a year. GDP can be measured according to the value of either what is produced and sold or what is demanded and bought. Because the total quantities of what is bought and what is sold need to be equal, by definition, these two ways of measuring GDP will lead to the same answer.

For example, in 2009 the total GDP of the United States was \$14.2 trillion. In terms of what was produced, this was 13.4 percent durable goods—things such as refrigerators and cars; another 13.4 percent nondurable goods, such as food and clothing; 66.2 percent services; and 7.7 percent structures. (This adds up to slightly more than 100 percent because about 1.1 percent of what was produced was added to inventories, and since it was not yet sold, it was not

counted as output.) When many people think about the economy, the first thing that leaps to mind is durable goods, hard stuff that comes out of factories. But the nondurable goods and services—such as health care, education, financial services, legal consulting, haircuts, car repair, lawn-mowing or home-cleaning services, and childcare—are well over half of all U.S. output. When people say we live in a service economy, that's what they mean. The share of services in the U.S. economy has been expanding for decades.

You can also measure GDP from the demand side. Demand from household consumption accounted for about 70 percent of GDP in 2009. Investment spending by businesses was running about 11 percent in 2009, but it varies considerably from year to year. Government spending is about 21 percent of GDP. That may seem a bit low when you recognize that federal, state, and local governments together collect about one-third of GDP in taxes in a typical year, but this figure represents only what government buys directly; government programs that pass money along to citizens, such as Social Security and welfare, show up as part of consumption. The final categories are exports and imports. Exports are demand from other countries for what is produced in the United States. Exports are added to total demand, but imports—U.S. demand going to products made abroad—are then subtracted from total demand. In recent years in the United States, imports have been a lot bigger than exports. That means the United States has a trade deficit, a subject we'll look at later on. Some quick shorthand is useful here. If you ask an economist what GDP is, sometimes the economist will answer that it is $C + I + G + X - M$: consumption plus investment plus government spending plus exports minus imports.

Who actually measures this stuff across the fourteen-figure U.S. gross domestic product? A branch of the U.S. Department of Commerce called the Bureau of Economic Analysis (BEA) collects data from all sorts of surveys and sources—some monthly, some quarterly, some annually. Sometimes the government statisticians need to extrapolate. Right after each quarter, they make a preliminary estimate;

they issue a final estimate later when all the data is in. Every five years, the BEA reviews all the results it has issued over that period and corrects them—sometimes substantially.

Economists often look next at per capita GDP, or GDP per person. This is a simple but rough way of estimating the standard of living at any given place and time. In 2009 that \$14.2 trillion in GDP would be divided by the population of 307 million, which equals about \$46,000 per person. Per capita GDP is useful as a comparison tool because it automatically adjusts for differences in population, either between countries or over time.

The next important calculation is called real GDP, meaning GDP adjusted for inflation. Say that in a given year, GDP increased 5 percent over the previous year. When the statisticians look at the change in prices, it turns out that 3 percent of the overall 5 percent rise in GDP is attributable to inflation—that is, to prices being higher. As a result, the other 2 percent is the real increase in goods and services being produced. Adjusting for inflation is important for short-run comparisons, and indispensable for long-run comparisons.

GDP has its share of conceptual imperfections. Indeed, *all* economic statistics have conceptual imperfections. Because GDP is a measure of what is bought and sold, things that affect the quality of life but aren't bought or sold aren't directly included in GDP. A classic example is home production. Around the late 1960s and the early 1970s, there was an enormous surge of women into the paid workforce. As a result, many goods and services that had previously been produced at home—meals, house cleaning, and childcare—were now more likely to be bought and sold in the market, and are now included in GDP. That's a significant change in GDP that does not reflect an actual change in the amount of those goods and services being produced in the economy—it only reflects whether these goods were produced at home or purchased in markets.

Many other things affect people's standard of living and their happiness but cannot be measured as things that are bought and sold. For example, if everyone worked ten hours a week less or had an extra

two weeks of vacation every year, but output remained the same, GDP would not show any overall gain. Greater or lesser pollution levels don't show up directly in GDP measures. Traffic congestion or length of commute doesn't show up as something bought and sold, except indirectly, such as in gasoline or cups of takeout coffee consumed. Negative events, such as a natural disaster, can lead to the rebuilding of a large part of a city, which makes short-term GDP look positive, but the locals have clearly suffered a lower standard of living. The costs of preventing crime count as part of GDP, but the costs of the actual crimes in terms of loss and violence are not part of GDP. The fact that people have longer life expectancies and are living healthier for longer doesn't show up in any direct way in GDP, although their spending on health care services is counted.

Even within GDP, you might ask—if you were in a skeptical frame of mind—if things that are priced the same really have the same value. For example, is a bag of potato chips the same value as an organic fresh apple? Does *Fortune* have the same value as a pornographic magazine? Is a ticket for admission to a shoot-'em-up movie the same value as a ticket for admission to an art museum? GDP is just about adding up what's bought and sold. It's not about value judgments.

One final note: GDP includes only finished goods, not the intermediate parts that go into making those goods—the steel that goes into the car, the lumber that goes into the chair, and so forth. If you added up the production of all the parts that go into making your car at each stage of production, you'd be counting the same steel or plastic over and over at each stage. Also, transfers of ownership do not show up in GDP; so, for example, it registers the new car you bought, but it doesn't register the used car you sold to your neighbor's son. It registers a new home you constructed or the costs of renovating an old home, but not an old home you bought. When you buy and sell shares of stock in a company, that doesn't register on the GDP, either, since nothing has been created, just exchanged. Only the stockbroker's fee, as a service, is included. When the stock market rises or falls substantially, it has no direct effect on GDP at all.

Despite these limits, the GDP is a worthwhile measure of the state of an economy. Societies with higher per capita GDP tend to be better off in lots of ways. They have more personal consumption—not just what we think of as luxuries but also things such as health care and education. Societies with high per capita GDP also tend to have cleaner air and water. They tend to have a greater degree of personal security. As a wise Nobel laureate named Robert Solow (1986) once said, “If you have to be obsessed by something, maximizing real national income is not a bad choice.”

Historically GDP shows a long-term upward trend: After adjusting for inflation, GDP in the mid-2000s was about 5.5 times as large as it was in 1950. The average growth of real GDP from 1950 through 2010 was about 3 percent per year. That's not to say that GDP rose in every single year. A significant and lasting downturn in GDP is called a recession. Some economists consider six months—two economic quarters—to be a lasting downturn, but there's nothing official about that span of time. The starting and ending points of recessions are not defined by any U.S. government agency; rather, they're defined by a committee of academic economists at a nonprofit research institution called the National Bureau of Economic Research (NBER).

The U.S. economy, according to NBER data, had twenty-three recessions from 1900 to 2010. Thus, there was a recession, on average, about every five years. In the aftermath of the 2001 recession, it was common for economists to note that the most recent recessions—in 2001 and in 1990–1991—had been relatively far apart and relatively mild. But the recession of 2007–2009 was brutally long and deep. Thus, a one-sentence summary of the U.S. macroeconomy, as captured by GDP, would report that the long-term trend is up, but this trend is sometimes interrupted, often once or twice a decade, by shorter periods of negative growth. Clearly, the policy goal in this setting should be to discourage or limit these downturns, while nourishing the roots of long-term growth.

CHAPTER 20

Economic Growth

Here's a question for you: Would you rather be a person with an average standard of living in the modern U.S. economy whose household earns about \$60,000 per year, or would you rather earn that same \$60,000 per year in 1925, which is the dollar equivalent of nearly \$800,000 in 2010? Before you choose too quickly, think about this: That \$60,000 income would make you extremely rich in 1925. You'd live in luxury, in a big house, with servants and the best of everything that was available. However, you'd also be living with the technological standards of 1925. No modern telecommunications. No modern transit. No modern health care.

Which would you prefer? When I pose this question to various groups, they tend to answer about two to one in favor of the present. Of course, this question has no right or wrong answer, but it illustrates some of the reasons why people value economic growth even over being relatively wealthy compared with others at a certain point in time.

Economic growth compounds over time. Small differences in annual growth rates work out to enormous differences in the standard of living over a generation or two. The formula for predicting the future value of an economy is to take its present value and multiply it by one plus the growth rate of the economy, raised to the power of the number of years into the future. Actually, it's essentially the same formula we used to calculate rates of return when people are saving for retirement. The process of compounding economic growth rates

is exactly the same as the process of compounding growth in interest rates.

Let's plug in some numbers to see how this works. Imagine you've got an economy that is growing at 1 percent a year. It starts off—just to keep the calculation simple—with a GDP of 100, measured in the imaginary currency of this hypothetical economy. After ten years, it would be just a bit over 110; the compounding doesn't do much over those ten years. After twenty-five years, 1 percent growth a year would give you a GDP of 128. After forty years, GDP would reach 149. Not bad, but not overwhelming, either.

Now let's run the numbers for a growth rate of 3 percent a year, which is more or less the average for the U.S. economy over the past few decades. After ten years of growth at 3 percent, GDP would go from 100 to 134; after twenty-five years, to 209; after forty years, to 326. In other words, over forty years of growth at 3 percent a year, an economy would more than triple in size. The power of compounding starts to make a real difference.

An Economy Starting with GDP of 100 and Growing for Ten, Twenty-five, and Forty Years

TIME	ANNUAL RATE OF REAL ECONOMIC GROWTH		
	3 percent	5 percent	8 percent
Ten years	134	163	216
Twenty-five years	209	339	685
Forty years	326	704	2,172

Note: The table is calculated using the formula $FV(1+r)^t = FV$, where FV is the present value of the size of the economy (in this example, a GDP of 100), r is the annual growth rate of the economy, t is time, and FV is the future value.

What if you had a growth rate of 5 percent a year? That's the rate the U.S. economy might have in a really good year, and many countries, including Brazil and Mexico, have had growth rates like that on

an ongoing basis for sustained periods of time. At 5 percent a year, an economy that starts out at 100 GDP reaches 163 after ten years. It reaches 339 after twenty-five years; it more than triples. And if you could sustain that growth rate for forty years, the GDP would be 704; it would increase more than sevenfold over forty years.

Finally, let's try an 8 percent rate of economic growth. To be honest, 8 percent is the far upper end of what's possible, at least over long periods. We're talking about the fastest periods of Japanese economic growth, back in the 1960s and '70s; we're talking about the growth rate of China in more recent decades. But let's look at the numbers anyway. With an 8 percent rate of economic growth, in just ten years, GDP would go from 100 to 216, so it would more than double in size in a single decade. After twenty-five years, the economy would go from 100 to 685—almost sevenfold in a time period well within the life span of a person; it's even within the working life of a person. If you could sustain that 8 percent rate of growth for forty years—which no country has done, but just for illustration's sake—then the GDP would grow from 100 to 2,172. In other words, over the course of a person's working life, from age twenty-five to sixty-five, they could see the economy expand by a multiple of 22. That's an extraordinary change in standard of living.

The lesson here is that seemingly small differences of a few percentage points in the annual growth rate have a huge effect. In the long run, you can argue that economic growth is the only thing that matters to the standard of living.

When a country starts off behind other nations in terms of the size of its economy, can it catch up? Definitely, if it can sustain a high rate of growth over time. In fact, some economists believe that countries that start off at lower levels of productivity might be able to take advantage of a period of so-called "catch-up growth." The country that is behind can copy and use technologies that have been invented in other places; they don't have to invent it all themselves. It seems as if the lower-income nations of the world should be able to diminish the gap in per capita GDP with higher-income nations gradually over time.

However, this didn't happen much in the twentieth century. The world economy, from about 1870 up to the early 2000s, has actually seen a divergence between the richest and poorest economies of the world, not a convergence. In 1870 the per capita GDP of the richest countries in the world was about nine times the size of the GDP of the poorest countries in the world. By 1960 the GDP of the richest countries in the world was about thirty-eight times greater. In 1990 it was forty-five times greater. Essentially what happened is that the poorest countries in the world were near a subsistence standard of living in 1870 and remained at a subsistence standard of living, while economic growth in the rest of the world has been compounding all that time.

The lowest-income countries are not poor because of globalization; they're poor because they are almost completely detached from the rest of the world economy. In fact, globalization and the international economy have been extraordinarily helpful in lifting countries such as Japan, China, and India out of poverty. Saying that globalization creates poverty is kind of like saying that exercise makes you overweight because you don't do it. If you're not participating, then the activity in which you are not participating is probably not causing the problem, either.

However, even though the twentieth century was a time of divergence between the high-income and the low-income countries, might the twenty-first century be a time of convergence? Economies of countries such as China and India, who were among the poorest in the world in the 1970s, have surged forward in recent decades. If their example can continue, and be followed, over the course of the next century the gap between rich and poor nations will diminish.

Predictions about catch-up growth are controversial. Some argue that we should expect continued divergence. We can't be sure, for example, when Africa will experience a productivity takeoff, if ever. There are several countries, particularly in Latin America, that seem to take a few leaps forward and then a few steps back. Also, if an economy is far behind, it can take an extraordinarily long time to catch up. We can demonstrate this by using our formula for the

effects of long-term growth. Let's say nation A has a real GDP of \$500 per year and nation B has a real GDP of \$30,000 per year. If nation A somehow sustained a remarkable 8 percent growth for forty years, its real GDP would reach \$10,862. Meanwhile, if nation B sustained a mediocre 2 percent growth over that same forty years, its real GDP would be up a little over \$66,000. So after forty years of explosive growth in one country and modest growth in the other, per capita GDP in nation B is only six or seven times higher than that in nation A, instead of being sixty times higher. Even under the most optimistic conditions, the world's poorest nations will remain behind in standard of living half a century from now.

The underlying cause of long-term economic growth is a rise in productivity growth—that is, higher output per hour worked or higher output per worker. The three big drivers of productivity growth are an increase in physical capital, that is, more capital equipment for workers to use on the job; more human capital, meaning workers who have more experience or better education; and better technology, that is, more efficient ways of producing things. In practice, these work together in the context of the incentives in a market-oriented economy. However, a standard approach is to calculate how much education and experience per worker have increased and how much physical capital equipment per worker has increased. Then, any remaining growth that cannot be explained by these factors is commonly attributed to improved technology—where “technology” is a broad term referring to all the large and small innovations that change what is produced.

When economists break down the determinants of economic growth for an economy such as the United States, a common finding is that about one-fourth of long-term economic growth can be explained by growth in human capital, such as more education and more experience. Another one-fourth of economic growth can be explained by physical capital: more machinery to work with, more places producing goods. But about one-half of all growth is new technology. If you do a similar breakdown of the reasons for growth in low-income countries, where education levels and physical capital are being updated more

rapidly, more of their productivity growth tends to come from gains in physical and human capital and less from new technology.

In the 1950s and '60s, U.S. productivity growth—as measured by the increase in output per hour worked—hovered near 3 percent a year. But starting in the 1970s, productivity growth sank to about 1.5 or 2.0 percent per year, where it stayed for approximately twenty-five years, rising again in the late 1990s. This twenty-five-year period suffered a lot of difficult economic adjustments: periods of high inflation and high unemployment; spikes in oil prices; a big defense spending drop and then a big defense spending buildup; a buildup of budget deficits. But of all those issues, the productivity slowdown was the most significant from a long-run perspective. Let's say the economy lost 1.5 percent productivity growth per year for thirty years. In rough terms, not worrying about compound growth rates, real GDP would be, as a result, 45 percent smaller at the end of that time than it would otherwise have been. Thus, in an alternate universe where that growth slowdown didn't happen, GDP would be nearly 50 percent larger. People would have had nearly 50 percent more money in their paychecks. Whether your goal for government is tax cuts or spending increases, having per capita GDP be almost 50 percent larger would make your goal a lot easier to achieve.

Starting in the late 1990s, the U.S. economy began to see a boost in productivity amid a lot of talk about a “new economy,” one built on the dramatic advances in information and communications technology. Productivity growth bounces up and down from year to year, depending on recessions and recoveries, but from 1996 to 2010, productivity growth averaged about 2.6 percent per year. The single most important factor for the long-term prospects of the U.S. economy is whether productivity growth falls back to the quarter century of slow growth from the 1970s into the early 1990s, or whether this higher level of productivity growth can be sustained for the long term.

