Chris Edmond ⟨cpedmond@unimelb.edu.au⟩

Introduction to consumption-based asset pricing

We will begin our brief look at asset pricing with a review of the essentials of Robert Lucas's (1978) representative agent approach to asset pricing. We will also use this model to introduce an equilibrium concept — recursive competitive equilibrium — that we will make use of when studying decentralized problems.

A. Lucas trees

Consider an endowment economy with a single type of **durable asset**. There are a large number of identical households each endowed with a single, identical, non-depreciating fruit tree, $s_0 = 1$. A fruit tree produces dividends (fruit), x_t, according to some **exogenous stochastic process**. Dividends cannot be stored, the only store of value are the trees. We are interested in pricing the assets (trees). Denote the price of a tree by p_t. The representative consumer's sequence problem is to choose a sequence of **shares** $\{s_{t+1}\}$ (claims to the ownership of the representative fruit tree) to maximize

$$\mathsf{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t U(c_t) \right\}, \qquad 0 < \beta < 1$$

subject to, for each $t = 0, 1, ...$

$$c_t + p_t s_{t+1} \leq (p_t + x_t) s_t$$

$$c_t, s_{t+1} \geq 0$$

$$\text{given } s_0, x_0$$

and an exogenous stochastic process for the dividends.

The household's expenditure is constrained by their **wealth**, $w_t \equiv (p_t + x_t)s_t$ (the value of the stock of trees held at t plus the value of the period flow of dividends per tree at t). Wealth can either be eaten (as a consumption good) or used to buy more trees for delivery next period.

The sequence problem can be reformulated as a dynamic programming problem

$$V(s, x) = \max_{s' \geq 0} \{ U(c) + \beta \mathsf{E}[V(s', x') \mid x] \}$$

subject to

$$c + p(x)s' \leq [p(x) + x]s$$

and a given Markov process that defines transition probabilities for the dividends. In this formulation, the **individual's state** is (s, x). A solution to this dynamic programming problem is a policy function g of the individual state that gives $s' = g(s, x)$. The price of the asset $p(x)$, is a function of just the **aggregate state** of the economy, namely x, and an individual take this function as given when solving their dynamic programming problems.

Notice that we can also think of wealth itself being an individual-specific state and instead write

$$V(w, x) = \max_{s' \geq 0}\{U(c) + \beta \mathsf{E}[V(w', x') \mid x]\}$$

subject to

$$s' = \frac{1}{p(x)}(w - c)$$
$$w' = \frac{p(x') + x'}{p(x)}(w - c)$$

where the household's budget constraint has been combined with the definition of next period's wealth, $w' = [p(x') + x']s'$, to construct a law of motion for wealth. The **gross rate of return** on the asset is given by $[p(x') + x']/p(x)$.

There is no production in this economy, so there are no firms and we can go straight to the equilibrium concept.

DEFINITION. A **recursive competitive equilibrium** is a collection of functions: (i) a value function V, (ii) a policy function g, and (iii) a pricing function p, such that:

1. Given the pricing function p, the value function V and the policy function g solve the household's dynamic programming problem, and

2. Markets clear

$$g(w, x) = 1 \qquad \text{for all } (w, x)$$

From the market clearing condition, $g(w, x) = 1$, the flow budget constraint requires

$$c + p(x) \cdot 1 = [p(x) + x] \cdot 1$$

or

$$c = x$$

This is the **resource constraint** for the economy. Since every household is the same, the equilibrium outcome (using the budget constraint) must be the **no-trade** outcome. Since every household wants to be on the same side of the market, the only outcome can be that every household simply retains ownership of its initially endowed tree. We are interested in finding the prices that **support** this no trade equilibrium.

Now let's think about how to characterize the price of the tree. The FOC for the household's problem is

$$U'(c)p(x) = \beta \mathsf{E}\left\{ \frac{\partial V(w', x')}{\partial w'}[p(x') + x'] \,\Big|\, x \right\}$$

and the envelope condition gives

$$\frac{\partial V(w, x)}{\partial w} = U'(c)$$

Combining these conditions, we can write

$$U'(c)p(x) = \mathsf{E}\left\{ \beta U'(c')[p(x') + x'] \mid x \right\}$$

Using the equilibrium condition that each period's consumption always equals that period's dividends, $c = x$, we can simplify to

$$U'(x)p(x) = \mathsf{E}\left\{ \beta U'(x')[p(x') + x'] \mid x \right\}$$

Now introduce a new function $\psi(x)$ defined by

$$\psi(x) \equiv U'(x)p(x)$$

which we can think of as the marginal utility loss incurred to obtain an extra unit of the tree. We have to solve a single functional equation in this new unknown

$$\psi(x) = \mathsf{E}\left\{ \beta\psi(x') + \beta U'(x')x' \mid x \right\}$$

If we can solve the functional equation for $\psi(x)$ then we have the price of the tree (since $U'(x)$ is known from the primitives of the model). This functional equation is typically easier to solve than

a Bellman equation because it does not involve a max operator, just first order conditions. Indeed, if x is defined by a Markov chain then solving this functional equation problem just involves solving a finite system of linear equations.

Closed form example

Let $U(c) = \log(c)$. The relevant operator, say T, is defined by

$$(T\psi)(x) = \mathsf{E}\left\{\beta\psi(x') + \beta \mid x\right\}$$

We want to find a ψ such that $T\psi = \psi$, which we can do by iterating until a functional form is preserved.

Step 1. Guess that $\psi_0(x) = 0$ all x. This implies

$$(T\psi_0)(x) = \mathsf{E}\left\{\beta 0 + \beta \mid x\right\} = \beta$$

(well, that was silly anyway — this is a model with an always positive marginal utility of consumption, so it would be surprising indeed if the asset price was zero which is what it would have to be if $\psi_0(x) = 0$ all x).

Step 2. OK, guess instead that $\psi_1(x) = A > 0$ all x. This implies

$$(T\psi_1)(x) = \mathsf{E}\left\{\beta A + \beta \mid x\right\} = \beta(A + 1)$$

This time the T operation preserves the functional form. Now let's use the method of undetermined coefficients to solve for this unknown coefficient A. If $A = \beta(A + 1)$, then $A = \beta/(1 - \beta)$, and we have our functional form for $\psi(x)$

$$\psi(x) = \frac{\beta}{1 - \beta} \qquad \text{all } x$$

and so

$$p(x) = \frac{\beta}{1 - \beta}x$$

We have found a function mapping the aggregate state of the economy x into a price for the tree, $p(x)$.

Notice that in this special example, the distribution over future dividends did not matter. In fact, future dividends don't matter at all in the logarithmic case: all that matters is the subjective discount rate and how much fruit there is today (the income and substitution effects just exactly offset each other). This log example is quite counterfactual, since it implies that the **price/dividend ratio** is a constant

$$\psi(x) = \frac{p(x)}{x} = \frac{\beta}{1 - \beta} \qquad \text{all } x$$

Once we have solved for the pricing function, we can compute other objects of interest. For example, with log utility the gross return on a Lucas tree is

$$R(x', x) \equiv \frac{p(x') + x'}{p(x)} = \frac{\frac{\beta}{1-\beta}x' + x'}{\frac{\beta}{1-\beta}x} = \frac{1}{\beta}\frac{x'}{x}$$

Hence the gross return depends on the pure rate of time discount through the constant β^{-1} term plus the (random) growth rate of the supply of fruit.

Euler equations

An alternative approach to asset pricing is to work more closely with the sequence formulation of the problem. We can write the fundamental functional equation of the problem, the equilibrium consumption Euler equation, as

$$U'(x_t)p_t = \mathsf{E}_t \left\{ \beta U'(x_{t+1})(p_{t+1} + x_{t+1}) \right\}$$

and so

$$p_t = \mathsf{E}_t \left\{ \beta \frac{U'(x_{t+1})}{U'(x_t)}(p_{t+1} + x_{t+1}) \right\}$$

or

$$1 = \mathsf{E}_t \left\{ \beta \frac{U'(x_{t+1})}{U'(x_t)} \left(\frac{p_{t+1} + x_{t+1}}{p_t} \right) \right\} = \mathsf{E}_t \left\{ \beta \frac{U'(x_{t+1})}{U'(x_t)} R_{t+1} \right\}$$

We can now recursively substitute forward using p_{t+j+1} to eliminate p_{t+j} and the **law of iterated expectations**, that $\mathsf{E}_t \{\mathsf{E}_{t+1} \{p_{t+j}\}\} = \mathsf{E}_t \{p_{t+j}\}$ to obtain

$$p_t = \lim_{T \to \infty} \mathsf{E}_t \left\{ \sum_{t=1}^{T} \beta^j \frac{U'(x_{t+j})}{U'(x_t)} x_{t+j} \right\} + \lim_{T \to \infty} \mathsf{E}_t \left\{ \beta^{T+1} \frac{U'(x_{t+T+1})}{U'(x_t)} p_{t+T+1} \right\}$$

As usual, we will presume — or make boundedness assumptions on the utility function sufficient to ensure — that the limiting price term goes to zero. We can then write the solution to the difference equation as

$$p_t = \mathsf{E}_t \left\{ \sum_{t=1}^{\infty} \beta^j \frac{U'(x_{t+j})}{U'(x_t)} x_{t+j} \right\}$$

The expression

$$M_{t+j} = \beta^j \frac{U'(x_{t+j})}{U'(x_t)}$$

is a random variable — often referred to as a **stochastic discount factor** — and reflects both the usual time preference of the household and its desire for smoothness in consumption (aversion to bumpy consumption paths).

B. Contingent claims markets

A one-period-ahead contingent claim is an asset that delivers **one** unit of the consumption good (fruit) if and only if a particular state is realized. Suppose that there is only a finite number of possible states with transitions given by a Markov chain. Let $\pi(x', x) = \Pr(x_{t+1} = x' | x_t = x)$ denote the transition probabilities for state x' given state x.

Denote by $q(x', x)$ the price in state x of a **contingent claim** to one unit of the consumption good if x' is realized in the next period. The price of a sure claim to a unit of the consumption good is the price of a claim to one unit of the consumption good in the next period no matter what state is actually realized next period. Thus, the **price of a sure claim** in state x is found by summing this **pricing kernel** $q(x', x)$ over all x'. We can think of a sure claim to a unit of the consumption good as a safe one-period asset (a one-period **pure discount bond**).

Consider an individual household that is endowed with a tree and can trade in markets for trees and markets for contingent claims for all possible states. Given wealth w and this period's state x, the individual household's dynamic programming problem is

$$V(x, w) = \max_{s' \geq 0, B(x') \geq 0} \left\{ U(c) + \beta \sum_{x'} V(x', w') \pi(x', x) \right\}$$

where the maximization is over choices of shares s' and a portfolio of contingent claims $B(x')$ and is subject to the budget constraint

$$c + p(x)s' + \sum_{x'} q(x', x) B(x') \leq w$$

with wealth

$$w = [p(x) + x]s + B(x)$$

As well as a quantity of trees, each household chooses a quantity of contingent claims for each future state x'. This vector of quantities is summarized by B with typical element $B(x')$. Policy functions are

$$s' = g_s(x, w)$$
$$B(x') = g_B(x', x, w)$$

DEFINITION. A **recursive competitive equilibrium** is a collection of functions: (i) a value function V, (ii) policy functions g_s, g_B and (iii) pricing functions p and q such that:

1. Given the pricing functions p and q, the value function v and the individual decision rules g_s and g_B solve the household's dynamic programming problem,

2. Markets clear

$$s' = g_s(x, w) = 1 \qquad \text{for all } (x, w)$$
$$B(x') = g_B(x', x, w) = 0 \qquad \text{for all } (x, w) \text{ and all } x'$$

Goods market clearing is then

$$c = x$$

The pricing functions are again characterized using the first order conditions for the household's choice of trees and contingent claims. Leaving aside the trees for the moment, we can find the contingent claims pricing kernel by examining the first order condition for $B(x')$ and the related envelope condition. The first order condition for some typical state x' and associated $B(x')$ is

$$B(x') : \qquad U'(c)q(x', x) = \beta \frac{\partial V(x', w')}{\partial w'} \pi(x', x)$$

The envelope condition requires

$$\frac{\partial V(x, w)}{\partial w} = U'(c)$$

which implies the Euler equation

$$U'(c)q(x', x) = \beta U'(c')\pi(x', x)$$

The **equilibrium Euler equation**, where $c = x$, is then just

$$U'(x)q(x', x) = \beta U'(x')\pi(x', x)$$

or

$$q(x', x) = \beta \frac{U'(x')}{U'(x)}\pi(x', x)$$

so the pricing kernel depends on the household's impatience to consume, attitude to bumpiness in consumption and the probability transition density between states. It's also pretty clear that the pricing kernel is intimately related to the stochastic discount factor defined above.

The price of a sure claim to a single unit of the consumption good is

$$\sum_{x'} q(x', x) = \sum_{x'} \beta \frac{U'(x')}{U'(x)}\pi(x', x)$$
$$= \mathsf{E}\left\{ \beta \frac{U'(x')}{U'(x)} \bigg| x \right\}$$

Indeed, the price of any complicated asset can be found in the same way. Let φ be the payoff function for an asset that pays $\varphi(x)$ in state x. Then the price of the asset, $p_\varphi(x)$, is

$$p_\varphi(x) = \sum_{x'} q(x', x)\varphi(x')$$
$$= \sum_{x'} \beta \frac{U'(x')}{U'(x)}\varphi(x')\pi(x', x)$$
$$= \mathsf{E}\left\{ \beta \frac{U'(x')}{U'(x)}\varphi(x') \bigg| x \right\}$$

Alternatively, we can write this in terms of a gross return $R_\varphi(x', x) = \frac{\varphi(x')}{p_\varphi(x)}$ as

$$1 = \mathsf{E}\left\{ \beta \frac{U'(x')}{U'(x)}R_\varphi(x', x) \bigg| x \right\}$$

Since shares in the tree have payoff $\varphi(x') = p(x') + x'$ we immediately have the consumption Euler

equation for returns on Lucas trees

$$1 = \mathsf{E}\left\{\beta\frac{U'(x')}{U'(x)}\frac{p(x') + x'}{p(x)}\,\bigg|\,x\right\}$$

In general, we can **synthesize** the price of other assets given the prices of all the contingent claims. Given the pricing kernel, the price of any payoff can be expressed as a linear combination of the form $p_\varphi = \mathsf{E}\{M\varphi\}$. The neat thing about this formula is that φ can describe the payoffs of an enormously complicated asset — say a particularly exotic option — but the price p_φ is straightforward to compute if we know how to calculate the pricing kernel.

A huge literature in economics and finance uses alternative assumptions about preferences (different specifications of U) and data on consumption, $c = x$, to estimate and test asset pricing models.

C. Consumption CAPM

We can obtain a capital asset pricing model (CAPM) relationship in the following way. Let $R_f(x)$ denote the gross return on a sure claim to a unit of consumption next period. The notation R_f is designed to indicate "risk free", but this is a little misleading in the sense that $R_f(x)$ is only risk free one-period-ahead. In general this return (equivalently, the price of a one-period-ahead risk free bond) does depend on the **current** state. More specifically, the return satisfies

$$1 = \mathsf{E}\left\{\beta\frac{U'(x')}{U'(x)}R_f(x)\,\bigg|\,x\right\}$$

and so the price is

$$\frac{1}{R_f(x)} = \mathsf{E}\left\{\beta\frac{U'(x')}{U'(x)}\,\bigg|\,x\right\}$$

Now consider a bona-fide risky asset with gross return $R_\varphi(x', x)$ contingent on next period's state. This return satisfies

$$
\begin{aligned}
1 &= \mathsf{E}\left\{\beta\frac{U'(x')}{U'(x)}R_\varphi(x', x)\,\bigg|\,x\right\} \\
&= \mathsf{E}\left\{\beta\frac{U'(x')}{U'(x)}\,\bigg|\,x\right\}\mathsf{E}\left\{R_\varphi(x', x)\,\big|\,x\right\} + \mathsf{Cov}\left\{\beta\frac{U'(x')}{U'(x)}, R_\varphi(x', x)\,\bigg|\,x\right\} \\
&= \mathsf{E}\left\{\frac{R_\varphi(x', x)}{R_f(x)}\,\bigg|\,x\right\} + \mathsf{Cov}\left\{\beta\frac{U'(x')}{U'(x)}, R_\varphi(x', x)\,\bigg|\,x\right\}
\end{aligned}
$$

or

$$E\left\{ \frac{R_\varphi(x', x) - R_f(x)}{R_f(x)} \,\middle|\, x \right\} = -\mathsf{Cov}\left\{ \beta \frac{U'(x')}{U'(x)}, R_\varphi(x', x) \,\middle|\, x \right\}$$

The expected **excess return** on an asset is negatively related to the covariance of that asset's gross return with the stochastic discount factor. This relationship is often written in time series notation as

$$E_t\left\{ \frac{R_{t+1}^\varphi - R_t^f}{R_t^f} \right\} = -\mathsf{Cov}_t\left\{ \beta \frac{U'(x_{t+1})}{U'(x_t)}, R_{t+1}^\varphi \right\}$$

The term on the left hand side is the expected excess return on the asset with payoffs φ over the risk free return (the **risk premium**). To see this, notice that $R_{t+1}^\varphi / R_t^f \cong 1 + r_{t+1}^\varphi - r_t^f$ where the little rs denote net returns. Hence

$$E_t\left\{ \frac{R_{t+1}^\varphi - R_t^f}{R_t^f} \right\} \cong E_t\{r_{t+1}^\varphi - r_t^f\} = E_t\{r_{t+1}^\varphi\} - r_t^f$$

The expected excess return is proportional to the **covariance** of the return of that asset with the stochastic discount factor. A high risk premium is demanded of an asset that is poor from an insurance perspective, i.e., an asset that has a high return only when the marginal utility of consumption is low (and hence consumption itself is relatively high) so that the covariance is a large negative number. Put differently, an asset is more risky the more its return co-varies negatively with the stochastic discount factor. Notice that in general the risk premium is time- and state-varying.

Chris Edmond

10 September 2004