Chapter 7

Uncertainty

Many aspects of economic life are not fully predictable. For example, at the aggregate level, technologies and policies can change in unpredictable ways. Life is even more uncertain at the individual level due to the unpredictability of income, health, and other events. We now introduce the main techniques that macroeconomists use to incorporate uncertainty into our analysis of the economy.

This chapter covers several issues. First, we introduce analytical tools that are helpful in analyzing stochastic economies. These tools include mathematical concepts related to stochastic processes as well as economic concepts related to decision making under uncertainty. We then use these tools to analyze the social planner's problem for a version of the neoclassical growth model with stochastic productivity. This important example is introduced in Section 7.3. We then discuss how agents trade with each other in an uncertain economic environment and present the competitive equilibrium of the stochastic neoclassical growth model. Finally, we briefly present an environment with incomplete insurance markets.

7.1 Stochastic processes

A **stochastic process** is a collection of random variables indexed by time. Suppose at each date t there is a random outcome X_t. The collection of these random variables $\{X_t : t \in \mathcal{T}\}$ is a stochastic process, where the set $\mathcal{T}$ is the span of time we are interested in. When modeling an economy with uncertainty, we typically assume that some fundamental features of the economy follow an exogenous stochastic process. For example, we often assume that the level of productivity fluctuates over time and is modeled as an exogenous stochastic process. This randomness in fundamentals generates randomness in endogenous variables. Our economic model determines the stochastic process these endogenous variables follow. We will now introduce some general properties of stochastic processes before turning to a few of the main types of processes used by macroeconomists.

7.1.1 Properties of stochastic processes

When working with stochastic processes we often need to take expectations of them. One way to think about the expectation of a random variable X_t that is part of a stochastic process is to imagine multiple realizations of the entire stochastic process $\left\{X_t^{(i)} : t \in \mathcal{T}, i \in 1, \cdots, I\right\}$, where i indexes the different realizations. Taking the expectation across i (i.e., averaging across realizations using probability weights) then gives the unconditional expectation. A related concept is a conditional expectation given information up to some point in time. For example we could suppose we have observed the realization of the stochastic process up to date t and then imagine different possible realizations for $t+1$ and later dates. Taking the expectation of X_{t+1} across these different realizations is then an expectation conditional on information through date t. We often use $\mathbb{E}_t[X_{t+1}]$ to denote such a conditional expectation. To be clear, $\mathbb{E}_t[X_{t+1}]$ means the expectation of X_{t+1} conditional on all information available at date t, not just the realization of X_t observed at date t. A very useful property of expectations is the **law of iterated expectations**, which for any dates $t < s < \tau$ says

$$\mathbb{E}_t\left[\mathbb{E}_s\left[X_\tau\right]\right] = \mathbb{E}_t\left[X_\tau\right].$$

A proof of the law of iterated expectations appears in the appendix.

In addition to expectations, we are often interested in the second moments of a stochastic process. These second moments are captured by the autocovariances. The j-th autocovariance is given by $\mathbb{E}\left[(X_t - \mu_t)(X_{t-j} - \mu_{t-j})\right]$ where μ_t is the unconditional expectation of X_t.

Many economic theories assume or imply stochastic processes that are stationary. A process is **covariance stationary** if neither the unconditional expectations nor the autocovariances depend on time t. In practice, stationarity means that the consequences of a shock to the process eventually fade. If this were not the case, as time goes by, the effects of all the shocks that have occurred would accumulate and the distribution of X_t would change, e.g., become more and more dispersed, as t increases.

In most cases we look at, a stationary process will be **ergodic**. For an ergodic process, observing a long time series allows us to understand the distribution of the stochastic process. For example, the unconditional expectation of X_t is the expectation across different possible realizations of X_t. For an ergodic process, an average of a long time series $(1/T)\sum_{t=1}^{T} X_t$ will converge to $\mathbb{E}[X_t]$ as $T \to \infty$.[1] Ergodicity is useful because in practice we often have data on a single long time series.

A stochastic process is (first-order) **Markov** if its current value summarizes all the available information that is useful for predicting its future realizations. For example, the conditional distribution of X_{t+1} conditional on X_t is the same as the conditional distribution of X_{t+1} conditional on $\{X_\tau\}_{\tau \le t}$, in which case knowledge of X_τ for $\tau < t$ does not add any relevant information beyond that contained in X_t. More formally, a Markov process satisfies $\Pr[X_{t+k} = x | X_\tau \forall \tau \le t] = \Pr[X_{t+k} = x | X_t]$ $\forall t, k \ge 0$. In dynamic programming, it is convenient if a single state variable summarizes the process. For this reason, macroeconomists often work with Markov processes.

Over time, a deterministic sequence might converge to a particular value. A stochastic process, on the other hand, might continually be subject to random shocks and therefore not

[1] See Hamilton (1994) for the conditions under which a stationary process is ergodic.

settle down to a particular value but nevertheless there is a sense in which it can converge. A stationary distribution (or invariant distribution) $\bar{\pi}(X)$ of a Markov process X has the property that if X_t has distribution $\bar{\pi}(X)$ then X_{t+1} also has distribution $\bar{\pi}(X)$.

7.1.2 Markov chains

Let x_t be a random variable that takes on values in the discrete set $\mathcal{X} = \{\bar{x}_1, \bar{x}_2, \cdots, \bar{x}_N\}$. If x_t follows a Markov process, the probability that a particular x_{t+1} occurs depends only on x_t and not on earlier values. A Markov process with a discrete state space is called a **Markov chain**. As x_t takes on discrete values, we can summarize the process in terms of a **transition matrix**. Suppose the probability of moving from state i to j is given by P_{ij}. We can then collect these transition probabilities in a matrix P. Row i represents the probabilities of states 1 through N occurring next period conditional on the current state being i. As exactly one of these states will occur, these probabilities must sum to one. By this logic, each row of P must sum to one.

In addition to the transition matrix, a full description of the stochastic process also requires knowing the initial distribution over the states. We will represent this distribution as a $1 \times N$ vector of probabilities π_0 such that the ith element of π_0 gives $\Pr[x_0 = \bar{x}_i]$.

Given the initial distribution π_0, the transition matrix determines the probability distributions for all x_t for $t \geq 1$. We have

$$\Pr[x_1 = \bar{x}_j] = \sum_{i=1}^{N} \Pr[x_1 = \bar{x}_j | x_0 = \bar{x}_i] \times \Pr[x_0 = \bar{x}_i] = \sum_{i=1}^{N} P_{ij} \times [\pi_0]_i$$

where $[\pi_0]_i$ is the ith element of π_0. Notice that this sum is the product of π_0 against the ith column of P. Repeating this logic for each $j = 1, \cdots, N$ we have $\pi_1 = \pi_0 \times P$. This relationship generalizes to $\pi_{t+1} = \pi_t \times P$ and by repeated substitution we have

$$\pi_{t+k} = \pi_t \times P^k.$$

This is a very useful property of Markov chains: the conditional distributions k periods ahead can be found by raising the transition matrix to the power k.

For a Markov chain, a stationary distribution $\bar{\pi}$ is one for which $\bar{\pi} = \bar{\pi} \times P$. If we transpose this definition we have $\bar{\pi}' = P' \times \bar{\pi}'$ and we can see that the stationary distribution is the eigenvector of P' associated with a unit eigenvalue.[2]

To give an example with an economic interpretation, suppose workers can be employed or unemployed. Unemployed workers find jobs with probability $f \in (0,1)$ and lose (or separate from) jobs with probability $s \in (0,1)$. Then we can represent the transitions across the employed/unemployed states by the transition matrix

$$P = \begin{pmatrix} 1-s & s \\ f & 1-f \end{pmatrix},$$

[2] A transition matrix will always have a unit eigenvalue. As the rows of P sum to one, we know $\mathbf{1} = P\mathbf{1}$, where $\mathbf{1}$ is a column vector of ones. This equation says P has a unit eigenvalue and the eigenvalues of P and P' are the same.

where the first state represents being employed and the second represents being unemployed. As the stationary distribution has two probabilities that sum to one, there is only one unknown. Let $\bar{u}$ be the stationary (steady state) unemployment rate so that we have $\bar{\pi} = \begin{bmatrix} 1-\bar{u} & \bar{u} \end{bmatrix}$. The eigenvector of P' associated with a unit eigenvalue solves

$$\begin{pmatrix} -s & f \\ s & -f \end{pmatrix} \begin{pmatrix} 1-\bar{u} \\ \bar{u} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

Or, equivalently,

$$\bar{u} = s(1-\bar{u}) + (1-f)\bar{u}.$$

The steady state unemployment rate is equal flow into unemployment (separations) plus the flow of unemployed who remain jobless. Solving this equation yields

$$\bar{u} = \frac{s}{s+f}.$$

Notice that there is a unique stationary distribution of this economy. Suppose we start our economy with an unemployment rate u_0. As time goes by, it is also straightforward to check that the unemployment rate will converge to $\bar{u}$ regardless of what u_0 we start with.[3]

When will a Markov chain more generally converge to a unique stationary distribution? It is not always the case, as the following examples demonstrate. Consider a Markov chain with a transition matrix equal to the identity matrix. No matter what initial distribution we start with, the distribution will forever remain stationary so there is not a unique distribution although it converges immediately. As another example, consider the Markov chain with transition matrix given by

$$P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

This Markov chain has a unique stationary distribution of [1/2, 1/2], but unless we start with that distribution, the Markov chain will never converge to it. To see this, suppose we start with probability mass $p \neq 1/2$ on the first state and $1-p$ on the second. The mass on the first state will oscillate between p and $1-p$ forever.

A simple necessary condition for convergence and uniqueness is that from each state there is a positive probability of moving to any other state. This condition is actually substantially stronger than we need. Weaker conditions are as follows. State j of a Markov chain is said to be reachable from state i if there is some n such that $\Pr(x_n = \bar{x}_j | x_0 = \bar{x}_i) > 0$. A Markov chain is said to be irreducible if all states are reachable from all other states. State i of a Markov chain is said to be aperiodic if there is some n such that for all $n' \geq n$ $\Pr(x_{n'} = \bar{x}_i | x_n = \bar{x}_i) > 0$. A Markov chain is aperiodic if all states are aperiodic. An irreducible Markov chain has a unique stationary distribution, $\bar{\pi}$, and if it is aperiodic then $\lim_{t\to\infty} \pi_0 P^t = \bar{\pi}$ for all π_0.

[3]We obtain $u_{t+1} = s(1-u_t) + (1-f)u_t = s + (1-f-s)u_t$. By repeated substitution we see that $u_t = s(1 + \lambda + \lambda^2 + \cdots + \lambda^t u_0)$, with $\lambda \equiv 1-f-s$, which will converge to $s/(s+f)$ since both s and f are strictly between 0 and 1.

7.1.3 Autoregressive processes

We now turn our attention to stochastic processes with continuous distributions. A very common formulation is the autoregressive process of order one or **AR(1) process** for short. The stochastic process x_t follows an AR(1) if it satisfies

$$x_t = \rho x_{t-1} + b\varepsilon_t + (1-\rho)\mu \tag{7.1}$$

where $x_t \in \mathbb{R}$, ρ, b, and μ are scalar coefficients, and ε_t is a stochastic process that satisfies $\mathbb{E}_{t-1}[\varepsilon_t] = 0$, $\mathbb{E}_{t-1}[\varepsilon_t^2] = 1$, and $\mathbb{E}_{t-1}[\varepsilon_t \varepsilon_{t+s}] = 0$ for all $s > 0$.[4] We call the process (7.1) an AR(1) because x_t depends on just one lagged value x_{t-1}.

We can calculate moments of this process by expressing x_t as a moving average of past ε's. Through repeated substitution we arrive at

$$\begin{aligned} x_t &= \mu + b\varepsilon_t + \rho b \varepsilon_{t-1} + \rho^2 b \varepsilon_{t-2} + \cdots \\ &= \mu + b \sum_{s=0}^{\infty} \rho^s \varepsilon_{t-s}. \end{aligned}$$

If $|\rho| < 1$, the effect on x_t of shocks, ε, in the distant past vanishes and the process is stationary. Taking an unconditional expectation we see $\mathbb{E}[x_t] = \mu$ since $\mathbb{E}[\varepsilon_{t-s}] = \mathbb{E}[\mathbb{E}_{t-s-1}[\varepsilon_{t-s}]] = 0$ for all s. The unconditional variance of x_t is given by

$$\operatorname{Var}[x_t] = \sum_{s=0}^{\infty} (b\rho^s)^2 \operatorname{Var}[\varepsilon_{t-s}] = \frac{b^2}{1-\rho^2},$$

where we have used the fact that the ε's have unit standard deviation and are uncorrelated across time. Similarly, the covariance of x_t and x_{t+j} is

$$\begin{aligned} \operatorname{Cov}(x_t, x_{t+j}) &= \mathbb{E}\left[(x_t - \mu)(x_{t+j} - \mu)\right] \\ &= \mathbb{E}\left[\left(b \sum_{s=j}^{\infty} \rho^{s-j} \varepsilon_{t-s+j}\right)\left(b \sum_{s=0}^{\infty} \rho^s \varepsilon_{t+j-s}\right)\right] \\ &= b^2\left(\rho^j + \rho^{j+2} + \rho^{j+4} + \cdots\right) \\ &= \frac{b^2 \rho^j}{1-\rho^2}. \end{aligned}$$

The correlation between x_t and x_{t+j} is therefore ρ^j. In summary, the parameter μ determines the level of the process, the parameter b determines the volatility of the process, and the parameter ρ determines the persistence of the process.

7.1.4 Linear stochastic difference equations

We now will generalize our autoregressive specification to allow for vector-valued random variables. Let x_t be a column vector in $\mathbb{R}^n$. Let ε_t be a random variable in $\mathbb{R}^m$. The stochastic process ε_t is assumed to satisfy

$$\mathbb{E}_t[\varepsilon_{t+1}] = 0, \tag{7.2}$$

[4] The assumption that ε_t has unit variance is a normalization as the parameter b scales the effects of ε_t on x_t.

$$\mathbb{E}_t\left[\varepsilon_{t+1}\varepsilon'_{t+1}\right] = I, \tag{7.3}$$

and

$$\mathbb{E}_t\left[\varepsilon_{t+1}\varepsilon'_{t+s}\right] = 0 \quad \forall s > 1. \tag{7.4}$$

The second condition says the elements of ε_t are uncorrelated with each other and have unit standard deviation. The third condition states that ε_t is uncorrelated across time. We assume that x_t follows a linear stochastic difference equation:

$$x_t = Ax_{t-1} + B\varepsilon_t + C. \tag{7.5}$$

The $n \times n$ matrix A controls how x_{t-1} affects x_t. If all the eigenvalues of A are smaller than 1 in absolute value, then the effects of past shocks will eventually fade and x_t will be a stationary process. The $n \times m$ matrix B captures the effects of ε_t on x_t.[5] Lastly, the $n \times 1$ vector C affects the mean of x_t as we describe next.

The unconditional expectation of x_t is

$$\mu \equiv \mathbb{E}\left[x_t\right] = A\mathbb{E}\left[x_{t-1}\right] + C = A\mu + C.$$

Solving this equation yields $\mu = (I - A)^{-1}C$. Similarly, let $\Gamma(0)$ be the unconditional covariance matrix of x_t. The definition of a covariance matrix gives us

$$\begin{aligned}\Gamma(0) &= \mathbb{E}\left[(x_t - \mu)(x_t - \mu)'\right] \\ &= \mathbb{E}\left[A(x_{t-1} - \mu)(x_{t-1} - \mu)'A' + B\varepsilon_t\varepsilon'_t B'\right] \\ &= A\Gamma(0)A' + BB',\end{aligned} \tag{7.6}$$

where we have used the fact that ε_t is independent of x_{t-1}.[6]

We are often interested in the behavior of a stochastic process following a particular event. For example, we might be interested in how the economy would behave following a TFP shock. In this thought experiment, we suppose no further shocks occur. Using $C = \mu - A\mu$, rewrite (7.5) as

$$x_t - \mu = A(x_{t-1} - \mu) + B\varepsilon_t.$$

Suppose there is a particular shock ε_t at t and then no future shocks. Repeated substitution yields

$$x_{t+h} - \mu = A^{h+1}(x_{t-1} - \mu) + A^h B\varepsilon_t.$$

The effect of ε_t on x_{t+h} is given by $A^h B\varepsilon_t$. This change in the future evolution of the process is called the impulse response of x_t to the shock ε_t. The function $\mathcal{F}(h) = A^h B\varepsilon$ is the **impulse response function** of x to the particular shock ε. The impulse response function tells us how x responds to the shock as a function of the time since the shock has occurred. We saw some examples of impulse responses in the deterministic, non-linear Solow model in Section 3.5.2. If the Solow model is extended to include stochastic shocks, the results in this section would apply only to a linear approximation to the Solow model.

[5]The assumption reflected in (7.3) that the ε_t are uncorrelated with each other and have unit standard deviations is a normalization since the matrix B can rescale their standard deviations and impart correlations across their effects.

[6]Equation (7.6) is a Lyapunov equation and can be used to solve for $\Gamma(0)$. Many software packages are available to solve such equations.

7.2 Choice under uncertainty

We will now begin our discussion of how agents make choices under uncertainty. In this section we start by introducing a framework for modeling uncertainty in a way that can introduce uncertainty without restricting ourselves to a specific stochastic process. We then discuss preferences over risky consumption outcomes. Risk aversion is an important aspect of preferences in an uncertain environment and we will demonstrate the implications of risk aversion through a portfolio choice problem.

7.2.1 Stochastic events

To incorporate uncertainty into economic theory, it is often convenient to define a stochastic event that determines all the risky outcomes. The idea here is that there are many different ways the world may take shape in the future and our uncertainty is that we do not know which of these "worlds" we live in. As our theories are typically dynamic, we need to allow for our uncertainty to resolve over time. Let $\omega_t \in \Omega_t$ be the stochastic event realized at date t and let $\omega^t = \{\omega_0, \omega_1, \cdots, \omega_t\} \in \Omega^t$ be the history of events up to date t. To give an example, suppose your income each month can either be high or low. The event ω_t determines whether your income is high or low in month t and ω^t gives a list of all the past events from which we can infer your past incomes. Figure 7.1 shows an example of how these stochastic events could unfold for $t = 0, 1, 2$ when there are two possible realizations of ω_t at each date: $\omega_t \in \{0, 1\}$.

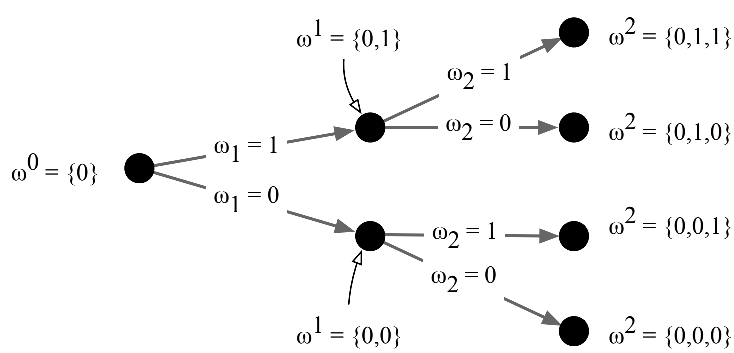

Figure 7.1: Example event tree.

The probability that history ω^t will be realized at date t from the perspective of date 0 is given by $\pi_t(\omega^t)$. The conditional probability of ω^t given ω^τ for $t > \tau$ is $\pi_t(\omega^t|\omega^\tau)$. An outcome at date t is a function of the history up to date t. For example, the balance in your bank account reflects not just the randomness in your current income, but also the fluctuations in your previous incomes (as well as the changes in spending they induce). We could write your assets as $a_t(\omega^t)$ to indicate that it depends on the whole history of events leading up to date t.

7.2.2 Expected utility and risk aversion

We now extend the preferences we introduced in Chapter 4 to incorporate uncertainty according to *expected utility theory.* Thus, utility over stochastic events is then a convex linear combination of a function $u(c)$, where c is random, with the linear coefficients being the probabilities with which the different outcomes for c are realized. Applied to the context of preferences over time, suppose $\{c_t(\omega^t) : \forall t, \omega^t\}$ and $\{\tilde{c}_t(\omega^t) : \forall t, \omega^t\}$ are two consumption processes. We will say the c process is preferred to the $\tilde{c}$ process if and only if

$$\sum_{t=0}^{\infty} \sum_{\omega^t \in \Omega^t} \pi_t(\omega^t)\beta^t u(c_t(\omega^t)) > \sum_{t=0}^{\infty} \sum_{\omega^t \in \Omega^t} \pi_t(\omega^t)\beta^t u(\tilde{c}_t(\omega^t)).$$

In many situations, we will not write the stochastic events explicitly because the time subscript on the variables is sufficient to keep track of which histories they depend on. Using that notational convention, the above statement would be

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t) > \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(\tilde{c}_t).$$

As we get started, however, it is useful to be explicit about the histories.

Risk aversion is a fundamental concept of choice under uncertainty that will allow us to explain, among other things, why consumers buy insurance and why savers hold diversified portfolios. Risk aversion is the idea that a less risky consumption stream is preferred to a more risky stream with the same expected consumption. Mathematically, risk aversion is implied by Jensen's inequality when $u(\cdot)$ is concave. If $u(\cdot)$ is linear, then the consumer ranks consumption streams by the expected level of consumption and is said to be risk neutral.

The curvature of $u(\cdot)$ determines the level of risk aversion. The coefficient of absolute risk aversion is defined as $-u''(c)/u'(c)$. A more concave utility function leads to higher risk aversion. The second derivative is normalized by the first derivative to capture the change in curvature as utility changes.[7] The coefficient of absolute risk aversion refers to the attitude towards changes in consumption of a given (absolute) size. The constant absolute risk aversion (CARA) utility function is given by $u(c) = -\exp(-\alpha c)$. For this utility function, the coefficient of absolute risk aversion is α at all levels of consumption.

Alternatively, the coefficient of relative risk aversion measures a consumer's attitude to proportional changes in consumption. It is defined as $-cu''(c)/u'(c)$. Our power utility function

$$u(c) = \frac{c^{1-\sigma} - 1}{1 - \sigma},$$

motivated by its consistency with exact balanced growth, has a constant coefficient of relative risk aversion. Taking derivatives, we see the coefficient of relative risk aversion is simply σ. As the coefficient of relative risk aversion is the same at all levels of consumption, these preferences are also known as the constant relative risk aversion (CRRA) utility function.

[7] In expected utility theory, a positive affine transformation of a utility function represents the same preferences. That is, the utility function $v(c)$ defined by $au(c) + b$ for constants $a > 0$ and b is equivalent to the utility function u. Normalizing the coefficient of risk aversion by $u'(c)$ makes the coefficient of absolute risk aversion invariant to such an affine transformation of u.

In most macroeconomic applications, the CRRA function is used because, as we have discussed, it is the only one consistent with balanced, long-run growth. From the perspective here, using a CRRA function means that decisions about risks today and one hundred years ago—when the level of consumption was much lower—would have been made the same way if the risks were the same in percentage terms.

Recall from Section 4.2.4 that the elasticity of intertemporal substitution is the elasticity of consumption growth with respect to the real interest rate. More precisely, the consumption Euler equation for power utility implies $d\log[c_{t+1}/c_t]/dR_{t+1} = 1/\sigma$, implying that the elasticity of intertemporal substitution actually equals the inverse of the coefficient of relative risk aversion. The intuition here is that both of these features of the utility function are determined by the curvature of the utility function. If the utility function displays strong diminishing marginal utility, then a consumer will be unwilling to accept higher consumption in one state of the world in exchange for low consumption in another state of the world (i.e., risk aversion is high). Similarly, the consumer will be unwilling to accept low consumption in period t in exchange for high consumption in period $t+1$ (i.e., is unwilling to substitute intertemporally). Intuitively put, one single parameter guides the desire for smoothing consumption across time and states of nature.[8]

7.2.3 Portfolio choice

A simple portfolio choice problem can help illustrate the differences between the two concepts of risk aversion introduced above. Suppose an individual has $W > 0$ units of wealth at date 0 to allocate between a risk-free asset and a risky asset. The assets will pay off at date 1 and the individual will consume the proceeds. The (gross) return on the risk-free asset is known to be R^f while the return on the risky asset is unknown and denoted Z. Let A be the assets invested in the risky asset while $W - A$ are invested in the risk-free asset. The investor's decision problem is

$$\max_A \mathbb{E}_0 \left[u\left(R^f(W-A) + ZA\right)\right].$$

The first-order condition of this problem is

$$\mathbb{E}_0 \left[u'\left(R^f(W-A) + ZA\right)\left(Z - R^f\right)\right] = 0.$$

Suppose the investor's utility function is the CRRA utility function $u(c) = c^{1-\sigma}/(1-\sigma)$. The first-order condition of the portfolio choice problem becomes

$$\mathbb{E}_0 \left[\left(R^f(W-A) + ZA\right)^{-\sigma}\left(Z - R^f\right)\right] = 0.$$

Rearranging we arrive at

$$\mathbb{E}_0 \left[\left(R^f\left(1 - \frac{A}{W}\right) + Z\frac{A}{W}\right)^{-\sigma}\left(Z - R^f\right)\right] = 0,$$

[8] There is a generalization of the power-function preferences that allows one to separate risk aversion from intertemporal substitution using two separate parameters. This case will be discussed in Chapter 16 below.

where we have brought $W^{-\sigma}$ outside the expectation because it is known at date 0. This equation determines A/W as a function of σ, R^f, and the distribution of the risky asset return. Notice that the solution for A/W does not depend on W, which means that at any level of wealth, the investor will allocate the same fraction of savings to risky assets: the rich and the poor choose the same risk exposure.

Now suppose the investor has the CARA utility function $u(c) = -\exp(-\alpha c)$. The first-order condition then becomes

$$\mathbb{E}_0\left[\alpha \exp\left\{-\alpha R^f W\right\} \exp\left\{-\alpha\left(Z - R^f\right)A\right\}\left(Z - R^f\right)\right] = 0.$$

We can bring $\exp\left\{-\alpha R^f W\right\}$ outside the expectation because it is known at date 0 to arrive at

$$\mathbb{E}_0\left[\exp\left\{-\alpha\left(Z - R^f\right)A\right\}\left(Z - R^f\right)\right] = 0.$$

This equation gives a solution for A that does not depend on W. In this case, the investor allocates a particular level of savings (an absolute number of goods or dollars) to risky assets regardless of their wealth. This means that the rich have a lower risky share than poorer households.

In the data, as we shall see in Chapter 21, the rich on average choose a higher risky share. Neither of these simple models can match this fact but the CRRA case is more in line with the data.

7.3 The stochastic growth model

We can now use the tools of choice under uncertainty and stochastic processes to analyze a stochastic version of the neoclassical growth model. As briefly discussed in Chapter 3 above, in this model, total factor productivity (TFP) is assumed to follow an exogenous stochastic process. The fluctuations in TFP then give rise to endogenous fluctuations in output, consumption, and investment; the model was first studied, as a planning problem, in Brock and Mirman (1972), and then became a workhorse framework in macroeconomics, and we devote Chapter 14 below to it. Through this important example, we will discuss how our methods for dynamic optimization can be extended to allow for uncertainty.

7.3.1 A two-period economy

To begin, suppose the economy exists for two periods. In period 0, the level of TFP in period 1 is not known. Let ω_1 be the stochastic event in period 1. Let $\pi_1(\omega_1)$ be the probability of ω_1. TFP at date 1 is given by $A_1(\omega_1)$.

The economy is inhabited by a representative household with expected utility preferences that ranks consumption streams according to

$$U = u(C_0) + \beta \sum_{\omega_1 \in \Omega_1} \pi_1(\omega_1) u\left(C_1(\omega_1)\right), \tag{7.7}$$

where $C_1(\omega_1)$ is the level of consumption if ω_1 occurs. In period 0, the economy is endowed with K_0 units of capital, which are used to produce $Y_0 = K_0^\alpha$ units of output. This output is then used for consumption and investment subject to the date-0 resource constraint

$$K_1 + C_0 = K_0^\alpha + (1-\delta)K_0. \tag{7.8}$$

In period 1, the value of ω_1 becomes known and the economy produces $Y_1(\omega_1) = A_1(\omega_1)K_1^\alpha$. As there are no further periods, there is no reason to invest in capital so the resource constraint in period 1 is

$$C_1(\omega_1) = A_1(\omega_1)K_1^\alpha + (1-\delta)K_1. \tag{7.9}$$

The first important thing to note is that A_1, Y_1, and C_1 are all functions of the event ω_1. From the perspective of date 0, agents do not know A_1 and therefore they cannot know how much will be produced or consumed. When we formulate a decision problem in date 0, the agents will not choose specific values for Y_1 and C_1, but rather they will choose a plan for how they will respond to each realization of ω_1. This is an important feature of optimization under uncertainty: the choice variable is a contingent plan for actions following each history of stochastic events.

The planner's problem for this economy is to choose C_0, K_1, and $\{C_1(\omega_1) : \forall\omega_1\}$ to maximize (7.7) subject to (7.8) and (7.9). We can form the Lagrangian as

$$\begin{aligned}\mathcal{L} = u(C_0) + \beta \sum_{\omega_1\in\Omega_1} \pi_1(\omega_1)u\left(C_1(\omega_1)\right) - \lambda_0\left[K_1 + C_0 - K_0^\alpha - (1-\delta)K_0\right] \\ - \sum_{\omega_1\in\Omega_1} \lambda_1(\omega_1)\left[C_1(\omega_1) - A_1(\omega_1)K_1^\alpha - (1-\delta)K_1\right],\end{aligned}$$

where λ_0 and the λ_1's are Lagrange multipliers on (7.8) and (7.9), respectively. As (7.9) must hold for each realization of ω_1 we treat that as a separate constraint for each ω_1. We therefore have separate Lagrange multipliers for each ω_1 and we use the sum to include all of them in the Lagrangian.

Taking the first-order conditions for this problem, we have

$$u'(C_0) = \lambda_0,$$

$$\lambda_0 = \sum_{\omega_1\in\Omega_1} \lambda_1(\omega_1)\left(\alpha A_1(\omega_1)K_1^{\alpha-1} + 1 - \delta\right),$$

and

$$\beta\pi(\omega_1)u'(C_1(\omega_1)) = \lambda_1(\omega_1). \qquad \forall\omega_1.$$

The second line is the first-order condition for K_1. On the right-hand side we have a sum over all possible realizations of ω_1 because when the planner chooses K_1 they do not know which ω_1 will occur so they need to take into account how K_1 affects output and consumption after each one. In contrast, the third line is the first-order condition with respect to $C_1(\omega_1)$ for a specific ω_1 and there is one such equation for each ω_1.

Combining the first-order conditions to eliminate the Lagrange multipliers we have

$$u'(C_0) = \beta \sum_{\omega_1 \in \Omega_1} \pi(\omega_1) u' \left(C_1(\omega_1)\right) \left(\alpha A_1(\omega_1) K_1^{\alpha-1} + 1 - \delta\right).$$

This is the stochastic consumption Euler equation for the planner. The left-hand side is the marginal utility loss in date 0 from saving one more unit. The right-hand side is the expected marginal utility gain in period 1 from saving one more unit. Notice that we sum over ω_1 and weight the outcomes by $\pi_1(\omega_1)$ so we are taking an expectation. The term $\alpha A_1(\omega_1) K_1^{\alpha-1} + 1 - \delta$ is the marginal increase in resources from increasing K_1 and $u'\left(C_1(\omega_1)\right)$ is the marginal utility of consuming more. The former is the return on saving, which is stochastic as it depends on TFP. The latter reflects the fact that the return on capital is valued differently after different realizations of ω_1. This is due to diminishing marginal utility—when TFP is high, consumption will be high and the marginal value of consuming more is low. The uncertainty in TFP generates uncertainty in $C_1(\omega_1)$, which in turn generates uncertainty in marginal utility.

The Euler equation brings us to another important point about dynamic optimization under uncertainty that turns out to be general: the return to saving is evaluated differently in different states of the world. We do not just focus on the expected return, but instead a weighted average that accounts for the different value of resources in different situations.

7.3.2 An infinite-horizon economy

We now consider an infinite-horizon version of the model. At each date t there is a realization ω_t and the date-0 probability of a history ω^t is given by $\pi_t(\omega^t)$. The representative household has preferences given by

$$U = \sum_{t=0}^{\infty} \sum_{\omega^t \in \Omega^t} \pi_t(\omega^t) \beta^t u\left(C_t(\omega^t)\right). \tag{7.10}$$

Unlike the two-period model, consumption at date t now depends on the whole history of stochastic events. At each date, a stochastic TFP is realized and used to produce output

$$Y_t(\omega^t) = A_t(\omega^t) F(K_t(\omega^{t-1}), L_t(\omega^t)),$$

where $L_t(\omega^t)$ is the labor input and $A_t(\omega^t)$ is total factor productivity. The production function is twice continuously differentiable in K and L, is strictly increasing and strictly concave in both arguments, and is constant returns to scale. Note that the capital that is used in production at date t, denoted K_t, is selected at date $t-1$ and therefore can only depend on the information that is available at the time it is selected. Therefore K_t is a function of ω^{t-1}, not ω^t.

The economy is endowed with one unit of labor each period and we assume, for simplicity, that there is no preference for leisure so that labor supply is inelastic and equal to one. The aggregate resource constraint at date t is

$$K_{t+1}(\omega^t) + C_t(\omega^t) = f\left(A_t(\omega^t), K_t(\omega^{t-1})\right), \tag{7.11}$$

where we have defined $f(A, K) \equiv AF(K, 1) + (1 - \delta)K$. The economy begins with an initial endowment of capital, K_0. Negative capital holdings are not possible.

The planner's problem for this economy is to maximize (7.10) subject to (7.11) where the constraint applies to each t and each ω^t. The Lagrangian of this problem is

$$\mathcal{L} = \sum_{t=0}^{\infty} \sum_{\omega^t \in \Omega^t} \left\{ \pi_t(\omega^t)\beta^t u\left(C_t(\omega^t)\right) - \lambda_t(\omega^t)\left[K_{t+1}(\omega^t) + C_t(\omega^t) - f\left(A_t(\omega^t), K_t(\omega^{t-1})\right)\right]\right\}.$$

When we take the first-order conditions of this problem a key point is that our choice of $K_{t+1}(\omega^t)$ will affect production in $t+1$ for all histories ω^{t+1} that are possible given that we have already reached ω^t. For example, refer back to Figure 7.1 and suppose we have reached $\omega^1 = \{0, 0\}$ at date 1 and we are choosing $K_2(\omega^1)$. This choice of capital will affect production at date 2 for histories $\{0, 0, 0\}$ and $\{0, 0, 1\}$ because these are possible following ω^1. This choice of capital will not affect production for histories $\{0, 1, 0\}$ or $\{0, 1, 1\}$ because these are not possible given ω^1. The first-order condition for $K_{t+1}(\omega^t)$ is therefore

$$\lambda_t(\omega^t) = \sum_{\{\omega^{t+1}|\omega^t\}} \lambda_{t+1}(\omega^{t+1}) f_2\left(A_{t+1}(\omega^{t+1}), K_{t+1}(\omega^t)\right), \tag{7.12}$$

where the notation $\{\omega^{t+1}|\omega^t\}$ indicates that we sum over the histories ω^{t+1} that are possible given ω^t and $f_i(\cdot, \cdot)$ represents the partial derivative with respect to ith argument. The first-order condition with respect to $C_t(\omega^t)$ is

$$\pi_t(\omega^t)\beta^t u'\left(C_t(\omega^t)\right) = \lambda_t(\omega^t).$$

Using this to eliminate the Lagrange multipliers in (7.12) we arrive at the consumption Euler equation for this problem

$$u'\left(C_t(\omega^t)\right) = \beta \sum_{\{\omega^{t+1}|\omega^t\}} \pi_{t+1}(\omega^{t+1}|\omega^t) u'\left(C_{t+1}(\omega^{t+1})\right) f_2\left(A_{t+1}(\omega^{t+1}), K_{t+1}(\omega^t)\right), \tag{7.13}$$

where we have defined the conditional probability $\pi_{t+1}(\omega^{t+1}|\omega^t) = \pi_{t+1}(\omega^{t+1})/\pi_t(\omega^t)$. The consumption Euler equation has the same interpretation as in the two-period economy. The right-hand side has a weighted average of the returns on savings with weights corresponding to the marginal utility of consumption in different states of the world.

In many applications, the time subscripts on variables are sufficient to indicate the history of events they depend on in which case we can rewrite (7.13) as

$$u'(C_t) = \beta \mathbb{E}_t\left[u'(C_{t+1}) f_2(A_{t+1}, K_{t+1})\right]. \tag{7.14}$$

The $\mathbb{E}_t$ indicates we are taking conditional expectation just as our sum over $\{\omega^{t+1}|\omega^t\}$ does.

We can now use equation (7.11), at (t, ω^t) and at all the nodes in the following period, to eliminate consumption from the Euler equation (7.13). In the deterministic model, we followed the corresponding method and arrived at a second-order difference equation in capital. Here, we obtain a second-order *stochastic* difference equation in capital, which holds at all nodes in the event tree. In the deterministic model, we also had a transversality condition—as a second-order difference equation and an initial value for capital leave one degree of freedom to choose capital—and we could then pin down a solution to the difference equation. With uncertainty, the situation has much of the same structure: a transversality condition must be added to determine a solution, and it has the same interpretation as before: it is a self-imposed constraint not to over-accumulate at infinity, expressed as an *expected* present value.[9]

Clearly, solving for a stochastic sequence of capital levels appears daunting. Closed-form solutions exist in special cases, but only under very special assumptions; one is when the utility function is logarithmic, the production function is Cobb-Douglas, and $\delta = 1$. For that case it is possible to verify that a constant saving rate is optimal. In other cases, one must apply numerical methods to solve the model. One such approach is to linearize the model around the steady state of the deterministic model; we show how to do this in Section 7.3.4. If non-linearities are believed to be important, a way forward is to solve the model numerically in a dynamic-programming version of the model. We now look at how it is formulated.

7.3.3 A recursive formulation

We will now analyze a recursive version of the same economy. To do so, we will assume that TFP follows a first-order Markov process so we only need to keep track of the most recent realization in order to know the distribution of its future realizations. As recursive modeling keeps track of the history of the economy explicitly through well-chosen state variables, it is customary not to use histories of stochastic events (ω^t) in this context. Following this convention, we let $\pi(A'|A)$ be the probability of A' occurring next period given the current TFP A. The planner's problem can now be expressed as the following Bellman equation

$$V(A, K) = \max_{C, K' \geq 0} \left\{ u(C) + \beta \sum_{A'} \left[\pi(A'|A) V(A', K') \right] \right\}$$

subject to

$$K' + C = f(A, K).$$

The difference compared to dynamic programming under certainty is that now the Bellman equation has an expectation over continuation values. In the next period, there will be a value of having states (A', K') but A' is not known yet. As in the case with certainty, the recursive formulation delivers the same solution as the sequential formulation of the problem. It may seem surprising that this works out since here we are only taking expectations one period into the future, while in the sequential formulation we take expectations of outcomes far in the future. However, the two sets of expectations are actually the same. In the recursive formulation $V(A', K')$ is a random variable that includes a term $\mathbb{E}[V(A'', K'')|A']$. By the law of iterated expectations, the expectation of this term conditional on A becomes $\mathbb{E}[V(A'', K'')|A]$. The same logic applies for the value function at all future dates.

[9] In the case with stochastic shocks, there is still just one degree of freedom given the first-order conditions. To see this intuitively, the Euler equation at any node (t, ω^t) can be used to solve for $k_{t+1}(\omega^t)$ as a function of $k_t(\omega^{t-1})$ and an expression involving future values $k_{t+2}(\omega^t, \omega_{t+1})$ and substituted into the previous Euler equation. When repeated infinitely many times, we have a first equation involving k_0 and $k_1(\omega_0)$ only—the latter is the only remaining degree of freedom.

To analyze the recursive economy, we can proceed with the same steps as we would for the recursive economy without uncertainty. We can substitute the constraint in to the Bellman equation

$$V(A,K) = \max_{K'} \left\{ u\left(f(A,K) - K'\right) + \beta \sum_{A'} \left[\pi(A'|A)V(A',K')\right] \right\}$$

and take the first-order condition with respect to K' to obtain

$$u'(C) = \beta \sum_{A'} \pi(A'|A) V_2(A',K'). \tag{7.15}$$

The envelope condition gives us

$$V_2(A,K) = u'(C) f_2(A,K).$$

Using this to eliminate the derivative of the value function in (7.15) we obtain a similar consumption Euler equation as we had before

$$u'(C) = \beta \sum_{A'} \pi(A'|A) \left[u'(C') f_2(A,K)\right].$$

We can then rewrite this Euler equation as a functional equation that determines the savings policy function. To do so, let $g(A,K)$ denote the choice of K' as a function of states (A,K). Then write the resource constraint as $C = f(A,K) - g(A,K)$. Substituting these definitions into the Euler equation we have

$$u'\left(f\left(A,K\right) - g\left(A,K\right)\right) = \beta \sum_{A'} \pi(A'|A) \left[u'\left(\underbrace{f(A', g(A,K)) - g(A', g(A,K))}_{=C'} \right) f_2(A', g(A,K)) \right], \tag{7.16}$$

where $f_2(A', g(A,K))$ denotes the derivative of the production function with respect to K, but evaluated at (A', K'). This equation must hold for all (A,K) and it implicitly defines the function $g(A,K)$ that is the solution to the planner's problem. In the next section, we will use this functional Euler equation to derive a complete solution to the planner's problem.

7.3.4 Solving the model via linearization

Section 7.3.3 derived a functional Euler equation that the solution to the planner's problem must satisfy. In this model, productivity follows an exogenous stochastic process. We will now show how we can use a linear approximation to the functional Euler equation to derive a linear stochastic difference equation that the endogenous variables in the model must follow. We will then use the properties of linear stochastic difference equations from Section 7.1.4 to derive properties of the planner's solution.

We will now assume that TFP follows an AR(1) process given by $A' = \rho A + (1-\rho)\bar{A} + \varepsilon'$ with $\mathbb{E}_t[\varepsilon'] = 0$. We will approximate the behavior of the economy around the deterministic steady state, which is the same notion of a steady state we have studied before. Now that we have shocks in the model, the interpretation of the steady state changes. The economy will never converge to the steady state if it is constantly hit by shocks.[10] The deterministic steady state is the point the economy would converge to if all shocks take their unconditional expectation forever and the agents in the model expect this.[11]

We take a linear approximation of (7.16) around the steady state. For a variable X we will use the notation $\hat{X}_t \equiv X_t - \bar{X}$, where $\bar{X}$ is the steady-state value.[12] Our linear approximation (which is tedious but straightforward to derive) is

$$(f_K - g_K)\hat{K} + (f_A - g_A)\hat{A} = \mathbb{E}\left[\begin{array}{c} f_K\left((f_K - g_K)\left(g_K\hat{K} + g_A\hat{A}\right) + (f_A - g_A)\hat{A}'\right) \\ +\frac{u'}{u''}\left(f_{KK}g_K\hat{K} + f_{KK}g_A\hat{A} + f_{KA}\hat{A}'\right) \end{array} \middle| A \right],$$

where f_K is the derivative of $f(A, K)$ evaluated at the steady state. Other abbreviations f_i and g_i $(i = A, K)$ represent corresponding derivatives. f_{ij} $(i, j = A, K)$ are second derivatives, also evaluated at the steady state. Notice that $\hat{A}' = \rho\hat{A} + \varepsilon'$ so the distribution of A' given A is determined by the distribution of ε. We will write the expectation as summing over ε' and substitute in for A'.

Because this is a linear equation, we can pass the expectation operator inside the right-hand side to obtain

$$(f_K - g_K)\hat{K} + (f_A - g_A)\hat{A} = \beta\left[\begin{array}{c} f_K\left((f_K - g_K)\left(g_K\hat{K} + g_A\hat{A}\right) + (f_A - g_A)(\rho\hat{A} + \mathbb{E}[\varepsilon'])\right) \\ +\frac{u'}{u''}\left(f_{KK}g_K\hat{K} + f_{KK}g_A\hat{A} + f_{KA}(\rho\hat{A} + \mathbb{E}[\varepsilon'])\right) \end{array} \right].$$

As $\mathbb{E}[\varepsilon'] = 0$, it is as if there is no uncertainty and $\hat{A}'$ is treated as if it is known to be at its expected value $\rho\hat{A}$. This is a general feature of analyzing a stochastic economy through linearization known as **certainty equivalence**—once one linearizes the economy, only expected values matter. In particular, the variance of the exogenous shock does not influence outcomes.

Recall that equation (7.16) must hold for all values of A and K. The equation above is a linear approximation to (7.16) and must hold for each $\hat{K}$ and each $\hat{A}$. The only way this can be true is if the coefficients on $\hat{K}$ on the left-hand side equal those on $\hat{K}$ on the right-hand side and, similarly, the coefficients on $\hat{A}$ match. Imposing that these coefficients match gives us two equations that allow us to solve for g_K and g_A. Conveniently, one equation contains g_K only. Therefore, starting with the coefficients on $\hat{K}$ we have

$$(f_K - g_K) = \beta\left[f_K(f_K - g_K)g_K + \frac{u'}{u''}f_{KK}g_K\right].$$

[10]In addition, the average value of variables will not coincide with those at the deterministic steady state if the model is non-linear. As the shock variance becomes smaller and smaller, however, they will become increasingly similar and, in the limit where the shock variance is zero, coincide.

[11]That the agents perceive the environment to be deterministic is a subtle but important point. There is an alternative notion of a steady state in which the agents in the model perceive there to be risk but ex post all the shocks are realized at their unconditional means. Such a steady state is sometimes called a "stochastic steady state" or a "risky steady state."

[12]Previously in the text, we used this notation for deviations in logs; both linear and log-linear deviations are used in practice.

Rearrange to obtain

$$g_K^2 - \left[1 + \beta^{-1} + \frac{u'}{u''}\frac{f_{KK}}{f_K}\right] g_K + \beta^{-1} = 0. \tag{7.17}$$

where we have used $\beta f_K = 1$, which follows from the steady state Euler equation. Equation (7.17) is a quadratic equation in g_K. The equation will have one root less than one and one root greater than one. To verify this, note that the coefficient on g_K^2 is positive so (7.17) is an upward facing parabola as shown in Figure 7.2; the quadratic intersects the y-axis at β^{-1}, which is positive; and at $g_K = 1$ the quadratic takes the value $-\frac{u'}{u''}\frac{f_{KK}}{f_K}$, which is negative if the utility and production functions are both strictly increasing and strictly concave. The quadratic therefore takes the form illustrated in Figure 7.2 and has one root between 0 and 1 and one greater than 1. The relevant root is the smaller one, because it is the one that

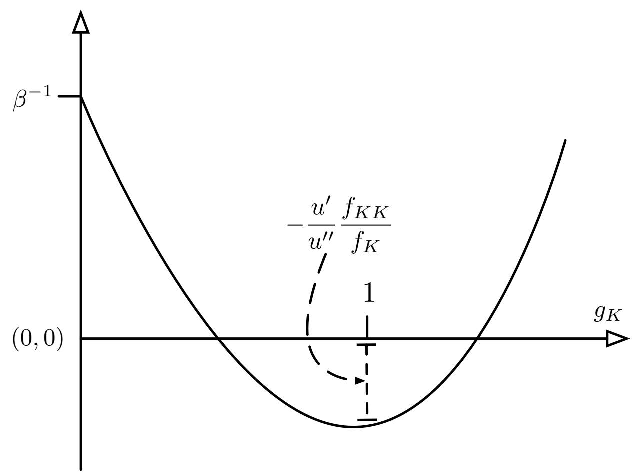

Figure 7.2: Quadratic equation to determine g_K in the linearized stochastic growth model.

is consistent with the transversality condition. The root above one is explosive in that any initial condition or shock has an increasing effect on the economy.

Solving for g_A is more straightforward because matching coefficients on g_A leads to a linear relationship with the solution

$$g_A = \left(f_K - g_K + 1 - \rho + \frac{u'}{u''}\frac{f_{KK}}{f_K}\right)^{-1}\left[(1-\rho)f_A - \rho\frac{u'}{u''}\frac{f_{KA}}{f_K}\right].$$

We have solved for the derivatives of the savings policy rule. The level of the policy rule is determined by the requirement that $(\bar{K}, \bar{A})$ is a steady state so we have

$$K' = \bar{K} + g_K(K - \bar{K}) + g_A(A - \bar{A}). \tag{7.18}$$

If we augment this equation with

$$A' = \bar{A} + \rho(A - \bar{A}) + \varepsilon', \tag{7.19}$$

we have a system of two linear stochastic difference equations and we can apply the techniques described in Section 7.1.4 to analyze the behavior of the economy.

For a numerical illustration of the properties of the stochastic growth model we will make some specific assumptions about the production function and the parameters of the model. We assume $f(A, K) = AK^{\alpha} + (1 - \delta)K$, with $\alpha = 0.3$, $\delta = 0.02$. The persistence of TFP is $\rho = 0.95$, the standard deviation of the innovations is 0.5%, and $\bar{A} = 1$. Finally, we assume $u(c) = \log(c)$ and $\beta = 0.99$.

With our parameterized solution to the model we can generate random draws of $\{\varepsilon_t\}_{t=0}^{T}$ and iterate equations (7.18) and (7.19) forward to simulate the behavior of the economy. Notice here that we are simulating the behavior of the state variables. Given the state variables it is straightforward to calculate the simulated path for output (using the production function) and the simulated path for consumption (using the aggregate resource constraint). Simulated paths for these variables are displayed in Figure 7.3(a). Notice how TFP and output exhibit much more high-frequency variation than do capital and consumption. Consumption is smooth because of the diminishing marginal utility of consumption—when output is high, it is preferable to save some of the extra resources to consume them later rather than consume all of them when marginal utility is low. Capital is smooth because it reflects the accumulation of savings over many periods—one period of high investment will not make a big percentage difference to the capital stock. Panel (b) of Figure 7.3 shows the impulse response functions following a one standard deviation shock to TFP. K_t is pre-determined so it does not respond in the period the shock occurs. Therefore, on impact of the shock, output increases by the same percentage amount as TFP. Consumption increases, but not as much as output as some of the increase in output is directed to increased savings. Over time, capital increases and the increase in output exceeds that of productivity.

Figure 7.3(c) shows the results from simulating the economy for a long time and forming a histogram with the simulated data on the capital stock. The solid line in the figure shows the theoretical unconditional distribution of the capital stock as calculated from equation (7.6).[13] As the figure shows, the economy fluctuates in the vicinity of the steady state. Sometimes capital drifts higher, sometimes lower, but it tends to return towards the steady-state level. Panel (d) of the figure shows why this is the case. The figure plots $g(A, K)$ as a function of K for two levels of A—one high and one low. The dashed line is the 45-degree line. For low K and high A, the savings policy rule is above the 45-degree line and the capital stock will increase. Similarly, for high K and low A, the savings policy rule is below the 45-degree line and the capital stock will decrease. As A fluctuates, the savings policy will shift up and down leading the capital stock to fluctuate. But note that for high A and high K, or low A and low K, the savings policy intersects the 45-degree line. These intersections imply that capital will not move out of this range (unless A gets even higher or even lower).

7.4 Competitive market trade under uncertainty

We will now begin to discuss market interactions in an uncertain economy and in the next section we will use these theoretical tools to analyze a decentralized equilibrium of the

[13]Specifically, we simulated the economy using Gaussian random variables for ε. As the dynamics of the economy are linear, the distribution of the state variables is also Gaussian. We use (7.6) to solve for the unconditional covariance matrix A and K. We have plotted a Gaussian distribution with the unconditional variance of K and a mean equal to the steady state capital stock.

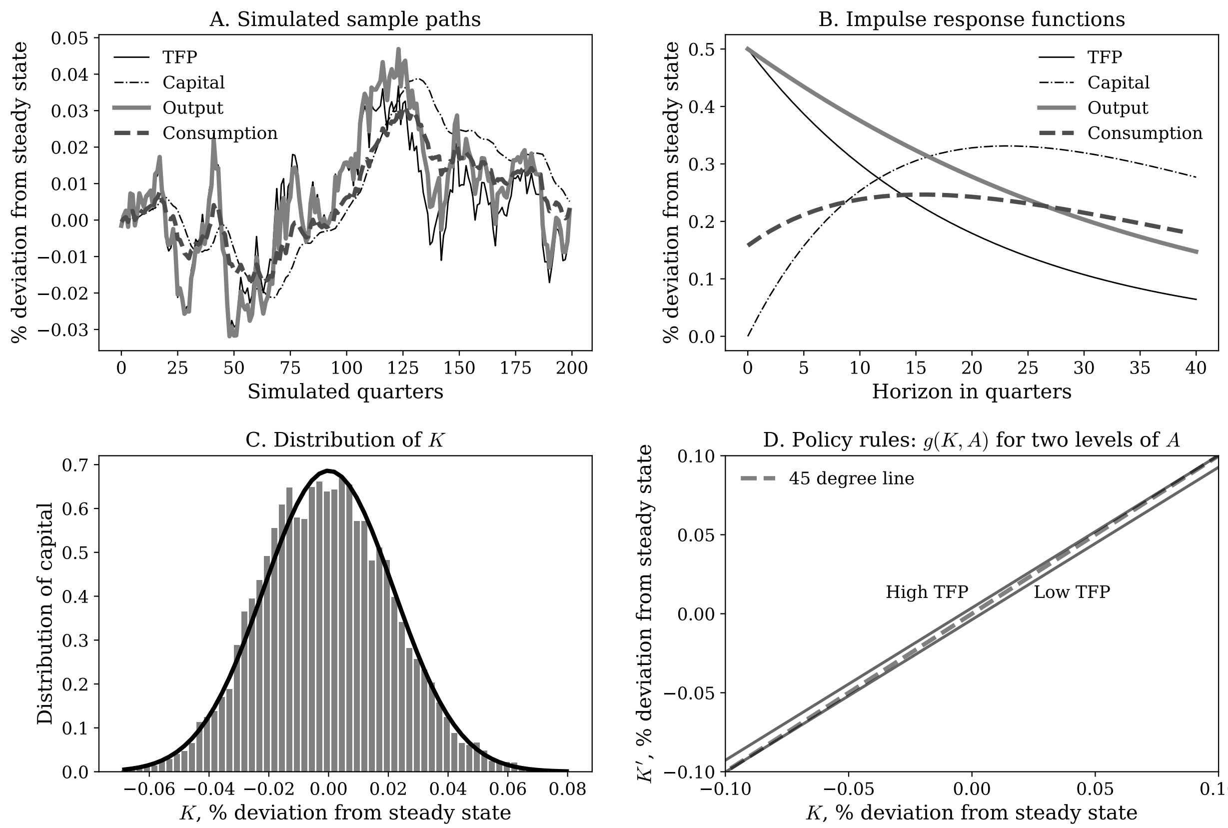

Figure 7.3: Numerical illustration of the stochastic growth model.

stochastic growth model. Before we get to that, we will first discuss how models of trade under uncertainty allow us to analyze the way agents insure themselves by sharing risks between them.

Broadly speaking, we can classify economic models of trade under uncertainty into two groups: complete markets models and incomplete markets models. In models with complete markets, agents can buy and sell goods with contracts tailored to every possible state of the world. They can write a contract for how they will behave after every possible history ω^t. This means that any possible risk can be insured at some price. In incomplete markets models, some of these contracts are not available—some risks can not be insured against. In this chapter, we will mostly focus on complete markets models. We start here because it is simpler, not because it is more realistic, but we will introduce an incomplete markets environment in Section 7.6.

Suppose the economy is populated by a set $\mathcal{I}$ of infinitely-lived households; this set could be the continuum [0,1], as in most of Chapter 5, or it could be a different, perhaps smaller set. Each household has expected utility preferences given by

$$\sum_{t=0}^{\infty} \sum_{\omega^t \in \Omega^t} \pi_t(\omega^t) \beta^t u\left(c_{i,t}(\omega^t)\right),$$

where $c_{i,t}(\omega^t)$ is the consumption of household $i \in \mathcal{I}$ after history $\omega^t \in \Omega^t$. We do not spell out the nature of ω_t just yet, as it depends on the population details (the nature of the set $\mathcal{I}$). We assume u is strictly increasing and strictly concave. Each household is endowed with a stochastic income stream that depends on the stochastic event. In particular let the income of household i at date t after history ω^t be $y_{i,t}(\omega^t)$. The good is perishable, meaning that it cannot be stored from one period to the next and must be consumed in the period it arrives in the economy. A consumption allocation is feasible if the total consumption of goods at t and ω^t is equal to the total endowment of goods

$$\sum_{i\in\mathcal{I}} c_{i,t}(\omega^t) \leq \sum_{i\in\mathcal{I}} y_{i,t}(\omega^t).$$

We will begin by considering a market structure in which there is trade only at date 0. For each date t and history ω^t, there is a contract that says the seller will pay the buyer one unit of good at that date if that history has been realized, and nothing otherwise. These are called *Arrow securities*. We will denote the date-0 price of obtaining one unit of goods after history ω^t as $p_t(\omega^t)$. This price is denominated in terms of date-0 consumption.

The decision problem of household i is to choose a contingent plan of $\{c_{i,t}(\omega^t) : \forall t, \omega^t\}$ to maximize the expected utility preferences subject to the budget constraint

$$\sum_{t=0}^{\infty} \sum_{\omega^t\in\Omega^t} p_t(\omega^t) c_{i,t}(\omega^t) = \sum_{t=0}^{\infty} \sum_{\omega^t\in\Omega^t} p_t(\omega^t) y_{i,t}(\omega^t).$$

This budget constraint says the date-0 cost of the consumption plan is less than the date-0 value of the income process. It is as if the agent, at date 0, sells claims to all their future income and then uses those funds to buy consumption goods to be delivered at future dates if particular histories occur.

Formally, we have the following.

Definition 13 *An **Arrow-Debreu competitive equilibrium** is a set of stochastic sequences* $\left\{c^\star_{i,t}(\omega^t) : \forall t, \omega^t\right\}$, *for each* $i \in \mathcal{I}$, *and* $\{p_t(\omega^t) : \forall t, \omega^t\}$ *such that*

1. *for each* i, $\left\{c^\star_{i,t}(\omega^t) : \forall t, \omega^t\right\}$ *solves*

$$\max_{\{c_t(\omega^t):\forall t,\omega^t\}} \sum_{t=0}^{\infty} \sum_{\omega^t\in\Omega^t} \beta^t \pi(\omega^t) u(c_t(\omega^t)) \quad \textit{subject to} \quad \sum_{t=0}^{\infty} \sum_{\omega^t\in\Omega^t} p_t(\omega^t) c_{i,t}(\omega^t) = \sum_{t=0}^{\infty} \sum_{\omega^t\in\Omega^t} p_t(\omega^t) y_{i,t}(\omega^t)$$

2. $\sum_{i\in\mathcal{I}} c^\star_{i,t}(\omega^t) di = \sum_{i\in\mathcal{I}} y_{i,t}(\omega^t) di$ *for all* (t, ω^t).

To characterize the equilibrium, the Lagrangian of household i is

$$\mathcal{L} = \sum_{t=0}^{\infty} \sum_{\omega^t\in\Omega^t} \beta^t \pi_t(\omega^t) u(c_{i,t}(\omega^t)) + \lambda_i \left[\sum_{t=0}^{\infty} \sum_{\omega^t\in\Omega^t} p_t(\omega^t) \left(y_{i,t}(\omega^t) - c_{i,t}(\omega^t) \right) \right],$$

where λ_i is the Lagrange multiplier on the date-0 budget constraint. The first-order condition of household i with respect to $c_{i,t}(\omega^t)$ is

$$\beta^t \pi_t(\omega^t) u'(c_{i,t}(\omega^t)) = \lambda_i p_t(\omega^t) \tag{7.20}$$

We can understand several properties of the equilibrium consumption allocation from the first-order condition.

Insurance Consider equation (7.20) for two different values of ω^t, call them ω^t and $(\omega^t)'$ and take the ratio of these two equations to arrive at

$$\frac{\beta^t \pi_t(\omega^t) u'(c_{i,t}(\omega^t))}{\beta^t \pi_t((\omega^t)') u'(c_{i,t}((\omega^t)'))} = \frac{\lambda_i p_t(\omega^t)}{\lambda_i p_t((\omega^t)')}.$$

If the prices satisfy $p_t(\omega^t) = \bar{p}_t \times \pi_t(\omega^t)$ with $\bar{p}_t > 0$, which we call **actuarially fair prices**, we have

$$u'(c_{i,t}(\omega^t)) = u'(c_{i,t}((\omega^t)'))$$

and

$$c_{i,t}(\omega^t) = c_{i,t}((\omega^t)'),$$

where the second line follows from $u(\cdot)$ being strictly concave.[14] With actuarially fair prices, the households will buy full insurance and consumption does not depend on ω^t. Full insurance is, of course, only feasible in equilibrium if total resources do not vary across the different values of ω^t so in general, prices have to adjust to reflect not just probabilities, but also relative scarcity, across states.

Risk sharing For ω^t, take the ratio of equation (7.20) for household i and household j:

$$\frac{u'(c_{i,t}(\omega^t))}{u'(c_{j,t}(\omega^t))} = \frac{\lambda_i}{\lambda_j}.$$

Now solve for the consumption of household i in terms of that of household j

$$c_{i,t}(\omega^t) = u'^{-1}\left(\frac{\lambda_i}{\lambda_j} u'(c_{j,t}(\omega^t))\right).$$

Goods market clearing requires $\sum_i c_{i,t}(\omega^t) = \sum_i y_{i,t}(\omega^t)$, so we obtain

$$\sum_i u'^{-1}\left(\frac{\lambda_i}{\lambda_j} u'(c_{j,t}(\omega^t))\right) = \sum_i y_{i,t}(\omega^t). \tag{7.21}$$

Equation (7.21) relates the consumption of household j to the aggregate supply of goods and the Lagrange multipliers of all the households. Importantly, those Lagrange multipliers are constant across time so $c_{j,t}(\omega^t)$ varies over time as a function of aggregate income not as a function of $y_{j,t}$. This is a very important result for complete markets models: all idiosyncratic risk is insured away and consumption fluctuations only reflect aggregate risks.

[14] In a static context, an actuarially fair gamble is one in which the expected payoff is equal to the cost of the gamble. To adapt this definition to a dynamic context, let's say prices are actuarially fair if the price of any payoff at date t is equal to the price of any other payoff at date t that has the same expected payoff. The condition $p_t(\omega^t) = \bar{p}_t \times \pi_t(\omega^t)$ imposes this definition. To see this, consider a set of payoffs at different histories that can arise at date t given by $x_t(\omega^t)$. The date-0 value of such a portfolio of claims is $\sum_{\omega^t} x_t(\omega^t) p_t(\omega^t) = \bar{p}_t \mathbb{E}[x_t(\omega^t)]$, which only depends on the date and the expected payoff.

Aggregate and idiosyncratic risks Aggregate risks lead to movements in aggregate variables such as aggregate income or aggregate consumption while idiosyncratic risks affect an individual's circumstances but do not affect the aggregate. One person becoming unemployed is an example of an idiosyncratic shock while an event that changes the unemployment rate is an example of an aggregate shock. If the set of consumers, $\mathcal{I}$, is finite then an individual shock, by definition, is an aggregate shock, albeit a small one if $\mathcal{I}$ has many elements. If $\mathcal{I} = [0, 1]$, then the set is (even uncountably) infinite and one individual's shock is truly idiosyncratic, unless it is synchronized/correlated with the shocks of others. To illustrate, an interesting case is precisely that where the shock is "employment" (say, with income y_e for the individual) or "unemployment" (with income $y_u < y_e$). The date-t shock ω_t would then specify a whole function taking, for each $i \in [0, 1]$, the value e or u. The whole event tree would even be hard to imagine. A special case is where individuals' employment outcomes are independent draws with probabilities π_e and $\pi_u = 1 - \pi_e$, respectively. Then if we can appeal to a law of large numbers there would be no aggregate uncertainty: aggregate resources would be a deterministic value $\pi_e y_e + \pi_u y_u$. However, each individual faces uncertainty, though in this case markets can allow full insurance. This kind of model, where the law of large numbers is assumed to hold, is often used in macroeconomics.[15] Now imagine that π_e is random: an aggregate shock, which itself could, e.g., take on two values (say, high or low) as well vary over time. Then individuals' shocks are correlated, though if one conditions on the aggregate shock (high or low unemployment), their shocks can be thought of as purely idiosyncratic and uncorrelated. Our notation involving $\mathcal{I}$ and ω^t is abstract and meant to capture all these possibilities.

Sequential trading We can implement a complete set of markets with an alternative trading arrangement in which agents only trade securities that pay off in the next period and then trade again every period. This parallels our two ways to define equilibrium in deterministic contexts in Sections 5.2–5.3: we now merely have stochastic sequences.

For each event ω_{t+1} that can occur at $t + 1$, there is an asset traded at t that pays one unit at $t + 1$ if that event occurs and zero otherwise. These are known as Arrow securities. Let $q_t(\omega_{t+1}|\omega^t)$ be the price at t of a unit of consumption at $t + 1$ if event ω_{t+1} occurs. This price can depend on the history leading up to date t, ω^t, and is denominated in terms of consumption after history ω^t. Let $a_{i,t+1}(\omega^{t+1})$ be the amount of this asset held by household i. The budget constraint of the household is then

$$c_{i,t}(\omega^t) + \sum_{\omega_{t+1}} q_t(\omega_{t+1}|\omega^t) a_{i,t+1}(\omega^{t+1}) \leq y_{i,t}(\omega^t) + a_{i,t}(\omega^t). \qquad (7.22)$$

Financial wealth, $a_{i,t}(\omega^t)$, becomes a state variable for the household's problem. This wealth allows the household to consume more than its income stream in the current period and future periods. When financial wealth is negative, the household must consume less than its income either now or in the future.

While agents only trade assets that pay off one period in the future, they can use these asset prices to value payoffs further in the future. One unit of goods at $t + 2$ after history ω^{t+2} has a value in date t of $q_{t+1}(\omega_{t+2}|\omega^{t+1}) \times q_t(\omega_{t+1}|\omega^t)$. In this product, the first term

[15] Such an assumption involves mathematical subtleties; see, e.g., Uhlig (1996).

discounts the unit of goods back to $t+1$ and the second term discounts it from $t+1$ to t. In general, we can define these discounts recursively as

$$\tilde{q}^t_{\tau+1}(\omega^{\tau+1}) = q_\tau(\omega_{\tau+1}|\omega^\tau)\tilde{q}^t_\tau(\omega^\tau)$$

with $\tilde{q}^t_t(\omega^t) = 1$. While the households do not trade assets for dates $\tau > t+1$ at date t, they do correctly anticipate the prices that will prevail in the future and the $\tilde{q}$ terms reflect these expectations.

At date 0, we assume that households have no financial wealth positive or negative because we assume that no trades have occurred prior to date 0 and so no household has a financial claim on any other household. Similar to models without uncertainty, the no Ponzi game constraint requires that the household could repay if it consumes nothing forever:

$$a_{i,t}(\omega^t) \geq -\sum_{\tau=t}^{\infty}\sum_{\omega^\tau}\tilde{q}^t_\tau(\omega^\tau)y_{i,\tau}(\omega^\tau). \tag{7.23}$$

Notice that this constraint rules out Ponzi games: it is the "natural borrowing limit," discussed in Section 4.3.1, now applying state by state.

We can now write the household's problem as

$$\max_{\{c_{i,t}(\omega^t),a_{i,t}(\omega_t)\}_{\forall t,\omega^t}}\sum_t\sum_{\omega^t}\beta^t\pi_t(\omega^t)u(c_{i,t}(\omega^t))$$

such that (7.22) and (7.23) hold for all t and ω^t.

A competitive equilibrium is a consumption allocation $c_{i,t}(\omega^t)$ for all i, t, and ω^t; asset positions $a_{i,t}(\omega^t)$ for all i, t, and ω^t; a price system $q_t(\omega_{t+1}|\omega^t)$ for all t, ω_{t+1} and ω^t such that (i) for all i, the consumption-savings plan is optimal taking the prices, borrowing constraints, and $a_{i,0} = 0$ as given; (ii) for all t and ω^t, the goods markets clear: $\sum_i(c_{i,t}(\omega^t) - y_{i,t}(\omega^t)) = 0$; (iii) for all t and ω^t, the asset market clears $\sum_i a_{i,t}(\omega^t) = 0$.

Definition 14 *A **sequential competitive equilibrium** is a set of stochastic sequences* $\left\{c^\star_{i,t}(\omega^t) : \forall t, \omega^t\right\}$ *and* $\left\{a^\star_{i,t+1}(\omega^t) : \forall t, \omega^t\right\}$ *for each* $i \in \mathcal{I}$*, and* $\{q_t(\omega^t) : \forall t, \omega^t\}$ *such that*

1. *for each* i, $\left(\left\{c^\star_{i,t}(\omega^t) : \forall t, \omega^t\right\}, \left\{a^\star_{i,t+1}(\omega^t) : \forall t, \omega^t\right\}\right)$ *solves*

$$\max_{\{c_t(\omega^t):\forall t,\omega^t\},\{a_{t+1}(\omega^t):\forall t,\omega^t\}}\sum_{t=0}^{\infty}\sum_{\omega^t\in\Omega^t}\beta^t\pi(\omega^t)u(c_t(\omega^t))$$

subject to

$$c_{i,t}(\omega^t) + \sum_{\omega_{t+1}}q_t(\omega_{t+1}|\omega^t)a_{i,t+1}(\omega^{t+1}) = y_{i,t}(\omega^t) + a_{i,t}(\omega^t) \tag{7.24}$$

and (nPg)

$$a_{i,t}(\omega^t) \geq -\sum_{\tau=t}^{\infty}\sum_{\omega^\tau}\tilde{q}^t_\tau(\omega^\tau)y_{i,\tau}(\omega^\tau) \tag{7.25}$$

2. $\sum_i \left(c^\star_{i,t}(\omega^t) - y_{i,t}(\omega^t)\right) = 0$ *and* $\sum_i a^\star_{i,t+1}(\omega^t) = 0$ *for all* t *and* ω^t.

Here, requirement 2 has two conditions (for each date and state); one of these implies the other.

As in the case of certainty, the equilibrium allocation under sequential trading is the same as the one that arises with date-0 trading. To prove this, the key step is to add the stochastic sequence of budget constraints, multiplied by the appropriate prices, to arrive at a time-zero consolidated constraint (after using the nPg condition). Then, after seeing how the prices in the two settings map into each other, it becomes clear that the consumers solve the same problems in the two equilibrium definitions.

Spanning and complete markets Arrow securities are a convenient modeling device, but they are not recognizable as assets that we normally trade. Most real-world assets pay off in more than one state of nature. A system of markets would still be complete if we can construct portfolios of the available assets that have payoffs equivalent to a full set of Arrow securities.

Suppose there are S states of the world that might be realized at date $t+1$ and there are N assets traded at each date. We can construct the $S \times N$ payoff matrix D that lists what each asset pays in each state. A portfolio is a vector $\theta \in \mathbb{R}^N$ that lists the weights on each asset. The $S \times 1$ vector listing the payoff of a portfolio θ in each state of the world is given by $D\theta$.

The range of the matrix D is the space of all possible payoffs that can be constructed by making portfolios of the N assets. Let

$$\mathcal{M} \equiv \left\{z \in \mathbb{R}^S : z = D\theta \text{ for some } \theta \in \mathbb{R}^N\right\}.$$

If $\mathcal{M} = \mathbb{R}^S$, the system of markets is complete. In this sense, a complete market means that one can construct a portfolio with any conceivable payoff vector. A system of markets will be complete if and only if $\text{rank}(D) = S$. If this rank condition is satisfied by the N assets, we say the assets **span the payoff space**. If a system of markets is complete, there are portfolios $\{\theta_j^{\text{A}}\}_{j=1}^S$ such that $D\theta_j^{\text{A}}$ pays one unit if state j occurs and zero units otherwise. The portfolios $\{\theta_j^{\text{A}}\}_{j=1}^S$ are the Arrow securities. For much more on this, see Chapter 16.

7.5 Competitive equilibrium in the growth model

We now use the framework of trade under uncertainty we just developed to define a competitive equilibrium for the stochastic growth model. We will state all of the assumptions here even though some of them were already introduced in Section 7.3.

We assume, for simplicity, that all individuals are identical. The representative household is endowed with k_0 units of capital and one unit of labor that is supplied inelastically. The intertemporal utility function is

$$\sum_t \sum_{\omega^t} \beta^t \pi_t(\omega^t) u(c_t(\omega^t)).$$

Output is produced according to

$$y_t(\omega^t) = A_t(\omega^t)F(k_t(\omega^{t-1}), 1),$$

where F has the usual neoclassical properties. The aggregate resource constraint is

$$k_{t+1}(\omega^t) + c_t(\omega^t) = (1-\delta)k_t(\omega^{t-1}) + y_t(\omega^t).$$

Turning to markets, the representative household accumulates capital and rents it to a representative firm in a spot market at price $r_t(\omega^t)$. The total, gross, return is $r_t(\omega^t)+1-\delta$, thus also inclusive of the undepreciated capital. Similarly, the household rents its labor to the firm in a spot market at a price $w_t(\omega^t)$.

The firm's problem is exactly as we have discussed before and results in the first-order conditions

$$r_t(\omega^t) = A_t(\omega^t)F_1(k_t(\omega^{t-1}), 1) \tag{7.26}$$

and

$$w_t(\omega^t) = A_t(\omega^t)F_2(k_t(\omega^{t-1}), 1). \tag{7.27}$$

The household's budget constraint is

$$c_t(\omega^t) + k_{t+1}(\omega^t) = (r_t(\omega^t) + 1 - \delta)k_t(\omega^{t-1}) + w_t(\omega^t).$$

We could allow the household to trade Arrow securities contingent on ω^{t+1} but, without another party to trade with, the representative household must have a zero position in each security in equilibrium. Therefore, although the consumer faces incomplete markets here, a full set of state-contingent assets would not change the equilibrium allocation. The household cannot hold a negative capital position, $k \geq 0$, but we will assume this constraint does not bind.

The Lagrangian of the household's problem is

$$\mathcal{L} = \sum_{t=0}^{\infty}\sum_{\omega^t\in\Omega^t}\left\{\beta^t\pi_t(\omega^t)u(c_t(\omega^t)) + \lambda_t(\omega^t)\left[(r_t(\omega^t)+1-\delta)k_t(\omega^{t-1}) + w_t(\omega^t) - c_t(\omega^t) - k_{t+1}(\omega^t)\right]\right\}.$$

Taking the first-order conditions with respect to $c_t(\omega_t)$ and $k_{t+1}(\omega^t)$ we have

$$\beta^t\pi_t(\omega^t)u'(c_t(\omega^t)) = \lambda_t(\omega^t)$$
$$\lambda_t(\omega^t) = \sum_{\omega^{t+1}|\omega^t}(r_{t+1}(\omega^{t+1}) + 1 - \delta)\lambda_{t+1}(\omega^{t+1})$$

and combining these we arrive at an Euler equation of

$$u'\left(c_t(\omega^t)\right) = \mathbb{E}_t\left[\beta u'\left(c_{t+1}(\omega^{t+1})\right)\left(r_{t+1}(\omega^{t+1}) + 1 - \delta\right)\right]. \tag{7.28}$$

A competitive equilibrium of this economy is a set of stochastic processes

$$\{r_t(\omega^t), w_t(\omega^t), c_t(\omega^t), k_{t+1}(\omega^t)\}_{\forall t,\omega^t}$$

such that $c_t(\omega^t)$ and $k_{t+1}(\omega^t)$ are optimal in the household's problem given the prices, the prices are set by competitive profit-maximizing firms in accordance with equations (7.26) and (7.27) and the resource constraint is satisfied.

Definition 15 *A **sequential competitive equilibrium** is a set of stochastic sequences* $\{c_t^\star(\omega^t) : \forall t, \omega^t\}$ *and* $\left\{k_{t+1}^\star(\omega^t) : \forall t, \omega^t\right\}$ *and* $\{(r_t(\omega^t), w_t(\omega^t) : \forall t, \omega^t\}$ *such that*

1. $\left(\{c_t^\star(\omega^t) : \forall t, \omega^t\}, \left\{k_{t+1}^\star(\omega^t) : \forall t, \omega^t\right\}\right)$ *solves*

$$\max_{\{c_t(\omega^t):\forall t,\omega^t\},\{k_{t+1}(\omega^t):\forall t,\omega^t\}} \sum_{t=0}^{\infty} \sum_{\omega^t \in \Omega^t} \beta^t \pi(\omega^t) u(c_t(\omega^t))$$

subject to the nPg condition and

$$c_t(\omega^t) + k_{t+1}(\omega^t) = (r_t(\omega^t) + 1 - \delta) k_t(\omega^{t-1}) + w_t(\omega^t) \tag{7.29}$$

2. *for all* t *and* ω^t,

$$r_t(\omega^t) = A_t(\omega^t) F_1(k_t^\star(\omega^{t-1}), 1) \quad \text{and} \quad w_t(\omega^t) = A_t(\omega^t) F_2(k_t^\star(\omega^{t-1}), 1)$$

3. *for all* t *and* ω^t,

$$k_{t+1}^\star(\omega^t) + c_t^\star(\omega^t) = (1-\delta) k_t^\star(\omega^{t-1}) + A_t(\omega^t) F(k_t^\star(\omega^{t-1}), 1).$$

In this definition of equilibrium the last condition is superfluous. This is, as in the deterministic case, because a constant returns to scale production function is homogeneous of degree one and by Euler's theorem we therefore have

$$r_t(\omega^t) k_t(\omega^{t-1}) + w_t(\omega^t) = A_t(\omega^t) F(k_t(\omega^{t-1}), 1) = Y_t(\omega^t).$$

If we substitute this into the household budget constraint we arrive at the aggregate resource constraint, which is the goods market clearing condition.

The competitive equilibrium allocation is the same as we obtain from the planner's problem. If we substitute equation (7.26) into (7.28) and we arrive at the same Euler equation as we obtained in the planner's problem, equation (7.13). Furthermore, if we substitute into (7.26) and (7.27) into the household budget constraint, we obtain the same resource constraint as applies to the planner's problem.

The equivalence between the planner's solution and the competitive equilibrium allocation is not affected by adding uncertainty to the model. Indeed, there is nothing fundamentally different about an economy with uncertainty as compared to a deterministic one. Instead of indexing goods just by time, we now index goods by time and histories. As a result, the first and second welfare theorems continue to apply.

In Section 7.3 we used the functional Euler equation to solve the planner's problem for the stochastic growth model. Due to the equivalence between the competitive equilibrium and the planner's problem, the solution we found is also the solution to the competitive equilibrium. In fact, in Appendix 7.A we formulate a recursive competitive equilibrium and derive the exact same functional Euler equation as we found for the planner's problem.

7.6 An incomplete-markets economy

In this chapter, we have mostly focused on competitive equilibria with complete markets. If a system of markets is incomplete, i.e., if some insurance contracts are not available, then the planner's problem and the competitive equilibrium will not coincide in general. In this case, the first welfare theorem fails to hold because some markets are missing (the markets for those insurance contracts). Here we will just give a brief introduction to incomplete-market models and return to this topic in later chapters.

Consider a consumer with preferences given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t).$$

The consumer receives a stochastic income stream y_t and can borrow or save in an asset that pays gross interest $1+r$, which is known and constant. Here we just describe the consumer's decision problem and we will take r as given.

Letting a_t be the agent's assets at the start of period t and $q = 1/(1+r)$, the budget constraint is

$$qa_{t+1} + c_t = a_t + y_t.$$

We assume there is some lower limit to how much the consumer can borrow, $a_{t+1} \geq \underline{a}$, where $\underline{a}$ could be the natural borrowing limit (the amount the consumer could repay if they received the lowest possible income realization in all future periods) or some more restrictive borrowing constraint.

If we assume the endowment process is a first-order Markov process, the consumer's problem can be stated recursively as

$$V(a, y) = \max_{a' \geq \underline{a}} \left\{ u(a + y - qa') + \beta \mathbb{E} \left[V(a', y') | y \right] \right\}.$$

The first-order condition of this problem is

$$u'(c)q = \beta \mathbb{E} \left[V_1(a', y') | y \right]$$

and the envelope condition is

$$V_1(a, y) = u'(c).$$

Combining these, and the fact that $q = 1/(1+r)$, we have the Euler equation

$$u'(c) = (1+r)\beta \mathbb{E} \left[u'(c') | y \right].$$

This is an example of an incomplete-market environment because the single asset cannot insure the consumer against the income risk. In other terms, using the single asset with a constant return, the consumer can borrow and save but the payoff of their portfolio is independent of their income in the next period. As a result, their consumption will fluctuate in response to the realizations in their individual endowment. However, the consumers can partially self insure. They can accumulate savings that they can spend down if they receive low endowments in the future. In this way the consumer can partially smooth their

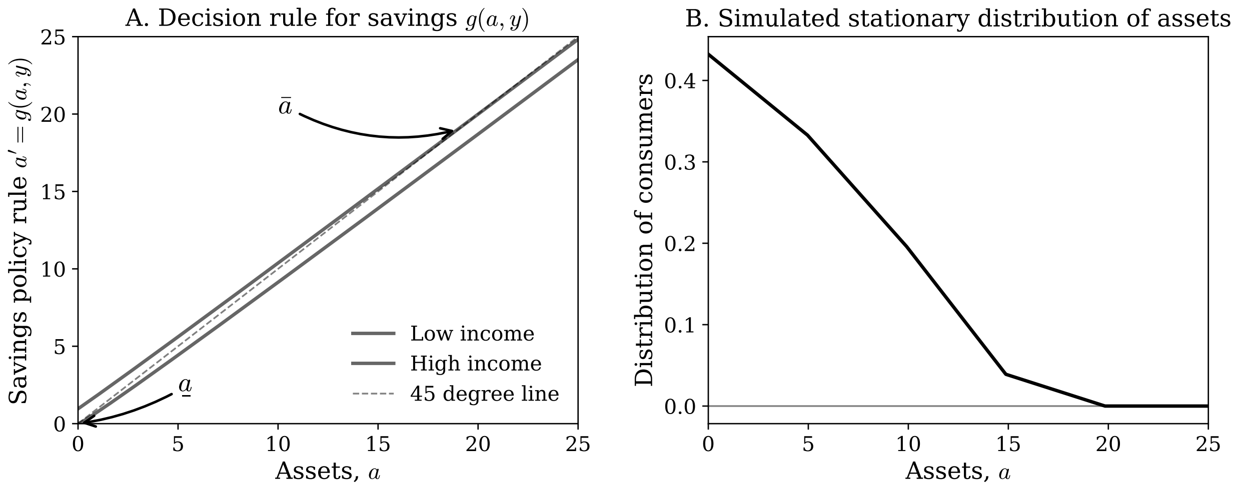

Figure 7.4: Decisions rules and stationary distribution for the incomplete-market consumption-savings problem.

consumption. However, they will not in general choose to fully smooth their consumption. To see this, suppose a consumer wished to have a perfectly smooth consumption path. The only way they could do this is to set their consumption to a level that would be sustainable if they received the lowest possible income realizations at all future dates. If they choose a higher level, there would be a chance that they would be unlucky and have to reduce their consumption. But this extremely conservative plan will not be optimal even for a consumer with very high risk aversion. Instead, the consumer will choose to adjust their consumption level in response to the endowments they receive. In Chapter 11, and later in Chapter 21, we study problems of self insurance in much more detail.

Panel (A) of Figure 7.4 plots the decision rule $a' = g(a, y)$ for a version of this model in which income can take two values each period. When income is high, the consumer accumulates savings up to the point $\bar{a}$ where the upper line crosses the 45-degree line. When income is low, the consumer spends down their savings until it reaches the borrowing constraint. If the consumer begins with initial assets above $\bar{a}$, they will continuously spend down their assets regardless of their income until their assets reach $\bar{a}$. So in the long run the consumer will have asset holdings on the ergodic set $[\underline{a}, \bar{a}]$. Within this set, however, the consumer will sometimes be moving towards higher asset levels (when they have high income) and will sometimes be moving towards lower levels (when they have low income). Contrast this saving rule with what we found in our deterministic steady state endowment economy in Section 5.4.1 where the decision rules were $a' = g(a)$ so they are simply the 45-degree line. In Figure 7.4, the decision rule of high-income consumers has a slope less than 45 degrees but an intercept above the 45 degree line while the decision rule of low-income consumers is on the 45-degree line at $\underline{a}$ but has a slope less than 45 degrees. The fact that the decision rules are above and below the 45-degree line is the source of partial insurance—those with high income put some of their "extra" resources into savings while those with low income draw down their savings.

Now suppose we simulate a large number of consumers each solving this decision problem and each receiving their own independent endowment process. The simulation produces a distribution of consumers over asset levels that reflects the fact that some consumers have been lucky and received many high endowments and some have been unlucky. Moreover, as consumption is a function of the consumer's assets and income, the simulation also produces a distribution of consumption levels. At a point in time, the distribution of asset holdings will depend in part on the distribution of asset holdings in the previous period. In the long run, the distribution will converge to a stationary distribution. Panel (B) of Figure 7.4 plots the stationary distribution of assets for our simulation.

In some contexts the complete market and representative agent assumptions are sensible simplifications of reality, but when they are not we often turn to incomplete-market models that build on the framework we have just presented. Such models can be used to study a wide range of issues and we will return to this topic in the chapters that follow.