CHAPTER ELEVEN: EQUILIBRIUM NATIONAL INCOME AND FINANCIAL BALANCES IN A KEYNESIAN WORLD

The question is simplicity itself. Consider, for instance, Figure 1, which shows us the ups and downs of US economic growth since1972:

Figure 1: Annualized Percentage Growth of US Real GDP, 1972-2011

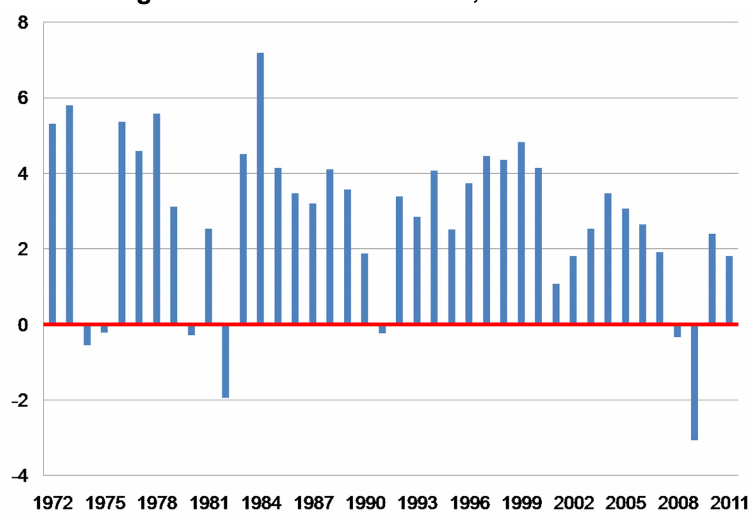

The US economy's growth has varied greatly from quarter to quarter, with several periods of shrinkage (recessions).

In some quarters the economy grew robustly, in others it slumped. The simple question is "Why?"

We would really like to know the answer. If we did, we could hope to control these fluctuations and engineer steady prosperity, now and forever. Economists would be held in very high esteem.

Unfortunately, this simple question does not have a simple answer, and we will devote much of the remainder of this book to investigations that delve into aspects of the puzzle. In this chapter, however, we *will* give a simple response, even though it is incomplete or misleading in several respects. It will provide a framework for thinking about the economy and a starting point for the detail yet to come. In a sense, it functions for macroeconomics in the way supply and demand diagrams functioned for micro: simplifications, but useful in organizing our thoughts.

Our strategy will be to build up the model step-by-step. To do this, we will begin with *extremely* simplified economies and then add one additional piece at a time. It will be somewhat repetitive, but at least we will have the entire sequence. You will note that each piece is based on essentially the same core idea, that it is the size of the national economy that adjusts to arrive at an equilibrium. When we are done, we will go one step further and ask what this model implies for the financial balances of households, businesses and government, linking the NIPA (National Income and Product Accounts) orientation of the model to the beginnings of a financial, Flow of Funds perspective.

Step 1: the consume everything economy

On the Friday after Thanksgiving, millions of American shoppers descend on the country's malls and big box stores to gobble up the sale items while they last. It is a frenzy of consumption that lasts for a few hours—and then it's over. But imagine that an economy operates this way all the time, with consumers spending every penny they have as soon as they get it. And even stranger, imagine that the entire output of the economy consists of consumer goods—nothing is produced for investment. That is the bizarre vision that lies behind the first step in our analysis.

The consumption-only economy, if it exists, adheres to the circular flow of income. Households buy goods from firms, who distribute their revenues back to households. The total income of the economy, which we will designate with the letter Y ("yincome") is the same, whether measured as the value of goods produced or the earnings received by households. Since only consumer goods are produced, we can also say that

(1) $C \equiv Y$

where C is the total value of consumption goods purchased by households. This is an identity, not an equation, as indicated by the use of an "≡" rather than an "=". It is always true because it is a matter of definition—in setting up the model we defined Y as consisting solely of C. Equations are only conditionally true; we solve them to find out what it means *if* they are true. As you saw in Chapter 4, this is an important distinction. Meanwhile, geometrically, the situation appears as Figure 2.

Figure 2: Consumption and Income in the Consume Everything Economy

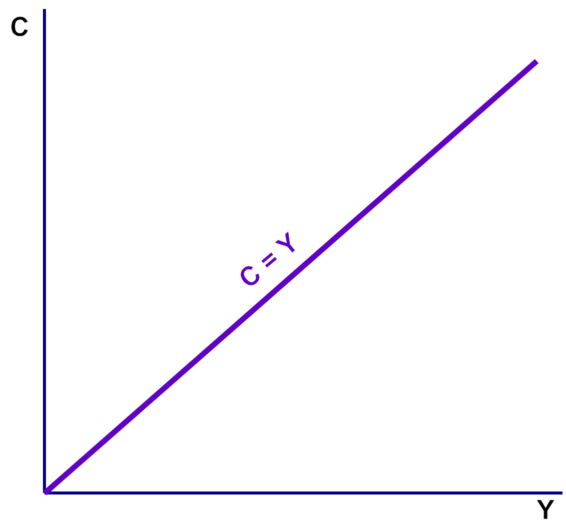

In an economy in which everything is consumed, total consumption and income are indeterminate.

We have drawn a 45º line from the origin to represent all the points at which total consumption in the economy equals total income. What the diagram is telling us is that the equality between income and consumption does not give us enough information to figure out how large either of these will be. Perhaps everyone is working overtime, producing a cornucopia of consumables, then racing to the stores after work to empty the shelves—and provide the revenues so that their employers can pay them to buy all this stuff. Or perhaps only a few people are working and consuming, and everyone else is waiting, waiting.... Either would be consistent with Figure 2: the first possibility would be far out to the northeast along the C=Y line, the second would be very close to the origin. The model isn't helping us much.

Step 2: add savings

So let us complicate matters just a little and give households a choice: they can either spend the money or save it. If they save, they tuck it into a secure place at home—under the mattress or, if they have a waterbed, behind a picture frame. Algebraically, we can write

(2) $C + S \equiv Y$

where S is the total volume of savings properly tucked away. We would like to draw this in a diagram, but we have to make further assumptions about C and S. Let's suppose that the amount households *wish* to spend on consumption is a linear function of income—that it depends on how much income they earn and looks like a straight line in a diagram like Figure 2. We know that the equation for a line consists of a slope and an intercept, so we could write, very generally

$$(3) \qquad C* = a + \frac{\Delta C}{\Delta Y} Y$$

where C* is the desired quantity of consumption, and $\frac{\Delta C}{\Delta Y}$ is the slope of the line—the change in consumption as income changes. We will translate this into economic language as the **consumption function**

$$(4) \qquad C^* = C_0 + MPC \bullet Y$$

where C_0 is called "autonomous consumption" and MPC the "marginal propensity to consume". It is a consumption function because it represents desired consumption as a function of household income. Formally, **autonomous consumption** refers to the amount households desire to consume irrespective of their income; you could imagine it as the total amount they would consume if their income were zero. (They would draw down savings, for instance, to finance a portion of their essential needs.) The **marginal propensity to consume** (MPC), meanwhile, means the proportion of any additional income consumers desire to spend. Note that both of these have the word "desire" in their definitions. There is no objective determination of either of them; they are purely subjective and represent the psychological underpinnings of consumption behavior. You can't be sure you are seeing them in the real world, if only because people sometimes make mistakes; they consume more than or less than they intended. Nevertheless, by looking at the behavior of large numbers of people over extended time periods, you may be able to come up with a reasonable estimate of them. (Your estimate will be more reasonable if you allow consumption to assume a more complex pattern than the simple linear model we are working with; if you continue in economics, you may get a chance to do this.) Note also that we use an equal sign rather than an identity (≡) since people may or may not manage to achieve their desired level of consumption.

So let us see how equation 4 looks in a diagram:

Figure 3: Equilibrium National Income in a Model with Consumption and Savings

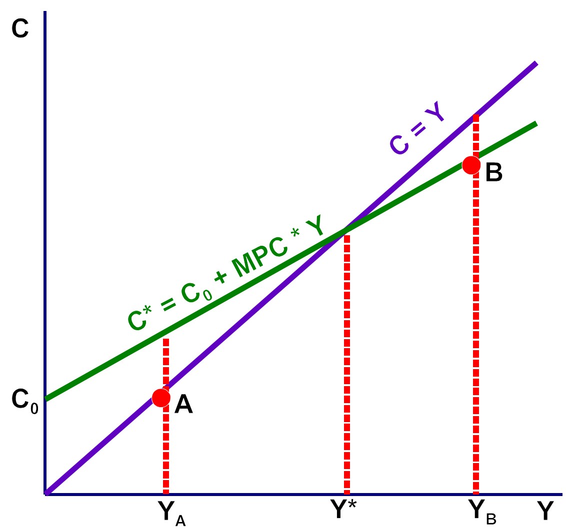

With autonomous consumption at C_0 and the marginal propensity to consume given by the slope of the line $C^* = C_0 + MPC \bullet Y$, equilibrium national income is determined to be at Y*.

The line $C^* = C_0 + MPC \bullet Y$ shows the desired level of consumption for each possible level of income Y (the consumption function). If Y = 0, households will still consume C_0; this is their autonomous consumption.

Then, as Y increases, so does C*, rising at the rate of MPC. Any point along this desired consumption line is a potential equilibrium, in the sense that it reflects what households *want* to do, given their income level. But the amount of consumption households engage in is also the amount of income they will earn, according to the circular flow. Thus $C \equiv Y$ is a *necessary* condition for the economy: the value of the goods they consume has to be equal to the income households will receive. Only the intersection of these two lines yields an outcome that is both *desired* and *feasible*. It represents a situation in which everyone is making the choices they prefer, given the choices made by everyone else (which generate their income). This is why we call it **equilibrium national income.** (Figure 3 also provides us with our first example of a **Keynesian Cross** diagram, which locates an equilibrium where desired spending "crosses" the circular flow constraint. The Keynesian Cross is also sometimes call the income-expenditure model, in honor of the quantities measured on its two axes.)

To see this in more detail, suppose that national income is below its equilibrium at Y_A. The circular flow requires that actual spending be equal to this income, so we are at point A along the line C = Y. But this is below the desired consumption level according to our simple consumption function: consumers are spending less than they want at this income level. So they will try to spend more. As they do this, they will also be raising their aggregate income, as given by the circular flow. This process will stop when income reaches Y*. By the same token, if income is initially above its equilibrium level at Y_B, the economy is at point B. Consumers realize they are spending more than they want, so they cut back, and again the process ends at Y*.

Here we come to an extremely important moment in our understanding of macroeconomics. Economics defines equilibrium as a state in which all parties are acting according to their preferences, given the actions taken by the other parties. It is a mutually consistent set of behaviors. The result is that there is no "inner" tendency for the situation to change once it is in equilibrium; no one wants to do anything differently. Any change must come from an external force. Students just beginning to learn economics should be cautioned against defining equilibrium in terms of supply equaling demand, and here is one of the reasons. Y* in Figure 3 is an equilibrium according to the general definition of equilibrium, but it is *not* based on any sort of supply and demand reasoning. Unless we stick close to a general definition of equilibrium that can apply to contexts like Figure 3, we lose the ability to use this important tool.

A second point is that equilibrium requires not only a state of "no tendency to change", but also an adjustment process that brings it about if it does not already exist. In our discussion of the implications of Y_A and Y_B we sketched just such an adjustment process. In other words, Y* is an equilibrium not only because it represents a mutually consistent set of desired actions, but also because there is a tendency to move toward it if it is not the initial state of affairs. But what exactly is this adjustment? *It is national income that adjusts to bring about the equality of desired and actual consumption.* This is a crucial point: you should stay with it and not read on until you are sure you understand it, since it is the linchpin for the entire Keynesian edifice. Again: *national income adjusts to bring about macroeconomic equilibrium.* This assertion will have to be strongly qualified later on, as we build more complex models of how the economy works, but to qualify it you have to understand it, so make sure you see the idea clearly.

Now get ready for the third implication, which is quite remarkable by any standard. Figure 3 was drawn rather arbitrarily. We did not do any research to determine what either autonomous consumption or the marginal propensity to consume are. (And there are no numbers on either axis either, so the entire diagram is in the realm of speculation.) Y* is the result of how we drew the consumption function, but perhaps it was drawn incorrectly and we should change it. For instance, what if C_0 is too low and should be moved further up the Y-axis? This would shift the entire consumption function upward, and it would then intersect C = Y at a higher level of national income. This is not particularly shocking. But what about changes in MPC? What if, instead of desiring to spend, say, 80% of their income on consumption, households decide they want to spend 90%? This shift is depicted in Figure 4.

Consumption function C^*_1 is our original desired consumption, and it yields the equilibrium national income Y^*_1. When MPC goes up, this means the slope of C* goes up, so we can draw a new consumption function C^*_2. This would result in an increase of equilibrium national income to Y^*_2. Consider what this means: the more of their income households want to spend, and the less of it they want to save, the more income they will have. What a result—the behavior that we are usually told will make us rich, saving, actually makes us poorer (in this model), while spending like there is no tomorrow means that we will have even more to spend. This is such a striking feature of Keynesian models like this one that it has been given its own name, the **paradox of thrift**. Where this paradox applies, it means that the members of an economy can raise their income by reducing their saving and increasing their spending. At first glance, it seems like a case of having your cake and eating it too (or maybe the other way around). How can it be explained?

The answer is that *there is actually no saving taking place in this model!* People may wish to save if their income is high enough (that is what an MPC < 1 means), but they never have a chance. The reason is that their savings are not spent on anything; they just disappear into various household hiding places. Since they don't make their way to firms as revenues for purchases, they can't be returned to households as income—the circular flow. Thus income cannot be high enough to generate savings: the money coming out of the firms and into the households can't be greater than the money going into the firms in the first place. This is the deeper meaning of the paradox of thrift: by desiring to save more, household do not actually save more, they only earn less.

Figure 4: The Effect of an Increase in MPC on Equilibrium National Income

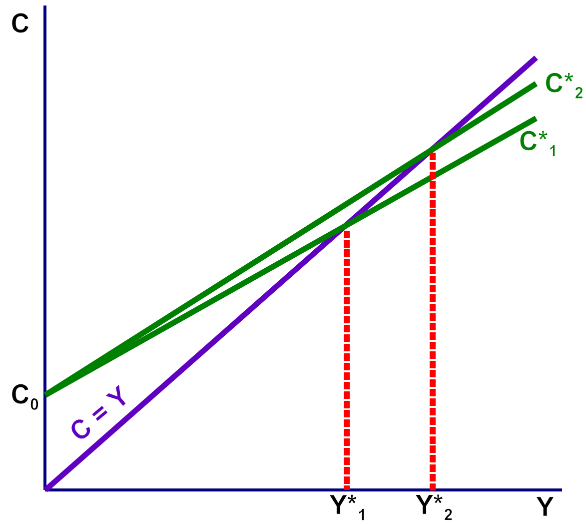

An increase in MPC is represented by the shift from C^*_1 to C^*_2 and results in an increase in equilibrium national income from Y^*_1 to Y^*_2.

Having seen how the model works graphically, let's see how it works algebraically. We have one equation and one identity:

(5a) $C^* = C_0 + MPC \cdot Y$
(5b) $Y \equiv C$

(5a) is the consumption function and solves for desired consumption; (5b) says that national income identically equals the national product, where the product consists only of (actual) consumption. Equilibrium occurs when desired consumption according to the consumption function equals actual consumption according to the circular flow; that is, where $C^* = C$. This means that in equilibrium we can combine these two relationships to get

(6) $Y^* = C_0 + MPC \cdot Y^*$

(We put the asterisk after Y to indicate that it is the value in equilibrium.) Solve for Y^* by consolidating terms and then dividing:

(7) $Y^* - MPC \cdot Y^* = Y^*(1 - MPC) = C_0$

(8) $$Y* = C_0 \frac{1}{1 - MPC}$$

This formula, simple as it is, already contains the two crucial elements found in even the most complex versions of Keynesian income analysis, autonomous spending and the **national income multiplier**. C_0 is autonomous spending; it sets a sort of baseline for the size of the economy. It is multiplied by a fraction whose numerator is 1 and whose denominator is (1 - MPC). Since MPC is the fraction of income people choose to spend, it has to be between 0 and 1. This means that the entire fraction must be greater than one, and it is larger as MPC is larger. Suppose, for instance, that MPC is .5, meaning that people wish to spend half their income (in addition to autonomous spending). Then the fraction has the value of 2. If it were .75, the fraction would equal 4. Since this fraction multiplies autonomous consumption to determine equilibrium national income, we call it the national income multiplier. You should think of it as literally multiplying some initial change in autonomous consumption to arrive at the ultimate effect on equilibrium income. If the multiplier is 2, for instance, and C_0 increases by $1 billion, then equilibrium national income increases by $2 billion. Go back to Figure 4 and see if you can visualize this multiplying effect by comparing the level of C_0 to the level of Y*—and how this effect increases when MPC goes up and the consumption function rotates counterclockwise.

Step 3: Add Investment

It was a bit artificial to introduce savings but give them nowhere to go, so let's also add investment to the mix. By investment, we mean what it means in the National Income and Product Accounts: spending on goods and services for the purpose of increasing production capacity in future periods or, equivalently, additions to the stock of capital. It is important to keep this in mind. Simply saving money and putting it into an interest-bearing account, or loaning it out at a rate of interest, doesn't register as investment; the investment occurs only if and when this money is actually spent.

Now that we are allowing for investment, we have a new expression for the national product:

(9) $$C + I \equiv Y$$

where I is aggregate investment spending. Firms take in revenue from both households (consumption) and other firms (investment), and pay income back to households. In this simple formulation, where do firms get the extra money to invest? They borrow it, either directly from households (by issuing bonds) or via the intermediation of banks, which accept savings in the form of deposits and then make loans.

Just as with consumption, we can speak of desired investment, which may or may not be the same as actual investment. By desired investment, we mean the amount that businesses wish to spend for this purpose, given all the relevant conditions in the economy. The most important factor in businesses' investment decision-making is the range of expectations they have for future profits: this is why they invest in the first place. There are other considerations as well, which we will take up in a later chapter; for now the point to emphasize is that investment, unlike consumption, has very little to do with *current* income. The simplest way to express this is to say that, for the time being, we will treat all desired investment as autonomous:

(10) $$I^* = I_0$$

By doing this, we are treating investment as essentially the same as autonomous consumption—there is no difference (so far) between an extra billion dollars of investment or an extra billion of autonomous consumption. They have exactly the same effect in our model. To see this, consider Figure 5.

Our new desired spending is the sum of I* and C*, where C* is still determined by the consumption function (with C_0 and MPC). The (C* + I*) line is therefore parallel to the old C* relationship, indicated by the dotted line. Equilibrium national income is at the new Y*, greater than what it was in Figure 3. As before, the circular flow requires that the economy be somewhere on the 45º line that equates national income (Y) with national product (actual C + I). Where desired spending is equal to actual spending is the equilibrium. If national income is too low, for instance at Y_A, spending is at point A, below its desired level—but whose spending is out of adjustment? Now that we have two sources of spending, households and firms, we don't know; it could be either. If it is the households that are spending less than they want, this is easy to understand. Perhaps they were following their old spending plans and were taken by surprise by greater-than-expected income. Adjustment takes the form of revising these plans (assuming that MPC remains unchanged). But it is also possible for firms to invest less or more than they intend. The main way this happens is through changes in inventories. In NIPA terms, the inventories carried by firms are counted as part of their investment. At point A, for instance, it may be the case that firms are experiencing a reduction in inventories, reducing their measured investment below the intended I_0. Their response is likely to be to increase production and build their inventories back up. This would take the form of increased spending, so it would be recycled back to households as increased incomes, moving the economy in the direction of Y*. Thus, whether the gap between desired and actual spending occurs on the side of households (consumption) or firms (investment), or some combination of the two, the result is a process that tends to restore equilibrium.

Figure 5: Equilibrium National Income in a Model with Consumption and Investment

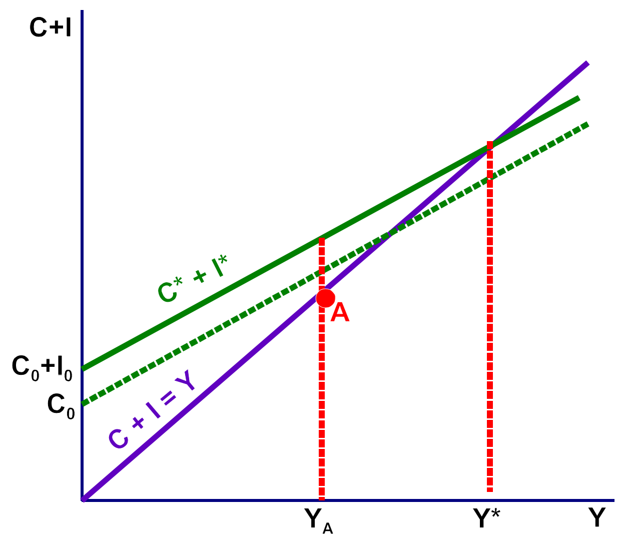

Desired spending C* + I* is consistent with the circular flow constraint C + I = Y at equilibrium national income Y*. If income is too low (Y_A), actual spending at A is below its desired level, so households and/or firms will alter their spending behavior.

An algebraic approach to the model will give us additional insight. In equilibrium, where Y = Y*:

(11) $Y^* = C^* + I^* = C_0 + MPC \cdot Y^* + I_0$

Again we collect terms and divide:

(12) $Y^* - MPC \cdot Y^* = Y^* \cdot (1 - MPC) = C_0 + I_0$

(13) $$Y^* = (C_0 + I_0)\frac{1}{1 - MPC}$$

Equation (13) is the same as equation (8), except that it now adds autonomous investment to autonomous consumption before multiplying by the multiplier. This means we can rewrite this formula as

(14) $Y^* = A \cdot NIM$

where A is the sum of autonomous spending, in this case autonomous consumption plus autonomous spending, and NIM is the national income multiplier. If you know the amount of A and can calculate NIM, you can also compute Y*.

This is a good place to stop and think about what this multiplier means. Suppose that someone comes up with a new invention, and businesses increase their investment to take advantage of it. If this increase is $1 billion, it means that an extra billion dollars of purchases will occur, generating income for new employees, owners of natural resources and other providers of inputs. But these people, having more income, will spend some of it on new consumption, and this in turn provides new employment and income, and new consumption, and so on. Since only some of the additional income is spent, each round of ripple effects becomes smaller, until the whole process is complete and we are at a new equilibrium income for the economy. This new equilibrium will be greater than $1B, and the multiplier tells us just how much greater. The more people increase their consumption as their income goes up, the higher MPC and the higher the multiplier. The point is that the national income multiplier is not just a creature of algebra; it has a clear meaning in real life, one we can see each day.

Meanwhile, to get a different take on how this model functions, consider the question of savings. Since all income is either spent or saved, we can write an identity for total savings S

(15) $S \equiv Y - C$

Since $C^* = C_0 + MPC \cdot Y$,

(16) $S^* = Y - C_0 - MPC \cdot Y = Y\,(1 - MPC) - C_0$

In equilibrium, drawing on equation (13),

(17) $$S^* = Y^* - C^* = (C_0 + I_0)\frac{1}{1 - MPC}(1 - MPC) - C_0$$

$$= C_0 + I_0 - C_0 = I_0 = I^*$$

In other words, when we solve for savings in equilibrium, we discover it is equal to equilibrium investment. Recall from the previous chapter that the classical argument regarding credit markets was that interest rates would adjust so that the amount of money that potential savers would save and lend would be equal to what potential investors would be willing to borrow. The model we are working through in this chapter, however, tells an entirely different story: *it is national income that adjusts so that, at its equilibrium level, savings will equal investment*. In fact, it does not require much imagination to see that, in the Keynesian model, investment plays the active role in the process and savings the passive one. If people get it into their heads to save more, they can decrease their marginal propensity to consume, but, as we have seen, this reduces national income and fails to increase savings at all—the paradox of thrift. On the other hand, if investors increase their desired autonomous investment, this will shift the desired spending line upward and equilibrium income outward. The result is more income out of which to save, so that, in the end, savings still equal investment, but at a higher level. This is such an important result, it is worth taking a few more minutes to let it sink in. To summarize:

- National income adjusts so that, in equilibrium, savings will equal investment.
- An increase in desired savings, for instance through a decrease in MPC, will defeat itself: savings will not rise, but national income will fall.
- An increase in autonomous investment will increase national income, thereby generating the savings that can be loaned out to finance it.

Of course, results as extreme as this would not occur in a more realistic portrayal, but that is exactly the reason for studying this highly simplified model in such detail: it reveals the Keynesian vision in its purest, most dramatic form. We will add many qualifications down the road, but for now it is important to see its core logic.

Step 4: Add Taxes and Government Spending

The key to understanding how government fits into the model is to recognize that its two directly economic functions, taxing and spending, are independent of one another. Taxes are determined by one set of policies, spending by another. Sometimes the two are related, but not always, and never in a strict manner. And if, at some point in the future, government were to link taxing and spending more tightly, this could be added as a further wrinkle. For now, however, we will treat money coming in and money going out as two separate phenomena.

Taxes take many forms. There are income taxes, excise taxes, value-added taxes, property taxes, and on and on. If this were a text in public finance, we would cover them all in detail. That would be a distraction here, so we will pick just one and represent it in the simplest possible way. Let's assume that all government revenue is collected in the form of an income tax, and that there is a single rate that applies to all individuals and income levels with no deductions or loopholes. Call this rate t; it will be a number between 0 and 1, but usually at the lower end of this spectrum, perhaps 15-20%. As we have defined it, then,

(18) $T \equiv t \bullet Y$

This means we can also define after-tax income as Y(1-t).

On the spending side it is important to remember that, in the world of the national income and product accounts, government spending consists only of actual purchases of goods and services (including labor) by the public sector; it does not include transfers, like benefit payments. In accounting terms, transfers should be deducted from tax payments to determine net taxes; in this chapter we will assume that there aren't any. So what determines the *desired* level of government spending? Here we are referring to government decision-makers, such as political leaders and higher-level officials, who set the spending policies for the various levels and branches of government. The short answer for now is, we don't know. What we can do, then, is treat government spending the way we treated investment, as a purely autonomous expenditure, unrelated to current income and otherwise unexplained:

(19) $G^* \equiv G_0$

With this in mind, we can see that adding government does two things to our model, one familiar, the other new. The familiar part is represented by G_0, which is treated just like C_0 and I_0; it shifts the desired spending line up and equilibrium income to the right. The tax rate is a bit more complicated, however. It enters through the use of after-tax income in the consumption function, since consumers presumably base their consumption decisions on their after-tax rather than pre-tax income:

(20) $C^* = C_0 + (1 - MPC)\ Y\ (1-t)$

Figure 6 shows what incorporating these two effects does to the determination of equilibrium income.

Figure 6: Equilibrium National Income in a Model with Consumption, Investment, Taxes and Government Spending

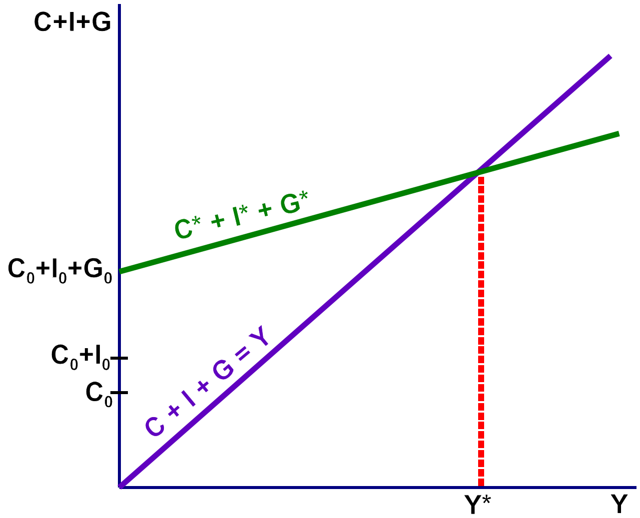

With autonomous spending set at $C_0+I_0+G_0$, and the national income multiplier determined by the marginal propensity to consume and the tax rate, equilibrium national income is at Y*.

The addition of government to the mix has the potential to both increase and decrease equilibrium income. On the one hand, government spending is spending, so it has the same effect as autonomous consumption (C_0) or investment (I_0). This shifts the desired spending curve, C*+I*+G*, upward and equilibrium income outward. On the other, by reducing the income available to households, taxation reduces the multiplier, which makes the slope of the desired spending curve flatter. This tends to produce an intersection with the 45º line further to the left. How the two effects net out depends on the relative size of spending versus taxes, as we will see in much greater detail in the next chapter.

The algebra is interesting mainly for the way it changes the calculation of the national income multiplier. In equilibrium we have

(21) $Y^* = C^* + I^* + G^* = C_0 + I_0 + G_0 + MPC \cdot Y^* (1-t)$

Collecting terms and solving for Y* gives

(22) $Y^* - MPC \cdot Y^* (1-t) = C_0 + I_0 + G_0$

(23) $Y^* [1 - MPC (1-t)] = C_0 + I_0 + G_0$

(24) $$Y^* = (C_0 + I_0 + G_0) \frac{1}{1 - MPC(1-t)}$$

or

(25) $Y^* = A \cdot NIM$, where $A = C_0 + I_0 + G_0$ and $NIM = \dfrac{1}{1 - MPC(1-t)}$

The big difference between (14) and (25) is that the national income multiplier has changed. The effect of introducing a tax rate like t is that MPC (1-t) is smaller than MPC, so the denominator of the fraction is bigger, and the fraction as a whole (the multiplier) is smaller. For instance, earlier we calculated the multiplier when MPC was .75; it was 4. Now, with t, it will be less. Suppose t = .2, meaning that there is a constant tax rate of 20% of income. Then

(26) $MPC\ (1\text{-}t) = .75\ (.8) = .6$

and

(27) $NIM = \frac{1}{1-.6} = 2.5$

Thinking back to the real-world meaning of the multiplier, what is going on is that less of the additional money entering the economy from an increase in autonomous spending is itself being spent. Recall the story about the invention that leads to an extra billion dollars of investment. Once government is added to the picture, people whose income goes up because of the surge in investment will spend less of it than before, because some of their earnings are now taken by taxes. These tax revenues are effectively lost to the recirculation process described earlier; they "disappear" into government receipts. If more tax revenues meant more government spending, this money would return to the economy again, but we are assuming that taxes and government spending are determined separately and have no effect on each other. We will have more to say about this in the next chapter, but for now it is as reasonable a simplifying assumption as any other.

Step 5: add imports and exports

The final step in constructing our simple Keynesian model is to recognize that our "cartoon" economy is part of a larger global economy, linked by international trade. As we did with government spending and tax revenues, we will assume that imports and exports are completely independent of one another, the result of entirely different factors. This seems even more plausible here, since, by purchasing an import, we have little or no effect on the desire of foreigners to purchase our exports, and vice versa.

Let's start with exports. The simplest representation is that the value of our goods that foreigners want to buy is the result of events taking place in their countries which are beyond our ability to influence. If so, they are autonomous in the sense we have been using this word, and $X^* \equiv X_0$.

Imports are a bit more complicated, even when we try to simplify them as much as possible. Perhaps the simplest approach is to assume that a constant proportion m of all purchases made by households, businesses and government, as well as goods produced for export, are imported from abroad. Of course, this is unlikely to be true; in particular, we should expect that government spending would have little import content, since the biggest single item in their budget is salary—but we are trying to keep things simple. So, based on this assumption, we can express imports this way:

(28) $M \equiv m\ (C + I + G + X)$

Note that exports are measured in the national income accounts as the value of goods shipped, not the value added, so it is entirely possible (in fact it is certain) that a portion of our exports represent inputs that were imported to make them. Meanwhile, since only domestically produced goods generate income for households within the country, the circular flow identity is

(29) $Y \equiv (1 - m)\ (C + I + G + X)$

We can see that international trade will have three effects on equilibrium income. First, it increases autonomous spending by the amount of exports, since we are assuming that exports are autonomous. Second, it *decreases* autonomous spending to the extent that autonomous consumption, investment, government spending and exports are composed of imports. We can't say without more information which of these two effects is stronger. Third, it decreases the national income multiplier because some of the increased income resulting from autonomous spending will leak out through imports, just as it does through taxes and savings. For example, if a new business investment or government program leads to more workers being hired, they will spend some of their new income on new domestically-produced goods, but not as much as in earlier versions of the model, since a portion of their spending will go to imports. Combining all these effects, we might see a picture like Figure 7.

Figure 7 is drawn to reflect the possibility that autonomous spending goes up with trade, because the boost derived from exports (X_0) outweighs the effect of reducing this and other autonomous spending by the import share (m). It could have been otherwise, however, in which case the green desired spending line might begin at a point below $C_0 + I_0 + G_0$. We have also drawn the slope of the green line much flatter, reflecting the impact of spending impetus lost to imports. To see how this works algebraically, let's solve for equilibrium Y* as before. (The algebraic expressions keep getting longer, but the method remains the same.)

(30) $$Y^* = (1-m)(C^* + I^* + G^* + X^*) = (1-m)[C_0 + I_0 + G_0 + X_0 + MPC \cdot Y^*(1-t)]$$

(31) $$\frac{Y^*}{1-m} = C_0 + I_0 + G_0 + X_0 + MPC \cdot Y^*(1-t)$$

Figure 7: Equilibrium National Income with Consumption, Investment, Taxes, Government Spending, Imports and Exports

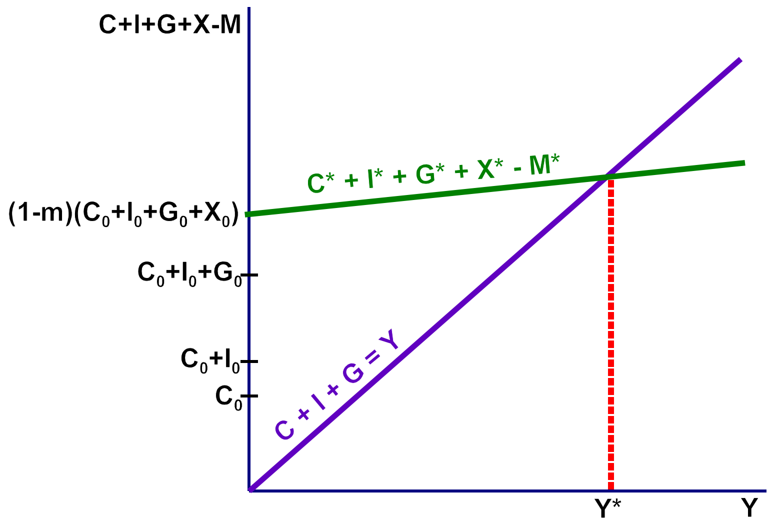

With autonomous exports plus the domestic share of autonomous domestic spending, plus the national income multiplier (based on MPC, the tax rate, and the import rate), equilibrium national income occurs at Y*.

(32) $$Y^* \left[\frac{1}{1-m} - MPC\,(1-t)\right] = C_0 + I_0 + G_0 + X_0$$

(33) $$Y^* = (C_0 + I_0 + G_0 + X_0)\,\frac{1}{\frac{1}{1-m} - MPC\,(1-t)}$$

Once again we have something that looks like autonomous spending times a multiplier, but with a somewhat more intricate formula for the multiplier. The only change is that, in the denominator, 1 has been replaced by $\frac{1}{1-m}$. Since m is between 0 and 1, this change will make the denominator larger and the whole multiplier smaller, which is what we expected based on the logic of diminishing ripple effects from spending: some of the money that might otherwise have generated further rounds of demand and production are escaping to other locations in the global economy. The upshot is that the C*+I*+G*+X*-M* line has a flatter slope in Figure 7 than the equivalent in Figure 6.

And this completes the entire stripped-down Keynesian model of equilibrium national income. To sum up where we have come to, the model illustrates several aspects of macroeconomic analysis:

- National income is *always* composed of consumption, investment, government spending and net exports. This is an identity: it is true by definition and in every moment.

- Equilibrium national income results from *desired* consumption, investment, government spending and net exports. It is an equilibrium because all the participants in the economy are making the choices they prefer to make, and their choices are mutually consistent—meaning that they are basing their decisions on the income level that results from everyone making these decisions. There is also a process that tends to bring the economy to its equilibrium income level if it is out of equilibrium.

- The Keynesian view is that national income adjusts to bring about equality between the supply and demand for credit, or savings and investment.

- Measures to increase savings directly are vulnerable to the paradox of thrift, according to which reductions in the marginal propensity to consume, rather than yielding more saving, only result in lower national income. Savings will increase, however, if investment demand increases, through an expansion of national income.

- Autonomous increases or decreases in spending are magnified by the national income multiplier. The multiplier is greater if the marginal propensity to consume is greater, but it is reduced by taxes and the share of spending going to imports.

Several of these generalizations will have to be amended as we introduce more complexity to our analysis, but all of them will survive to some extent. They represent the core of the new thinking Keynes brought to economics 70 years ago.

A word on macroeconomic models

Equation 30 gives us a simple macroeconomic model, stripped down enough to write in a single line of algebra, but sufficiently detailed to be fitted to an actual economy without nonsensical results. It is not remotely sophisticated enough to use in real economic work, but it has everything it needs to illustrate what macroeconomic models are and how they work.

The most basic distinction to make in a model is between the **variables** and the **parameters**. Both are represented by algebraic symbols, but they mean very different things. A variable is an economic outcome of interest. It might be observable in advance, like exports in the way they have been depicted above, or it might be an outcome we are trying to predict by using the model, such as equilibrium income, tax revenues or imports. Variables, like their name suggests, will vary from one application of the model to the next—for instance, from one year to the next, based on the information specific to that year. A parameter is a value used in calculating the model that is believed to remain constant over multiple applications; examples in Equation 30 include MPC and m. The tax rate t might be a parameter as well, if it is not changing every year. Practically speaking, we first try to figure out what the parameters are, and then we plug in specific numbers for some variables in order to compute the rest.

Where do the parameters come from? In textbook examples we are often told what to use for them, but in real life—when economists use models to make forecasts—they are deduced by **fitting** the model to the available data. Understanding what “fitting” means is crucial; the usefulness of a model depends largely on how well it can be fitted. Suppose we have a model like the one represented in Equation 30. In itself it doesn’t tell us very much; unless we know what MPC, m and t are, for instance, we can’t say what the multiplier is. Unfortunately, it’s not as though you could go to some government source, like the US Bureau of Economic Analysis or the World Bank, and look them up. What you can do is get many years’ worth of observable economic data for things like consumption, investment, imports, national income and so on. From this you can use statistical techniques to estimate the values of the parameters that do the best job of making a model like Equation 30 fit the data. For instance, perhaps an MPC of .8 enables you to use the Keynesian consumption relation in Equation 20 to account for both total consumption and national income, given the level of taxes. This value of MPC is not selected simply because someone tells you to use it; it is “right” because, if you select it, the model works in the sense that it fits the actual data. If you selected some other MPC, .6 for example, it might not be possible to use Equation 20 at all, because either the level of total consumption or the level of national income would have to be wrong; you couldn’t use the actual values for both in the same equation.

What this means is that the first test of a model is whether it can be fitted. If it can, it means that parameter values can be selected that permit the model to correctly compute the variables that are observed in the real world. Of course, the model is unlikely to produce an exact fit: there may well be error. But if the error is small enough, meaning not enough to seriously impair the usefulness of the model, we can live with it.

Once the model has been fitted, the next step is to use it for forecasting. This means plugging one or more new variables into it, while continuing to use the same (fitted) parameters. For instance, suppose the you are interested in how a change in government spending will affect equilibrium national income according to the model presented in Equation 30. You would leave all the parameters the same, so that MPC, m and t would remain at whatever their fitted values required. Some of the variables would also remain the same, like exports. (In this model we assume that the level of exports depends entirely on what people in other countries want to buy, and this is presumably not affected by changes in the spending of our own government.) But government spending would indeed change, so we would use some new value for it. By computing the formula using existing parameters but one or more new variables, we can calculate a new equilibrium income. That’s how one uses a model to do forecasting.

Which brings us to the second test of a model: how well does it forecast? If parameters are fitted to be consistent with past or current data, can the model use them to reliably predict future outcomes? To find out, try fitting the model and see what happens. If the model proves that it can do this job well, it’s worth using.

Of course, matters are not quite this simple. (They never are.) A big problem is that, even though parameters, to be useful, have to be held constant (this is what makes them parameters), in real life they change continuously. MPC, the tendency of consumers to spend out of additional income, is not really constant over time; it goes up and down, even if only slightly. This is one reason models generate error. You can’t allow parameters to change with each run of the model, since calculation would be impossible: you need to be given some algebraic values in order to compute the others, where “given” in this context means “holding fitted parameters as constant”. You can update them with new information, but for each new forecast you have to keep them at their most recent fitted level.

And what about the simple Keynesian model contained in Equation 30? Let’s see how well it does with real-world data. Here we will focus on just the consumption function piece of it:

(20) $C^* = C_0 + (1 - MPC)\ Y\ (1-t)$

Table 1 gives us information on disposable income and total consumption in the US for the years 2010-2012.

Table 1: Income, Consumption and Net Taxes, United States (in billions of current dollars)

Year	Y-T	Δ(Y-T)	C	ΔC	MPC	C_0
2010	11,127		10,216			
2011	11,549	422	10,729	513	1.22	-3,312
2012	11,931	382	?			

Y-T is disposable (after-tax) income. Changes in disposable income and consumption (C) are calculated from the previous year. The marginal propensity to consume (MPC) and level of autonomous consumption (C_0) are calculated using the formulas below.

The consumption function has two parameters, MPC and C_0. To fit it, we need to calculate the parameters that fit the data. Suppose we have only the two years 2010 and 2011. To get the MPC we use the formula implied by its definition:

(34) MPC = ΔC ÷ Δ(Y-T) = 513/422 = 1.22

This is odd! It is telling us that for any change in their after-tax income, consumers will increase their spending by 22% *more* than that amount! The numbers aren't lying: from 2010 to 2011 consumers *did* increase their spending at a faster rate than their available income grew. Of course, we are looking at only one pair of years, and it is unlikely that we would find such eagerness to spend across a wider time frame. We can interpret this result by noting that consumers were emerging from a devastating economic downturn during which their spending was depressed, so they responded to the modest recovery of 2010-11 by returning to their previous spending habits. In other words, not only disposable income changed during this period; psychology did too. The consumption function is not sophisticated enough to incorporate the sort of variables that might be related to this psychological change.

Since the MPC is greater than one, autonomous consumption, according to Equation 20, has to be negative for 2011. Its formula, which comes from rearranging terms, is

(35) C_0 = C – MPC (Y – T) = 10,729 – 1.22 (11,549) = -$3,312B

Negative consumption is absurd; what this result is telling us is that our fitted consumption function can't possibly work at incomes much below what was experienced during 2010-11. To be precise, it calls into question the *linearity* of the function—its representation as a straight line rather than one whose slope changes over different values of (Y-T). This is another way in which Equation 20 is unsophisticated.

So now we can see that our simple consumption function passes the first test, sort of: it can be fitted, although the parameter values we get when we do this, the calculated values for MPC and C_0 are implausible over a larger historical range.

The second test is whether the model has predictive value. Let's use it to predict the level of consumption, based on disposable income data for 2012. Prediction, we should remember, means keeping the parameters the same, but inserting new data for the variables. In the case of Equation 20 there is only one variable for which we need data, disposable income. Its value is given in Table 1. Simply inserting the numbers into Equation 20 gives us our prediction for C:

(36) C* = C_0 + MPC (Y - T) = -3312 + 1.22 (11,931) = $11,193B

How did we do? Well, the actual level of consumption in 2012 was $11,121B. We were off by $72B. Is that a lot? It's almost 20% above the actual increase in consumption recorded over that time period. All in all, it's what we would expect from a primitive model: it goes in the right direction, but its precision leaves something to be desired. Graphically, our model-fitting exercise looks like this:

Figure 8: US Consumption, 2010-12, Fitted and Actual

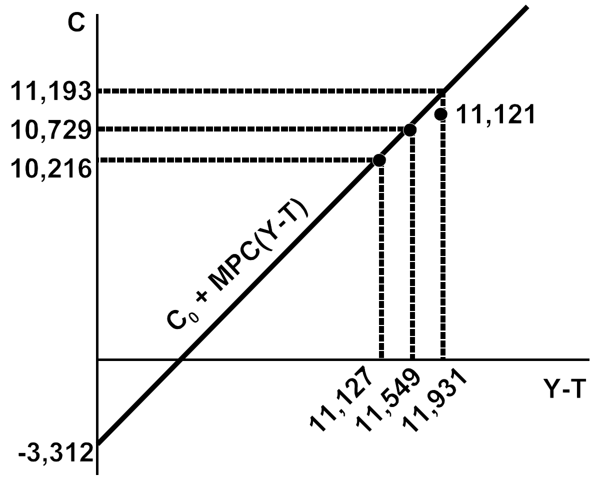

The simple Keynesian consumption function is fitted to the first two points, 2010 and 2011, but it overestimates the actual consumption level in 2012.

The first two points represent actual disposable income and consumption for 2010 and 2011. To fit a simple linear consumption function to it, we constructed a straight line whose slope is 1.22 (MPC) and vertical intercept is -3,312 (C_0). Prediction means extrapolating this line to estimate a level of consumption corresponding to the level of disposable income for 2012. The line proves to be somewhat above the true level of consumption in 2012.

The lesson we can learn from this is that it is not enough for a model to make intellectual sense: it also has to be capable of being fitted to existing data and do a reasonable job of predicting new data. A large portion of debate in the sciences takes the form of dueling models: whose model fits and predicts better, yours or mine? Researchers are constantly tweaking their models, trying to show that theirs beats the competition. If nonspecialists—most of us—want to follow the game, we need to pay attention to the models and the empirical hoops they are made to jump through. As we will see, much of the controversy in macroeconomics can be interpreted as a contest between models.

As for our simple Keynesian model, we can say this: while it passes the first test—it is capable of being fitted—it doesn't do very well on the second. No professional forecaster would choose to use it. But the simplicity that makes it unsuitable for high-level work also makes it useful for conveying the core logic of Keynesian macroeconomics. We will let the matter rest now, then take it up again in Chapter 15.

Equilibrium national income and financial balances

Except for the rather sketchy depiction of savings and investment, we haven't considered the relationship between our model of national income and the financial position of households, businesses or government. This is a major gap: if equilibrium income, as we have defined it, leads to a rapid accumulation of debt it won't be an equilibrium for very long. Fortunately, we have tools to study this problem that fit nicely into the framework we've already developed.

Recall that there is an aggregate relationship between net exports and the net saving or borrowing of the entire economy:

(37) $C + I + G \equiv E$, where E is **aggregate expenditure**

(38) $C + I + G + NX \equiv Y$

So therefore

(39) $NX \equiv Y - E$

In other words, the sum of all domestic budgets of every sort—every household, business and government budget—is captured in the expression Y-E, and it is equal to the country's net exports. If exports exceed imports, then there is more income being earned than households, firms and government are spending; if imports exceed exports, the opposite must be true. Note that this relationship is based on identities, not equilibrium behavior. It is not that people's saving and spending are adapting somehow to trade or that trade is somehow responding to spending and saving; rather, the trade balance is exactly the same as the balance between production and consumption in the economy as a whole. These are two ways of describing the same thing.

What identity (39) tells us is that the *sum* of the various budgets equal the trade balance, but it doesn't say which budgets play which role. Suppose, for instance, that an economy has a trade deficit, as the United States does today. According to (39), if we add up the borrowing and lending of households, firms and government, the total will be negative. In fact, all three sectors of the economy have accumulated debt: households borrowed extensively via mortgages and other credit instruments, firms borrowed to finance investment (and build up their cash position), and governments ran fiscal deficits. If the trade deficit narrows, this means there will be some combination of more domestic production (which generates domestic income) and fewer purchases of imports (which means, in itself, less spending), so budgets will be closer to breaking even. Again, however, we don't know in advance *which* budgets will improve.

This connection between trade balances and aggregate budgets comes into sharper focus if we consider the relationship between trade and national income. Figure 8 depicts a typical situation in which autonomous exports exceed the autonomous component of imports, but additional imports are generated as income increases, since households in particular spend a portion of their earnings on imported goods. The slope of the line, then, depends on the import share m: the bigger m is, the steeper the slope will be.

In interpreting Figure 9, the most important point to remember is that Y_A is *not* an equilibrium. It does not reflect anyone's preferences, nor is there a process that leads the economy to it; it is simply the level of national income at which the trade balance is exactly zero. Equilibrium income, determined through a process like that in Figure 7, could be either less than or more than Y_A. If it is less, there is a tendency for this country to run trade surpluses; if it is more there is a tendency to run trade deficits. These tendencies are based on a large number of factors we haven't yet considered—but will in a future chapter. For the present, the point to keep in mind is that greater national income, holding all else constant, leads to either a smaller trade surplus or a larger trade deficit, and the trade balance in turn is equivalent to the overall balance of borrowing and lending in the economy as a whole.

Figure 9: Net Exports as a Function of National Income

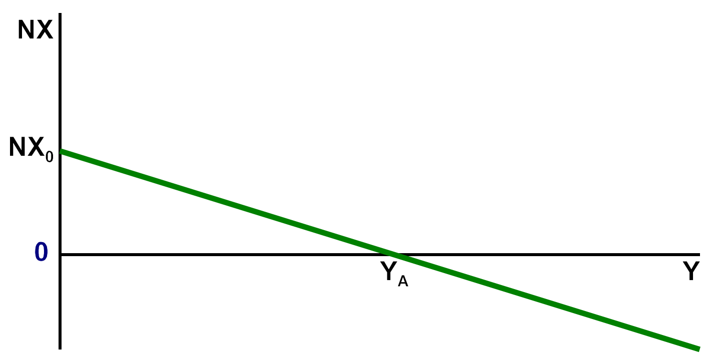

Net exports decline with national income. At any national income less than Y_A net exports will be positive; if national income is more than Y_A there will be a trade deficit.

Thus far we have looked at the net borrowing or lending of the economy as a whole, but of course beneath this aggregate there are deficits and surpluses appearing in different industries, regions, social classes, etc. The Federal Reserve's Flow of Funds Accounts, for instance, records all the lending that occurs in the economy, even the debts incurred by individuals and institutions that might be lending at the same time. (It is not uncommon, for instance, for a household to borrow by taking out a mortgage, while also saving by accumulating balances in bank accounts and pension funds.) What can we say about these gross flows of credit, compared to the net flows after debits and credits cancel out, in relation to national income?

Figure 10 provides one possible scenario. It assumes there are two types of households in the economy, lower-income (L) and upper-income (U). Each has its desired consumption schedule, C^*_L and C^*_U respectively, made up of autonomous consumption (C_{0-L} and C_{0-U}) and addition consumption based on the group's MPC. The total income in the economy is determined in the same way as in Figure 7, but it is not shown here. (Equilibrium national income would be well to the right of the levels depicted in Figure 9, being the sum of Y_L and Y_U.)

Suppose the lower income group receives income Y_L. The height of its consumption function C^*_L gives the level of consumption it will choose, which exceeds its income by $-S_L$, as measured by the distance from the 45º line at which $C_L = Y_L$. By -S we refer to dissaving—drawing down previous savings and/or borrowing. The upper income group, meanwhile, might receive income Y_U, for which its consumption function indicates net savings S_U. It is possible, in fact, that this state of affairs could be financed by the L group borrowing from the U group to close its financing gap.

For our purposes, what is most relevant about this simple model is its implicit relationship between income, saving and dissaving. If L's income goes up, its need to borrow or dissave will go down; similarly, if U's income goes down, it will save less or even pass the break-even point after which it begins to dissave. Two broad generalities emerge: (1) A general increase in income that affects both groups will reduce the total amount of dissaving and increase the total amount of saving. (2) Holding combined income constant, a redistribution from U to L—greater income equality—will tend to reduce both saving and dissaving. It might have little effect on the net savings of both groups combined, but it will reduce the amount of borrowing and lending between them. Greater equality, in other words, is normally associated with fewer and smaller pockets of negative financial balances in the economy.

Figure 10: Saving and Dissaving Across Two Types of Households

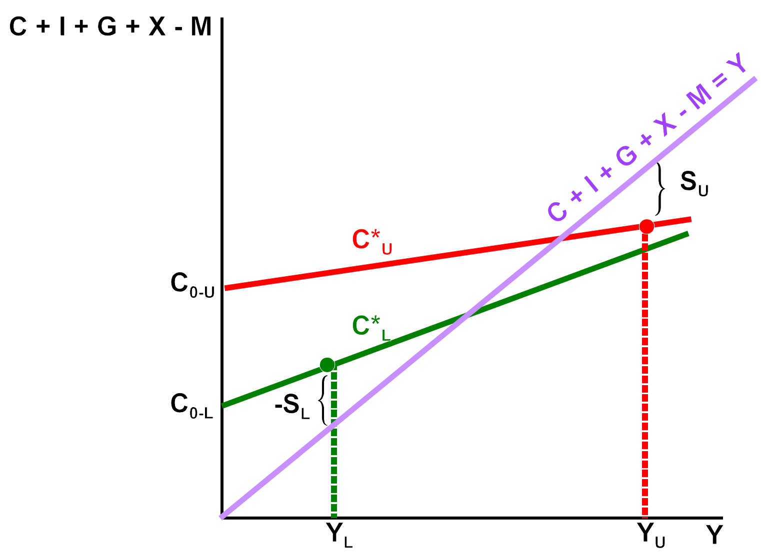

There are two income groups, lower (L) and upper (U). When L receives income Y_L its desired consumption exceeds its income by $-S_L$ (dissaving). When U receives Y_U its income exceeds its desired consumption by S_U (saving).

One final word about the two simple pictures of financial surpluses and deficits captured in Figures 9 and 10: they pull in opposite directions. The first story about net exports tells us that at higher levels of national income, the aggregate budgets of households, firms and government will become negative or at least less positive. The second, about two income groups, says that more income for both groups will reduce the total amount of borrowing. How can this be? The answer is that the first story is necessarily correct for the entire economy, since it is based on the national income and product identities. The second *could* be true for some portion of the economy, but it is incomplete, since it doesn't tell us where the incomes for these two groups came from or how the consumption decisions they make will affect the incomes of other participants in the economy. So: overall, in an economy whose international trade is determined in roughly the way outlined above, it must be true that aggregate savings fall as the economy expands, but there are likely to be eddies in the larger stream, sectors in which savings rise, offset by declines elsewhere.

The main points

1. This chapter presents the core ideas of Keynes' approach to equilibrium national income by building up, step-by-step, a simple algebraic and diagrammatic model.

2. The starting point is a consumption-only economy. This model is indeterminate: it does not give us enough information to ascertain how much or little income there will be.

3. Next we add savings. This implies a consumption function, which tells us how much people wish to consume out of their income. Combining a consumption function with the circular flow identity establishes an equilibrium level of national income. It is an equilibrium because it embodies the desired choices of each individual in the economy, it is feasible, and there is an adjustment process that leads to it if the economy falls out of equilibrium for a time. It presents the simplest version of the general Keynesian formulation $Y^* = A * NIM$, where Y^* is equilibrium income, A is autonomous spending, and NIM is the national income multiplier. It also illustrates the paradox of thrift, according to which an increased desire to save does not result in greater savings, but only a lower level of income—since income is equal to what is spent.

4. By adding investment we make it possible to have savings as well, since now money not spent in consumption can be loaned to firms to make investments, which also add to the size of the economy. This allows us to see how, in stripped-down Keynesian formulation, it national income, not interest rates (or other prices) that adjust in order to achieve equality between desired savings and desired (credit-financed) investment. It is clear from the characteristics of this model that it is the level of desired investment that determines the amount of savings in the economy and not the other way around.

5. The treatment of government's role in the economy rests on the assumption that taxation and spending are entirely independent choices. Taxes lower national income by reducing the national income multiplier: for any given infusion of new spending in the economy, less is recycled into additional consumption because some of it is diverted to tax revenues. Government spending increases national income by augmenting desired spending overall. (Government spending helps expand the economy in exactly the same way that any other form of spending does.) At this level of generality we don't know which effect is greater.

6. The last step involves adding international trade. Exports are taken to be autonomous; they depend on choices made in other countries and do not vary with domestic income. Imports, on the other hand, are represented as a fixed percentage of total spending: the more income and spending there are in the economy, the more imports it will consume. Exports increase equilibrium national income, while imports reduce it by reducing the national income multiplier.

7. Even though it is too simple to use for professional purposes, the algebraic model developed in this chapter has the essential features of any macroeconomic model. It has variable and parameters. The parameters can be fitted by using real economic data, and this allows the model to forecast future economic variables based on assumptions fed into the model. The model passes the first empirical test—its parameters can be fitted—but performs poorly on the second, forecasting future outcomes in the economy with acceptable precision.

8. The sum of all budgets in an economy is identically equal to the trade balance (or, more precisely, the current account). The trade balance in turn depends on the level of national income. Since exports do not depend on domestic income but imports increase with it, as the level of income rises the trade balance falls. There is a particular level of national income at which trade will be exactly in balance, but this is not an equilibrium—it is a coincidence.

9. Borrowing and lending between different groups within the economy has some relation to national income: as income rises, and with constant marginal propensities to consume, we expect less borrowing. At the same time, greater equality of income distribution should be associated with less borrowing and lending between income groups.

Terms to define
aggregate expenditures
autonomous consumption
equilibrium national income
fitting a model
Keynesian Cross
marginal propensity to consume
national income multiplier
paradox of thrift
variables vs parameters

1. Think about your own consumption decisions. How well does Equation 20 represent *you*? Do you have an MPC? If you were to write your own consumption equation, what might it look like?

2. How well does Figure 9 describe your own saving or dissaving choices? Are they a function of your income in the way that this Figure represents? If not, how are they different? Is there a relationship to income that you could graph?