Chapter 8

Intertemporal valuation

The equilibrium models we have seen so far are static models—they determine securities prices at a reference time on the basis of expected future revenues to which they provide a claim. They ignore a key fact: These revenues themselves depend on the securities' future prices. To account for this we need to consider an economy in which prices are set on markets that are open during consecutive periods over an unlimited horizon.

We will begin with the *representative agent* model, which will serve as a benchmark. The economy consists of a single agent who consumes the available resources at each point in time and in every state of nature. It is easy to define an (implicit) price for every good and *all* assets as the quantity of the numeraire the consumer is prepared to surrender at the margin in exchange for one unit of the good or security under consideration. This, in particular, will yield the dynamics of the spot curve and link it to the evolution of resources, as well as allowing us to examine the relationship between the forward rate curve, say for one year, and the spot curve that will prevail then.

The representative agent hypothesis is simplistic. Nonetheless, the conclusions this model yields apply when consumers are numerous but identical (in terms of tastes, resources, and expectations). They also partially extend to an economy with complete markets.

After having presented the model, we will examine the spot curve and its evolution along with the valuation of securities with risky returns and a finite lifespan. Subsequently, we will discuss one of the empirical paradoxes of the valuation of securities, the *equity premium puzzle*. We will conclude with an examination of assets with an infinite lifespan in connection with the phenomenon of bubbles.

1 The representative agent model

1.1 The economy

Consider an intertemporal economy with an infinite time horizon, in which time is discrete, $t = 0, \ldots, \infty$. There is a single consumer in this economy, whose lifespan is also infinite.

From the perspective of time 0, future consumption may be uncertain. We describe it with a consumption program, $\mathbf{c} = (\tilde{c}_t)_{t=0,\ldots,\infty}$, where $\tilde{c}_t$ is the (possibly stochastic) forecasted level of consumption at time t.

The consumer's preferences are defined over these consumption plans. We represent them with a von Neumann Morgenstern utility function U, which is intertemporally additive, $U(\mathbf{c}) = \sum_{t=0}^{\infty} \delta^t u(c_t)$, where u is an increasing, concave, and continuously differentiable function from $\mathbb{R}$ into $\mathbb{R}$. The *psychological discount factor*, δ, is positive and less than one, capturing the consumer's preference for the present. We denote i the *psychological discount rate*, which is connected to δ by the relationship:

$$\delta = \frac{1}{1+i}.$$

The utility level associated with $\mathbf{c}$ is given by

$$E[U(\mathbf{c})] = E_0 \left[\sum_{t=0}^{\infty} \delta^t u(c_t) \right],$$

where the mathematical expectation is conditional on the information available at time 0. Note that the consumer's tastes are stable over time: At time t preferences are represented with the same utility function as that given above, the only difference being that the mathematical expectation is taken conditional on the information available at time t, and not that available at time 0. This analysis can easily be extended to the case in which intertemporal utility is the sum of utility levels that may vary between dates, $E_0\left[\sum_{t=0}^{\infty}\delta^t u_t(c_t)\right]$.

The economy's resources, which our agent receives, are given by $\tilde{\omega}_t$, $t = 0, \ldots$ They are exogenous and fixed. The good cannot be stored, nor invested or transferred from one period to the next. Thus, the only resource allocation that is realizable is $\tilde{c}_t = \tilde{\omega}_t$, for all t.

Before proceeding, let us recall several definitions.

1.2 The spot curve: a review

Assume we are at time 0. Recall that the spot curve can easily be obtained from the prices of zero coupons (see Chapter 1). A zero coupon maturing at τ provides a claim on the delivery of one unit of the good at time τ, $\tau = 1, \ldots$

The *interest rate*, $r(\tau)$, on a loan taken at time 0 and due at τ, is defined from the price of the zero coupon, $q(\tau)$, maturing at τ, by

$$\frac{1}{[1+r(\tau)]^{\tau}} = q(\tau). \qquad (8.1) \quad \{\texttt{taux0}\}$$

All rates correspond to loan operations with a duration of a unit of time, and are all measured in terms of the same unit, regardless of their maturity.

The *spot curve* at time 0 is the curve $\tau \to r(\tau)$, for $\tau = 1, \ldots,$ that compares the yields, per unit of time, of loans with different maturities.

The forward interest rate at t and maturing at τ is defined as

$$\frac{1}{[1+f_t(\tau)]^{\tau}} = \frac{q(\tau+t)}{q(t)}. \qquad (8.2) \quad \{\texttt{terme}\}$$

By arbitrage, $f_t(\tau)$ is the prevailing rate at time 0 for loan operations contracted at time t with a maturity of τ periods.

The *spot curve at t* is the curve f_t: $\tau \to f_t(\tau)$.

Spot curves and forward rate curves are observable at present. Typically, the forward rate curve evolves over time. We denote the spot curve that will materialize at t, r_t: $\tau \to r_t(\tau)$ (with this notation, we have $r = r_0 = f_0$). Studies of rate dynamics focus specifically on the relationship between forward rates, $f_t(\tau)$, that are observable today and future rates $r_t(\tau)$ that will be realized at t. In other words, what information about the realization at time t of spot curve $r_t(\cdot)$ can we glean from current observations on the forward curve, t, $f_t(\cdot)$?

The representative agent model allows this type of question to be addressed once we have defined expectations on resources and their evolution.

In the next section, resources are sure (though not necessarily constant) and perfectly anticipated. Subsequently we will abandon that unrealistic assumption and assume that resources follow a Markovian process. This analysis will be conducted assuming that *expectations are rational*: Future events are drawn from a distribution that is consistent with the agent's expectation on the probabilities.

2 Risk-free aggregate resources

2.1 The interest rate curve and its evolution

In the representative agent model, calculating the prices of zero coupons and the associated rates—in relationship with the fundamental characteristics of the economy—is simple when resources are sure. Consider a consumer with preferences $E_0\left[\sum_{t=0}^{\infty} \delta^t u(c_t)\right]$, who can buy a zero coupon maturing at τ at a price $q(\tau)$. He adjusts his portfolio to satisfy the marginal condition

$$q(\tau)u'(c_0) = \delta^t E_0[u'(\tilde{c}_t)].$$

In an economy with a single agent, this consumption satisfies

$$\tilde{c}_\tau = \omega_\tau$$

at equilibrium. Thus, we obtain:

{p0c} $$q(\tau) = \frac{\delta^{\tau} u'(\omega_{\tau})}{u'(\omega_0)}. \tag{8.3}$$

In other words, the price of a unit of the good available at τ in terms of the good available today is equal to the marginal rate of substitution between these two periods.

The spot curve is thus entirely determined by the evolution of resources and preferences:

$$\frac{1 + r(\tau)}{1 + i} = \left[\frac{u'(\omega_0)}{u'(\omega_{\tau})}\right]^{1/\tau}. \tag{8.4}$$ {taux}

It immediately follows that:

When resources are expected to be constant over time the rates are all equal to the consumer's psychological discount rate. Rates are constant for all maturities and we say that the spot curve is flat.

When resources are sure but vary over time, we have

$$\omega_{\tau} > \omega_0 \Longleftrightarrow r(\tau) > i.$$

This property follows from the consumer's declining marginal utility. For him to accept a lower level of consumption at time 0 than at time τ, the price of the good at time τ must be sufficiently low, implying that the return is greater than the psychological discount factor.

Let us finally note that an increase in resources does not necessarily mean that the spot curve is increasing, as the following example illustrates.

Example 1

Assume that the function u is isoelastic:

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \quad \gamma > 0,\ \gamma \neq 1 \text{ or } u(c) = \ln c.$$

We obtain

$$1 + r(\tau) = (1 + i) \left(\frac{\omega_\tau}{\omega_0} \right)^{\gamma/\tau}.$$

If resources increase at a constant rate g, $\omega_\tau = \omega_0 g^\tau$, then the rate curve is flat and equal to $(1 + i)g^\gamma - 1$. ■

We will now examine the forward rate curve for t and the curve that will materialize at t. The forward rate at t maturing at τ, defined by (8.2), is here equal to

$$[1 + f_t(\tau)]^\tau = \frac{q(t)}{q(\tau + t)} = \frac{\delta^\tau u'(\omega_t)}{u'(\omega_{t+\tau})}. \tag{8.5}$$

The rate for maturities at τ that will be realized at t is given by the same expression as (8.4), but offset by t periods.

$$\frac{1 + r_t(\tau)}{1 + i} = \left[\frac{u'(\omega_t)}{u'(\omega_{t+\tau})} \right]^{1/\tau}.$$

If there is no uncertainty concerning aggregate resources, the forecast of ω_t is independent of the current date and equal to the value observed at t. Thus, we obtain:

If there is no uncertainty regarding aggregate resources, the forward rate curve for t: $\tau \to f_t(\tau)$ coincides with the spot rate that will materialize at t.

Recall that, by definition, the identity

$$\frac{1}{q(\tau)} = [1 + r(\tau)]^\tau = \frac{q(0)}{q(1)} \prod_{t=1}^{\tau-1} \frac{q(t)}{q(t+1)} = [1 + r(1)] \prod_{t=1}^{\tau-1} [1 + f_t(1)]$$

obtains. Moreover, since it is the case that in a model without uncertainty and with perfect foresight, $f_t(1) = r_t(1)$, the preceding equation can be written

$$\frac{1}{q(\tau)} = [1 + r(\tau)]^\tau = [1 + r(1)][1 + r_1(1)] \cdots [1 + r_{\tau-1}(1)].$$

The rate maturing at τ is the geometric mean of the successive short-term rates. Thus, a long-term rate exceeding today's short-term rate implies that the latter will rise and, in fact, eventually assume values greater than that of the current long-term rates.

2.2 The valuation of risky assets

Even if resources are certain, securities may pay stochastic dividends. The same valuation principle applies. We will model uncertainty using states of nature. During each period, we will find ourselves in one of a finite number of states of nature, e_t in $\mathcal{E}_t$, which determine the dividends to be distributed. We denote the probability, at time 0, that state e_t will materialize at time t with $\pi(e_t)$. Clearly, the values of risky assets at time 0 depend crucially on the expected distribution of future states. We continue to assume that expectations are *correct*: The agent has perfect knowledge of the distribution π.

As we saw in Chapter 2, a convenient procedure consists of evaluating the *state prices*, since they allow existing assets with a finite lifespan to be valued by arbitrage.

An Arrow-Debreu, or contingent, asset—defined as a security that provides a claim to the good at time τ if state e_τ materializes—can be associated with each date τ and each event e_τ. Its price, $q(e_\tau)$, is determined in a manner analogous to that of zero coupons. Let this security be tradable. The representative agent's optimization yields the first-order condition:

$$q(e_\tau)u'(c_0) = \delta^\tau \pi(e_\tau)u'[c(e_\tau)].$$

At equilibrium we must have

$$c_0 = \omega_0 \text{ and } c(e_\tau) = \omega_\tau.$$

This directly yields the state prices $q(e_\tau)$ in terms of the good available at time 0:

$$q(e_\tau) = \frac{\delta^\tau \pi(e_\tau)u'(\omega_\tau)}{u'(\omega_0)},$$

or, using the price of the zero coupon from (8.3),

$$q(e_\tau) = \pi(e_\tau)q(\tau).$$

In other words, the price contingent on a state at time τ equals the price of the zero coupon maturing at τ multiplied by the probability of that state. Consequently, the risk-adjusted probability is here identical to the objective probability. Of course, this is because there is no aggregate risk: The interest rate is determined by the (sure) marginal utilities of resources at the times in question, and under the von Neumann Morgenstern assumptions the risk-adjusted probability coincides with the (subjective and objective) probability of the occurrence of the states. Using the definition of the interest rate, the preceding equality can be written

$$q(e_\tau) = \frac{1}{[1 + r(\tau)]^\tau}\pi(e_\tau).$$

Now consider a security with a finite lifespan and paying stochastic dividends: The owner of one unit of this asset receives $d(e_\tau)$ at time τ if the state of nature e_τ materializes; dividends are nil beyond some time T (in the last section this assumption is abandoned). Using the principle of the absence of arbitrage opportunities, this security's price, p_0, expressed in terms of the good today, follows from the state prices:[1]

$$p_0 = \sum_{t=1}^{T} \sum_{e_t \in \mathcal{E}_t} q(e_t)d(e_t),$$

[1] We can also reason directly. The additional utility that this security contributes at the margin equals:

$$\sum_{t=1}^{T} \delta^t u'(\omega_t) \sum_{e_t \in \mathcal{E}_t} \pi(e_t)d(e_t).$$

Its price, p_0, expressed in today's good, is thus

$$p(e_0) = \frac{1}{u'(\omega_0)} \sum_{t=1}^{T} \delta^t u'(\omega_t) \sum_{e_t \in \mathcal{E}_t} \pi(e_t)d(e_t),$$

yielding the result we seek.

or

$$p_0 = \sum_{t=1}^{T} \frac{1}{[1+r(t)]^t} \left[\sum_{e_t \in \mathcal{E}_t} \pi(e_t) d(e_t) \right].$$

The probability of the occurrence of the states $\pi(e_t)$ is conditional on the information available at present. Analogously, at time τ, the price is a function of the realized state e_τ:

$$p(e_\tau) = \sum_{t=\tau+1}^{T} \frac{1}{[1+r_\tau(t)]^{t-\tau}} \sum_{e_t \in \mathcal{E}_t} \pi(e_t|e_\tau) d(e_t). \qquad (8.6)$$ {fondamental}

The following proposition captures these results.

Theorem 8.1

Assume that there is no uncertainty regarding the economy's total resources.

1. *The forward rate curve at t: $\tau \rightarrow f_t(\tau)$ coincides with the spot curve that will materialize at t.*

2. *The value of the risky security with a finite lifespan equals the discounted sum, using the term structure of interest rates, of the expected dividends it will yield conditional on the available information.*

■

3 Risky future resources

Though it serves a pedagogical purpose, the assumption that resources are risk-free is clearly too restrictive. The previous analysis can easily be extended to the case in which resources follow a dynamic stochastic process, provided it is known. We then obtain a stochastic model of the evolution over time of the interest rate curve.[2] From here on, resources $\tilde{\omega}_t$ are random and measurable with respect to the state e_t. Also, we can write $\tilde{\omega}_t = \omega(e_t)$.

[2]We here apply a simplified version of the model proposed by Lucas (1978).

3.1 The interest rate curve

We can apply a reasoning analogous to the preceding to the evaluation of zero coupons. At equilibrium, the agent expects to consume 0 at time τ:

$$\tilde{c}_\tau = \tilde{\omega}_\tau.$$

Asset prices must be such that there is no incentive to deviate from this consumption. A marginal increase in the unconditional consumption of the good at time τ, made possible by a zero coupon, increases future utility by $\delta^\tau E_0[u'(\tilde{\omega}_\tau)]$ while decreasing current utility by $q(\tau)u'(\omega_0)$. Thus:

$$q(\tau) = \frac{\delta^\tau E_0[u'(\tilde{\omega}_\tau)]}{u'(\omega_0)},$$

and, for the interest rate, using the psychological discount rate, i, defined by $\delta = 1/(1+i)$:

{ta}

$$\frac{1+r(\tau)}{1+i} = \left\{ \frac{u'(\omega_0)}{E_0[u'(\tilde{\omega}_\tau)]} \right\}^{1/\tau}. \tag{8.7}$$

In order to examine the impact of uncertainty on the spot curve, let us refer to a situation with no uncertainty in which the resource at time τ equals $\overline{\omega}_\tau$. Now consider an alternative in which resources are stochastic, but with the same mathematical expectation as in the reference situation:

$$E_0(\tilde{\omega}_\tau) = \overline{\omega}_\tau.$$

According to Equation (8.7), the interest rate with respect to the sure reference situation[3]

- increases if marginal utility is concave;
- remains unaltered if utility is quadratic;
- decreases if marginal utility is convex.

[3] The interest rate in the case of uncertainty is greater than in the case of certainty if $E_0[u'(\tilde{\omega}_\tau)] < u'[E_0(\tilde{\omega}_\tau)]$, which obtains if u' is concave.

The impact of uncertainty on the interest rate curve thus depends on very specific features of agents' preferences and is ambiguous even in the simplistic model of the representative agent. However, the third case is frequently deemed the most plausible. It captures what is called the *precautionary* effect: Faced with a higher future risk, the individual prefers to increase savings today and transfer wealth to the future. This leads to a decline in rates in order to equilibrate the market.

We can also look at the links between the forward prices of zero coupons and their prices in the future. The forward price for t, $q(\tau + t)/q(t)$, and the future price at t, $q_t(\tau)$, of a zero coupon maturing at τ are respectively given by:

$$\begin{aligned} \frac{q(\tau+t)}{q(t)} &= \delta^\tau \left\{ \frac{E_0[u'(\tilde{\omega}_{\tau+t})]}{E_0[u'(\tilde{\omega}_t)]} \right\}, \quad \text{and} \\ q_t(\tau) &= \delta^\tau \left\{ \frac{E_t[u'(\tilde{\omega}_{\tau+t})]}{u'(\omega_t)} \right\}. \end{aligned}$$

Seen from time 0, the future price, $q_t(\tau)$, is stochastic: It depends partly on the realization of wealth, ω_t, (through the denominator), and partly on information concerning future wealth that may change the conditional expectation in the numerator. Unless we assume risk neutrality, there is little likelihood that the forward price will be an unbiased estimator of the price in the future. In general, $E_0[q_t(\tau)]$ differs from the forward price $q(\tau+t)/q(t)$.

Now assume the simple (and unrealistic) situation in which there is no information between 0 and τ on the available resources at time $\tau + t$. In this event $E_t[u'(\tilde{\omega}_{\tau+t})]$, which is random from the perspective of time 0 *a priori*, is constant and equal to its current value $E_0[u'(\tilde{\omega}_{\tau+t})]$. Thus:

$$E_0[q_t(\tau)] = \delta^\tau E_0 \left\{ \frac{E_0[u'(\tilde{\omega}_{\tau+t})]}{u'(\omega_t)} \right\} = \delta^\tau E_0[u'(\tilde{\omega}_{\tau+t})] E_0 \left[\frac{1}{u'(\tilde{\omega}_t)} \right].$$

The inequality[4]

$$\frac{1}{E_0[u'(\tilde{\omega}_t)]} \leq E_0 \left[\frac{1}{u'(\tilde{\omega}_t)} \right]$$

[4]Since the inverse function is convex, this follows form Jensen's inequality.

implies

$$E_0(q_t) \leq \frac{q(\tau + t)}{q(t)}.$$

Thus, in the absence of an information effect, the future price will, on average, be below the forward price. This is interpreted as follows. Consider an investor wishing to take out a loan in t periods maturing at $t + \tau$. The forward price at t fixes the cost of the loan today and thus allows her to hedge against variations in this price at time t. If this price is below that expected to prevail at t, and if waiting will not bring the benefit of further information, she has every reason to commit to the operation today. This is not true if waiting an additional period will yield a clearer picture of future needs.

This type of analysis is often conducted on the spot rate, as in the example below. We thus wish to know whether there is a systematic difference between the forward rate curve and the expectation of curves that will materialize.

3.2 An example

We revert to the situation in Example 1, where preferences are of the form

$$\frac{1}{1-\gamma} E\left[\sum_{t=0}^{\infty} \delta^t (c_t)^{1-\gamma}\right],$$

with γ strictly positive and not equal to one.

We assume that the distribution of resources (i.e. national output) is lognormal: The joint distribution of the $\log(\tilde{\omega}_t)_{t=1,\ldots,T}$-s is normal for all T.

Taking logs in (8.7), we obtain

$$\log[1 + r(\tau)] = \log(1 + i) - \frac{1}{\tau}\left[\log E_0(\tilde{\omega}_\tau^{-\gamma}) - \log(\omega_0^{-\gamma})\right]$$

Note that the distribution of $\tilde{\omega}_t^{-\gamma}$ is lognormal: $\log \tilde{\omega}_t^{-\gamma}$ is normal with expectation $-\gamma E(\log \omega_t)$ and variance $\gamma^2 \mathrm{var}(\log \tilde{\omega}_t)$. Now, if X is distributed lognormally:

$$\log E(X) = E(\log X) + \frac{1}{2}\text{var}(\log X).$$

From this we can derive the following expression for the future interest rate at τ:

$$\log[1 + r(\tau)] = \log(1 + i) + \frac{\gamma}{\tau}\left\{E_0\left[\log(\tilde{\omega}_\tau/\omega_0)\right] - \frac{\gamma}{2}\text{var}_0(\log\tilde{\omega}_\tau)\right\}, \quad (8.8)$$ {ltau}

which only depends on the values of $E_0(\log\tilde{\omega}_\tau)$ and $\text{var}_0(\log\tilde{\omega}_\tau)$.

Owing to its dual role in the representative agent's preferences, the parameter γ has two effects on the interest rate curve.

As in the case with certainty, it indicates the degree of intertemporal substitution the individual will accept. If he expects an increase in wealth, $E_0[\log(\tilde{\omega}_\tau/\omega_0)] > 0$, he will wish to borrow, and all the more so as γ increases, reflecting a higher interest rate. But γ also measures relative risk aversion at a given point in time—and so the rate depends on the future variance, $\text{var}_0(\log\tilde{\omega}_\tau)$. Here, rates are decreasing with this variance. This is the precautionary effect mentioned above, which is stronger as γ rises.

We can study the relationship between the spot curve maturing at 1 and the curve that will materialize in one period. According to (8.2), we have

$$[1 + f_1(\tau)]^\tau = \frac{[1 + r(\tau+1)]^{\tau+1}}{1 + r(1)}$$

which, using (8.8) yields

$$\log[1 + f_1(\tau)] =$$
$$\log(1 + i) + \frac{\gamma}{\tau}\left\{E_0(\log\tilde{\omega}_{\tau+1} - \log\tilde{\omega}_1) - \frac{\gamma}{2}[\text{var}_0(\log\tilde{\omega}_{\tau+1}) - \text{var}_0(\log\tilde{\omega}_1)]\right\}.$$

The spot curve that will materialize at $t = 1$ is given by:

$$\log[1 + r_1(\tau)] = \log(1 + i) + \frac{\gamma}{\tau}\left\{E_1\left[\log(\tilde{\omega}_{\tau+1}/\omega_1)\right] - \frac{\gamma}{2}\text{var}_1(\log\tilde{\omega}_{\tau+1})\right\}.$$

Seen from time 0, it is random, depending on the information available at time 1, which clearly includes ω_1. To examine the difference between forward rate curves and the expectation of curves that will materialize in the future,[5] notice that

$$\log[1 + f_1(\tau)] - E_0\{\log[1 + r_1(\tau)]\} = \\ -\frac{\gamma^2}{2\tau}\left\{[\mathrm{var}_0(\log\tilde{\omega}_{\tau+1}) - \mathrm{var}_1(\log\tilde{\omega}_{\tau+1})] - \mathrm{var}_0(\log\tilde{\omega}_1)\right\}.$$

We once again find the two previously mentioned effects: The term between square brackets captures the impact of new information at time 1 on future resources, and the second the impact of realized wealth at time 1. There is no reason why this expression should be equal to zero.

Thus, there is generally a bias, the sign of which depends on the precision of the information that will be available at maturity. If, for example, this information is not good quality (the variances of $\log\tilde{\omega}_{\tau+1}$ are nearly equal at time 0 and time 1), the bias will be positive. On average, the forward rate will be higher than the future rate.

Let us develop these results for when resources follow a first-order autoregressive process:

$$\log\omega_t = g + \rho\log\omega_{t-1} + \varepsilon_t,$$

where g is a real number determining the long-term level of resources, ρ falls in the interval $]-1, +1[$, and the ε_t-s are independent normal variables with mean zero and variance σ^2.

A simple calculation yields

$$\begin{aligned} E_0(\log\omega_t) &= g\frac{1-\rho^t}{1-\rho} + \rho^t\log(\omega_0) \quad \text{and} \\ \mathrm{var}_0(\log\omega_t) &= \sigma^2\frac{1-\rho^{t+2}}{1-\rho^2}. \end{aligned}$$

[5] To be rigorous, we will work with $\log(1+r)$, and not r. We obtain analogous results when we directly study the rates.

If resources are independent and identically distributed, $\rho = 0$, and

$$\log[1 + f_1(\tau)] - E_0 \{\log[1 + r_1(\tau)]\} = \frac{\gamma^2}{2\tau}\sigma^2.$$

The forward rate is consistently above the expectation of the rate that will materialize, and this bias declines with the maturity. For non-nil ρ:

$$\log[1 + f_1(\tau)] - E_0 \{\log[1 + r_1(\tau)]\} = \frac{\gamma^2}{2\tau}\frac{\sigma^2}{1-\rho^2}(1 - \rho^3 - \rho^{\tau+2} + \rho^{\tau+3}).$$

Changes to the shape of the spot curve are not so simple any more. For the usual, positive, values of ρ, the increase in rates for maturities that are near is less than for more remote ones. The expected future spot curve has a steeper slope than the current spot curve.

3.3 The dynamics of securities prices

As previously, it is convenient to begin by computing the prices of Arrow-Debreu assets:

$$q(e_\tau) = \frac{\delta^\tau \pi(e_\tau) u'[\omega(e_\tau)]}{u'[\omega(e_0)]},$$

since we can calculate the price of any other security with a finite lifespan from the state prices:

$$p_0 = \sum_{t=1}^{T} \sum_{e_t \in \mathcal{E}_t} q(e_t) d(e_t).$$

This equality is often written in different ways, using the *risk-adjusted probability* or the *stochastic discount factor*, or by incorporating yields.

Let us begin with the first formulation. By definition, the sum of the state prices for time τ is equal to the price of a zero coupon for the same date:

$$\sum_{e_\tau} q(e_\tau) = q_\tau.$$

Dividing the state prices by the price of the zero coupon, we obtain a probability. More precisely, we can write

$$q(e_\tau) = \frac{\delta^\tau E_0[u'(\tilde{\omega}_\tau)]}{u'[\omega(e_0)])} \frac{\pi(e_\tau)u'[\omega(e_\tau)]}{E_0[u'(\tilde{\omega}_\tau)]} = \frac{1}{[1+r(\tau)]^\tau}\pi^*(e_\tau)$$

where:

$$\pi^*(e_\tau) = \frac{\pi(e_\tau)u'[\omega(e_\tau)]}{E_0[u'(\tilde{\omega}_\tau)]}.$$

By construction, π^* is, in fact, a measure or probability. It is equal to the objective probability if expectations on resources are certain or if the individual is risk neutral. The price of an Arrow-Debreu asset as a function of interest rates has exactly the same expression as in the previous section when the objective probability is replaced by the risk-adjusted probability. It directly follows that:

$$p_0 = \sum_{t=1}^{T} \frac{1}{[1+r(t)]^t} \left[\sum_{e_t \in \mathcal{E}_t} \pi^*(e_t) d(e_t) \right].$$

or:

The value of a risky asset is equal to the discounted sum, using the term structure of interest rates, of the mathematical expectation of the dividends it will yield, computed with the risk-adjusted probability.

These equations can also be written in terms of yields. Let $R_k(e_t, e_{t+1})$ represent the gross yield of asset k at time t in state e_t. If the state at $t+1$ is e_{t+1}, then

$$R_k(e_t, e_{t+1}) = \frac{p_k(e_{t+1}) + d(e_{t+1})}{p(e_t)}.$$

The fundamental price-setting relationship is thus

{equity} $$1 = E\left[\delta \frac{u'(\tilde{\omega}_{t+1})}{u'(\omega_t)} \tilde{R}_k \,\middle|\, e_t \right]. \tag{8.9}$$

Let $\tilde{\delta}_{t+1}$ be the *stochastic discount factor*[6] defined by

[6] This formulation is often used in econometric tests, where the "true," historical, probability is preferred.

$$\tilde{\delta}_{t+1} = \delta \frac{u'(\tilde{\omega}_{t+1})}{u'(\omega_t)}.$$

According to (8.9), if consumption were risk-free, all securities should have the same expected yield, equal to $1/E(\tilde{\delta})$, regardless of their variance. Thus, there should not be a risk premium. However, consumption varies over time—and these variations cannot be forecast perfectly.

It is within this variation, and in the correlation between movements in consumption and asset yields, that the risk premium arises. To see this, let us rewrite (8.9)

$$1 = E_t(\tilde{\delta}_{t+1})E_t(\tilde{R}_k) + \mathrm{cov}_t(\tilde{\delta}_{t+1}, \tilde{R}_k).$$

Applying the formula to the risk-free asset (indexed 0) and the stocks (security 1), and eliminating element by element, we find:

$$E_t(\tilde{R}_1) - R_0 = -\frac{\mathrm{cov}_t(\tilde{\delta}_{t+1}, \tilde{R}_1)}{E_t(\tilde{\delta}_{t+1})}.$$

It is the differences between the covariances of the yields and the stochastic discount rate that underlies the expected yields in our model. The correlation is essential: The expected yield of a security whose yield is uncorrelated with forecasted consumption, sometimes called a zero-β, is always equal to $1/E(\tilde{\delta})$.

Example 2

Let us illustrate with the isoelastic utility function:

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \quad \gamma \geq 0.$$

Assume that all variables are lognormal. Taking the log of (8.9), we obtain:[7]

$$0 = \log\delta + E(\log R_k) - \gamma E\left[\log\left(\frac{\tilde{c}_{t+1}}{c_t}\right)\right] + \frac{1}{2}\mathrm{var}\left[\log\tilde{R}_k - \gamma\log\left(\frac{\tilde{c}_{t+1}}{c_t}\right)\right],$$

We can rewrite this letting σ_k represent the standard error of the log of security k's yield, σ_c the standard error of the growth rate of consumption, and σ_{kc}, the covariance of the log of k's yield with the log of the growth rate of consumption:

{prime} $$0 = \log\delta + E(\log R_k) - \gamma E\left[\log\left(\frac{\tilde{c}_{t+1}}{c_t}\right)\right] + \frac{1}{2}\left(\sigma_k^2 + \gamma^2\sigma_c^2 - 2\gamma\sigma_{kc}\right). \quad (8.10)$$

Applying the formula to the risk-free asset (indexed 0) and the stock (security 1), and eliminating element by element, we find:

{equity1} $$E(\log R_1) - \log R_0 + \frac{1}{2}\sigma_1^2 = \gamma\sigma_{1c}. \quad (8.11)$$

which yields, as the risk premium:

$$E(R_1) - R_0 = R_0\,\exp(\gamma\sigma_{1c}).$$

■

4 Empirical Verification

It is reasonable to seek to test the foregoing equations. We know the real yields of bonds and shares over long periods. For example, the log of the annual yield of the Standard and Poor's 500 index from 1889 to 1994 shows a mean of 6.0 per cent, and a standard error of 16.7 per cent. The same calculation applied to six-month commercial debt obligations—the best approximation to the risk-free rate available for long periods—reveals a mean

[7] We are again using the fact that, if X is lognormal:

$$\log E(X) = E(\log X) + \frac{1}{2}\mathrm{var}(\log X).$$

of 1.8 per cent. Is the 4.2 per cent average premium on shares justified by the risks that they impose on stockholders. If we use the representative agent approach, the risk premium depends on the shape of the utility function and the consumption process.

4.1 Isoelastic utilities

The first empirical studies were conducted by Mehra and Prescott on U.S. data. Their results, which were subsequently confirmed, led to what is known as the *equity premium puzzle*. They assumed an isoelastic utility function as in the previous examples, and sought to verify whether (8.10) and (8.11) are compatible with the orders of magnitude suggested by the statistics. We have already described the yields of stocks and bonds. It remains to specify the evolution of consumption. In practice, we consider the purchases of nondurables and services. [8] For the United States, the mean of the log of the ratio c_{t+1}/c_t over the period 1889–1994 was 1.7 per cent and the standard error was 3.3 per cent. We can use this historical data, if we assume that intertemporal variability throughout the past century equaled the conditional variance during the period (bear in mind that we are really only looking at orders of magnitude: more precise calculations have revealed that the phenomenon persists when the evolution of the variance over time is accounted for). The first element of (8.11), evaluated with the aforementioned data, yields a little over 6 per cent (the risk-free rate, 1.8 per cent, is barely greater than the mean of the variance of the securities' yields, $(0.167)^2/2$). Thus, the key term is the covariance between the rate of growth of consumption and the yield on the stock exchange. A high correlation (agents consume more when the market is high) indicates that the stock exchange is poorly suited for providing investors with insurance to smooth shocks to their consumption profile: Investors demand a risk premium that increases with the correlation. Indeed, over the period studies,

[8] Purchases of durables fluctuate considerably with the business cycle, styles, and expected price rises. In all likelihood, the services which these goods provide to consumers, which enter into the utility function but are not directly observable by statisticians, evolve much more smoothly than the purchases.

the correlation is high (0.49), resulting in a covariance of 0.27 per cent. Nonetheless, it remains too low: Equation (8.11) assumes a coefficient γ greater than 20.

This value is highly implausible. Experiments that deal with risky choices tend to yield values for γ that are below 4 or 5. Given these values, an agent would purchase high-yield securities and sell sure assets.

There is a further problem related to the risk-free interest rate. Equation (8.10) applied to the risk-free rate allows estimation of the rate of preference for the present when γ and the growth rate of consumption are known:

$$\log \delta = -\log R_0 + \gamma E\left[\log\left(\frac{c_{t+1}}{c_t}\right)\right] - \frac{\gamma^2 \sigma_c^2}{2}.$$

Since the mean rate of growth of consumption is approximately equal to the risk-free rate (1.7 or 1.8 per cent), and the standard error of consumption is 3.3 per cent (variance 0.001), we see that a reasonable value for γ, between 2 and 10, implies a psychological discount factor greater than 1. This is the *risk-free rate paradox*. Given the low interest rate, our theoretical consumers (with $\delta < 1$) will wish to incur debt and consume more today than tomorrow. Only if future consumption is very uncertain and/or they are highly risk averse can the mean growth rate of their consumption be equal to the risk-free interest rate.

4.2 Beyond the representative agent

The preceding analysis requires many ancillary assumptions: the choice of the form of the utility function, lognormality of yields and consumption growth, and a representative agent. It is reasonable to wonder whether these assumptions may lie at the root of the incompatibility between the theory and the observations, or whether the whole portfolio choice model must be scrapped.

We will here draw on the approach in Hansen and Jagannathan. The central idea is to retain the rationality of agents while assuming as little as possible regarding the unobservable, in particular the marginal rate of substitution and the stochastic discount factor. We begin with observations on securities yields and seek to derive a range of admissible values for mathematical expectations on the variance (or volatility) of the stochastic discount factor of an investor operating on the markets for these assets.

Applying the same reasoning as above, a marginal investment in security k that leaves the investor's expected utility unchanged yields

$$u'[\omega(e_t)] = \delta E\left\{u'[\omega(e_{t+1})[R_k(e_t, e_{t+1})|e_t]\right\}.$$

We point out that this equality, which we found in the context of the representative agent economy (where consumption is, in fact, macroeconomic consumption) also applies to any individual market participant, provided purchases and sales of k are not constrained (in this case aggregate values for consumption and utility are replaced with the corresponding amounts for individual agents).

Let us identify some forms for the stochastic discount factor that are compatible with the statistical observation on yields. We can write Equation (8.9) in terms of the expectation and covariance:

$$1 = E(R_k)E(\tilde{\delta}) + \text{cov}(R_k, \tilde{\delta}).$$

Stacking these equalities for all securities yields the following equation, in matrix notation:

$$\mathbb{1}_K = E(\tilde{\delta})\ E(R) + \text{cov}(R, \tilde{\delta}). \qquad (8.12)$$ {margin}

From this we can derive a lower bound for $\text{var}(\tilde{\delta})$ given a value of $E(\tilde{\delta})$. If $\tilde{\delta}$ is in the L^2-space of measurable random variables defined on the fundamental space, then $\tilde{\delta}$ can always be decomposed into its projection $x'\tilde{R}$ onto the subspace generated by the securities' yields and an orthogonal element of that projection. We have, surely: $\text{var}(\tilde{\delta}) \geq \text{var}(x'\tilde{R})$. If we denote the yields' variance-covariance matrix Γ, this inequality becomes

$$\text{var}(\tilde{\delta}) \geq x'\Gamma x,$$

and the first-order condition (8.12) is written

$$\mathbb{1}_K = E(\tilde{\delta})\ E(R) + \Gamma x.$$

Whence, taking x out of the f.o.c. and substituting it into the inequality:

$$\mathrm{var}(\tilde{\delta}) \geq x'\Gamma x = [\mathbb{1}_K - E(\tilde{\delta})\ E(R)]'\Gamma^{-1}[\mathbb{1}_K - E(\tilde{\delta})\ E(R)],$$

or:

$$\mathrm{var}(\tilde{\delta}) \geq [E(R')\Gamma^{-1}E(R)][E(\tilde{\delta})]^2 - 2[\mathbb{1}_K'\Gamma^{-1}E(R)]E(\tilde{\delta}) + [\mathbb{1}_K'\Gamma^{-1}\mathbb{1}_K].$$

Thus, we have identified a lower bound on the volatility of the stochastic discount factor. When this inequality holds as a strict equality, the set of all points $1/E(\tilde{\delta})$ as a function of $\sigma_{\tilde{\delta}}/E(\tilde{\delta})$ describes a hyperbola (just as in the case of the mean-variance efficient portfolios we saw in Chapter 4, Section 2.2). A stochastic discount factor is admissible if, when the values of $[\sigma_{\tilde{\delta}}/E(\tilde{\delta}),\ 1/E(\tilde{\delta})]$ have been inserted into the efficient portfolios, the representative point yields a standard error that is greater than the minimum and thus located to the right of the hyperbola.

The lower bound improves (i.e. becomes more binding) as the variety of securities under consideration increases. If a single security were to represent the entire market (security 1), the hyperbola would collapse into its two asymptotes and the inequality reduce to

$$\frac{\mathrm{var}(\tilde{\delta})}{[E(\tilde{\delta})]^2} \geq \frac{[E(R_1) - 1/E(\tilde{\delta})]^2}{(\sigma_1)^2}.$$

If the expectation of the stochastic discount factor is equal to the reciprocal of the expected yield of the asset, then the volatility is unbounded. Otherwise, there is a strictly positive lower bound to this discount factor's volatility or, more precisely, to the ratio of its standard error to its mathematical expectation, which increases with the distance between $1/E(R_1)$ and $E(\tilde{\delta})$ and decreases with σ_1.

When we introduce a representative set of financial assets we are again confronted with the problems from the previous section: With an isoelastic utility function and data on fluctuations in consumption, Hansen and Jagannathan demonstrate that the stochastic discount factor does not satisfy the inequality given above.

5 Fundamental value and bubbles

Let us return to the case with certainty in resources. In the framework of the current model, when expectations are rational, Equation (8.6) is often called the *fundamental value* of the security. Its observed value may deviate from the fundamental value if expectations are incorrect. However, this may also occur when the asset's lifespan is not finite, even if expectations are assumed to be correct. In this case we are observing a *bubble*.

To illustrate this vocabulary and simultaneously examine the dynamics of asset prices, we only need to look at the case of resources that are constant and sure over time. Then the interest rate curve is flat and all rates, at all times, are equal to the psychological discount rate i (see Section 2). Consider a dividend-paying asset linked to the evolution of the state of nature e_t. We can write:

$$p(e_0) = \sum_{t=1}^{\infty} \frac{1}{(1+i)^t} \sum_{e_t \in \mathcal{E}_t} \pi(e_t) d(e_t),$$

and similarly for time τ. If the state e_τ materializes, the price (after the dividend has been paid) will be

$$p(e_\tau) = \sum_{t=\tau+1}^{\infty} \frac{1}{(1+i)^{t-\tau}} \sum_{e_t \in \mathcal{E}_t} \pi(e_t | e_\tau) d(e_t).$$

A fundamental property of this process results from the relationship between two successive prices, which we can see by using the identity of the conditional probabilities for $\tau \geq 0$:

$$p(e_\tau) = \frac{1}{1+i} \sum_{e_{\tau+1} \in \mathcal{E}_{\tau+1}} \pi(e_{\tau+1} | e_\tau) [p(e_{\tau+1}) + d(e_{\tau+1})].$$

The price of the asset today equals the discounted value of the mathematical expectation of its resale price tomorrow, increased by the dividends. Indeed, this is the only property that can be directly inferred from the rationality of the agent's behavior. This is often assimilated into the notion of the *efficiency* of financial markets: Markets completely reflect all available information. However, in our model they do not capture the fundamental value of a security having an infinite lifespan. This is easy to see, for example if we consider an asset that does not yield any dividends regardless of the state of nature. This is a property of money. Its fundamental value is equal to zero. However, the equation

$$p(e_\tau) = \frac{1}{(1+i)} E[p(e_{\tau+1})|e_\tau]$$

is solved by many series of prices. If we limit ourselves to non-stochastic solutions, they are given by

$$p(\tau) = (1+i)^\tau p(0),$$

for any nonnegative $p(0)$. *A priori*, all these solutions are acceptable. The null solution is the only one that, like the security and the resources, is stationary. Non-stationary solutions are called *bubbles*. The price of the security today is only justified by the fact that it is expected to increase at the rate i tomorrow, and so on into the future. This price rises exponentially until expectations collapse and... the bubble bursts!

Bibliographical Note

The representative agent model was introduced by Lucas in 1978 and provides the basis for many asset valuation models. In 1986, Campbell contributed a useful extension. We have focused our description on consistency relationships between the stochastic features of securities' yields and movements in consumption. Mehra and Prescott identified and popularized the *equity premium puzzle*. Hansen and Jagannathan proposed a nonparametric test for it.

We have been working with yields rather than prices. The earliest empirical studies sought to test for equality of the price of a security with the discounted expectation of the future income it will provide its owner. Shiller (1981) rejected that property, called the *efficiency of markets*: For plausible discount rates, prices variations appear too great compared to the variance in dividends. The literature subsequently turned its attention to the choice of discount rate, which led to the CCAPM and the stochastic discount factor.

Campbell, J.Y. (1986) : "Bond and Stock Returns in a Simple Exchange Model", *Quarterly Journal of Economics*, 101, 785-803.

Hansen, L., and R. Jagannathan (1991) : "Restrictions on Intertemporal Rates of Substitution Implied by Asset Returns", *Journal of Political Economy*, 99(2), 225-262.

Lucas, R.E., (1978) : "Asset Prices in an Exchange Economy", *Econometrica*, 46(6), 1429-1435.

Mehra, R. and E.C. Prescott (1985) : "The Equity Premium : A Puzzle", *Journal of Monetary Economics*, 15, 145-151.

Shiller, R. J. (1981) : "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends ?", *American Economic Review*, 71, 421-436.

Exercises

Exercise 18: *Recursive preferences*

1. Consider preferences represented by

$$\frac{1}{1-\gamma}\left[c_1^{1-\gamma} + \delta E\left(c_2^{1-\gamma}\right)\right]$$

   We assume that there is no randomness and are interested in the choice of a consumer at two dates given the budget constraint

$$p_1c_1 + p_2c_2 = p_1\omega_1 + p_2\omega_2.$$

(we have $p_2/p_1 = 1/(1+r)$). Show that the ratio of consumption at the two times, c_2/c_1, only depends on the price ratio p_2/p_1.

2. We call the relative variation in the consumption ratio caused by a change in the price ratio the intertemporal elasticity of substitution

$$e = -\frac{\partial(c_2/c_1)}{c_2/c_1}\frac{\partial(p_2/p_1)}{p_2/p_1}.$$

In other words, if the price ratio increases by one per cent, then consumption will fall by e per cent. Compute this elasticity. Compare it to the risk aversion.

3. Now assume that preferences are given by

$$\frac{1}{1-\gamma}\left[c_1^{1-\phi} + \delta E(c_2^{1-\gamma})^{\frac{1-\phi}{1-\gamma}}\right]$$

Answer the same questions again. Draw conclusions.

■

Exercise 19: *Growth of the spot curve*

We seek a very simple model to study the growth of the spot curve. We are particularly interested in the impact of expected future rates and the supply of securities.

Consider two securities:

- a zero coupon with a "close" maturity of 1 and costing $\frac{1}{1+r}$ today, at $t = 0$,
- a zero coupon with a "remote" maturity of 2 and costing p_0 today. Let $\tilde{p}_1^a$ represent the expected value of its price at time 1. This expectation is assumed identical for all investors.

There are n investors at time $t = 0$. Investor i, $i = 1, \ldots n$, has wealth ω_i and preferences over wealth at time 1. They are represented by a mean-variance function. i seeks to maximize

$$E(\tilde{c}) - \frac{a_i}{2} var(\tilde{c})$$

subject to the budget constraint.

1. Compute the quantity of zero coupons demanded by investor i.

2. The government supplies these zero coupons: The zero coupon maturing at 2 is fixed at M, while the one maturing at 1 is adjusted to make the short-term rate equal to r.

   a) Derive the equilibrium price p_0 as a function of the anticipated variance of the expectation $\tilde{p}_1^a$ and the supply of securities M.

   b) Comment.

3. a) Express the relationship linking p_0 to the "remote" rate R with maturity 2, and that linking $\tilde{p}_1^a$ to the "close" rate $\tilde{r}_1^a$ expected at time 1. Derive the equilibrium relationship between the rates.

   b) Using the fact that the rates are small relative to 1, linearize the foregoing expression. Under what conditions will the "remote" rate be greater than the "close" rate? Can the spot curve be increasing?

■