Chapter 4
Inflation and Interest Rates in the Consumption-Savings Framework

The lifetime budget constraint (LBC) from the two-period consumption-savings model is a useful vehicle for introducing and analyzing the important macroeconomic relationships between inflation, nominal interest rates, real interest rates, savings, and debt. Before doing so, we present definitions of these terms and a basic relationship among them.

The Fisher Equation

Inflation is a general rise in an economy's price level over time. Formally, an economy's rate of inflation is defined as the percentage increase in the price level from one period of time to another period of time. In any period t, the inflation rate relative to period $t-1$ is defined as

$$\pi_t = \frac{P_t - P_{t-1}}{P_{t-1}},$$

where π denotes the inflation rate.[32] As a matter of terminology, a **deflation** (negative inflation) occurs when $\pi < 0$, and a **disinflation** occurs when π decreases over time (but is still positive at every point in time). For example, if in four consecutive years, inflation was 20%, 15%, 10%, and 5%, we say that disinflation is occurring – even though the price level increased in each of the four years.

In our consideration of the consumption-savings model, we defined the nominal interest rate as the return on each dollar kept in a bank account from one period to the next. For example, if your savings account (in which you keep dollars) pays you \$3 per year for every \$100 you have on balance, the nominal interest rate on your savings account is three percent.

Because of inflation, however, a dollar right now is not the same thing as a dollar one year from now because a dollar one year from now will buy you less (generally) than a dollar right now. That is, the purchasing power of a dollar changes over time due to inflation. Because it is goods (i.e., consumption) that individuals ultimately care about and not the dollars in their pockets or bank accounts, it is extremely useful to define another kind of interest rate, the **real interest rate. A real interest rate is a return that is measured in terms of goods rather than in terms of dollars.** Understanding the

[32] Not to be confused with profits, which is what π often represents in microeconomics. The usage is almost always clear from the context.

difference between a nominal interest rate and a real interest rate is important. An example will help illustrate the issue.

Example:
Consider an economy in which there is only one good – macroeconomics textbooks, say. In the year 2012, the price of a textbook is $100. Wishing to purchase 5 textbooks (because macroeconomics texts are so much fun to read), but having no money with which to buy them, you borrow $500 from a bank. The terms of the loan contract are that you must pay back the principal plus 10% interest in one year – in other words, you must pay back $550 in one year. After one year has passed, you repay the bank $550. If there has been zero inflation during the intervening one year, then the purchasing power of that $550 is 5.5 textbooks, because the price of one textbook is still $100. Rather than thinking about the loan and repayment in terms of dollars, however, we can think about it in terms of real goods (textbooks). In 2012, you borrowed 5 textbooks (what $500 in 2012 could be used to purchase) and in 2013, you paid back 5.5 textbooks (what $550 in 2013 could be used to purchase). Thus, in terms of textbooks, you paid back 10% more than you borrowed.

However, consider the situation if there **had** been inflation during the course of the intervening year. Say in the year 2013 that the price of a textbook had risen to $110, meaning that there had been 10% inflation during the year. In this case, the $550 repayment can be used to purchase only 5 textbooks, rather than 5.5 textbooks. So we can think about this case as if you had borrowed 5 textbooks and repaid 5 textbooks – that is, you did not pay back any additional textbooks, *even though* you repaid more dollars than you had borrowed.

In the zero-inflation case in the above example, the nominal interest rate is 10% and the real interest rate is 10%. In the 10%-inflation case, however, the nominal interest rate was still 10% but the real interest rate (the extra textbooks you had to pay back) was zero percent. This relationship between the nominal interest rate, the real interest rate, and the inflation rate is captured by the **Fisher equation,**

$$r_t = i_t - \pi_t \tag{15}$$

where r is the real interest rate, i is the nominal interest rate, and π is the inflation rate. Although almost all interest rates in economic transactions are specified in nominal terms, we will see that it is actually the real interest rate that determines much of macroeconomic activity.

Actually, however, the Fisher equation as stated in expression (15) is a bit of a simplification. The **exact Fisher equation** is

$$(1+i_t) = (1+r_t)(1+\pi_t)\,, \tag{16}$$

the details of which we will not describe here. This more accurate form of the Fisher equation turns out to be more convenient than its simplification in thinking about our two-period consumption-savings model. Before we analyze the topics of inflation, nominal interest rates, and real interest rates in the consumption-savings model, let's quickly see why expression (15) is in fact an approximation of expression (16). Multiplying out the terms on the right-hand-side of expression (16), we get

$$1+i_t = 1+\pi_t + r_t + r_t\pi_t . \tag{17}$$

If both r and π are small, which they usually are in developed economies (e.g., the U.S., Europe, Japan, etc.), then the term $r\pi$ is very close to zero. For example, if $r = 0.02$ and $\pi = 0.02$, then $r\pi = 0.0004$, which is essentially zero. So we may as well ignore this term. Dropping this term and then canceling the ones on both sides of expression (17) immediately yields the "casual" Fisher Equation of expression (15). The simplified Fisher equation of (15) is useful for quick analysis, but for our consumption-savings model it will almost always be more useful to think in terms of the exact Fisher equation (16).

For the two-period analysis below, the only economically meaningful inflation rate is that occurs between period 1 and period 2. According to our definition of inflation above, the inflation rate between period 1 and period 2 is

$$\pi_2 = \frac{P_2 - P_1}{P_1} . \tag{18}$$

So π_2 measures the percentage change in the price level (here, the nominal price of the consumption basket) between period 1 and period 2. For use below, it is helpful to re-arrange expression (18). First, separate the two terms on the right-hand side to get

$$\pi_2 = \frac{P_2}{P_1} - 1;$$

next, add 1 to both sides, which gives

$$1+\pi_2 = \frac{P_2}{P_1}.$$

Finally, taking the inverses of both sides leads to

$$\frac{1}{1+\pi_2} = \frac{P_1}{P_2} . \tag{19}$$

Consumption-Savings Model in Real Units

Recall the **nominal LBC** of the two-period model,

$$P_1 c_1 + \frac{P_2 c_2}{1+i} = Y_1 + \frac{Y_2}{1+i} + (1+i)A_0, \tag{20}$$

where the notation is exactly as we have already developed. Each term is in **nominal** units in this expression. As shown in the diagram, we can recast the framework into purely **real (goods-denominated) units** and re-do the entire analysis.

Dividing the nominal LBC by P_1 is the first step in re-casting the analysis in real units:

$$c_1 + \frac{P_2}{P_1}\cdot\frac{c_2}{1+i} = \frac{Y_1}{P_1} + \frac{Y_2}{P_1\cdot(1+i)} + (1+i)\cdot\frac{A_0}{P_1}.$$

The "labor income" terms Y_1 and Y_2 are nominal income. Define **real income** in period 1 and period 2, respectively, as

$$y_1 \equiv \frac{Y_1}{P_1} \tag{21}$$

and

$$y_2 \equiv \frac{Y_2}{P_2}. \tag{22}$$

Notice now we have to be careful in distinguishing upper-case Y from lower-case y!

Substituting y_1 into the LBC gives

$$c_1 + \frac{P_2}{P_1}\cdot\frac{c_2}{1+i} = y_1 + \frac{Y_2}{P_1\cdot(1+i)} + (1+i)\cdot\frac{A_0}{P_1}.$$

To substitute y_2, observe that we can multiply and divide the second term on the right-hand side by P_2, which gives

$$c_1 + \frac{P_2}{P_1}\cdot\frac{c_2}{1+i} = y_1 + \frac{Y_2}{P_2}\cdot\frac{P_2}{P_1}\cdot\frac{1}{1+i} + (1+i)\cdot\frac{A_0}{P_1}$$

(all we have done is multiply by "1," which is always a valid mathematical operation). Now using the definition y_2, we have

$$c_1 + \frac{P_2}{P_1}\cdot\frac{c_2}{1+i} = y_1 + y_2\cdot\frac{P_2}{P_1}\cdot\frac{1}{1+i} + (1+i)\cdot\frac{A_0}{P_1}.$$

The definition of inflation allows us to replace the $\frac{P_2}{P_1}$ terms to obtain

$$c_1 + c_2\cdot\left(\frac{1+\pi_2}{1+i}\right) = y_1 + y_2\cdot\left(\frac{1+\pi_2}{1+i}\right) + (1+i)\cdot\frac{A_0}{P_1}$$

Next, using the exact Fisher expression $\frac{1+i}{1+\pi_2} = 1+r$, rewrite the LBC once again as

$$c_1 + \frac{c_2}{1+r} = y_1 + \frac{y_2}{1+r} + (1+i)\cdot\frac{A_0}{P_1}.$$

What's left to deal with is the seemingly complicated term at the far right-hand side. In terms of economics, it represents the **nominal receipts from the A_0** wealth with which the consumer began period 1, **stated in terms of period-1 purchasing power,** hence the appearance of P_1 in the denominator.

Using the same procedure as before, let's multiply and divide this term by P_0 (the nominal price level in period zero, or more generally stated, the nominal price level "in the past"), which gives us

$$c_1 + \frac{c_2}{1+r} = y_1 + \frac{y_2}{1+r} + (1+i)\cdot\frac{P_0}{P_1}\cdot\frac{A_0}{P_0}$$

Using the definition of inflation allows us to rewrite this as

$$c_1 + \frac{c_2}{1+r} = y_1 + \frac{y_2}{1+r} + \frac{1+i}{1+\pi_1}\cdot\frac{A_0}{P_0}.$$

Two steps remain. First, invoke the exact Fisher relationship. Second, define $a_0 \equiv \frac{A_0}{P_0}$ as the **real net wealth** of the consumer at the very end of period 0 and hence, equivalently and as shown in the timeline, at the very start of period 1. Finally, **the LBC in real terms is**

$$c_1 + \frac{c_2}{1+r} = y_1 + \frac{y_2}{1+r} + (1+r)\cdot a_0,$$

which is highly analogous to the LBC in nominal terms. Indeed, the two describe the same exact budget restriction on consumer optimization.

The real form of the LBC emphasizes that consumption (which is a real variable! Nobody eats dollar bills or sits down in front of a dollar bill to watch a baseball game!) decisions over time are ultimately dependent on real factors of the economy: the real interest rate and real ("labor") income.

It is true that in modern economies with developed monetary exchange and financial markets, dollar prices and nominal interest rates are the objects people seem to think in terms of when making consumption and savings decisions. This facet of reality is indeed why our analysis so far has been framed in nominal units

But we can boil these dollar prices and nominal interest rates down to real interest rates and describe much of consumer theory solely in terms of real factors.

None of this is to say, though, that consideration of currencies, dollar prices, and nominal interest rates are unimportant or uninteresting topics. Indeed, the whole field of "monetary economics" is primarily concerned with these issues, and we will have a lot to say later about monetary economics. Depending in which issues we are analyzing, we will use either the LBC in real terms or the LBC in nominal terms. If we are considering issues of inflation, for example, then the nominal LBC will typically be more appropriate.

We proceed now with the nominal LBC. For diagrammatical purposes, it will be, just as before, easier to assume that $a_0 = 0$ (that is, the individual has no initial wealth). Rearranging the real LBC into the ready-to-be-graphed "slope-intercept" form, we have

$$c_2 = -(1+r)c_1 + (1+r)y_1 + y_2. \tag{23}$$

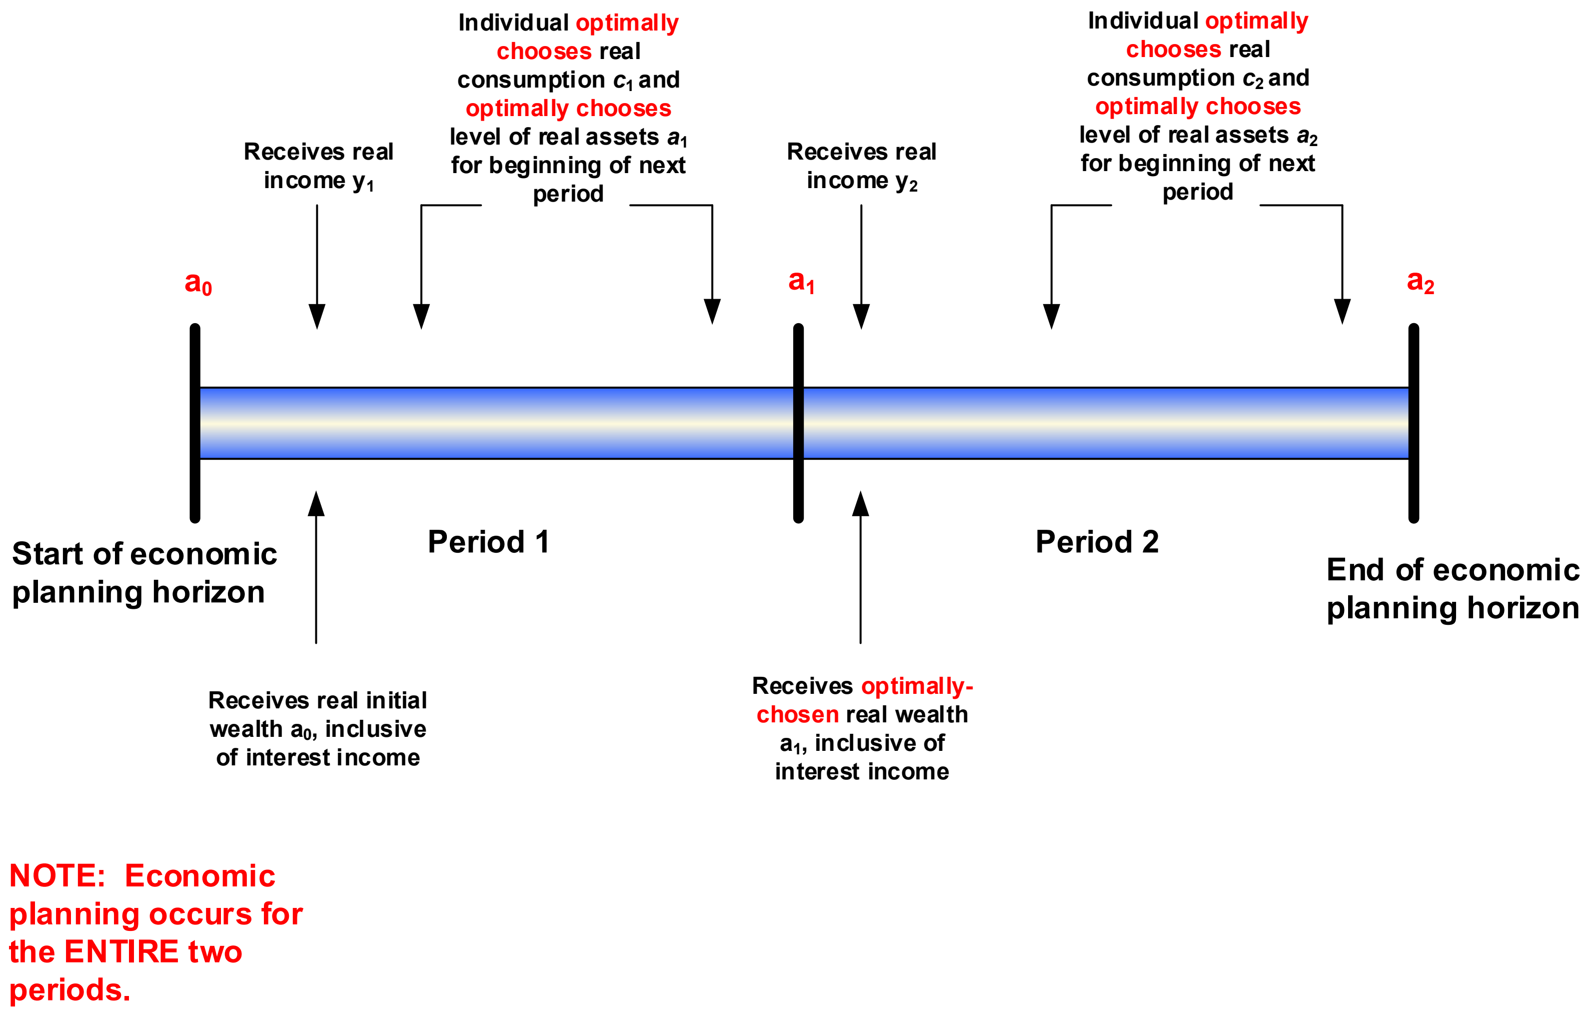

Figure 23. Timing of events in two-period consumption-savings framework, stated in real units.

The utility function $u(c_1, c_2)$ is unaffected by all of these manipulations of the LBC, meaning the indifference map is unaffected – as it must be, since budget constraints and indifference curves are two completely independent concepts.

Graphically, then, an example of an individual's optimal choice is shown in Figure 24 (which takes as given $a_0 = 0$). In this example, the individual consumes more than his real income in period 1, leading him to be in debt at the end of period 1; in period 2, he must repay the debt with interest and therefore consume less than his period-2 income. The definition of **real private savings** during the course of period 1 can be stated as

$$s_1^{priv} = ra_0 + y_1 - c_1,$$

which is quite analogous to one statement of **nominal private savings** during the course of period 1 (which, recall, was $S_1^{priv} = iA_0 + Y_1 - P_1c_1$

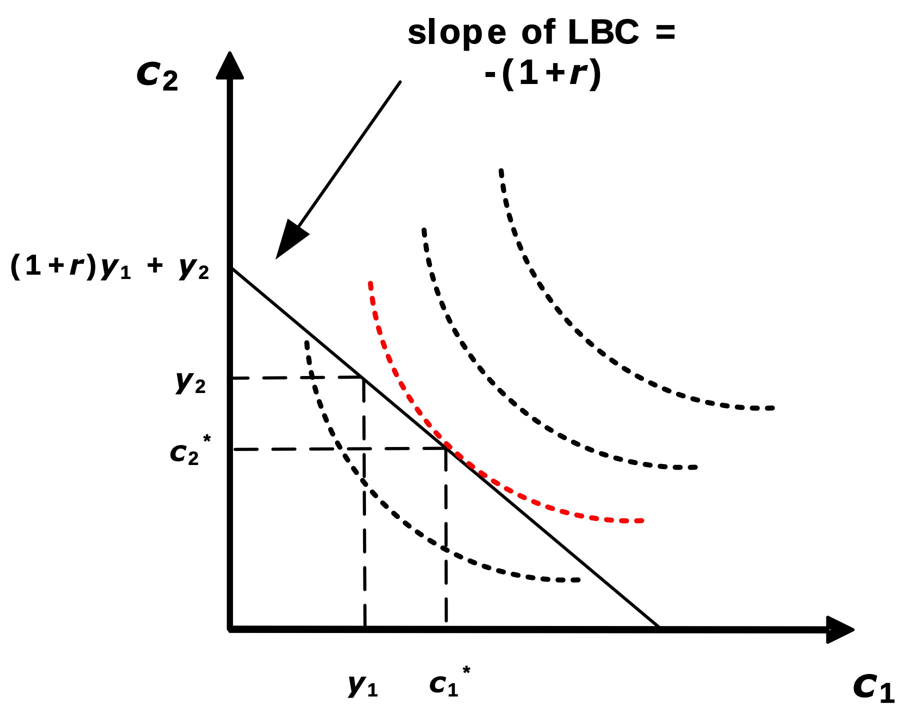

Figure 24. The interaction of the individual's LBC (here presented in real terms) and his preferences (represented by the indifference map) determine the individual's optimal consumption over time, here c_1* in period 1 and c_2* in period 2. The individual begins period 1 with $a_0 = 0$.

The Aggregate Private Savings Function

With the aid of Figure 24, we will now consider how changes in the real interest rate affect savings decisions of individuals. In our two-period model, there is only one time that the individual actually makes a decision about saving/borrowing: in period 1, when he must decide how of his period-1 labor income to save for period 2 or how much to borrow so that he can consume more than his period-1 labor income. As such, what we are exactly interested in is how S_1^{priv} (the same notation as before – private savings in period 1) is affected by r. Put more mathematically, what we are interested in is what the **private savings function** looks like.

Let us begin by supposing that the initial situation is as shown in Figure 24, in which the individual is a debtor at the end of period 1. Consider what happens to his optimal choice if the real interest rate r rises, while his real labor income y_1 and y_2 both remain constant. Such a rise in the real interest rate causes the LBC to both become steeper and have a higher vertical intercept, which we can see by analyzing the LBC (23). **In fact, the new LBC must still go through the point (y_1, y_2) because that is still a possible consumption choice for the individual.** That is, regardless of what the real interest rate is, it is always possible for the individual to simply not borrow or save in period 1 and simply consume his real labor income in each period. Because this is always possible, the point (y_1, y_2) must always lie on the LBC. Thus, the new LBC at the higher real interest rate is as shown in Figure 25. Also shown in Figure 25 are the new optimal consumption choices of the individual at the new higher interest rate. Specifically, notice that consumption in period 1 has decreased.

Because labor income in period 1 is unchanged, this means that his savings in period 1 has risen. Recall that private savings in period 1 is

$$S_1^{priv} = Y_1 - P_1 c_1 \tag{24}$$

in nominal terms. We can divide this expression through by P_1 to get savings in real terms,

$$s_1^{priv} = y_1 - c_1, \tag{25}$$

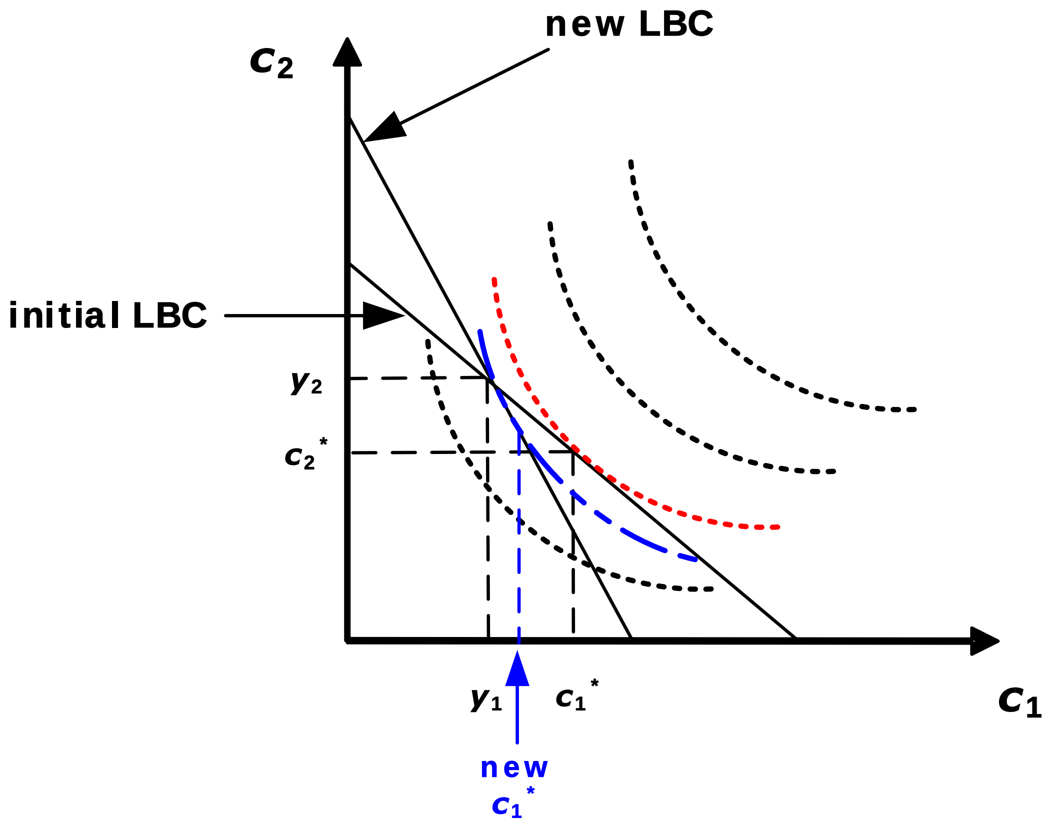

Figure 25. If at the initial real interest rate the individual chose to be a debtor at the end of period 1, then a rise in the real interest rate necessarily lowers consumption in period 1, implying that savings during period 1 has increased (or, equivalently, as shown, dissaving has decreased).

Notice the distinction between lower-case s_1^{priv}, which denotes real savings, and our earlier upper-case S_1^{priv}, which denotes nominal savings. The relationship is simply that $s_1^{priv} = S_1^{priv} / P_1$.[33] Thus, with unchanged y_1 and a decreased c_1^*, s_1^{priv} has increased. Actually, in Figure 25, savings is still negative after the rise in the real interest rate – but it is less negative, so indeed private savings has increased.

The preceding analysis seems to suggest that there is a positive relationship between the real interest rate and private savings. However, the conclusion is not so straightforward because we need to consider a different possible initial situation. Rather than the initial situation depicted in Figure 24, suppose instead that Figure 26 depicted the initial situation of the individual. In Figure 26, the optimal choice of the individual is such that

[33] By now, you should be noticing how to convert any nominal variable into its corresponding real variable – simply divide by the price level. The one slight exception is the nominal interest rate – to convert to the real interest rate requires use of the inflation rate (which itself depends on price levels, so the idea is still the same).

he consumes less in period 1 than his labor income in period 1, allowing him to accumulate positive wealth for period 2. That is, he saves during period 1.

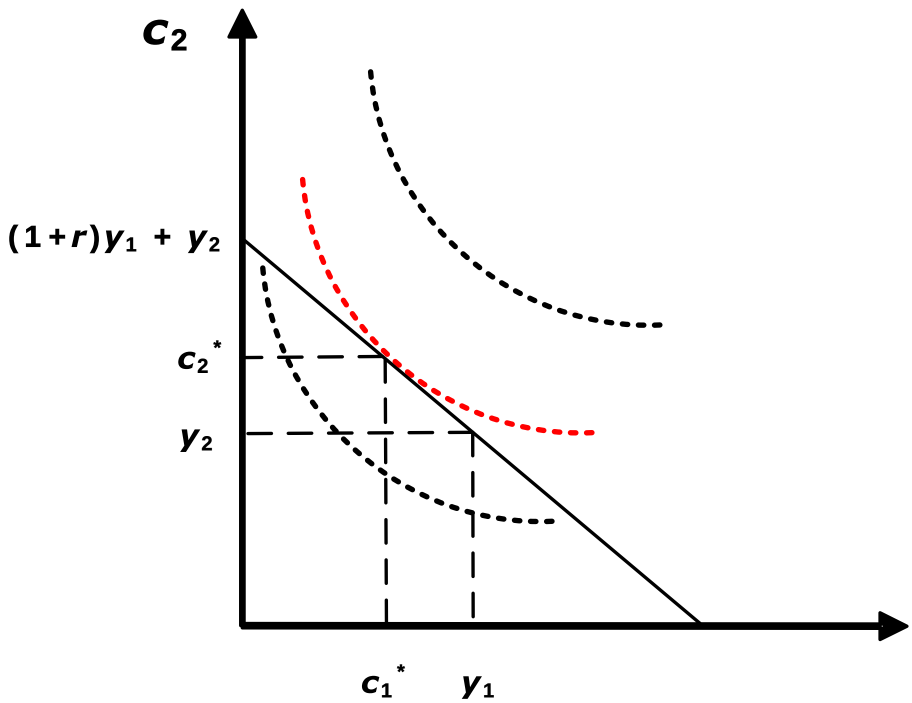

Figure 26. At the initial real interest rate, the individual's optimal choice may be such that he is **not** a debtor at the end of period 1 but rather a saver. This is because he chooses to consume less in period 1 than his labor income in period 1, which allows him to consume more in period 2 than his labor income in period 2.

Now suppose the real interest rate rises, with labor income y_1 and y_2 both held constant. The budget line again becomes steeper by pivoting around the point (y_1, y_2), as shown in both Figure 27 and Figure 28. However, depending on the exact shapes of the individual's indifference curves, the individual's consumption in period 1 may fall (shown in Figure 27) or rise (shown in Figure 28). In terms of his savings in period 1, then, a rise in the real interest rate may induce either a rise in savings (shown in Figure 27) or a fall in savings (shown in Figure 28).

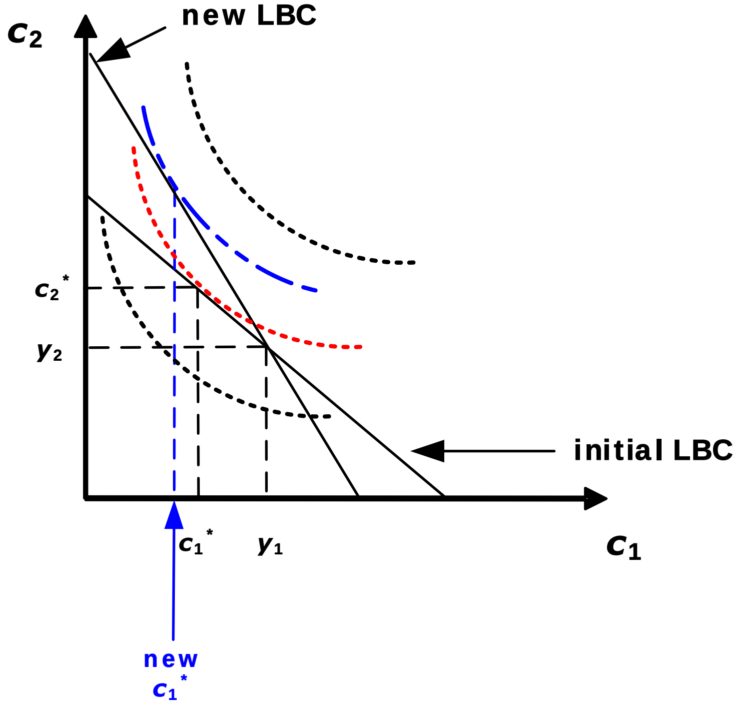

Figure 27. If the initial situation is such that the individual optimally chose to be a saver at the end of period 1, then a rise in the real interest rate may cause his savings in period 1 to increase…..

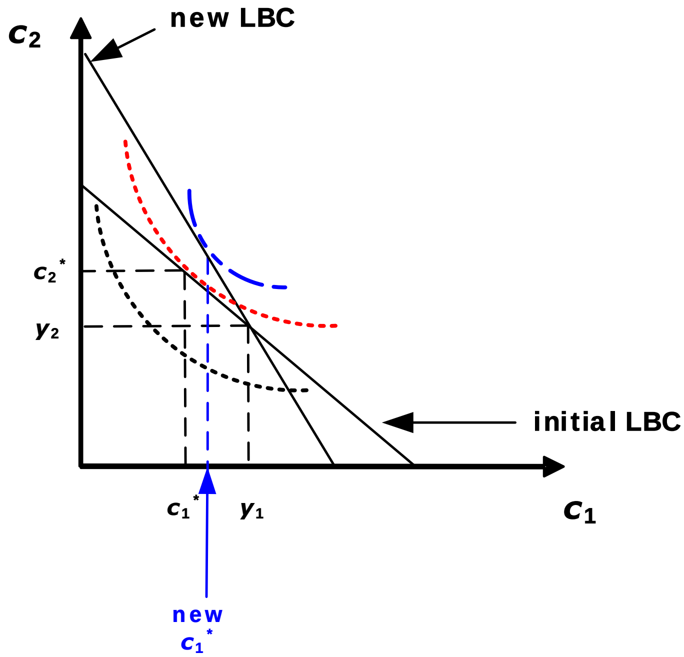

Figure 28. … or decrease, depending on the shape of his indifference map (i.e., depending on exactly what functional form his utility function has). Thus, for an individual who optimally initially chooses to be a

saver during period 1, it is impossible to determine theoretically in which direction his savings changes if the real interest rate rises.

Where does this leave us in terms of our ultimate conclusion about how private savings reacts to a rise in the real interest rate? Not very far theoretically, unfortunately. The summary of the above analysis is as follows. If an individual is initially a debtor at the end of period 1, then a rise in the real interest rate necessarily increases his savings during period 1. On the other hand, if an individual is initially a saver at the end of period 1, then a rise in the real interest rate may increase or decrease his savings during period 1. Overall, then, theory cannot guide us as to how private savings at the macroeconomic level responds to a rise in the real interest rate!

Where theory fails, we can turn to data. Many empirical studies conclude that the real interest rate in fact has a very weak effect, if any effect at all, on private savings behavior. The studies that do show that real interest rates do influence savings almost always conclude that a rise in the real interest rate leads to a rise in savings. The interpretation of such an effect seems straightforward: if all of a sudden the interest rate on your savings account rises (and inflation is held constant), then you may be tempted to put more money in your savings account in order to earn more interest income in the future.

We will adopt the (somewhat weak) empirical conclusion that the real interest rate has a positive effect on private savings – thus we will proceed with our macroeconomic models as if Figure 25 and Figure 27 are correct and Figure 28 is incorrect.[34]

This leads us to graph the upward-sloping aggregate private savings function in Figure 29.

Stocks vs. Flows

Let's return to the critical difference between stock variables and flow variables. Stated in terms of real goods (and as Figure 23 displays), the **stock (or, equivalently, accumulation) variables** are a_0, a_1, and a_2; and the **flow variables** are c_1, c_2, y_1, y_2; s_1, and s_2.

It is hard to emphasize how much the distinction between **stock variables** and **flow variables matters for all of macroeconomic analysis!** As our multi-period frameworks soon begin to include more and more time periods, the critical concepts of stocks vs. flows will continue to help us think about various economic events play out. So you are highly encouraged to understand the difference right away.

[34] Though debate amongst macroeconomists over this issue is not yet settled, this seems to be the most commonly-accepted interpretation of the results.

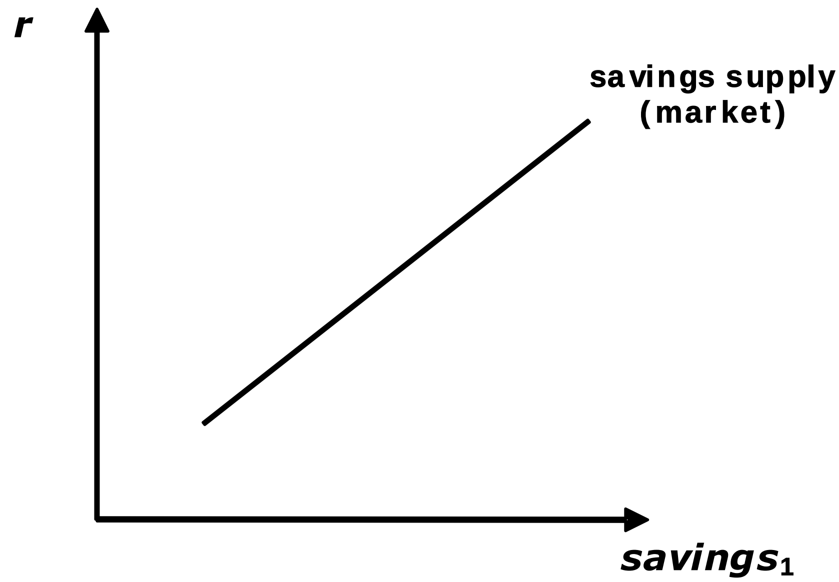

Figure 29. The upward-sloping aggregate private savings function**.**

Lagrange Characterization – the Consumption-Savings Optimality Condition

As we did with the consumption-leisure model, it is useful to work through the mechanics of analyzing the two-period model using our Lagrange tools. In analyzing multi-period models using Lagrangians, it turns out we have two alternative and distinctly useful ways of proceeding: an approach we will refer to as a **lifetime Lagrange formulation** and an approach we will refer to as a **sequential Lagrange formulation.**

These ideas will hopefully become clear as we describe how to pursue these two different Lagrange approaches, but the advantages and disadvantages of the two approaches can be summarized as follows. For a simple two-period model, the lifetime Lagrange formulation is essentially nothing more than a formal mathematical statement of the graphical analysis we have already conducted. It emphasizes, as the terminology suggests, that consumers can be viewed as making *lifetime* choices. The sequential Lagrange formulation, on the other hand, emphasizes the unfolding of economic events and choices over time, rather than starting from an explicitly lifetime view. In the end, the sequential approach will bring us to exactly the same conclusion(s) as the lifetime approach; the sequential approach will thus seem like a more circuitous mode of analysis.

We introduce the sequential Lagrangian approach, however, for two reasons. One reason is that when we soon extend things to an infinite-period model, in which graphical analysis becomes quite infeasible, the lifetime Lagrangian formulation (which, as just stated, is really just a mathematical formulation of analysis that can otherwise be carried out purely graphically) inhererently becomes a bit less interesting.

A second, and quite related, reason that sequential Lagrangian analysis is of interest is that it will allow us to explicitly track the dynamics of *asset prices* over time as macroeconomic events unfold over time. In the lifetime view of the two-period model, we effectively end up removing from our analysis the "intermediate asset position" A_1. In the richer infinite-period models to come, we will offer quite specific various interpretations of what A_1 "is," and we will naturally end up being concerned with "its price." Here, we have been loosely speaking of A as the "amount of money in the bank." This is a fine enough interpretation for now, but we will develop the concept of "A" much further in the chapters ahead, and the sequential Lagrangian approach will prove extremely useful in thinking about specific instantiations of A.

In what follows, we will formulate both the lifetime and sequential Lagrangians in nominal terms, but one could easily pursue either in real terms, as well – a useful exercise for you to try yourself.

Lifetime Lagrangian Formulation

To construct the lifetime Lagrangian for the two-period model, the general strategy is just as we have seen several times already: sum the objective function together with the constraint function (with a Lagrange multiplier attached to it) to form the Lagrangian, compute first-order conditions, and then conduct relevant analysis using the first-order conditions. The objective function to be maximized is obviously the consumer's lifetime utility function $u(c_1, c_2)$. The relevant constraint – recall we are pursuing the *lifetime* Lagrangian here – is the consumer's LBC. Associating the multiplier λ with the LBC, the lifetime Lagrangian for the two-period model is

$$u(c_1, c_2) + \lambda\left[Y_1 + \frac{Y_2}{1+i} - P_1c_1 - \frac{P_2c_2}{1+i}\right].$$

Note for simplicity we have dropped any initial assets, just as we did in our graphical analysis, by assuming $A_0 = 0$; none of the subsequent analysis depends on this simplifying assumption.

It should be clear by now that, apart from the first-order condition on the Lagrange multiplier, the two relevant first-order conditions that we need to compute are those with respect to c_1 and c_2. Indeed, these are the formal objects we need to compute. However, before simply proceeding to the mathematics, let's remind ourselves of what it means conceptually when we construct these objects. A first-order condition with respect to any particular variable (think in terms of basic calculus here) mathematically describes *how* a maximum is achieved by optimally setting/choosing that particular variable, taking as given the settings/choices for all other variables. In terms of the economics of our model, the consumer is *optimally* choosing *both* c_1 *and* c_2 (in order to maximize utility), which, from the formal mathematical perspective, requires computing first-order conditions of the Lagrangian with respect to both c_1 and c_2. Keep this discussion in mind when we consider the sequential Lagrangian.

The first-order conditions with respect to c_1 and c_2 (we'll neglect here the first-order condition with respect to λ, which, as should be obvious by now, simply returns to us the LBC) thus are:

$$\frac{\partial u}{\partial c_1} - \lambda P_1 = 0$$

$$\frac{\partial u}{\partial c_2} - \lambda \frac{P_2}{1+i} = 0$$

The next step, as usual, is to eliminate λ from these two conditions. From the first expression, we have $\lambda = \frac{\partial u / \partial c_1}{P_1}$; inserting this into the second expression gives us $\frac{\partial u}{\partial c_2} = \frac{\partial u}{\partial c_1} \frac{P_2}{P_1(1+i)}$. From earlier, we know that $\frac{P_2}{P_1(1+i)} = \frac{1+\pi_2}{1+i}$, which in turn, from the exact Fisher equation, we know is equal to $\frac{1}{1+r}$. Slightly rearranging the resulting expression $\frac{\partial u}{\partial c_2} = \frac{\partial u}{\partial c_1} \frac{1}{1+r}$ gives us

$$\frac{\partial u / \partial c_1}{\partial u / \partial c_2} = 1 + r,$$

which is our two-period model's **consumption-savings optimality condition**. The consumption-savings optimality condition describes what we saw graphically in **Figure 24**: when the representative consumer is making his optimal intertemporal choices, he chooses c_1 and c_2 in such a way as to equate his MRS between period-1 consumption and period-2 consumption (the left-hand-side of the above expression) to (one plus) the real interest rate (the right-hand-side of the above expression). The real interest rate (again, more precisely, one plus the real interest rate) is simply the slope of the consumer's LBC. The two-period model's consumption-savings optimality condition will be present in the richer infinite-period model we will build soon.

Sequential Lagrangian Formulation

We can alternatively cast the representative consumer's choice problem in the two-period world on a period-by-period basis. That is, rather than take the lifetime view of the consumer's decision-making process, we can take a more explicitly sequential view of events. A bit more precisely, we can think of the consumer as making optimal decisions for period 1 and *then* making optimal decisions for period 2. If there were more than just two periods, we could think of the consumer as *then* making optimal decisions for period 3, and *then* making optimal decisions for period 4, and *then* making optimal decisions for period 5, and so on.

In this explicitly sequential view of events, the consumer, in a given period, *chooses consumption for that period along with an asset position to carry into the subsequent period.* That is, in period t (where, in the two-period model, either t =1 or t = 2), the consumer chooses consumption c_t and asset position A_t; note well the time-subscripts here. Also, crucially, note that in the sequential formulation, we are thinking explicitly of the consumer as making an optimal choice with regard to intermediate asset positions; in the lifetime formulations of the two-period model, whether graphical or Lagrangian, we effectively removed intermediate asset positions from the analysis, as we have noted a couple of times. In the sequential formulation, we do not remove intermediate asset positions from the analysis; think of this as the consumer deciding how much to put in (or borrow from) the bank.

Formally, in order to construct the sequential Lagangian, we must, as always, determine what the relevant objective function and constraint(s) are. The objective function, as usual, is simply the representative consumer's utility function. In terms of constraints, in the sequential formulation we will impose **all of the period-by-period budget constraints,** rather than the LBC. In our two-period model, we obviously have only two budget constraints, one describing choice sets in period 1 and one describing choice sets in period 2.

Almost all of our Lagrangian analyses thus far have used only one constraint function. But recall from our review of basic mathematics that it is straightforward to extend the Lagrangian method to handle optimization problems with multiple constraints. All we need to do, once we have identified the appropriate constraints, is associate *distinct* Lagrange multipliers with each constraint and then proceed as usual.

To construct the sequential Lagrangian, then, associate the multiplier λ_1 with the period-1 budget constraint and the multiplier λ_2 with the period-2 budget constraint – note that λ_1 and λ_2 are distinct multipliers, which in principle have nothing to do with each other. The sequential Lagrangian is thus

$$u(c_1,c_2)+\lambda_1[Y_1-P_1c_1-A_1]+\lambda_2[Y_2+(1+i)A_1-P_2c_2].$$

In writing this Lagrangian, we have used our assumption that $A_0 = 0$ and our result that $A_2 = 0$. The sequential analysis then proceeds as follows. Compute the first-order conditions for the consumer's choice problem in period 1: recall from our discussion above that in period 1, the consumer optimally chooses c_1 and A_1. Mathematically, this requires us to compute the first-order conditions of the Lagrangian with respect to these two variables; they are

$$\frac{\partial u}{\partial c_1} - \lambda_1 P_1 = 0$$

$$-\lambda_1 + \lambda_2(1+i) = 0$$

Next, compute the first-order conditions for the consumer's choice problem in period 2: in period 2, the consumer optimally chooses c_2 and A_2. Mathematically, this requires us to compute the first-order conditions of the Lagrangian with respect to these two variables. Of course, in the two-period model, we have that $A_2 = 0$, so due solely to the artifice of the two-period model, we actually do not need to compute the first-order condition with respect to A_2; only if we had more than two periods in our model would we need to compute it. Thus, all we need from the period-2 optimization is the first-order condition with respect to c_2, which is

$$\frac{\partial u}{\partial c_2} - \lambda_2 P_2 = 0.$$

Let's proceed to eliminate multipliers from the three first-order conditions we just obtained (and note that we'll skip considering the first-order conditions with respect to the two multipliers – as should be obvious by now, they simply deliver back to us the period-1 budget constraint and the period-2 budget constraint). Note that we now have *two* multipliers to deal with. From the first-order condition on A_1, we have $\lambda_1 = \lambda_2(1+i)$. We'll have much more to say about this type of relationship between multipliers – this expression that links multipliers across time periods – when we study the infinite-period model; for now, let's just exploit the mathematics it provides. Take this expression for λ_1 and insert it in the first-order condition on c_1, yielding $\frac{\partial u}{\partial c_1} = \lambda_2(1+i)P_1$. We've gotten rid of the multiplier λ_1 but are still left with λ_2. Fortunately, we can use the first-order condition on c_2 to obtain an expression for the period-2 multiplier: $\lambda_2 = \frac{\partial u / \partial c_2}{P_2}$. Now, insert this expression into the previously-obtained condition to get

$$\frac{\partial u}{\partial c_1} = \frac{\partial u}{\partial c_2}\frac{(1+i)P_1}{P_2},$$

in which we finally have eliminated all multipliers. Rearranging this expression a bit,

$$\frac{\partial u / \partial c_1}{\partial u / \partial c_2} = \frac{(1+i)P_1}{P_2}.$$

We have seen the right-hand-side of this expression a couple of times already, and we know that we can transform it (using the definition of inflation and the Fisher relationship) into $1 + r$. Thus, the last expression becomes

$$\frac{\partial u / \partial c_1}{\partial u / \partial c_2} = 1 + r,$$

which clearly is simply the consumption-savings optimality condition we derived above in the lifetime formulation of the problem. Because we have already derived and discussed it, there of course is no reason to discuss the economics of it again.

The idea to really understand and appreciate here is that, whether we pursue the lifetime Lagrangian approach or the sequential Lagrangian approach, we arrive at exactly the same prediction regarding how consumers optimally allocate their intertemporal consumption choices: they do so in such a way as to equate the MRS between period-1 consumption and period-2 consumption to (one plus) the real interest rate.

The mathematical difference between the two approaches is that in the sequential approach we had to proceed by explicitly considering the first-order condition on the intermediate asset position A_1, which generated a relationship between Lagrange multipliers over time. Through the optimal decision on A_1, the consumer **does** take into account future period events, even though the mathematics may not make it seem apparent. In the lifetime approach, no such relationship had to formally be considered because there was, by construction, only one multiplier.

In the end, we should not be surprised that we reached the same conclusion using either approach – indeed, they are simply *alternative* approaches to the *same* problem, the problem being the representative consumer's utility maximization problem over time.

Optimal Numerical Choice

Regardless of a lifetime or sequential analysis, the same exact consumption-savings optimality condition arises: $\frac{\partial u / \partial c_1}{\partial u / \partial c_2} = 1 + r$. This expression is part of the heart of macroeconomic analysis.

However, if we actually wanted to solve for numerical values of the optimal choices of period-1 and period-2 consumption, the consumption-savings optimality condition is not enough. Why? **Because the consumption-savings optimality condition is one equation in two unknown variables.** A simple way to see this is to take the case of $u(c_1, c_2) = \ln c_1 + \ln c_2$. The consumption-savings optimality condition is thus $c_2/c_1 = 1+r$ (which at this point you should be able to obtain yourself). Even though the market real interest rate r is taken as given, it is clearly impossible to solve for **both** c_2 and c_1 from this **one** equation.

This might be obvious by this point (especially given all of the indifference-curve/budget constraint diagrams in Figure 25, Figure 26, Figure 27, and Figure 28!), but to complete the numerical solution of the two-period framework requires us to use **both the consumption-savings optimality condition and the budget constraint** to pin down the optimal numerical choices of consumption across time. In other words, there are two equations in the two unknowns, period-1 consumption and period-2 consumption. The ensuing example takes us step-by-step though the analysis, and it also raises an important economic interpretation of the optimal consumption choices across time that arise.

Consumption Smoothing

The concept of **"consumption smoothing"** is an important underlying theme of the results that emerge from multi-period representative consumer utility maximization. This powerful and intuitive economic result arises not just in the two-period framework, but also in the progressively richer models we will construct later.

An example using the two-period model sheds light on the idea of consumption smoothing.

Consumption-Smoothing Example

Suppose the lifetime utility function is $u(c_1, c_2) = \ln c_1 + \ln c_2$. And also assume that P_1 = 1, P_2 = 1, A_0 = 0, r = 0.10.

Case 1: Suppose the lifetime stream of nominal income is concentrated in the "later" period of the consumer's economic planning horizon – for example, Y_1 = 2 and Y_2 = 11.

To solve for the optimal numerical values of c_1 and c_2 requires use of the pair of expressions

$$\left(\frac{\partial u / \partial c_1}{\partial u / \partial c_2} = \right)\frac{c_2}{c_1} = 1 + r$$

and

$$P_1 c_1 + \frac{P_2 c_2}{1+i} = Y_1 + \frac{Y_2}{1+i}.$$

The few steps of algebra are left for you to go through (which is good reinforcement of basics). The numerical values of the optimal choices of consumption across time turn out to be

$$c_1^* = 6, \quad c_2^* = 6.6.$$

Case 2: Suppose instead the lifetime stream of nominal income is more evenly spread through the "early" period and the "later" period of the consumer's economic planning horizon – for example, $Y_1 = 7$ and $Y_2 = 5.5$. Once again, the optimal numerical values of consumption are determined by the consumption-savings optimality condition and the budget constraint. And also once again leaving the few steps of algebra for you to verify, optimal choices of consumption across time turn out to be

$$c_1^* = 6, \quad c_2^* = 6.6.$$

Clearly, the lifetime path of optimal consumption is the same, despite the large difference between the Case 1 lifetime income path ($Y_1 = 2$, $Y_2 = 11$) and the Case 2 lifetime income path ($Y_1 = 7$, $Y_2 = 5.5$).

This example demonstrates the **two different facets of consumption smoothing.** The first aspect is that individuals prefer their consumption across time to not vary very much. This result arises due to strictly increasing and strictly concave lifetime utility, which is part of the **preference** side of the framework.

The second aspect arises from the **constraint** side of the framework. Despite the two very different income scenarios in the example, optimal c_1 and c_2 are identical. The identical optimal consumption streams, despite the very different income streams, is due to the ability of the individual to borrow (in Case 1) as much as he or she wants during period one, and hence be in debt at the very beginning of period two. This is highlighted in the **negative value of the** A_1 term that arises in Case 1:

$$\begin{aligned} A_1 &= Y_1 - P_1 c_1 + (1+i)A_0 \\ &= 2 - 6 + 0 \\ &= -4 \end{aligned}$$

If, counter to the example, the individual faced **another** constraint, in addition to the budget constraints, that allowed no borrowing at all during period one, the Case 1 consumption outcomes would be quite different: we would have $c_1 = 2$ and $c_2 = 11$ as the **"credit-constrained"** optimal choices for Case 1.

Without the credit constraint, the Case 1 individual is borrowing (that is, **dissaving**) during period one, and repaying the accumulated debt, inclusive of interest payments, in period two. In Case 2, the individual is saving during period one, and using the accumulated wealth (inclusive of interest earnings) for consumption in period two. Using all of the terminology and definitions of the two-period consumption-savings framework, you should be able to verify all of this for yourself.