Chapter 8
Infinite-Period Representative Consumer and Asset Pricing

Macroeconomic models, at least those based on micro-foundations, used in applied research and in practice assume that there is an infinite number of periods, rather than just two as we have been for the most part assuming. A two-period model is usually sufficient for the purpose of illustrating intuition about how consumers make intertemporal choices, but in order to achieve the higher quantitative precision needed for many research and policy questions, moving to an infinite-period model is desirable.

Here we will sketch the problem faced by an infinitely-lived representative consumer, describing preferences, budget constraints, and the general characterization of the solution. In sketching the basic model, we will see that in its natural formulation, it easily lends itself to a study of asset-pricing. Indeed, this model lies at the intersection of macroeconomic theory and finance theory and forms the basis of consumption-based asset-pricing theories. We will touch on some of these macro-finance linkages, but we really will only be able to whet our curiosity about more advanced finance theory. For the most part, we will index time by arbitrary indexes $t-1, t, t+1$, etc., rather than "naming" periods as "period 1," "period 2," and so on. That is, we will simply speak of "period *t*," "period *t*+1," "period *t*+2," etc.

Before we begin, we again point out that "the consumer" we are modeling is a stand-in for the economy as a whole. In that sense, we of course do not literally mean that a particular individual considers his intertemporal planning horizon to be infinite when making choices. But to the extent that "the economy" outlives any given individual, an infinitely-lived representative agent is, as usual, a simple representation.

Preferences

The utility function that is relevant in the infinite-period model in principle is a lifetime utility function just as in our simple two-period model. As before, suppose that time begins in period one but now never ends. The lifetime utility function can thus be written as

$$v(c_1, c_2, c_3, c_4, c_5, ...).$$

This function describes total utility as a function of consumption in every period 1, 2, 3, … and is the analog of the utility function $u(c_1, c_2)$ in our two-period model. The function v above is quite intractable mathematically because it takes an infinite number of arguments. Largely for this reason, in practice an **instantaneous utility function** that describes how utility in a given period depends on consumption in a given period is typically used. The easiest formulation to consider is the additively separable function,

$$v(c_1, c_2, c_3, c_4, c_5, ...) = u(c_1) + \beta u(c_2) + \beta^2 u(c_3) + \beta^3 u(c_4) + \beta^4 u(c_5) + ...,$$

where $u(.)$ is the instantaneous utility function. As written, period-*t* utility depends only on period-*t* consumption.[53] We will discuss the term β that we have introduced into this utility function below.

There is nothing special about a "period one." It is just as informative to assume that decisions occur in period t, meaning that decisions about t-1, t-2, etc. quantities cannot be undone. Thus, at the beginning of period t, the planning horizon remaining in front of the consumer is $t, t+1, t+2, ...$. In our infinite-period model, we will thus adopt the convention that decision-making in period t is under consideration. Thus the relevant lifetime utility function for the representative consumer when making decisions in period t is

$$u(c_t) + \beta u(c_{t+1}) + \beta^2 u(c_{t+2}) + \beta^3 u(c_{t+3}) + ... = \sum_{s=0}^{\infty} \beta^s u(c_{t+s}). \qquad (1.1)$$

The summation operator on the right-hand-side is a useful way of representing the utility function.

Impatience

We have also introduced a **time discount factor,** denoted β, in the above formulation to represent the idea that utility further out in the future is not as valuable as utility closer in time to the present moment. The discount factor β is a value between zero and one. The way we have written the above lifetime utility function, we are assuming we are currently in period one, because period-one utility is not discounted at all by β.

The parameter β is meant to be a crude way of modeling the idea of "impatience." It probably strikes us all as generally reasonable to think of humans as impatient beings: all else equal, most of us (all of us?) would prefer to have *x* units of goods right this instant rather than one year from now, and we would probably also prefer to have those *x* units one year from now rather than two years from now. The time discount factor β gets at this idea: because $\beta < 1$, a given quantity of period-2 consumption does not generate as much utility as does the same quantity of period-1 consumption **when viewed from the perspective of period 1.** Furthermore, when viewed from the perspective of period 1, a given quantity of period-3 consumption does not generate as much utility as does the

[53] This itself may strike you as an unnecessary assumption. Indeed it is unnecessary, except that until recently computational limitations made this assumption an often practically necessary one. More recently, time-non-separable preferences, in which an instantaneous utility function of the form $u(c_t, c_{t-1})$ have gained increasing popularity, mostly because they have proven useful in resolving some anomalous predictions of first-generation representative consumer macro models.

same quantity of period-2 consumption. To capture this idea, we have introduced the β^2 term in front of $u(c_3)$: because $\beta < 1$, $\beta^2 < \beta$; this gets at the latter idea. By analogy, we have introduced β^3 in front of $u(c_4)$, β^4 in front of $u(c_5)$, and so on.

Whether the idea of impatience can be modeled simply as a "number between zero and one" is obviously quite debatable. Furthermore, whether impatience "builds up" over time by simply raising β to successively higher powers is obviously quite debatable. Crude or not, it does at least allow us to start getting at the idea of impatience. As we will see more often as we build ever-richer models, even making a start on formally modeling an idea is often great progress.

Assets and Budget Constraints

As in the two-period model, the consumer faces period-by-period budget constraints. Rather than just two, the consumer here faces an infinite number of budget constraints, one for each period. The general idea behind the flow budget constraints is just as in our basic two-period model. However, here we'll take a more concrete stand on what the assets are that consumers trade. Rather than just an ambigious, catch-all *"A"* as in our basic two-period model, let's suppose here that the assets that consumers buy and sell are "shares in the stock market" – as in, the Dow or S&P 500. Arguably, the most salient characterisitics of shares of stock (be it Microsoft stock, General Motors stock, or a broad Dow or S&P 500 index) are the price of one of share, and any potential dividends that ownership of a share entitles one to receive. We will model these features of stock.

Our infinite-period model's period-*t* **flow budget constraint** is thus

$$P_t c_t + S_t a_t = S_t a_{t-1} + D_t a_{t-1} + Y_t, \tag{1.2}$$

in which c_t is consumption in period *t*, P_t is the price level in period *t*, a_t is the consumer's holdings of real assets – shares of stock – at the end of period *t*, S_t is the nominal price in period *t* of one share, D_t is a nominal dividend paid by each share, and Y_t is nominal income of the consumer in period *t*, which we will assume the consumer has no control over. Note the terms involving assets. In period *t*, the consumer begins with asset holdings a_{t-1}. In period *t*, each unit of these assets has some value S_t, and each unit of these assets carried into *t* pay a dividend D_t. Each unit of asset (share of stock) the consumer wishes to carry into period *t*+1, denoted by a_t also has a unit price of S_t. In more formal-sounding language, S_t is an asset price – it is the price of each share of stock.

An analogous flow budget constraint holds in each period $t, t+1, t+2, \ldots$ In principle we could combine all these flow budget constraints into a single lifetime budget constraint, as we did in the two-period model. However, it seems more natural in the infinite-period model to work with the flow budget constraint, which acknowledges that the decision-making happens sequentially (ie, period-by-period), rather than once-and-for-all like we implicitly assumed in the two-period model; recall our discussion of the sequential (Lagrangian) approach to the two-period model.

Optimal Choice

In order to consider optimal choices, then, we must formulate a Lagrangian. Specifically, the problem of the representative consumer in period *t* is to choose consumption c_t and asset holdings a_t to maximize lifetime utility (1.1) subject to the flow budget constraint (1.2), taking as given the price P_t of consumption, the price S_t of assets, the per-unit dividend D_t, and income Y_t. The Lagrangian must take into account the fact that the flow budget constraint applies in every period from *t* onwards. Thus we have

$$\begin{aligned} L(c_t, a_t, \lambda_t; c_{t+s}, a_{t+s}, \lambda_{t+s}) &= \sum_{s=0}^{\infty} \beta^s u(c_{t+s}) \\ &+ \lambda_t \left[Y_t + S_t a_{t-1} + D_t a_{t-1} - P_t c_t - S_t a_t \right] \\ &+ \beta\lambda_{t+1} \left[Y_{t+1} + S_{t+1} a_t + D_{t+1} a_t - P_{t+1} c_{t+1} - S_{t+1} a_{t+1} \right] \\ &+ .... \end{aligned}$$

where λ_t is the **multiplier on the period-*t* budget constraint,** and the ellipsis indicate that technically the Lagrangian has an infinite number of terms corresponding to the infinite number of future flow budget constraints. As we will see, in the current problem it is sufficient to write out just the t and t+1 flow budget constraints. Also note carefully that the *t*+1 budget constraint in the Lagrangian is discounted by β. This is because *everything* about period *t*+1 is discounted when viewing from the perspective of time *t*, including income and expenditures. If we had written out the period *t*+2 budget constraint in the Lagrangian, it would be discounted by β^2, just as instantaneous utility in period *t*+2 is discounted by β^2. As should be clear by recalling our study of the two-period model, what we have formulated is a sequential Lagrangian, in which each distinct flow budget constraint receives its own distinct Lagrange multiplier.

The objects of choice in period *t* are c_t and a_t. In line with how a sequential Lagrangian analysis proceeds, the first-order conditions of the Lagrangian with respect to these objects are

$$u'(c_t) - \lambda_t P_t = 0$$

and

$$-\lambda_t S_t + \beta\lambda_{t+1}\left(S_{t+1} + D_{t+1}\right) = 0 .$$

These first-order conditions should make clear why we did not need to write out explicitly the *t*+2 budget constraint in the Lagrangian: neither c_t nor a_t, the objects of choice in period t, appear in the *t*+2 budget constraint. In contrast, a_t *does* appear in the *t*+1 flow budget constraint, so it is useful to write that term out. The reason a_t appears in the *t*+1 budget constraint is that assets carried from period *t* into period *t*+1 have some value and pay some dividend in period *t*+1.

From the first-order condition on consumption in period *t*, we have that $\lambda_t = \frac{u'(c_t)}{P_t}$. Also, from the **first-order condition on consumption in period *t*+1** (construct this term yourself)**,** we would have the very similar condition $\lambda_{t+1} = \frac{u'(c_{t+1})}{P_{t+1}}$. Inserting these expressions for both λ_t and λ_{t+1} into the first-order condition on shares of stock, we have

$$\frac{u'(c_t)S_t}{P_t} = \beta \frac{u'(c_{t+1})(S_{t+1}+D_{t+1})}{P_{t+1}}. \tag{1.3}$$

There are several alternative useful ways of looking at this expression. First, we can rearrange it to highlight the intertemporal marginal rate of substitution:

$$\frac{u'(c_t)}{\beta u'(c_{t+1})} = \frac{S_{t+1}+D_{t+1}}{S_t} \cdot \frac{P_t}{P_{t+1}}. \tag{1.4}$$

The left-hand-side is the intertemporal marginal rate of substitution – after all, it is simply a ratio of marginal utilities – between consumption in period t and t+1. This is simply the analog of our condition u_1/u_2 in the two-period economy. Turning to the right-hand-side, the term P_t/P_{t+1} is the inverse of the gross inflation rate between period t and t+1, that is, $1/(1+\pi_{t+1})$. The term $\frac{S_{t+1}+D_{t+1}}{S_t}$ is the **holding period return** of the asset a_t -- it measures the gain (or loss…) of holding the asset from period t to t+1. This gain is higher the higher is the period t+1 price and/or dividend, $S_{t+1}+D_{t+1}$, and is lower the higher is the current (period t) price S_t. Finally, also note that the discount factor β appears in the denominator of the left-hand-side of (1.4). This is because, from the perspective of period t, the marginal utility of period-t+1 consumption is discounted due to impatience.

The right-hand-side of (1.4) is the analog of the term $(1+r)$ from our two-period model. The reason $(1+r)$ does not appear explicitly is simply because of the assumption about the available assets we have made here. Later, when we study monetary models, we will assume there are assets in the environment that pay a nominal interest rate, as in our simple two-period model, which will allow us to regenerate that term. To aid us in thinking about some other issues, below, though, sometimes it will be useful to represent (1.4) as

$$\frac{u'(c_t)}{\beta u'(c_{t+1})} = 1 + r_t, \tag{1.5}$$

where the term 1 + r_t hides all of the details we see in equation (1.4); hiding these details can sometimes be useful.

Consumption-Based Asset Pricing

Expression (1.4) highlights optimal choices from a macroeconomic perspective, putting things into "MRS equals price ratio" form. Alternatively, and especially given our specific interpretation of *a* here as shares of stock, we can view things from a more finance-oriented perspective, by focusing on the asset price S_t. More precisely, we can think about what sorts of factors are relevant for determining what the price of a share of stock in any time period.

Return to the first-order condition on assets, which we reproduce here for convenience,

$$-\lambda_t S_t + \beta\lambda_{t+1}(S_{t+1} + D_{t+1}) = 0. \tag{1.6}$$

From this expression, we can solve for the period-t stock price,

$$S_t = \frac{\beta\lambda_{t+1}}{\lambda_t}(S_{t+1} + D_{t+1}). \tag{1.7}$$

In finance theory, one would identify two distinct components on the right-hand-side of (1.7): the term $\frac{\beta\lambda_{t+1}}{\lambda_t}$ is the **pricing kernel,** and the term $(S_{t+1} + D_{t+1})$ is the **future return.** Thus, what the asset-price equation (1.7) states is that the period-t price of a share of stock depends on the future return and a pricing kernel. The future return has two components, arising from any future dividends that buying a share of stock in period t entitles one to and any change in the share price itself between period t and period t+1.

The pricing kernel seems a bit more esoteric, being a function of the period-t and period-t+1 Lagrange multipliers. But here is where the link between finance and macroeconomics emerges. We know from our macroeconomic analysis that $\lambda_t = u'(c_t)/P_t$ and $\lambda_{t+1} = u'(c_{t+1})/P_{t+1}$. Inserting these expressions into (1.7) allows us to express the stock price S_t as

$$S_t = \beta\frac{u'(c_{t+1})}{u'(c_t)}(S_{t+1} + D_{t+1})\frac{P_t}{P_{t+1}}. \tag{1.8}$$

Further more, we know that $\frac{P_t}{P_{t+1}} = \frac{1}{1+\pi_{t+1}}$, where π_{t+1} is the rate of inflation between period t and period t+1. Rewriting one more time, we have that the stock price S_t is

$$S_t = \frac{\beta u'(c_{t+1})}{u'(c_t)} \left(\frac{S_{t+1} + D_{t+1}}{1+\pi_{t+1}} \right). \tag{1.9}$$

Now we can begin to more fully appreciate the linkages between macroeconomic events and asset (stock) prices. The asset-price equation (1.9) shows that stock prices in period t depend on what the future inflation rate will be and how consumption will change over time. For example, all else equal, the higher is $u'(c_{t+1})/u'(c_t)$, the higher will be S_t. And, all else equal, the higher is π_{t+1}, the lower will be S_t. We will explore such issues in more depth, but the broad point to appreciate here is that things such as monetary policy (which impinges in what inflation rate occurs in the economy) and how aggregate consumption evolves over time (recall that consumption makes up about 70% of total GDP) affect stock prices.

Steady-State

Our infinite-period model allows us to explore yet another issue, one that will be important to understand when we study business cycle issues as well as monetary policy issues. We have an infinite number of periods in our model, and in principle all variables – consumption, interest rates, asset prices, etc. – can be moving around over time. Indeed, in a dynamic economy, they inevitably do all move around over time, and understanding how and why certain variables evolve over time as they do is a broadly-defined goal of macroeconomics. But suppose for a moment that eventually the **real** variables in our infinite-period model "settle down" to some constant values.

Let's formally define a **steady state** of an economy as a situation in which **all *real* variables stop fluctuating over time.** Note the emphasis on the word *real* here. In our infinite-period model, a steady-state would involve consumption (which is a *real* variable) becoming constant over time, asset holdings *a* becoming constant over time, and the real interest rate becoming constant over time. Variables such as S_t, D_t, and P_t, because they are *nominal* variables, need *not* become constant over time in order to fit into our definition of steady-state, although they could become constant as well. To introduce more terminology, the steady-state of an economy is often referred to as the **long-run equilibrium** of an economy – think of it, if you will, as the "average" or "potential" performance of the economy (to invoke loose terms you likely encountered in basic macroeconomics).

To provide ourselves some more notation, suppose that the constant level of consumption to which the sequence of c_t eventually converges is $\bar{c}$; hence, we can think of the steady-state as a state of affairs in which $c_t = c_{t+1} = c_{t+2} = ... = \bar{c}$. Similarly, suppose that the constant level of *real* interest rate to which the sequence of real interest rates eventually converges is $\bar{r}$; hence, we can think of the steady-state as a state of affairs in which $r_t = r_{t+1} = r_{t+2} = ... = \bar{r}$. And so on for all real variables of our model.

Impatience and the Real Interest Rate

Consider expression (1.5), which is nothing more than the infinite-period model's **consumption-savings optimality condition.** Indeed, it is no different from our two-period model's consumption-savings optimality condition, apart from the introduction of the time discount factor. In a steady-state, the consumption-savings optimality condition can be expressed as

$$\frac{u'(\bar{c})}{\beta u'(\bar{c})} = 1 + \bar{r}. \tag{1.10}$$

Clearly, the $u'(\bar{c})$ terms cancel, leaving us with

$$\frac{1}{\beta} = 1 + \bar{r}. \tag{1.11}$$

Expression (1.11) captures an extremely critical idea embedded in virtually all of modern macroeconomic theory and thus is at the root of a wide range of both academic and policy discussions of macroeconomics. What it says is that, in the steady-state – alternatively, "in the long run," or "on average" – the real interest rate of the economy is fundamentally tied to the degree of impatience of consumers in the economy. The theoretical upper end of β is $\beta = 1$; if $\beta = 1$, then expression (1.11) immediately tells us that the long-run real interest rate equals zero. That is, if consumers are perfectly patient (which is what $\beta = 1$ means) there is no net real return from savings.

Suppose instead, for the sake of numerical illustration, that $\beta = 0.95$, meaning that consumers are somewhat impatient. Expression (1.11) then immediately allows us to conclude that the steady-state real interest rate in the economy is roughly $\bar{r} = 0.0526$. Suppose instead that, $\beta = 0.9$, meaning that consumers are somewhat more impatient. Expression (1.11) then immediately allows us to conclude that the steady-state real interest rate in the economy is roughly $\bar{r} = 0.11$.

To cast these conclusions in very broad perspective, the most primitive, fundamental source of "interest rates" in the economy is human impatience. If human beings were always infinitely-patient creatures ($\beta = 1$), (real) interest rates would be zero. Thus, the mere presence of impatience at all ($\beta < 1$) is the fundamental source of positive interest rates in the world. Not Wall Street; not central banks – the primitive reason for the general existence of positive interest rates is human impatience, however crudely we have modeled it. Expression (1.11) then also shows us that the more impatient consumers are (remember, we are always speaking of the representative consumer) the higher are real interest rates.

This deep connection between interest rates and people's inclination towards impatience cannot be overemphasized in its central importance to macroeconomic theory. It is a deceptively simple idea – equation (1.11) obviously looks simple enough, but the idea it captures will continue to be at the root of the richer models we'll continue building. As such, it is useful to wrap your mind around this idea as well as possible now.