

# Very brief mathematical appendix

---

Throughout this book, we make use of a few key mathematical tools that allow us to handle the dynamic problems that arise when dealing with issues of macroeconomic policy. In this appendix, we briefly go over the key solution techniques that we use as a simple user guide. Note that our focus here is not on rigor, but rather on user-friendliness, which may lead to some hair-raising moments for those more familiar with the formalism. For a more thorough (yet still economist-friendly) presentation, the reader can consult a number of textbooks, such as Acemoglu (2009) or Dixit and Pindyck (1994).

We now go over three key areas: (i) Dynamic optimisation in continuous time, (ii) Dynamic optimisation in discrete time, and (iii) Differential equations.

## A.1 | Dynamic optimisation in continuous time

We have described macroeconomic policy problems in discrete and continuous time at different points, depending on convenience. In continuous time, we can solve these problems using the optimality conditions from *optimal control* theory.<sup>1</sup>

What kinds of problems fit the optimal control framework? The idea is that you choose a certain path for a choice variable – the *control* variable – that maximises the total value over time of a function affected by that variable. This would be relatively easy, and well within the realm of standard constrained optimisation, if whatever value you chose for the control variable at a certain moment in time had no implication for what values it may take at the next moment. What makes things trickier, and more interesting, is when it is not the case. That is to say, when what you do now affects what your options are for tomorrow – or, in continuous time, the next infinitesimal moment. That’s what is captured by the *state* variable: a variable that contains the information from all the previous evolution of the dynamic system. The evolution of the state variable is described by a dynamic equation, the *equation of motion*.

The simplest way to see all of this is to look at a simple example. Consider a simplified consumer problem:<sup>2</sup>

$$\max_{\{c_t\}_{t=0}^T} \int_0^T u(c_t) e^{-\rho t} dt, \quad (\text{A.1})$$

---

### How to cite this book chapter:

Campante, F., Sturzenegger, F. and Velasco, A. 2021. *Advanced Macroeconomics: An Easy Guide*.

Appendix A. ‘Very brief mathematical appendix’, pp. 363–370. London: LSE Press.

DOI: <https://doi.org/10.31389/lsepress.ame.w> License: CC-BY-NC 4.0.

subject to the budget constraint  $c_t + \dot{a}_t = y_t + ra_t$  and to an initial level of assets  $a_0$ . In words, the consumer chooses the path for their consumption so as to maximise total utility over their lifetime, and whatever income (labour plus interest on assets) they do not consume is accumulated as assets. The control variable is  $c_t$  that is, what the consumer chooses in order to maximise utility and the state variable is  $a_t$  that is, what links one instant to the next, as described by the equation of motion:

$$\dot{a}_t = y_t + ra_t - c_t. \quad (\text{A.2})$$

The maximum principle can be summarised as a series of steps:

*Step 1 - Set up the Hamiltonian function:* The Hamiltonian is simply what is in the integral – the instantaneous value, at time  $t$ , of the function you are trying to maximise over time plus, “the right-hand-side of the equation of motion” multiplied by a function called the co-state variable, which we will denote as  $\lambda_t$ . In our example, we can write:

$$H_t = u(c_t) + \lambda_t[y_t + ra_t - c_t]. \quad (\text{A.3})$$

(This is the current-value version of the Hamiltonian because utility at time  $t$  is being evaluated at the current  $t$ , that is, without the time discounting represented by the term  $e^{-\rho t}$ . The present-value version, where we would add that discounting term and write  $u(c_t)e^{-\rho t}$ , works just as well, with some minor adaptation in the conditions we will lay out below. We believe the current-value Hamiltonian lends itself to a more natural economic interpretation.)

This looks a lot like the Lagrangian function from static optimisation, right? Well, the co-state variable has a natural economic interpretation that is analogous to the familiar Lagrange multiplier. It is the marginal benefit of a marginal addition to the stock of the state variable – that is, of relaxing the constraint. In economic parlance, it is the shadow value of the state variable.

The key idea behind the maximum principle is that the optimal trajectory of control, state, and co-state variables must maximise the Hamiltonian function. But what are the conditions for that?

*Step 2 - Maximise the Hamiltonian with respect to the control variable(s):* There is no integral in the Hamiltonian, it's just a function evaluated at a point in time. So this is just like static optimisation! The intuition is pretty clear: if you were not maximising the function at all instants considered separately, you probably could be doing better, right? For our purposes, this will boil down to taking the first-order condition with respect to the control variable(s). In our specific example, we would write

$$H_c = 0 \Rightarrow u'(c_t) = \lambda_t. \quad (\text{A.4})$$

This has your usual FOC interpretation: the marginal utility gain of increasing consumption has to be equal to the marginal cost, which is not adding to the stock of your assets, and is thus given by the co-state variable.

Importantly, you could have more than one control variable in a problem. What do you do? Well, as in static optimisation, you take FOCs for each of them.

*Step 3 - Figure out the optimal path of the co-state variable(s):* The Hamiltonian is static, but the problem is dynamic. This means that, at any given instant, you must figure out that whatever you leave to the next instant (your state variable) is consistent with optimisation. This is a key insight. Intuitively, maximising an infinite-dimensional problem (i.e. what's the right value for your control variable at every instant in a continuum) can be broken down into a sequence of choices between the current instant and the (infinitesimally) next.

Fair enough, but how can we guarantee that? The maximum principle tells you: it's about satisfying the co-state equations. These are basically about optimising the Hamiltonian with respect to the state

variable(s). The FOC here is a bit different from what you are used to, as you won't set derivatives equal to zero. Instead, you set them equal to  $-\dot{\lambda}_t + \rho\lambda_t$ .<sup>3</sup> In our example

$$H_a = -\dot{\lambda}_t + \rho\lambda_t \Rightarrow r\lambda_t = -\dot{\lambda}_t + \rho\lambda_t. \quad (\text{A.5})$$

It seems like this does not have much of an economic intuition, but think again. Consider your state variable as a (financial) asset. (It is literally true in our specific example, but also more broadly.) This condition can be rewritten as

$$H_a + \dot{\lambda}_t = \rho\lambda_t. \quad (\text{A.6})$$

The LHS is the total marginal return of holding this asset for an additional instant: the “dividend” from the flow of utility coming from it (measured by the marginal impact on the Hamiltonian,  $H_a$ ) plus the “capital gain” coming from the change in its valuation in utility terms ( $\dot{\lambda}_t$ ). The RHS is the required marginal payoff for it to make sense for you to hold the asset this additional instant; it has to compensate for your discount rate  $\rho$ . If the LHS is greater (resp. smaller) than the RHS, you should hold more (resp. less) of that asset. You can only be at an optimum if there is equality. In other words, you can think about this as an “asset pricing” condition that is necessary for dynamic optimality.

What if there were more than one state variable? Well, then you will have more than one equation of motion, and you will need to impose one such “asset pricing” dynamic optimality condition for each of them.

Equations (A.4) and (A.5), put together, yield a first-order differential equation that contains the information for the static and dynamic requirements for optimisation. As we will see in more detail later in the Appendix, a first-order differential equation allows for an infinite number of solutions, up to two constants. How do we pin down which solution is the true optimum?

*Step 4 - Transversality condition:* We need two conditions to pin down the constants that are left free by the differential equation. One of them is the initial condition: we know the state variable starts off at a value that is given at  $t = 0$  – in our example,  $a_0$ . The second is a terminal condition: how much should we have left at the end, to guarantee that we have indeed maximised our objective function? This is what the transversality condition gives us in the example:

$$a_T\lambda_T e^{-\rho T} = 0. \quad (\text{A.7})$$

Intuitively, as long as our state variable has any positive value in terms of generating utility (and that shadow value is given by the co-state variable,  $\lambda$ ), you should not be left with any of it at the end. After all, you could have consumed it and generated additional utility!

These optimality conditions fully characterise the solution for any dynamic problem we will have encountered in this book.

## A.2 | Dynamic optimisation in discrete time

In dynamic problems, it is sometimes just more convenient to model time as evolving in discrete intervals, as opposed to continuously. This doesn't make a difference for the economic intuition, as we will see, but does require different techniques.

These techniques come from *dynamic programming* theory.<sup>4</sup> The key (and truly remarkable) insight behind these techniques is to recognise the iterative (or recursive) nature of a lot of dynamic problems. Their structure essentially repeats over and over again through time. This means that – not coincidentally, echoing the lessons from optimal control theory in a different context – you can break

down such problems into a sequence of smaller ones. Optimizing through time is achieved by making sure that you choose properly between today and tomorrow, while accounting for the fact that tomorrow you will face a similar choice again.

This insight is beautifully encapsulated in the *Bellman equation*. To see it in action, let's consider the discrete-time version of the consumer problem we have seen in the previous section

$$\max_{\{c_t\}_{t=0}^T} \sum_0^T \beta^t u(c_t), \quad (\text{A.8})$$

subject to the equation of motion,

$$a_{t+1} - a_t = y_t + ra_t - c_t. \quad (\text{A.9})$$

The information in this recursive problem can be summarised using the concept of the *value function*,  $V(a_t)$ : it is the value of total utility at a given point in time (as a function of the state variable), conditional on optimal decisions being made over the entire future path. In other words, we can rewrite the problem as

$$V(a_t) = \max_{c_t} \{u(c_t) + \beta V(a_{t+1})\}. \quad (\text{A.10})$$

Here is the intuition: choosing optimal consumption means maximising current utility, while also leaving the amount of the state variable that lets you make the optimal choice at  $t + 1$ . Picking today's consumption is a much simpler task than choosing an entire path. If you do it right, it leads you to the same solution. That's the beauty of dynamic programming.

That seems reasonable enough, but how do you solve it?

*Step 1 - Maximise the value function with respect to the control variable:* Well, the first thing is, naturally enough, to take the first-order condition with respect to the control variable. In our example, where the control variable is  $c_t$ , we get

$$u'(c_t) - \beta V'(a_{t+1}) = 0, \quad (\text{A.11})$$

where the second term on the LHS is using the fact that, as per the equation of motion,  $a_{t+1}$  is a function of  $c_t$ . The intuition is the same as ever: the optimal choice between today and tomorrow equates the marginal gain of additional consumption to the marginal cost of leaving a marginal unit of assets for tomorrow. The latter is that it detracts from the future choice possibilities encapsulated in the value function.

Just as with optimal control theory, this summarises the fact that dynamic optimisation requires static optimisation – otherwise, you could have been doing better!

*Step 2 - Figure out the optimal path for the state variable:* Again, by analogy with the intuition from optimal control, we also need to figure out that the choice we are making today leaves the right amount of the state variable for an optimal decision tomorrow. This means figuring out the value function.

To see how that works, let's look at our example again. Since our problem has a nice, continuously differentiable value function, we can differentiate the value function with respect to the state variable:

$$V'(a_t) = (1 + r)u'(c_t). \quad (\text{A.12})$$

It's easy enough to see why we get this marginal utility term. Consumption is a function of the state variable, as per the equation of motion. But shouldn't there be a term on  $V'(a_{t+1})$  somewhere? After all,  $a_{t+1}$  is also a function of  $a_t$ . The trick is that we used an *envelope condition*.<sup>5</sup> Intuitively, as with any so-called envelope theorem, if you are optimising the value function, you should set the path of the state variable such that you cannot get additional utility from a marginal change, at the optimum. This means that the term on  $V'(a_{t+1})$  simplifies to zero.

In short, at the optimum, the only impact of an additional marginal unit of your state variable is the additional utility brought by converting it into consumption – the impact on the future value is, at the optimum, equal to zero.

*Step 3 - Putting it all together:* You will have noticed that, once you know  $V(a_t)$ , you also know  $V(a_{t+1})$ : just stick a  $t + 1$  wherever you see a  $t$ . We can put (A.11) and (A.12) together and obtain

$$u'(c_t) = \beta(1 + r)u'(c_{t+1}), \tag{A.13}$$

that is to say: our familiar Euler equation.

That will do it for our purposes here, though you should keep in mind that to nail the full path of your control variable, you still need the initial and terminal conditions. (In our specific example, it is pretty obvious that the consumer would choose to consume all of their assets in the last period of their life: a transversality condition.)

### A.3 | First-order differential equations

#### A.3.1 | Integrating factors

Typically, the solution to a dynamic problem will be a system of differential (or difference) equations, describing the evolution of key variables of interest over time. In the main text of the book, we have introduced *phase diagrams* as a tool for analysing the behaviour of such systems.

Oftentimes, though, we were interested in finding an analytical solution for a variable whose behaviour is described by a dynamic equation. In these cases, we used a method of solution involving integrating factors, which we will elaborate on now.

As usual, this is easiest to motivate in the context of a specific example. Let's take the consumer problem discussed above, and in particular the equation of motion described by (A.2), which we rewrite here, for convenience, in slightly different form:

$$\dot{a}_t - ra_t = y_t - c_t. \tag{A.14}$$

This illustrates a kind of first-order differential equation that can generally be written as

$$\frac{dz}{dx} + Pz = Q, \tag{A.15}$$

where  $P, Q$  are functions of  $x$  only. This is the kind that can be solved using integrating factors. The integrating factor is defined as:

$$I = e^{\int P dx}. \tag{A.16}$$

The trick is to multiply both sides of A.15 by the integrating factor, which yields

$$\frac{dz}{dx} e^{\int P dx} + Pz e^{\int P dx} = Q e^{\int P dx}. \tag{A.17}$$

You will notice that the LHS of this equation is what you get from differentiating  $ze^{\int P dx}$  with respect to  $x$ , using the product rule of differentiation. In other words, we can integrate both sides of A.17, using the Fundamental Theorem of Calculus, and obtain

$$\int \left( \frac{dz}{dx} e^{\int P dx} + Pz e^{\int P dx} \right) dx = \int Q e^{\int P dx} dx \Rightarrow z e^{\int P dx} = \int Q e^{\int P dx} dx + \kappa \tag{A.18}$$

This allows us to find a general solution for  $z$ , up to the constant  $\kappa$  (recall that any constant term would not affect the derivative, so any solution has to be up to a constant).

Let's look at that in the context of our consumer problem in (A.14). You can see that  $z$  is  $a$ ,  $x$  is  $t$ ,  $P$  is  $-r$ , and  $Q$  is  $y_t - c_t$ . So the integrating factor is  $e^{-\int r dt}$ , which can be simplified to  $e^{-rt}$ . Multiplying both sides by that integrating factor yields

$$\dot{a}_t e^{-rt} - r_t a_t e^{-rt} = y_t e^{-rt} - c_t e^{-rt}. \quad (\text{A.19})$$

Applying (A.18) allows us to find a general solution for  $a_t$ :

$$a_t e^{-rt} = \int_0^t (y_s e^{-rs} - c_s e^{-rs}) ds + \kappa \Rightarrow a_t = e^{rt} \left( \int_0^t (y_s e^{-rs} - c_s e^{-rs}) ds + \kappa \right) \Rightarrow a_t = \int_0^t (y_s - c_s) e^{r(t-s)} ds + \kappa e^{rt}, \quad (\text{A.20})$$

where we are using  $s$  to denote each instant over which we are integrating, up to  $t$ . How do we pin down the constant? This equation must hold for  $t = 0$ , which entails that  $\kappa = a_0$ . In other words,

$$a_t = \int_0^t (y_s - c_s) e^{r(t-s)} ds + a_0 e^{rt} \quad (\text{A.21})$$

This tells us that the consumer's assets, at any given point in time, equal the initial value of assets (compounded by the interest rate), plus the compound sum of their savings over time.

The differential equations we deal with can all be solved using this method, which is quite convenient. Depending on the specific problem, this solution then allows us to figure out the optimal level of consumption, or the path of the current account, or whatever else that interests us in the case at hand.

### A.3.2 | Eigenvalues and dynamics

In other occasions, we do not need an analytical solution, but want to have a way of figuring out the dynamic properties of a system of differential equations. Here's where linear systems are very convenient, because we can use tools of linear algebra to come to the rescue. This helps explain why we often focus on linear approximations (around a BGP, typically). Note that this focus entails important consequences: a linear approximation is good enough when you are sufficiently close to the point around which you are doing the approximation. If a shock gets you far from the BGP, then maybe the approximation is not going to work that well as a description of your economy!

Let's talk about the tools, and especially one concept that we mention quite a bit in the book: eigenvalues.

Consider a system of differential equations that we have studied, describing the solution of the basic AK model. In its simplest version, we can write it as

$$\dot{c}_t = (A - \rho)c_t \quad (\text{A.22})$$

$$\dot{k}_t = Ak_t - c_t \quad (\text{A.23})$$

The nice thing is that this system is already linear in  $c_t$ ,  $k_t$ , meaning that we can write it in matrix form:

$$\begin{bmatrix} \dot{c}_t \\ \dot{k}_t \end{bmatrix} = \begin{bmatrix} A - \rho & 0 \\ -1 & A \end{bmatrix} \begin{bmatrix} c_t \\ k_t \end{bmatrix} \quad (\text{A.24})$$

Let us denote the vector  $[c_t, k_t]$  as  $x_t$ ; we can write this more concisely as

$$\dot{x}_t = \Gamma x_t. \quad (\text{A.25})$$

Here's where the trick comes in: solutions to a system like (A.25) can be written as

$$x_t = \kappa_1 e^{\lambda_1 t} \eta^1 + \kappa_2 e^{\lambda_2 t} \eta^2, \quad (\text{A.26})$$

where  $\lambda_1$  and  $\lambda_2$  are the eigenvalues of the matrix  $\Gamma$ ,  $\eta^1$  and  $\eta^2$  are the corresponding eigenvectors, and  $\kappa_1$  and  $\kappa_2$  are the standard constants, of which there are two because this is a system of two differential equations.<sup>6</sup>

This means that the dynamic behaviour of the system will be closely tied to the eigenvalues. Why so? Imagine that both  $\lambda_1, \lambda_2 > 0$ .<sup>7</sup> It is easy to see that the solution (A.26) will behave explosively: as  $t$  grows,  $x_t$  will grow without bound, in absolute value, for any (nonzero)  $\kappa_1, \kappa_2$ . Such a system will not converge.

What if  $\lambda_1, \lambda_2 < 0$ ? Then, eventually, the solution will converge to zero (which means that the solution to the general differential equation will converge to the particular solution). This is a *stable* system; it will converge no matter where it starts.

Economists like stable systems that converge – but not quite so stable. After all, it seems natural to think that there is room for human choices to make a difference! Particularly since, in economic terms, there will often be the jumpy variables that we've been alluding to through the book – that is to say, those that need not follow an equation of motion.

You will have noticed, though, that there is a third case:  $\lambda_1 > 0, \lambda_2 < 0$ . (This is without loss of generality, of course!) In that case, the system will converge only in the case where  $\kappa_1$  is exactly equal to zero. Such a system is, technically speaking, not stable – it generally will not converge. But we refer to it as saddle-path stable. These are the most interesting, from an economic perspective, as convergence depends on purposeful behaviour by the agents in the model.

How do we know if a system is saddle-path stable without having to compute the eigenvalues? It suffices to recall, from linear algebra, that the determinant of a matrix is equal to the product of its eigenvalues. It immediately follows that, if  $\det(\Gamma) < 0$ , we are in a situation of saddle-path stability.<sup>8</sup> Such a system will converge only if the initial choice of the jumpy variable is the one that puts the system on the saddle path – that is to say, the one that delivers  $\kappa_1 = 0$  in the notation above.

## Notes

<sup>1</sup> These optimality conditions are synthesised in Pontryagin's maximum principle, derived in the 1950s from standard principles from the classical calculus of variations that goes back to Euler and Lagrange and others.

<sup>2</sup> Again, we are using subscripts to denote time, as opposed to parentheses, as more often done in continuous-time settings. We are iconoclasts!

<sup>3</sup> With the present-value Hamiltonian, this becomes  $-\dot{\lambda}_t$ . It is easy to check that they are equivalent.

<sup>4</sup> These techniques were developed by Richard Bellman in the 1950s.

<sup>5</sup> The relevant envelope theorem here is due to Benveniste and Scheinkman (1979).

<sup>6</sup> Recall that the eigenvalues of matrix  $\Gamma$  can be computed as the solution  $\lambda$  to the equation  $\det(A - \lambda I) = 0$ , where  $I$  is the identity matrix. The eigenvectors are defined as vectors  $\eta$  such that  $\Gamma\eta = \lambda\eta$ .

<sup>7</sup> More precisely, we are referring here to the *real* part of the eigenvalues – they can be complex numbers with an imaginary part.

<sup>8</sup> It is easy to check that the system in (A.24) does not converge if  $A > \rho$ . You will recall that this is exactly our conclusion from studying the AK model in Chapter 5.