Chapter 9

Continuous-time analytical techniques

In most chapters of this book, we develop models using a discrete-time framework. Discrete time is intuitive, aligns well with how we often approach data, and is typically more convenient for computation—especially when implementing numerical methods that require discretizing variables (see Chapter 10). However, in some cases, working in continuous time provides particular advantages. Certain economic problems become mathematically more tractable in continuous time, and key relationships—especially those involving differential equations—can be derived more transparently with "paper and pencil" techniques. Continuous-time models are especially helpful when we want to analyze economic variables that evolve smoothly over time, rather than in discrete jumps.

While discrete time remains the primary analytical tool throughout this textbook, the usefulness of continuous-time models will become clear in specific applications (such as those in Chapter 13). This chapter introduces the core analytical techniques specific to continuous-time models. Although most macroeconomic research and applications remain grounded in discrete time, it is important to become comfortable with continuous-time methods, as they offer valuable intuition and tools in particular research settings.

Discrete vs. continuous time: an intuitive overview Before introducing the formal notation and equations, it is helpful to clarify the distinction between discrete- and continuous-time models. In discrete time, we track economic variables—such as consumption, capital, or employment—at specific intervals (for example, monthly, quarterly, or yearly). This approach fits naturally with how economic data are typically measured and reported, and it facilitates computational work. Continuous-time models, in contrast, conceptualize time as a continuous flow, rather than a sequence of distinct intervals. Here, economic variables evolve smoothly at every instant, and their behavior is described by differential rather than difference equations.

To illustrate, consider the example of capital accumulation in an economy. In a discrete-time model, we observe the capital stock at the end of each period and model its change from one period to the next. With a continuous-time model, we track the capital stock's evolution at every moment, capturing the effects of investment and depreciation instantaneously. This provides a more precise and realistic depiction of gradual economic adjustments, especially when changes are ongoing rather than occurring in abrupt steps.

The rationale for adopting continuous-time frameworks arises from several key considerations. For some research problems—particularly those involving high-frequency dynamics, instantaneous decision-making, or complex intertemporal optimization—continuous-time models can substantially simplify the mathematics. Concepts such as instantaneous rates of change, derivatives, and integral calculus become central, often leading to elegant closed-form solutions that might be difficult or impossible to obtain in a discrete-time setting. This mathematical elegance, in turn, can foster deeper intuition and facilitate a more precise characterization of economic mechanisms.

9.1 Basic tools and notation in continuous time

To begin, let's clarify how we represent variables in continuous time. In discrete-time models, we use subscripts to indicate time—for example, X_t for $t = 0, 1, 2, \ldots$. In continuous-time models, however, we write variables as *functions* of time: $X(t)$, where t can take any non-negative real value ($t \in \mathbb{R}_+$). This change reflects the fact that, in continuous time, variables can evolve at every instant, not just at fixed intervals.

The key mathematical concept in continuous-time analysis is the *time derivative*, which measures the instantaneous rate of change of a variable. For a function $X(t)$, the time derivative is written as $dX(t)/dt$ or, more compactly, $\dot{X}(t)$. This "dot notation," introduced by Newton (and hence also called "Newton notation"), is commonly used in economics to denote time derivatives.

Let's start by connecting these ideas to something familiar: growth rates. In discrete time, if X_t grows at a constant rate γ, we can write the relationship as a *difference equation*:

$$\frac{X_{t+1} - X_t}{X_t} = \gamma, \tag{9.1}$$

which has the solution

$$X_t = (1 + \gamma)^t X_0. \tag{9.2}$$

A difference equation relates the change in a variable from one period to the next to its current level, and its solution is a sequence of values at discrete points in time. In continuous time, we express the same idea using a *differential equation*. If the *instantaneous* growth rate of $X(t)$ is γ, we write:

$$\frac{\dot{X}(t)}{X(t)} = \gamma. \tag{9.3}$$

Here, $\dot{X}(t)$ is the instantaneous rate of change of $X(t)$, and dividing by $X(t)$ gives the proportional (or percentage) rate of change at each moment. The solution to this differential equation can be written using the exponential function:

$$X(t) = e^{\gamma t} X(0), \tag{9.4}$$

where e is Euler's number (approximately 2.71828). In this form, $X(t)$ grows smoothly at rate γ for every instant $t \geq 0$. You can check this solution by substituting it back into equation (9.3). This can be verified by differentiating (9.4) with respect to time: using the properties of exponential functions and their derivatives (see Appendices 9.A.1 and 9.A.2), we have $\dot{X}(t) = \gamma e^{\gamma t} X(0) = \gamma X(t)$, which confirms that (9.4) solves the differential equation (9.3). It is important to recognize that the difference between discrete and continuous time extends to both notation and the types of equations used to describe dynamics. Difference equations (discrete time) generate sequences, while differential equations (continuous time) generate continuous functions of time.

A particularly useful property in continuous time is that the growth rate of $X(t)$ can be found by taking the natural logarithm and then differentiating (see Appendices 9.A.1 and 9.A.2):

$$\frac{\dot{X}(t)}{X(t)} = \frac{d}{dt} \log(X(t)). \tag{9.5}$$

This tells us that the derivative of the log of a variable gives its (instantaneous) growth rate. Intuitively, if a variable grows at a constant rate, its logarithm traces a straight line over time. This is why economists often plot time series—such as GDP or capital—on a log scale: constant growth appears as a straight line, making trends easier to spot. As an example, see Figure 2.1.

This property is particularly powerful when dealing with functions that involve products or powers, which are common in macroeconomics. For example, consider the Cobb-Douglas production function:

$$Y(t) = z(t) K(t)^{\alpha} L(t)^{1-\alpha},$$

where $Y(t)$ is output, $z(t)$ is total factor productivity (TFP), $K(t)$ is capital, and $L(t)$ is labor. Taking logs,

$$\log(Y(t)) = \log\left(z(t) K(t)^{\alpha} L(t)^{1-\alpha}\right) = \log(z(t)) + \alpha \log(K(t)) + (1-\alpha)\log(L(t)),$$

where the second equality uses properties of natural logs (see Appendix 9.A.1). Taking the time derivative on both sides and applying (9.5) gives:

$$\frac{\dot{Y}(t)}{Y(t)} = \frac{\dot{z}(t)}{z(t)} + \alpha \frac{\dot{K}(t)}{K(t)} + (1-\alpha)\frac{\dot{L}(t)}{L(t)}.$$

This is the growth accounting equation (2.1) in Chapter 2, where we also derived this equation in continuous time—precisely because it is particularly convenient. The difference here is that we are using a specific production function that makes the shares constant over time.

The formula (9.5) can also be used to derive (9.3) from (9.4). For example, if $X(t) = e^{\gamma t} X(0)$, taking logs gives $\log(X(t)) = \gamma t + \log(X(0))$. Differentiating with respect to t brings us back to γ, showing the connection between exponential growth and constant growth rates. Similarly, in discrete time, taking logs and differences of $X_t = (1+\gamma)^t X_0$ reveals that the growth rate is approximately γ when γ is small.[1]

[1]It's also helpful to see how these continuous-time formulas relate to the discrete-time case. For example, suppose $X_t = (1+\gamma)^t X_0$. Taking logs of both sides delivers $\log(X_t) = t\log(1+\gamma) + \log(X_0)$. Computing the difference, we obtain $\log(X_{t+1}) - \log(X_t) = \log(1+\gamma)$. Because this equation can be rewritten as $\log\left(1 + \frac{X_{t+1} - X_t}{X_t}\right) = \log(1+\gamma)$ and $\log(1+x) \approx x$ for a small x, it is approximately the case that $\frac{X_{t+1} - X_t}{X_t} = \gamma$.

9.2 Optimization in continuous-time models: Maximum principle

Having established the key notation and mathematical tools, we now turn to optimization in continuous time. In this section, we introduce the maximum principle, also known as the optimal control or Hamiltonian method. To motivate the continuous-time approach, we begin with a familiar benchmark: the finite-horizon consumption-saving problem, similar to Section 4.1.1.

The consumer solves the problem:

$$\max_{\{c_t, a_{t+1}\}_{t=0}^{T}} \sum_{t=0}^{T} \beta^t u(c_t),$$

subject to

$$a_{t+1} = w + (1+r)a_t - c_t, \qquad a_0 \text{ given,} \tag{9.6}$$

and a terminal (no-borrowing) condition

$$a_{T+1} \geq 0. \tag{9.7}$$

Here, $\beta \in (0,1)$ is the discount factor, $u(\cdot)$ is a strictly increasing, strictly concave utility function satisfying Inada conditions, c_t is consumption at period t, a_t is an asset at the beginning of period t, and $T > 0$ is the final period. The labor earnings w and interest rate r are given to the consumer.

We form a Lagrangian, slightly different from Section 4.1.1, as

$$L \equiv \sum_{t=0}^{T} \beta^t u(c_t) + \sum_{t=0}^{T} \mu_t((1+r)a_t + w - c_t - a_{t+1}) + \lambda a_{T+1},$$

where μ_t $(t = 0, 1, ..., T)$ are Lagrange multipliers (or costate variables) and λ enforces the terminal constraint. As in Section 4.1.1, we can proceed by applying the Kuhn-Tucker theorem to this problem directly, but here, we take a bit of a detour. The Lagrangian can be rewritten by expressing the budget constraint (9.6) in difference form,

$$a_{t+1} - a_t = ra_t + w - c_t. \tag{9.8}$$

Define the discrete-time *Hamiltonian* as

$$H_t \equiv \beta^t u(c_t) + \mu_t(ra_t + w - c_t).$$

Note H_t is a function of c_t and a_t (and μ_t). With this, the Lagrangian is:

$$L = \sum_{t=0}^{T} H_t - \sum_{t=0}^{T} \mu_t(a_{t+1} - a_t) + \lambda a_{T+1}.$$

Then, the first-order conditions for c_t, $t = 0, 1, 2, ..., T$ are

$$\frac{\partial H_t}{\partial c_t} = 0 \tag{9.9}$$

and the first-order conditions for a_{t+1}, $t = 0, 1, 2, ..., T-1$ are

$$\frac{\partial H_{t+1}}{\partial a_{t+1}} + \mu_{t+1} - \mu_t = 0. \tag{9.10}$$

In this context, *control variables* (like c_t) are chosen at time t, while *state variables* (like a_t) capture the system's current position, carried over from previous choices. This distinction mirrors what was introduced in Chapter 4, Section 4.1. State variables cannot be changed at time t, but some state variables at period $t+1$ can be decided on at period t. The control variable maximizes the Hamiltonian directly, while the state variable's condition accounts for changes in the costate variable μ_t over time (e.g., the fact that μ_t can change over time has to be taken into account). This consideration results in the extra term $\mu_{t+1} - \mu_t$.

The terminal condition (e.g., the FOC with respect to a_{T+1}) yields

$$-\mu_T + \lambda = 0. \tag{9.11}$$

The Kuhn-Tucker conditions are $\lambda \geq 0$, $a_{T+1} \geq 0$, and

$$\lambda a_{T+1} = 0. \tag{9.12}$$

Condition (9.9) for $t = T$ is $\beta^T u'(c_T) = \mu_T$, and combining this relationship with (9.11) and (9.12), we obtain

$$\beta^T u'(c_T) a_{T+1} = 0. \tag{9.13}$$

This condition is often called the *transversality condition* (TVC). We have discussed this concept in Chapter 4. In the finite-horizon case, because $\beta^T > 0$ and $u'(c_T) > 0$, the implication is that $a_{T+1} = 0$. This result does not necessarily apply when the planning horizon is infinite.

The condition (9.9) for time t is

$$\beta^t u'(c_t) = \mu_t$$

and the condition (9.10) for time $t+1$ is

$$\mu_t = (1+r)\mu_{t+1}.$$

The standard Euler equation follows by combining FOCs:

$$u'(c_t) = \beta(1+r)u'(c_{t+1}). \tag{9.14}$$

Given the initial condition a_0, the budget constraint (9.6), and the TVC (9.13), this characterizes the optimal paths of c_t and a_t. Using the Euler equation repeatedly from time 0 to $T-1$,

$$\beta^T u'(c_T) = \frac{u'(c_0)}{(1+r)^T}$$

holds, and therefore (together with the fact that $u'(c_0)$ is a positive number), (9.13) is equivalent to

$$\frac{a_{T+1}}{(1+r)^T} = 0.$$

Once again, in the finite-horizon case here, this condition is equivalent to $a_{T+1} = 0$.

For an infinite horizon, the setup is analogous:

$$\max_{\{c_t, a_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t),$$

subject to

$$a_{t+1} = w + (1+r)a_t - c_t,$$

a_0 is given, and

$$\lim_{T\to\infty} \frac{a_{T+1}}{(1+r)^T} \geq 0. \tag{9.15}$$

We assume r is not too large, so that the present value of utility is finite. Note the last constraint looks somewhat different from the finite-horizon case (9.7). In the finite horizon, $a_{T+1}/(1+r)^T \geq 0$ is equivalent to $a_{T+1} \geq 0$ because $(1+r)^T > 0$ for any T, so it does not make a difference. In the infinite horizon, $\lim_{T\to\infty} a_{T+1}/(1+r)^T \geq 0$ is not equivalent to $\lim_{T\to\infty} a_{T+1} \geq 0$ (consider the case where $a_t = a > 0$ for all t; this satisfies the former inequality but not the latter inequality). See Section 4.3.1 for the discussion of why (9.15), called the no-Ponzi game (nPg) condition, is the appropriate one in this case.

As in the finite-horizon case, rewrite the budget constraint as $a_{t+1} - a_t = ra_t + w - c_t$ and define the Hamiltonian

$$H_t \equiv \beta^t u(c_t) + \mu_t(ra_t + w - c_t),$$

where μ_t is the Lagrange multiplier.

Then, the first-order conditions are

$$\frac{\partial H_t}{\partial c_t} = 0$$

for $t = 0, 1, ...$ and

$$\frac{\partial H_{t+1}}{\partial a_{t+1}} + \mu_{t+1} - \mu_t = 0.$$

These two equations can be combined to obtain the same Euler equation as in the finite-horizon case (9.14). The TVC is

$$\lim_{T\to\infty} \beta^T u'(c_T) a_{T+1} = 0.$$

This condition is analogous to the finite-horizon expression (9.13). As in the finite-horizon case, this condition can also be expressed as

$$\lim_{T\to\infty} \frac{a_{T+1}}{(1+r)^T} = 0.$$

See Section 4.3.1 for further discussions on this condition.

The continuous-time version is formulated by switching notation from X_t to $X(t)$ and replacing sums with integrals:

$$\max_{c(t), a(t)} \int_0^{\infty} e^{-\rho t} u(c(t)) dt$$

subject to

$$\dot{a}(t) = w + ra(t) - c(t), \tag{9.16}$$

$a(0)$ is given, and

$$\lim_{T\to\infty} e^{-rT} a(T) \geq 0. \tag{9.17}$$

Let's begin by examining the objective function. The period utility $u(\cdot)$ (often called the "instantaneous utility" or "moment utility" in the continuous-time setting) is identical to the one in discrete time. However, instead of summing utility across periods, we now take an integral to capture utility received at every instant in time. One key difference in the continuous-time formulation is how we discount future utility. Here, $\rho > 0$ is known as the *discount rate*, and discounting is accomplished using the exponential term $e^{-\rho t}$. This serves a similar purpose to β^t in discrete time: it reduces the weight of future utility relative to the present. Note that a larger discount rate ρ means the consumer places less value on future utility. By contrast, in discrete time, a larger discount factor β (closer to 1) means the consumer cares more about the future. For a formal explanation of the connection between ρ and β, see the box below.

The discount rate ρ

Suppose that, in a discrete-time setup, we discount the next period using the discount factor $\beta \in (0, 1)$. We define the *discount rate* ρ by the value that satisfies

$$\beta = \frac{1}{1+\rho}. \tag{9.18}$$

When we discount t periods ahead, we multiply β^t, or equivalently, $[1/(1+\rho)]^t$. Now, imagine we split one period into many subperiods, each of which has a length Δ. For example, when one period is one year, we can think of Δ being one month and $1/\Delta = 12$. Let us discount one subperiod by the rate $\rho\Delta$. Thus, when we discount a variable for t periods (which is equal to t/Δ subperiods), we multiply

$$f(\Delta, t) = \left(\frac{1}{1+\rho\Delta}\right)^{\frac{t}{\Delta}}.$$

When we take a limit $\Delta \to 0$,

$$\lim_{\Delta\to 0} f(\Delta, t) = \lim_{\Delta\to 0} \left(\frac{1}{1+\rho\Delta}\right)^{\frac{t}{\Delta}} = \lim_{\Delta\to 0} \left[(1+\rho\Delta)^{1/(\rho\Delta)}\right]^{-\rho t} = e^{-\rho t},$$

where we used the fact $e = \lim_{x\to 0}(1+x)^{1/x}$. Thus, the continuous-time discount rate ρ that shows up in $e^{-\rho t}$ can be interpreted as analogous to ρ in (9.18) when a period can be divided into many subperiods so that we can discount "continuously."

Constraint (9.16) is similar to (9.8), except that the change in the asset is expressed as $\dot{a}(t)$ instead of $a_{t+1} - a_t$. The nPg condition (9.17) is similar to (9.15), except that the discounting is e^{-rT} instead of $1/(1+r)^T$. We can map these two with each other using the same procedure as in the above box. When r varies over time (denote it r_t or $r(t)$), we can generalize the nPg condition for the discrete-time case to

$$\lim_{T\to\infty} \frac{a_{T+1}}{\prod_{t=1}^{T}(1+r_t)} \geq 0$$

and the continuous-time case to

$$\lim_{T\to\infty} e^{-\int_0^T r(t)} a(T) \geq 0.$$

See Appendix 9.A.3 for detailed exposition.

Now, let us solve the problem. The steps are the same as in the discrete-time case. Define the Hamiltonian

$$H(t) \equiv e^{-\rho t} u(c(t)) + \mu(t)(ra(t) + w - c(t)), \tag{9.19}$$

where $\mu(t)$ is the costate variable. The first-order conditions become:

$$\frac{\partial H(t)}{\partial c(t)} = 0 \tag{9.20}$$

and

$$\frac{\partial H(t)}{\partial a(t)} + \dot{\mu}(t) = 0, \tag{9.21}$$

with the transversality condition

$$\lim_{T\to\infty} e^{-\rho T} u'(c(T)) a(T) = 0.$$

The only differences from the discrete-time case are that (i) the change in $\mu(t)$ is expressed as $\dot{\mu}(t)$ (instead of $\mu_{t+1} - \mu_t$) in (9.21), (ii) in (9.21), the relevant derivative of Hamiltonian is $\partial H(t)/\partial a(t)$ instead of $\partial H_{t+1}/\partial a_{t+1}$, because continuous-time setting has no "next period," and (iii) the discounting for the TVC is now $e^{-\rho T}$.

One can heuristically derive (9.20) and (9.21) in a similar manner as in the discrete-time case. Construct the Lagrangian

$$L = \int_0^\infty e^{-\rho t} u(c(t)) dt + \int_0^\infty \mu(t)(ra(t) + w - c(t) - \dot{a}(t)) dt.$$

(We ignore the terminal-condition issues here for the sake of exposition.) By rewriting the Hamiltonian as (9.19), the Lagrangian can be rewritten as

$$L = \int_0^\infty H(t) dt - \int_0^\infty \mu(t)\dot{a}(t) dt.$$

Applying integration by parts to the second term,

$$L = \int_0^\infty H(t) dt - \left[\lim_{T\to\infty} \mu(T)a(T) - \mu(0)a(0)\right] + \int_0^\infty \dot{\mu}(t)a(t) dt.$$

Taking the first-order conditions on $c(t)$ and $a(t)$ leads to (9.20) and (9.21).

Present-value Hamiltonians versus current-value Hamiltonians

The Hamiltonian that is used in this chapter is often called the *present-value Hamiltonian*. An alternative formulation is called the *current-value Hamiltonian*. The current-value Hamiltonian evaluates the utility at the current value, without discounting. In the context of the above consumption-saving problem, the current-value Hamiltonian is

$$\hat{H}(t) \equiv u(c(t)) + \hat{\mu}(t)(ra(t) + w - c(t)).$$

With the current-value Hamiltonian, the first-order conditions are modified to

$$\frac{\partial \hat{H}(t)}{\partial c(t)} = 0 \tag{9.22}$$

and

$$\frac{\partial \hat{H}(t)}{\partial a(t)} + \dot{\hat{\mu}}(t) - \rho\hat{\mu}(t) = 0. \tag{9.23}$$

Equation (9.22) corresponds to (9.20), and (9.23) corresponds to (9.21). One can also see the relationship $\mu(t) = e^{-\rho t}\hat{\mu}(t)$.

Using (9.20) and (9.21), we can derive the continuous-time version of the Euler equation. First, (9.20) can be written as

$$e^{-\rho t}u'(c(t)) = \mu(t). \tag{9.24}$$

The first-order condition for the state variable, (9.21), can be calculated as

$$r\mu(t) + \dot{\mu}(t) = 0. \tag{9.25}$$

To eliminate $\mu(t)$, first rewrite (9.25) as

$$\frac{\dot{\mu}(t)}{\mu(t)} = -r.$$

Applying the growth trick (9.5) to (9.24),

$$-\rho + \frac{u''(c(t))}{u'(c(t))}\dot{c}(t) = \frac{\dot{\mu}(t)}{\mu(t)}.$$

Combining these two, we obtain

$$-\frac{u''(c(t))c(t)}{u'(c(t))}\frac{\dot{c}(t)}{c(t)} = r - \rho. \tag{9.26}$$

This equation is the continuous-time version of the Euler equation. The term $-u''(c(t))c(t)/u'(c(t))$ is called the *coefficient of relative risk aversion*, and because $u''(\cdot) < 0$, $c(t) > 0$, and $u'(\cdot) > 0$, the coefficient is always positive. In a class of utility functions, called the CRRA utility (introduced in Section 4.2.4), where $u(c) = (c^{1-\sigma} - 1)/(1 - \sigma)$, the coefficient is constant: $-u''(c(t))c(t)/u'(c(t)) = \sigma$.

Discrete-time Euler equations versus continuous-time Euler equations

Although the continuous-time Euler equation (9.26) and the discrete-time Euler equation (9.14) may look different, their economic interpretation and intuition are fundamentally the same. In the continuous-time case, equation (9.26) shows that the growth rate of consumption, $\dot{c}(t)/c(t)$, is determined by the balance between the benefit of saving (measured by the interest rate r) and the desire to consume sooner rather than later (measured by the discount rate ρ). If the consumer discounts the future heavily (i.e., ρ is large), then the growth rate of consumption will be low or even negative—meaning the consumer prefers to enjoy more consumption today. Conversely, when the discount rate ρ is small, the consumer is more willing to postpone consumption, leading to faster consumption growth over time.

A similar relationship is reflected in the discrete-time Euler equation (9.14). Here, a small discount factor β (corresponding to strong discounting) implies that marginal utility today, $u'(c_t)$, must be relatively low compared to marginal utility tomorrow, $u'(c_{t+1})$. This suggests that the consumer allocates more consumption to the present than to the future. To see the connection more clearly, consider a first-order Taylor expansion: $u'(c_{t+1}) = u'(c_t + [c_{t+1} - c_t]) \approx u'(c_t) + u''(c_t)(c_{t+1} - c_t)$ and recall that $\beta = 1/(1+\rho)$. Substituting this into (9.14), we obtain:

$$-\frac{u''(c_t)c_t}{u'(c_t)}\frac{c_{t+1} - c_t}{c_t} = \frac{r - \rho}{1 + r},$$

which closely parallels the structure of the continuous-time Euler equation (9.26). This demonstrates that, despite the difference in notation, both approaches capture the same underlying trade-off between impatience and the rewards to saving.

9.3 Continuous-time growth models

In earlier chapters—specifically, Chapters 3 and 4—we introduced and analyzed the Solow model and the neoclassical growth model in discrete time. Those chapters provided the foundation for understanding long-run economic growth and dynamic optimization using the language of difference equations. Here, we revisit these classic models in continuous time. The purpose is twofold: first, to show how the same economic mechanisms can be formulated and analyzed using differential equations; and second, to provide hands-on examples that illustrate the tools and intuition developed earlier in this chapter. The continuous-time approach offers a more streamlined and, in some cases, more transparent characterization of the transition dynamics and steady states of growth models. For students, working through these continuous-time versions also clarifies the connections—and the distinctions—between discrete- and continuous-time techniques.

We begin with the continuous-time version of the Solow model, followed by the neoclassical growth model. In each case, we highlight both the economic interpretation and the formal differences that arise when moving from discrete to continuous time.

9.3.1 The Solow model

Let us begin by revisiting the Solow model with economic growth, first introduced in Chapter 3.2. The Solow model remains a cornerstone of growth theory, illustrating how capital accumulation and technological progress drive the expansion of output per worker over time. By expressing the model in continuous time, we can see clearly how the economy's capital stock evolves at each instant, and how the interplay of savings, depreciation, population growth, and technological progress shape the long-run steady state.

The aggregate production function takes the form

$$Y(t) = F(K(t), A(t)L(t)),$$

where $Y(t)$ denotes aggregate output (GDP), $K(t)$ is aggregate capital, $A(t)$ represents the level of technology, and $L(t)$ is aggregate labor. The assumptions on the aggregate production function $F(\cdot,\cdot)$ are the same as in Chapter 3 (see Section 3.1). We assume that both labor and technology grow at constant rates:

$$\frac{\dot{L}(t)}{L(t)} = n \tag{9.27}$$

and

$$\frac{\dot{A}(t)}{A(t)} = \gamma. \tag{9.28}$$

The capital stock evolves according to the law of motion:

$$\dot{K}(t) = I(t) - \delta K(t), \tag{9.29}$$

where $I(t)$ is aggregate investment and $\delta > 0$ is the depreciation rate. As in Chapter 3, we assume that a constant fraction $s \in (0,1)$ of aggregate income (which is the same as $Y(t)$) is saved and therefore invested:

$$I(t) = sF(K(t), A(t)L(t)). \tag{9.30}$$

To analyze the system's dynamics, it is helpful to express everything in terms of capital per effective unit of labor. Define

$$\tilde{k}(t) \equiv \frac{K(t)}{A(t)L(t)}.$$

Using the technique we learned in Section 10.1, we can obtain the growth rate of $k(t)$ as

$$\frac{\dot{\tilde{k}}(t)}{\tilde{k}(t)} = \frac{\dot{K}(t)}{K(t)} - \frac{\dot{A}(t)}{A(t)} - \frac{\dot{L}(t)}{L(t)}. \tag{9.31}$$

From (9.27), (9.28), (9.29), and (9.30), we can rewrite (9.31) as

$$\frac{\dot{\tilde{k}}(t)}{\tilde{k}(t)} = \frac{sF(K(t), A(t)L(t)) - \delta K(t)}{K(t)} - \gamma - n.$$

The left-hand side, $\frac{\dot{\tilde{k}}(t)}{\tilde{k}(t)}$, represents the proportional rate of change in capital per effective worker. The first term on the right is the net rate at which new capital is being created relative to the current capital stock. The subtracted terms, γ and n represent the rates of technological progress and population growth, respectively. To simplify further, divide the numerator and denominator of the first term by $A(t)L(t)$. Then, using the constant-returns property of the production function $F(\cdot,\cdot)$, we have:

$$\frac{F(K(t), A(t)L(t))}{A(t)L(t)} = F\left(\frac{K(t)}{A(t)L(t)}, 1\right) = f(\tilde{k}(t)),$$

where the first equality uses the constant-returns assumption on $F(\cdot,\cdot)$, and the second equality uses the definition $f(k) \equiv F(k,1)$ from Chapter 3. Using these definitions, we obtain

$$\frac{\dot{\tilde{k}}(t)}{\tilde{k}(t)} = \frac{sf(\tilde{k}(t))}{\tilde{k}(t)} - (\delta + \gamma + n).$$

Multiplying both sides by $\tilde{k}(t)$ yields a differential equation for the evolution of capital per effective worker that reads as follows.[2]

$$\dot{\tilde{k}}(t) = sf(\tilde{k}(t)) - (\delta + \gamma + n)\tilde{k}(t). \tag{9.32}$$

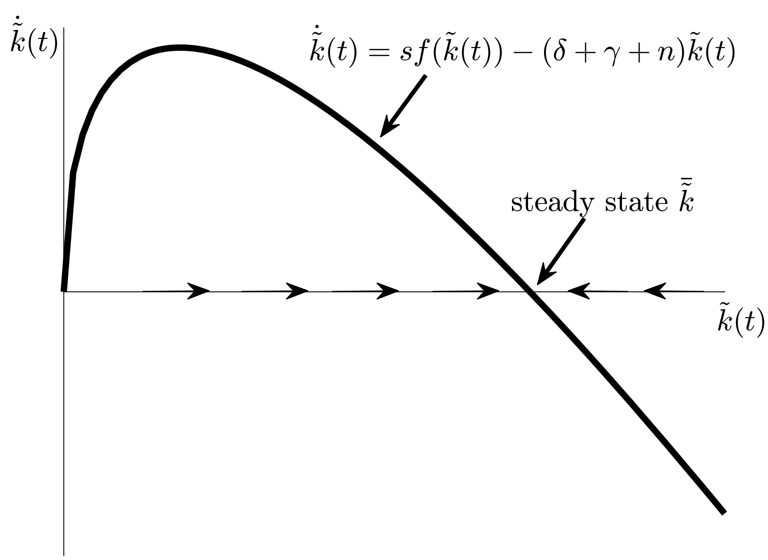

Figure 9.1: Solow model in continuous time

The dynamics described by equation (9.32) can easily be analyzed using the diagram depicted in Figure 9.1 (which is analogous to Figure 3.1 in the discrete-time version). The vertical axis measures the rate of change of capital per effective worker, $\dot{\tilde{k}}(t)$—that is, the

[2]This expression corresponds to the difference equations (3.5) and (3.7) in Chapter 3.

change of $\tilde{k}(t)$ over time—while the horizontal axis shows its current level of capital per effective worker, $\tilde{k}(t)$. The curve plotted in the figure represents the right-hand side of equation (9.32), incorporating the variables that drive the evolution of capital over time. The point where the curve crosses the horizontal axis corresponds to the steady state, denoted by $\bar{\tilde{k}}$. At this point, $\dot{\tilde{k}}(t) = 0$, implying that capital per effective worker remains unchanged over time. The behavior of the economy around this steady state reveals an important property of the Solow model: *convergence*. When $\tilde{k}(t)$ is to the left of the steady state ($\tilde{k}(t) < \bar{\tilde{k}}$), the value of $\dot{\tilde{k}}(t)$ is positive. This means that capital per effective worker is increasing, so the economy moves to the right in the diagram, gradually approaching the steady state. Conversely, when $\tilde{k}(t)$ exceeds the steady-state value ($\tilde{k}(t) > \bar{\tilde{k}}$), the rate of change $\dot{\tilde{k}}(t)$ becomes negative. In such case, capital per effective worker declines, and the economy moves to the left, once again approaching the steady state.

This dynamic adjustment process means that regardless of the economy's initial position, it will tend to move toward the steady-state level of capital per effective worker over time: from either direction, $\tilde{k}(t)$ gradually approaches $\bar{\tilde{k}}$. In the long run, the economy approaches the steady state, regardless of where it starts. The steady state itself is sometimes referred to as the balanced growth path, because along this path, the key economic variables—such as output per worker and capital per worker—grow at constant rates, determined by the rate of technological progress. In the long run, as the economy converges to the steady state, output per worker increases at the same rate as technology, and the overall path of the economy becomes predictable and stable. For example, because the per-capita GDP, $Y(t)/L(t)$, can be rewritten as

$$\frac{F(K(t), A(t)L(t))}{L(t)} = f(\tilde{k}(t))A(t)$$

and $f(\tilde{k}(t))$ is constant in the steady state, the growth rate of $Y(t)/L(t)$ in the long run is identical to the growth rate of $A(t)$, which is γ.

9.3.2 The neoclassical growth model

Let us now turn to the neoclassical growth model, building on the analysis introduced in 4.1.2. To simplify the exposition, suppose there is no population growth or technological progress; these features can be incorporated using the same techniques as discussed previously and further developed in Section 4.3.3.

The planner's problem in continuous time

Consider the social planner's problem, formulated here in continuous time. The planner seeks to maximize the present discounted value of aggregate utility, choosing time paths for consumption and capital to solve

$$\max_{c(t),k(t)} \int_0^\infty e^{-\rho t} u(c(t))dt,$$

subject to the dynamic resource constraint

$$\dot{k}(t) = f(k(t)) - \delta k(t) - c(t), \tag{9.33}$$

$$c(t), k(t) \geq 0, \quad \forall t,$$

and

$$k(0) \text{ given.}$$

Unlike in household problems with access to credit markets, the social planner in this setting cannot borrow against the future: the planner is bound by the non-negativity of capital at every instant (since this is a closed economy), so there is no need for an explicit no-Ponzi game condition. Following the same steps as in the previous consumption-saving problem, we can construct a Hamiltonian

$$H(t) \equiv e^{-\rho t} u(c(t)) + \mu(t)(f(k(t)) - \delta k(t) - c(t)).$$

where $\mu(t)$ is the costate variable associated with capital. The necessary first-order conditions are

$$\frac{\partial H(t)}{\partial c(t)} = 0$$

$$\frac{\partial H(t)}{\partial k(t)} + \dot{\mu}(t) = 0.$$

In addition, the transversality condition ensures that resources are not wasted in the limit,

$$\lim_{T \to \infty} e^{-\rho T} u'(c(T)) k(T) = 0. \tag{9.34}$$

For further tractability, let us focus on the widely used CRRA utility function, $u(c) = (c^{1-\sigma} - 1)/(1 - \sigma)$ with $\sigma > 0$ and $\sigma \neq 1$. Using the first-order conditions and following steps similar to those used in deriving equation (9.26) above, we can derive the Euler equation

$$\frac{\dot{c}(t)}{c(t)} = \frac{1}{\sigma} \left(f'(k(t)) - (\delta + \rho) \right). \tag{9.35}$$

The evolution of the economy is governed by the system of two differential equations, (9.33) and (9.35), together with the initial condition $k(0)$ and the TVC (9.34). These equations jointly determine the dynamic paths of capital and consumption over time. The trajectory of $(k(t), c(t))$ can be analyzed with a phase diagram (similar to the one in the Appendix of Chapter 4), illustrated in Figure 9.2.

The horizontal axis corresponds to the stock of capital $k(t)$ and the vertical axis to consumption $c(t)$. From equation (9.33), the sign of $\dot{k}(t)$—which captures how the capital stock evolves—depends on the relationship between consumption and output net of depreciation. Specifically, capital grows ($\dot{k}(t) > 0$) whenever $c(t) < f(k(t)) - \delta k(t)$. In the phase diagram, the curve $c(t) = f(k(t)) - \delta k(t)$, shown as a dash-dot line, marks the boundary where capital is neither rising nor falling ($\dot{k}(t) = 0$). Below this line, where consumption is lower, capital increases over time (illustrated by rightward arrows). Conversely, above this line, higher consumption leads to a decline in capital. Turning to equation (9.35), we see that consumption rises over time ($\dot{c}(t) > 0$) only when $f'(k(t)) - (\delta + \rho) > 0$. This condition depends solely on the level of capital, $k(t)$. Because $f'(\cdot)$ function is decreasing, this inequality implies $\dot{c}(t) > 0$ if and only if $k(t) < k^{**}$, where k^{**} satisfies $f'(k^{**}) = \delta + \rho$. The dotted line, $\dot{c}(t) = 0$, represents $k(t) = k^{**}$. To the left of this line, $c(t)$ increases over time (represented by the

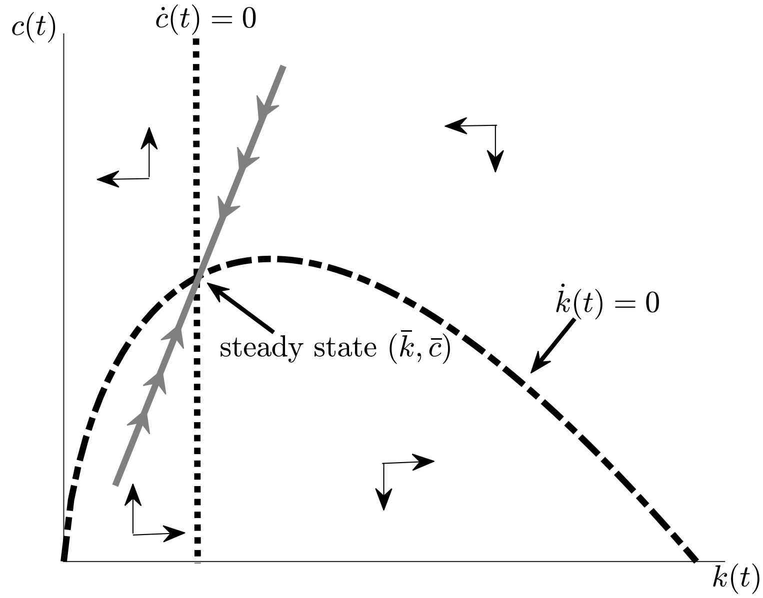

Figure 9.2: Phase diagram for Ramsey model in continuous time

upward arrows). Similarly, $c(t)$ decreases over time on the right side of the dotted line. The value k^{**} is called the *modified golden rule* capital stock. This value is smaller than the *golden rule* capital stock, which maximizes steady-state consumption.[3]

Starting from $k(0)$, the social planner has to choose the value of initial consumption $c(0)$ that is on the saddle path that is represented by the solid gray line. On the saddle path, $(k(t), c(t))$ gradually approaches the steady state $(\bar{k}, \bar{c})$ (here, $\bar{k} = k^{**}$) over time. This dynamic also satisfies the TVC (9.34), because $k(t)$ and $c(t)$ become constant in the long run. Thus, once again, the Ramsey model exhibits *convergence*. To see that the other values of $c(0)$ cannot be chosen, consider first the case when a larger value of $c(0)$ is chosen. One can see from the phase diagram that eventually $(k(t), c(t))$ reaches the situation $k(t) = 0$. At that point, the economy cannot produce anything, which implies $c(t)$ following the Euler equation is not feasible. Thus, this value of $c(0)$ does not satisfy the optimality conditions. If $c(0)$ is smaller than the one on the saddle path, eventually $(k(0), c(0))$ reaches the region where the value of $k(t)$ is large and $c(t) \to 0$. Appendix 9.A.4 shows this path does not satisfy the TVC. The consumers over-accumulate capital and do not consume much, which is a waste, as can be seen from the fact that the path of $c(t)$ is lower than the one on the saddle path.

[3]In Figure 9.2, the golden rule capital stock is the value of k that maximizes the dash-dot curve, since this curve represents steady-state consumption $c = f(k) - \delta k$. The maximizing condition, $f'(k) = \delta$, defines the golden rule. In contrast, the modified golden rule capital stock, k^{**}, solves $f'(k^{**}) = \delta + \rho$. Because $\rho > 0$, k^{**} is smaller than the golden rule level: consumers' impatience leads to less capital accumulation in the steady state.

The market equilibrium in continuous time

The analysis above characterizes the Pareto-optimal allocation, where all individuals are treated identically by the social planner. We now show that this allocation coincides with the outcome of the market equilibrium. While this equivalence follows from the first and second welfare theorems, it is nevertheless instructive to see explicitly how the decentralized market solution mirrors the planner's problem in this setting—offering another application of continuous-time methods.

Consider a continuum of households indexed by $i \in [0, 1]$, each with identical preferences and the same initial capital stock $k_i(0)$. The problem faced by a representative household is:

$$\max_{c_i(t), k_i(t)} \int_0^\infty e^{-\rho t} u(c_i(t)) dt,$$

subject to

$$\dot{k}_i(t) = r(t)k_i(t) + w(t) - \delta k_i(t) - c_i(t), \tag{9.36}$$

$$c_i(t), k_i(t) \geq 0, \quad \forall t,$$

and

$$k_i(0) \text{ given.}$$

Here, $r(t)$ denotes the rental rate of capital and $w(t)$ the wage rate. Applying the same optimization steps as before, the household's problem delivers the Euler equation

$$\frac{\dot{c}_i(t)}{c_i(t)} = \frac{1}{\sigma} \left(r(t) - (\delta + \rho) \right). \tag{9.37}$$

The representative firm maximizes profit. The firm's problem is static:

$$\max_{K(t), L(t)} F(K(t), L(t)) - r(t)K(t) - w(t)L(t),$$

where $K(t)$ is the factor demand for capital and $L(t)$ is the labor demand. The first-order conditions are

$$F_1(K(t), L(t)) = r(t)$$

and

$$F_2(K(t), L(t)) = w(t),$$

where $F_i(K(t), L(t))$ is the partial derivative with respect to input i. As in Section 3.3, we can show $F_1(K(t), L(t)) = f'(\tilde{k}(t))$ and $F_2(K(t), L(t)) = f(\tilde{k}(t)) - \tilde{k}(t)f'(\tilde{k}(t))$, where $f(k) \equiv F(k, 1)$ and $\tilde{k}(t) \equiv K(t)/L(t)$.

Let us now turn to the market equilibrium. There are three markets in this economy: the product market, the capital (rental) market, and the labor market. By Walras' law, we only need to consider two markets. In the capital market, the market-clearing condition is

$$\int k_i(t) di = K(t).$$

The right-hand side is capital demand, and the left-hand side is capital supply. In the labor market, because the aggregate supply of labor is 1, $1 = L(t)$. Let the aggregate capital in this economy be denoted as

$$k(t) = \int k_i(t) di.$$

Then, from the firm's optimization, the equilibrium prices are

$$r(t) = f'(k(t)) \tag{9.38}$$

and

$$w(t) = f(k(t)) - k(t)f'(k(t)). \tag{9.39}$$

Replacing (9.38) and (9.39) in (9.36) and (9.37) and using symmetry ($c_i(t) = c(t)$ and $k_i(t) = k(t)$), we obtain

$$\dot{k}(t) = f(k(t)) - \delta k(t) - c(t)$$

and

$$\frac{\dot{c}(t)}{c(t)} = \frac{1}{\sigma} \left(f'(k(t)) - (\delta + \rho) \right),$$

which are identical to (9.33) and (9.35).

9.4 Uncertainty: the Poisson process

Macroeconomic environments are inherently uncertain. Throughout this text, we have seen how uncertainty shapes savings decisions, labor supply, investment, and policy—whether in the form of aggregate shocks (see Chapters 7 and 14), income risks (Chapter 11), or events such as innovation and sovereign default (Chapters 13 and 24). In discrete-time models, it is natural to introduce uncertainty as a "shock in each period." For example, by specifying that output, productivity, or income can jump to a new value each quarter or year, according to some probability distribution. However, continuous-time modeling, which is essential in many applications (such as those in Chapter 13), does not come with a natural unit of time. So, how should we represent uncertainty in a continuous world?

There are two broad approaches for incorporating randomness in continuous time. The first captures the notion of uncertainty as a constant flow of small, frequent shocks, leading to what is called a *diffusion process*. Brownian motion, the prototypical example, is widely used in finance and has become increasingly prominent in macroeconomics, particularly in the analysis of asset prices, precautionary saving, and risk. However, a rigorous treatment of diffusion processes requires mathematical tools beyond the scope of this book.

The second approach focuses on the idea that randomness manifests as rare but discrete, often substantial, changes—events that occur unpredictably at random points in time (e.g., both the size and the frequency of shocks can be random). This class of models is built on the mathematics of *jump process*, the most fundamental of which is the Poisson process. The Poisson process is especially useful for modeling situations where events—such as technological breakthroughs, job loss, firm entry or exit—occur sporadically but with a well-defined average frequency. This structure underlies many applications in macroeconomics.

For instance, the Poisson process plays a central role in the analysis of job search and unemployment dynamics (see Chapter 20), and it forms the backbone of models of innovation and endogenous growth (see Chapter 13). By focusing on the Poisson process, we can introduce a rich form of uncertainty into continuous-time macroeconomic models in a way that remains mathematically accessible and tightly connected to a wide range of economic phenomena[4].

To consider a Poisson process, let us start with Bernoulli trials in discrete time. Consider a Bernoulli trial that can result in a "success" or a "failure." The situation can be finding a job, losing a job, succeeding in innovation, etc. Suppose the probability of success is $\lambda \in [0,1]$ for one trial. Suppose that during each unit of time (say, a year), one trial is made. From these assumptions, the expected number of successes during this one year is λ. Now, suppose that we divide the period into two and have one trial every six months. Each trial is independent. If we adjust the success probability of each trial to $\lambda/2$, the expected total number of successes during the one year is still λ. In general, if we divide the period (one year) into n subperiods and make the success probability λ/n in each subperiod, we still keep the expected total number of successes λ, although now we may have many successes during one year. The expected number of successes can be computed as $\underbrace{\frac{\lambda}{n} + \cdots + \frac{\lambda}{n}}_{n \text{ times}} = \lambda$.

With n trials, the distribution of the number of successes follows a binomial distribution

$$b(k,n) = \binom{n}{k} \left(\frac{\lambda}{n}\right)^k \left(1 - \frac{\lambda}{n}\right)^{n-k}, \tag{9.40}$$

where

$$\binom{n}{k} = \frac{n!}{k!(n-k)!}$$

is the binomial coefficient (the number of different ways one can have k successes out of n trials); $b(k,n)$ represents the probability of k successes during this time period.

Note that, in (9.40),

$$b(0,n) = \left(1 - \frac{\lambda}{n}\right)^n \tag{9.41}$$

and

$$\frac{b(k,n)}{b(k-1,n)} = \frac{n-(k-1)}{k} \frac{\lambda/n}{1-\lambda/n} = \frac{\lambda - (k-1)\lambda/n}{k - \lambda k/n} \tag{9.42}$$

for $k \geq 1$ hold. Let us consider a situation where $n \to \infty$. That is, the length of each subperiod approaches zero. In (9.41),

$$p(0) \equiv \lim_{n\to\infty} b(0,n) = e^{-\lambda}. \tag{9.43}$$

The second equality follows from the definition of e. Taking the limit of (9.42), we obtain

$$\lim_{n\to\infty} \frac{b(k,n)}{b(k-1,n)} = \frac{\lambda}{k}.$$

[4]A portion of the exposition that follows draws on Feller (1968).

Applying this formula for $k = 1, 2, 3, ...$ sequentially, we obtain

$$p(k) = e^{-\lambda} \frac{\lambda^k}{k!}.$$

This $p(k)$ represents the probabilities for the *Poisson distribution*. That is, $p(k)$ is the probability of k successes when the number of trials n approaches infinity (and the probability of success in each trial, λ/n, approaches zero).

The *Poisson process* is a stochastic process where the distribution of the number of successes between any time interval $(t, t+T]$, for $T > 0$, follows a Poisson distribution with

$$p(k) = e^{-(\lambda T)} \frac{(\lambda T)^k}{k!}.$$

The expected number of successes during the time interval $(t, t+T]$ is λT, and we saw that the Poisson process in continuous time can be understood as the limit of the discrete-time situation where repeated independent Bernoulli trials are conducted in every small time interval $\Delta = T/n$, where $\Delta \to 0$. Each trial has the probability of success $\lambda\Delta$, and the total number of trials is $n = T/\Delta$, thus keeping the expected number of successes as λT during the entire time period. Because the probability of success in each trial $\lambda\Delta \to 0$ and the total number of trials $n = T/\Delta \to \infty$ as $\Delta \to 0$, successes occur "infrequently" relative to the total number of trials.

A few notes are in order. First, by construction, the outcomes during $(t, t+T]$ and $(t+T+s, t+T+s+S]$ are independent of each other for any $s \geq 0$ and $S \geq 0$. Thus, the Poisson process is a memoryless process. Second, the initial restriction $\lambda \leq 1$ is not necessary; we can always start from a small-enough time interval (a large-enough n) such that $\lambda/n \leq 1$, and proceed with the above construction. Third, the probability that no success occurs during the time interval $(t, t+T]$ is, from (9.43), $e^{-\lambda T}$. If λ is time variant (denote as $\lambda(s)$ for time s), the same probability is expressed by $e^{-\int_t^{t+T} \lambda(s) ds}$.

As an example, consider the dynamics of employment and unemployment for a group of workers. Suppose that, for an individual employed worker, the event of job separation follows a Poisson process with parameter $\sigma > 0$. Then, the probability of keeping the job up to time t (i.e., the probability that no separation occurs up to time t) is, from the above formula, $e^{-\sigma t}$. Because the probability of separation during the time interval of length dt starting from time t is $e^{-\sigma t} \times \sigma dt$ (the probability of keeping the job up to time t times the probability of a separation event during the time period dt), the expected length of the remaining duration at the job is

$$\int_0^\infty t e^{-\sigma t} \sigma dt = \frac{1}{\sigma}.$$

Consider a group of many workers with total population N. Let the number of employed worker at time t be $e(t)$ and the number of unemployed worker be $u(t)$. Assume that a worker is either employed or unemployed, that is, $e(t) + u(t) = N$. Also assume that during any (short) time interval dt, an unemployed worker finds a job with probability λdt, where $\lambda > 0$. The total number of new job matches during dt time interval is $u(t) \times \lambda dt$ from the law of large numbers. Similarly, $e(t) \times \sigma dt$ is the total number of separation. Thus $u(t)$ follows the differential equation

$$du(t) = e(t)\sigma dt - u(t)\lambda dt,$$

where $du(t)$ is the change of $u(t)$ during the dt time interval, or, using the dot notation and $e(t) = N - u(t)$,

$$\dot{u}(t) = (N - u(t))\sigma - u(t)\lambda. \tag{9.44}$$

The steady-state mass of unemployment, $\bar{u}$, is (by setting $\dot{u}(t) = 0$)

$$\bar{u} = \frac{\sigma N}{\sigma + \lambda}.$$

By examining (9.44) closely, one can see the unemployment converges to the steady-state value, regardless of the initial unemployment. The unemployment rate in the steady state is

$$\frac{\bar{u}}{N} = \frac{\sigma}{\sigma + \lambda},$$

which is increasing in σ and decreasing in λ. A similar dynamics with discrete time (and $N = 1$) is presented in Chapter 20.