A Theory of Slack

How Economic Slack Shapes Markets, Business Cycles, and Policies

Pascal Michaillat

Draft version: April 2026
Draft URL: pascalmichaillat.org/18/

CHAPTER 14.

Slackish business cycle model

Chapter 13 showed that the unemployment gap in US data is generally positive and sharply countercyclical: it widens in recessions and narrows in expansions. The first task of part IV is to explain which aggregate shocks can generate these fluctuations. The second task is to clarify whether and how monetary policy—the main stabilization policy in the United States—can partially or completely undo such shocks so as to stabilize unemployment fluctuations. The final task is to assess whether fiscal policy can support monetary policy in stabilizing unemployment.

This chapter develops the simplest possible slackish business cycle model for that purpose. The model clarifies how aggregate demand and supply determine unemployment, and how monetary policy steers aggregate demand to stabilize unemployment in response to shocks. The model covers both normal times, when monetary policy operates in a conventional way by setting the nominal interest rate, and zero-lower-bound episodes, when monetary policy operates by forward guidance. The model also briefly studies fiscal policy in the form of public spending.

14.1. Overview of the model

The model is organized around one slackish market in which self-employed workers sell their production to other households.

In the model, households produce services and not goods. This simplifies the analysis greatly because if a service is available and no customer purchases it, the service is just lost and wasted—services cannot be stored. This simplifies the model a lot: we do not need to worry about inventory of unsold goods and the depreciation of unsold goods.

One might wonder how realistic it is to assume that households produce only services and not goods. It turns out that it is not such a bad assumption. In US data, there have always been more workers in service-producing industries than in goods-producing industries, but the gap between these two types of industries has widened significantly. From 1943 to today, the number of US workers in goods-producing industries has not changed much: 18.8 million workers in 1943, peaking at 25.0 million workers in 1979, and falling slightly to 21.7 million workers in 2024 (BLS 2025a). By contrast, the number of workers in service-providing industries has increased dramatically, from 23.8 million workers in 1943 to 136.3 million in 2024 (BLS 2025b). So the share of US workers in service-providing industries rose from 56% in 1943, to 72% in 1979, and to 86% in 2024. That is, of all workers, only 14% are producing goods in 2024, whereas 44% were producing goods in 1943. That is why modelling the US economy as a service economy seems like an accurate simplification.

We assume that households cannot produce services for themselves, so households also purchase and consume services produced by other households. The key idea is that a matching function mediates all the trades on the market for services, which produces slack.[1]

Another assumption here is that households produce services directly—there are no firms. Hence, we don't need to distinguish between a labor market and a product market: we only have the service market.

In the following chapters, we expand this baseline model to refine our understanding of unemployment fluctuations—both their macroeconomic causes and macroeconomic consequences. For instance, in chapter 15, we include separate product and labor markets to trace how slack propagates across markets.

14.2. Slackish market for services

The model is dynamic and organized around a single market for services. The structure of the market is similar to that of the dynamic slackish market model of chapter 9, and of the slackish labor market model of chapter 11.

We assume that there are many households, each indexed by $i \in [0,1]$. In each household, $h_i > 0$ workers are part of the labor force: either working or searching for a job. The aggregate number of workers in the labor force is $h = \int_0^1 h_i \, di$. The labor force comprises l employed workers and $U = h - l$ unemployed workers.

Hence the economy is composed of a labor force of size $h > 0$. Each worker in the labor force has the capacity to produce $a > 0$ services per unit time. The parameter a describes

[1]This is very much in the spirit of Diamond (1982), who develops a model of the product market in which self-employed workers must trade because they do not consume their own production.

labor productivity. The economy's aggregate capacity is ah services per unit time.

Services are sold through long-term worker-household relationships. Once workers have matched with a household, they become full-time employees of the household. They remain so until they separate, which occurs at rate $\lambda > 0$.

To find new employees, each household j posts V_j help-wanted ads. The reason why households must constantly hire help is that they constantly lose existing workers. The aggregate number of help-wanted ads is $V = \int_0^1 V_j \, dj$.

From the unemployment level U and vacancy level V, we also define the unemployment rate $u = U/h$ and the vacancy rate $v = V/h$.

The rate at which help-wanted ads and workers match is given by a matching function $m(U, V)$ that satisfies the assumptions laid out in chapter 4.

The aggregate tightness is the ratio of the two arguments in the matching function: $\theta = V/U = v/u$.

The labor market tightness determines the rate at which unemployed workers find jobs, $f(\theta) = m(1, \theta)$, and the rate at which help-wanted ads are filled, $q(\theta) = m(1/\theta, 1)$.

14.3. Equilibrium aggregate supply

The equilibrium aggregate supply is the output level when employment flows are balanced:

$$y^s(\theta) = [1 - u(\theta)] \, ah, \tag{14.1}$$

where the equilibrium unemployment rate is given by

$$u(\theta) = \frac{\lambda}{\lambda + f(\theta)}. \tag{14.2}$$

As we saw in chapter 11, the convergence to the equilibrium unemployment rate is rapid. Hence, the actual and equilibrium unemployment rates are indistinguishable: this is why we abstract from the law of motion of the unemployment rate (given by (9.3)) and consider the unemployment rate as a function of tightness (given by (14.2)).

14.4. Inflation rate and interest rates

Before studying the aggregate demand, we describe the prices of services and assets in the model.

14.4.1. Inflation rate

At time t, services are sold at a unit price $p(t)$. As in any slackish model, we need to specify a price norm to determine $p(t)$. In a dynamic model, the price norm determines the rate of inflation, $\pi(t) = \dot{p}(t)/p(t)$.

We saw in chapter 6 that in reality market prices respond very sluggishly to changes in market conditions. This appears to be true in the aggregate too. Of course, prices always rise, but the pace at which they do so does not seem to be affected by economic conditions. In figure 1, we saw that the US inflation rate barely fell when unemployment spiked during the Great Recession, and then that it barely rose when the unemployment rate recovered. More systematically, the US rate of inflation does not seem to respond much to shocks to monetary policy or to fluctuations in unemployment.[2]

Hence, we simply assume that the price of services grows at a fixed rate. This price norm could well be the result of constant communication from the central bank, which always tries to convince people that they will do anything to keep inflation at a fixed level—2% in the United States.

Formally, we assume that prices grow at a constant inflation rate π:

$$p(t) = p(0)\exp(\pi t).$$

In chapter 16, we will move away from fixed inflation and introduce endogenous inflation and a Phillips curve to study how fluctuations in slack might affect prices. For now, we study the model with fixed inflation.

14.4.2. Nominal interest rate

Next, we need to specify how nominal interest rates are set. We assume that the central bank follows an interest rate peg: they maintain the nominal interest rate $i(t)$ at some specific level i at all times in the absence of shocks. We also assume that nominal interest rates are subject to the zero-lower-bound constraint: $i(t) \geq 0$.

14.4.3. Real interest rate

The key price in our model is the real interest rate, because it is the relative price of government bonds relative to services:

$$r(t) = i(t) - \pi(t).$$

Because we have a fixed inflation rate and a fixed nominal interest rate, the real interest rate is also fixed at $r = i - \pi$. Note that the real interest rate must be less than the time discount rate ($r < \delta$) because if not, no household would want to consume anything and aggregate demand would collapse.

[2] For evidence, see Christiano, Eichenbaum, and Evans (1999, figure 2), Stock and Watson (2010, figure 1), and Stock and Watson (2020, figure 1).

14.5. Preference for wealth

In the following sections, we make the unusual assumption that wealth enters the utility function. Thanks to this assumption, the aggregate demand is nondegenerate and the model is therefore much more robust. How can we justify such an assumption?

The standard model assumes that people save to smooth consumption over time, but it has long been recognized that people seem to enjoy accumulating wealth irrespective of future consumption. Instead, we can argue that it has always been true that people value holding wealth in a few different ways. First, by introspection, we know that wealth defines social status, which is something that people value. Thus, people try to accumulate wealth not just because they want to use it for future consumption (like in a typical economic model), but also because they enjoy the lifestyle and social status that comes with having more money. This is something that a lot of economists have noticed over the years.

Describing the European upper class of the early 20th century, Keynes (1919, chapter 2) noted that

> The duty of saving became nine-tenths of virtue and the growth of the cake the object of true religion.... Saving was for old age or for your children; but this was only in theory—the virtue of the cake was that it was never to be consumed, neither by you nor by your children after you.

Fisher (1930, p. 17), who invented the model of consumption-saving that we use today, added that

> A man may include in the benefits of his wealth ... the social standing he thinks it gives him, or political power and influence, or the mere miserly sense of possession, or the satisfaction in the mere process of further accumulation.

Essentially, Keynes and Fisher, as well as many other economists, noted that by introspection, people enjoy accumulating wealth because it offers social status and power. Today, we have much stronger scientific evidence to back this claim: neuroscientists have actually confirmed this intuition we get by introspection—that wealth itself provides utility independent of consumption. After reviewing the literature, Camerer, Loewenstein, and Prelec (2005, p. 32) conclude that neuroscientific evidence confirms that wealth itself provides utility, independently of the consumption it can buy:

> Brain-scans conducted while people win or lose money suggest that money activates similar reward areas as do other 'primary reinforcers' like food and drugs, which implies that money confers direct utility, rather than simply being valued only for what it can buy.

A survey of 2000 Americans aged 62–75 conducted by the Employee Benefit Research Institute in 2020 and analyzed by Coy (2021) confirms these ideas. The survey first asks respondents what they plan to do over the course of their retirement with the wealth that they hold in their financial accounts. About 15% plan to grow their wealth, 1% plan to spend down none of their wealth, and another 35% plan to spend down a little, small portion of their wealth. So about 60% of the respondents do not plan to spend down their wealth in retirement.

Asked why they do not plan to consume their wealth in old age, about 35% of respondents say that spending down isn't necessary, and another 30% say that not spending makes them feel better. The survey explores that finding further and asks respondents whether they agree with the statement that "saving as much as I can makes me happy and fulfilled." About 65% of the respondents strongly or somewhat agree with the statement—a strong validation of the wealth-in-the-utility assumption.

Since people do value holding wealth, it is not inaccurate to assume that it enters the utility function. Among all the reasons why people may value wealth, we focus on social status: we postulate that people enjoy wealth because it provides social status. We therefore introduce relative (not absolute) wealth into the utility function. The assumption is convenient: in equilibrium everybody is the same, so relative wealth is zero. And the assumption seems plausible. Adam Smith, Ricardo, John Rae, J.S. Mill, Marshall, Veblen, and Frank Knight all believed that people accumulate wealth to attain high social status (Steedman 1981). More recently, a broad literature has documented that people seek to achieve high social status, and that accumulating wealth is a prevalent pathway to do so.[3]

14.6. Equilibrium aggregate demand

We now turn to the aggregate demand in the model, which we derive by solving the consumption-saving problem of households.

14.6.1. Household's utility function

To study the aggregate demand side of the model, we need to introduce the utility function for households. Much like in the basic model, households value two things: services and real wealth. Here, real wealth is held as real bonds instead of real money balances, but the idea is the same. The utility from services, $c(t)$, is

$$c(t)^{1-\alpha},$$

where $\alpha \in (0, 1)$ governs the concavity of the utility.

[3] See for instance Frank (1985), Cole, Mailath, and Postlewaite (1995), Weiss and Fershtman (1998), Heffetz and Frank (2011), Fiske (2010), Anderson, Hildreth, and Howland (2015), Cheng and Tracy (2013), Ridgeway (2014), Mattan, Kubota, and Cloutier (2017).

The utility from real wealth, $b(t)$, is

$$\mathcal{W}(b(t) - \bar{b}(t)),$$

where $\bar{b}(t)$ is the average real wealth, so that $b(t) - \bar{b}(t)$ is the relative real wealth. $\mathcal{W}$ is a strictly increasing and strictly concave function.

Then, the aggregate utility that a household maximizes is the discounted sum of flow utilities:

$$\int_0^{\infty} e^{-\delta t} \left[c(t)^{1-\alpha} + \mathcal{W}(b(t) - \bar{b}(t)) \right] dt.$$

14.6.2. Household's budget constraint

The next step to deriving the aggregate demand curve is to set up the budget constraint of the household. First, we need to figure out the income of the representative household in our model. There are two components to income: labor income and saving income. Labor income takes into account the number of people that are working in a given household: $[1 - u(t)]\, h$. Then, $a\,[1 - u(t)]\, h$ is the number of services sold by the household per unit time. Assuming all services have a price $p(t)$, labor income for a household is

$$p(t)a\,[1 - u(t)]\, h.$$

We know that households can also save using government bonds, which have an interest rate. This is investment/saving income:

$$i(t)B(t),$$

where $i(t)$ is the nominal interest rate and $B(t)$ is the number of bonds held by a household.

Next, we need to look at the sources of expenditure. First, to consume $c(t)$ services, a household must purchase more than $c(t)$ services—indeed, some services are used only to fill help-wanted ads and do not result in consumption. The amount of services paid for at any time is

$$y(t) = [1 + \tau(\theta(t))]\, c(t) \tag{14.3}$$

where $\tau(\theta)$ is the matching wedge:

$$\tau(\theta) = \frac{\kappa\lambda}{q(\theta) - \kappa\lambda}. \tag{14.4}$$

A household spends $p(t)\,[1 + \tau(\theta(t))]\, c(t)$ on purchasing services at time t.

Additionally, to finance the interest payments on the bonds from the government, there is a lump-sum tax, $T(t)$, that each household has to pay.

Now that we have our sources of income and expenditure, we can build the household's nominal budget constraint:

$$\dot{B}(t) = i(t)B(t) + p(t)a\left[1 - u(t)\right]h - p(t)\left[1 + \tau(\theta)\right]c(t) - T(t), \tag{14.5}$$

where $\dot{B}(t)$ is the change in savings, or nominal wealth, at time t. To obtain the budget constraint in real terms, we introduce the household's real stock of bonds,

$$b(t) = \frac{B(t)}{p(t)},$$

and the real interest rate,

$$r(t) = i(t) - \pi(t),$$

where $\pi(t)$ is the inflation rate:

$$\pi(t) = \frac{\dot{p}(t)}{p(t)}.$$

A useful relationship is $\ln(b(t)) = \ln(B(t)) - \ln(p(t))$. Taking the derivative of this relationship with respect to time, we get

$$\frac{\dot{b}(t)}{b(t)} = \frac{\dot{B}(t)}{B(t)} - \frac{\dot{p}(t)}{p(t)} = \frac{\dot{B}(t)}{B(t)} - \pi(t).$$

Thus, the temporal change in real wealth is related to the temporal change in nominal wealth:

$$\dot{b}(t) = b(t) \cdot \frac{\dot{B}(t)}{B(t)} - \pi(t)b(t) = \frac{\dot{B}(t)}{p(t)} - \pi(t)b(t), \tag{14.6}$$

because $b(t)/B(t) = 1/p(t)$.

By combining the nominal budget constraint (14.5) and the link between real and nominal wealth (14.6), we can finally write the real budget constraint:

$$\dot{b}(t) = r(t)b(t) + a\left[1 - u(t)\right]h - \left[1 + \tau(\theta)\right]c(t) - \frac{T(t)}{p(t)}.$$

14.6.3. Household's problem

We are now ready to solve the household's problem, which brings us closer to understanding the aggregate demand side of the model. The household's problem is to maximize utility subject to the budget constraint, taking all aggregate variables as given:

$$\max_{c(t),b(t)} \int_0^\infty e^{-\delta t}\left[c(t)^{1-\alpha} + \mathcal{W}(b(t) - \bar{b}(t))\right] dt$$

subject to

$$\dot{b}(t) = r(t)b(t) + a\left[1 - u(t)\right]h - \left[1 + \tau(\theta)\right]c(t) - \frac{T(t)}{p(t)}$$

and the no-Ponzi condition that households are not allowed to run a Ponzi scheme. The household takes as given $\theta(t)$, $u(t)$, $p(t)$, $T(t)$, $i(t)$, and the endowment of wealth ($B(0)$).

To solve this optimal control problem, we set up a Hamiltonian (see result F.24):

$$\begin{aligned}\mathcal{H}(t, c(t), b(t)) = c(t)^{1-\alpha} + \mathcal{W}(b(t) - \bar{b}(t))\\ + \psi(t)\left\{r(t)b(t) + (1 - u(t))ah - \left[1 + \tau(\theta)\right]c(t) - \frac{T(t)}{p(t)}\right\}.\end{aligned}$$

The variable $\psi(t)$ is the costate variable on the real budget constraint.

The two necessary conditions for optimality are the following:

$$\frac{\partial \mathcal{H}}{\partial c} = 0,$$
$$\frac{\partial \mathcal{H}}{\partial b} = \delta\psi(t) - \dot{\psi}(t).$$

In addition to these conditions, we also need the appropriate transversality condition. This can take different forms depending on the specific problem, but is usually something like: $\lim_{t\to\infty} e^{-\delta t}\psi(t)b(t) = 0$.

Notice that our utility function is concave in both consumption and real wealth, and that our budget constraint is linear in both consumption and real wealth. Thus, any interior solution to the necessary conditions is the global maximum of the household's problem (result F.26).

Let us now rework the necessary conditions to characterize the household's optimal behavior. As usual in consumption-saving problems, the behavior that maximizes the utility function follows an Euler equation. First, the derivative of the Hamiltonian with respect to consumption must equal zero:

$$\frac{\partial \mathcal{H}}{\partial c} = (1 - \alpha)c(t)^{-\alpha} - \psi(t)\left[1 + \tau(\theta(t))\right] = 0,$$

so that

$$(1 - \alpha)c(t)^{-\alpha} = \psi(t)\left[1 + \tau(\theta(t))\right]. \tag{14.7}$$

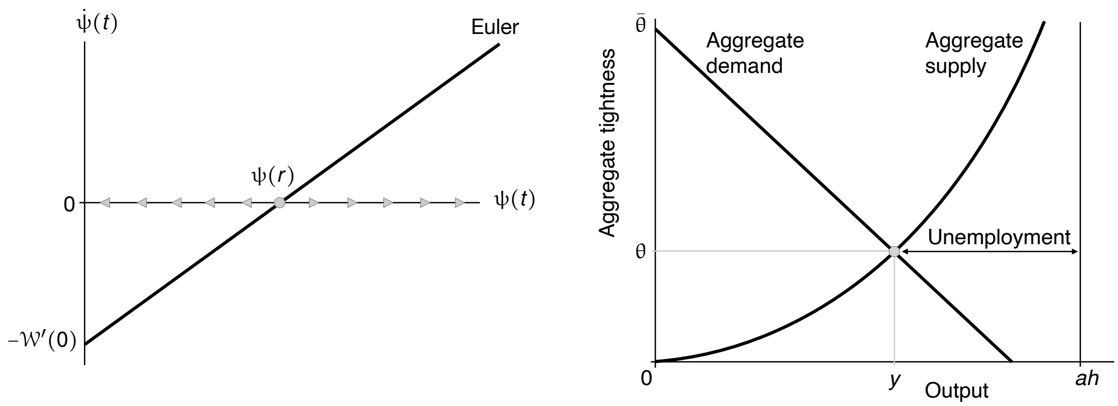

A. Phase line of the Euler equation

B. Aggregate demand and supply

FIGURE 14.1. Solution of the slackish business cycle model

The variable $\psi(t)$ is the costate variable associated with the household's real wealth. The Euler equation is given by equation (14.9). The equilibrium point $\psi(r)$ is given by equation (14.10). The aggregate supply curve is given by equation (14.1). The aggregate demand curve is given by equation (14.12).

From the second necessary condition,

$$\frac{\partial \mathcal{H}}{\partial b} = \mathcal{W}'(b(t) - \bar{b}(t)) + \psi(t)r(t) = \delta\psi(t) - \dot{\psi}(t)$$

which gives the differential equation describing how households save and spend:

$$\dot{\psi}(t) - [\delta - r(t)]\,\psi(t) - \mathcal{W}'(b(t) - \bar{b}(t)) = 0. \tag{14.8}$$

14.6.4. Euler equation

The differential equation (14.8) implicitly describes consumption and saving over time. From it, we obtain the aggregate demand. But to do that we need to understand its dynamical properties.

Since r is fixed in the model, and we have homogeneous households (so that $b(t) = \bar{b}(t)$), we can simplify (14.8) to

$$\dot{\psi}(t) = (\delta - r)\psi(t) - \mathcal{W}'(0), \tag{14.9}$$

where $\mathcal{W}'(0)$ is households' marginal utility of wealth. We now have our Euler equation: a first-order linear differential equation, whose properties are plotted on the phase line 14.1A.

We can see that the equilibrium point is at

$$\psi(r) = \frac{\mathcal{W}'(0)}{\delta - r}, \tag{14.10}$$

and that our system is a source.

The variable ψ is a costate variable, so it is not a predetermined variable at time t: it can jump at any time. If ψ jumps above the equilibrium point, it would start increasing over time and would diverge to ∞. This means that consumption would decrease and eventually become negative, violating one of the problem's feasibility constraints (see equation (14.7)). If ψ jumps below the equilibrium point, it would decrease over time, which means that consumption would increase over time and would diverge to ∞, eventually violating the resource constraint in the economy. Therefore, ψ must jump to the equilibrium point at $t = 0$—transition to the equilibrium point is immediate. At $t = 0$, $\psi(t)$ jumps to $\psi(r)$, so that at all times, $\psi(t) = \psi(r)$.

Therefore, using (14.7), we see that consumption is at all times

$$c = \left[\frac{(1-\alpha)(\delta - r)}{\mathcal{W}'(0)} \cdot \frac{1}{1+\tau(\theta)}\right]^{1/\alpha}, \tag{14.11}$$

where $0 < c < \infty$. This interior solution satisfies all the necessary conditions, so this is the unique global maximum of the household's problem.

14.6.5. Recovering the standard Euler equation

Our Euler equation looks a bit different from the standard Euler equation in macroeconomic models—which might be worrisome. But there is no need to worry: it is easy to recover the standard form by eliminating the matching wedge and the wealth in the utility function.

Indeed, setting $\tau = 0$ and $\mathcal{W}' = 0$ in (14.7) and (14.8), we find

$$(1-\alpha)c(t)^{-\alpha} = \psi(t), \qquad \frac{\dot{\psi}(t)}{\psi(t)} = \delta - r(t).$$

Next, we take the logarithm of the first equation and differentiate it with respect to time. We obtain

$$-\alpha \cdot \frac{\dot{c}(t)}{c(t)} = \frac{\dot{\psi}(t)}{\psi(t)}.$$

Combining the two results, we recover the standard Euler equation:

$$\frac{\dot{c}(t)}{c(t)} = \frac{1}{\alpha} \cdot \left[r(t) - \delta\right].$$

As usual, the equation says that the growth rate of consumption is determined by the gap between the real interest rate and the time discount rate.

14.6.6. Computing the equilibrium aggregate demand

We've seen that although our model is dynamic, consumption jumps directly to the equilibrium point at $t = 0$ and stays there. Thus, we can solve the business cycle model by supply-demand analysis, using tightness-quantity diagrams, with a very similar solution method to what we did for the market models.

The first step is to write down the aggregate demand curve, which gives us the amount of output that is demanded by households according to their Euler equation. The amount of output demanded is $y = [1 + \tau(\theta)]\, c$, where the amount of consumption demanded is given by (14.11). This gives the aggregate demand:

$$y^d(\theta, i) = \left[\frac{(1-\alpha)(\delta + \pi - i)}{\mathcal{W}'(0)}\right]^{1/\alpha} \cdot \frac{1}{[1+\tau(\theta)]^{1/\alpha - 1}}. \tag{14.12}$$

We can see that the aggregate demand is decreasing in aggregate tightness, from $y^d(0, i) > 0$ to $y^d(\bar{\theta}, i) = 0$ as tightness rises from $\theta = 0$ to $\theta = \bar{\theta}$. The aggregate demand is also decreasing in the nominal interest rate i and in fact in the real interest rate $r = i - \pi$.

14.7. Solution of the model

We now solve the slackish business cycle model. We distinguish normal times, when the nominal interest rate is positive, and zero-lower-bound episodes, when the nominal interest rate is at zero.

14.7.1. Solution in normal times

We know that aggregate tightness has to satisfy:

$$y^d(\theta, i) = y^s(\theta). \tag{14.13}$$

This equation implicitly defines tightness as a function of the nominal interest rate, $\theta(i)$. Then, from θ, we can back out all the other variables. We can plot the aggregate demand and supply curves to find the intersection, as shown in figure 14.1B.

The aggregate supply curve captures the capacity of the economy plus the matching process. The aggregate demand curve captures the fact that households maximize their utility given their budget constraint. At the intersection of the aggregate supply and aggregate demand curves, the matching process is respected (including both hiring and separations), as well as households' utility maximization process.

14.7.2. Solution at the zero lower bound

Our model behaves the same at the zero lower bound and away from it. In both situations, the aggregate demand and supply curves are identical. Indeed, in equations (14.1) and (14.12), nothing changes whether the economy is at the ZLB ($i = 0$) or not ($i > 0$). The model behaves exactly the same. Hence, the model accommodates permanent zero-lower-bound episodes, and business-cycle shocks and policies have identical effects then.

This property of the slackish business cycle model is entirely different from what occurs in a New Keynesian model. The New Keynesian model predicts that when the zero lower bound lasts sufficiently long, output and inflation collapse to implausibly low levels, and policies such as public spending and forward guidance have extraordinarily large effects (Michaillat and Saez 2021). More generally, in New Keynesian models, the zero lower bound is "a topsy-turvy world, in which many of the usual rules of macroeconomics are stood on their head" (Eggertsson and Krugman 2012, p. 1472). The zero lower bound is not topsy-turvy in slackish models.

14.8. Effects of aggregate demand and supply shocks

Now that we've characterized the solution in our dynamic model of slack, we can use comparative statics to look at the effect of aggregate demand and supply shocks on aggregate tightness, output, and employment.

14.8.1. Qualitative effects of aggregate demand shocks

The most natural source of a negative aggregate demand shock is an increase in the marginal utility of wealth, $\mathcal{W}'(0)$. Another possible negative aggregate demand shock is a decrease in the time discount rate, δ. After such a shock, households become more thrifty: they desire to save more and consume less, which depresses aggregate demand (equation (14.12)). Graphically, the shock triggers an inward rotation of the aggregate demand curve (figure 14.2A).

We can see that as the aggregate demand curve rotates downward, both aggregate tightness θ and output y decrease. The unemployment rate, $u(\theta)$, increases due to a decrease in tightness, as shown by (14.2). A negative demand shock means people want to consume less, so there is less output produced and a lower tightness. There is more slack in the economy, as a larger share of the economy's productive capacity remains unsold and unused.

Looking at the effect of an increase in the marginal utility of wealth, $\mathcal{W}'(0)$, is interesting because it produces the Keynesian paradox of thrift. The paradox states that in a situation where everyone wants to save more, because savings in the economy are fixed, everyone cuts their consumption but the amount of savings remains unchanged. When

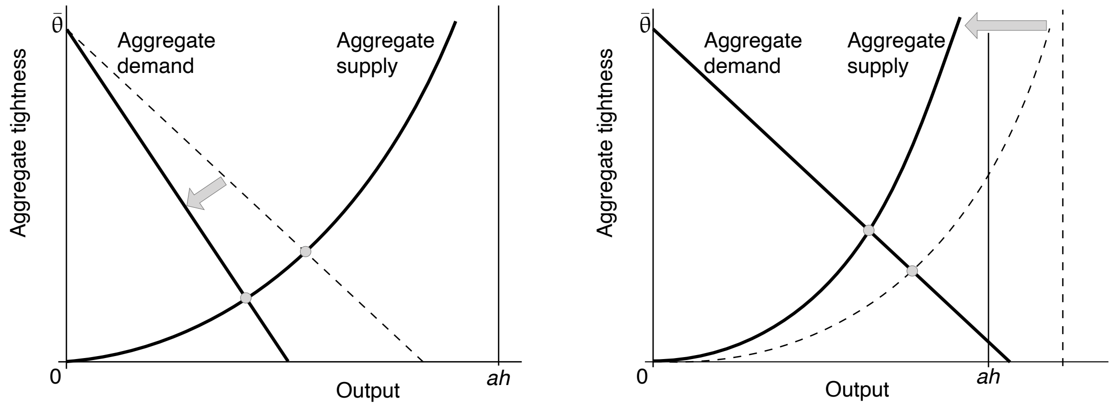

A. Negative aggregate demand shock

B. Negative aggregate supply shock

FIGURE 14.2. Aggregate demand and supply shocks in the slackish business cycle model

The aggregate supply curve is given by equation (14.1). The aggregate demand curve is given by equation (14.12).

the marginal utility of wealth is higher, people want to increase their wealth relative to their peers, so they favor saving over consumption. But in the aggregate relative wealth is fixed at 0 because everybody is the same; the only way to increase saving relative to consumption is to consume less. Thus, the only effect an increase in the marginal utility of wealth has is a decrease in output and consumption—savings remain the same in absolute and relative terms.

The efficiency analysis of chapter 10 continues to apply to this model, as the economy is organized around a single market that features a Beveridge curve with two sources of welfare cost: unemployed labor services and labor services devoted to recruiting instead of consumption. Hence, the efficient aggregate tightness θ^* is given by (10.10): $\theta^* \beta(\theta^*) \kappa = 1$, where $\beta(\theta)$ is the Beveridge elasticity, given by (10.9). The key is that the efficient tightness is not affected by aggregate demand shocks, so the efficient unemployment rate $u^* = u(\theta^*)$ is not affected either. This implies that the unemployment gap $u - u^*$ moves together with the unemployment rate. In particular, the unemployment gap widens when the unemployment rate rises after a negative aggregate demand shock.

14.8.2. Quantitative effects of aggregate demand shocks

To quantify the effect of aggregate demand shocks on the unemployment rate, we derive the semielasticity of unemployment with respect to the marginal utility of wealth, $\varepsilon^u_{\mathcal{W}'}$. The semielasticity is defined as follows:

$$\varepsilon^u_{\mathcal{W}'} = \mathcal{W}'(0) \cdot \frac{du}{d\mathcal{W}'(0)},$$

TABLE 14.1. Comparative statics in the slackish business cycle model

	Tightness θ	Output y	Employment l	Unemployment u	Gap $u - u^*$
A. Aggregate demand shocks					
Decrease in discount rate	–	–	–	+	+
Increase in marginal utility of wealth	–	–	–	+	+
B. Aggregate supply shocks					
Decrease in labor productivity	+	–	+	–	–
Decrease in labor force	+	–	–	–	–
C. Policy shocks					
Decrease in nominal interest rate	+	+	+	–	–

The comparative statics are obtained by comparative statics in figures 14.2A, 14.2B, and 14.4A. A variable's increase is denoted by "+" and a decrease by "-".

and it is related to the elasticity of unemployment with respect to the marginal utility of wealth by $\varepsilon^u_{\mathcal{W}'} = u \cdot \epsilon^u_{\mathcal{W}'}$. Semielasticities are typically easier to interpret than elasticities for rate variables that take small values close to zero (such as the unemployment rate or slack rate or interest rate). In terms of notation, we use a curly ε instead of the regular ϵ to denote semielasticities.

Using the analysis of demand shocks in slackish markets, it is straightforward to obtain the elasticity of tightness with respect to the marginal utility of wealth. Formula (7.10) applies here too, so

$$\epsilon^\theta_{\mathcal{W}'} = \frac{\epsilon^d_{\mathcal{W}'}}{\epsilon^s_\theta - \epsilon^d_\theta},$$

where $\epsilon^d_{\mathcal{W}'}$ is the elasticity of aggregate demand with respect to the marginal utility of wealth, $\epsilon^d_{\mathcal{W}'} = -1/\alpha$, ϵ^d_θ is the elasticity of aggregate demand with respect to tightness, given by (7.13), and ϵ^s_θ is the elasticity of aggregate supply with respect to tightness, given by (9.14). Combining all these elements, we find that the elasticity is

$$\epsilon^\theta_{\mathcal{W}'} = \frac{-1}{\alpha\left[1-\eta(\theta)\right]u(\theta) + (1-\alpha)\eta(\theta)\tau(\theta)}.$$

Next, we compute the elasticity of the unemployment rate with respect to tightness from equation (14.2) (using again various results from appendix B):

$$\epsilon^u_\theta = -\frac{f(\theta)}{f(\theta)+\lambda}\left[1-\eta(\theta)\right] = -\left[1-\eta(\theta)\right]\left[1-u(\theta)\right]. \tag{14.14}$$

Accordingly, the elasticity of the unemployment rate with respect to the marginal utility of wealth is

$$\epsilon^{u}_{\mathcal{W}'} = \epsilon^{u}_{\theta} \cdot \epsilon^{\theta}_{\mathcal{W}'} = \frac{1 - u(\theta)}{\alpha u(\theta) + (1-\alpha) \cdot \frac{\eta(\theta)}{1-\eta(\theta)} \cdot \tau(\theta)}.$$

Last, the semielasticity of the unemployment rate with respect to the marginal utility of wealth is

$$\varepsilon^{u}_{\mathcal{W}'} = u \cdot \epsilon^{u}_{\mathcal{W}'} = \frac{1 - u(\theta)}{\alpha + (1-\alpha) \cdot \frac{\eta(\theta)}{1-\eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)}}. \tag{14.15}$$

The semielasticity $\varepsilon^{u}_{\mathcal{W}'}$ gives the percentage-point increase in the unemployment rate when the marginal utility of wealth rises by 1%. We see that it is positive, confirming that the unemployment rate rises when the marginal utility of wealth rises (a negative aggregate demand shock). As usual in slackish models, the effect is sharply dependent on tightness—so the slackish business cycle model is state-dependent.

It's not easy to establish the monotonicity of the semielasticity because both the numerator and denominator are increasing in tightness. But we can obtain a few results, and guess a few properties of the semielasticity too—which we will validate numerically. First, when the economy is operating efficiently, we learn from (10.11) that the semielasticity is just $1 - u^* \approx 1$. Because $u(\theta)$ is small compared to 1, while the two terms in the denominator are commensurate, it is quite likely that the change in the denominator term in θ dominates the change in $u(\theta)$ in the numerator, which would then imply that the semielasticity is decreasing in θ around efficiency.

Figure 14.3A validates this conjecture when the slackish business cycle model is calibrated just like the labor market of chapter 11.[4] We see that the semielasticity is around 1 when the economy is efficient (tightness around 1). The semielasticity is for the most part decreasing in tightness. The logic is the same as in figure 9.3A: when the economy is slack, changes in aggregate demand have large effects on quantities and thus the unemployment rate; conversely, when the economy is tight, changes in aggregate demand have small effects on quantities and thus the unemployment rate. The monotonicity breaks down only when tightness is very low and the unemployment rate very high, because then the movements in the numerator of the semielasticity, $1 - u(\theta)$, take over.

14.8.3. Qualitative effects of aggregate supply shocks

The next step is to look at the effects of aggregate supply shocks on the economy and unemployment. A negative aggregate supply shock could be caused by a decrease in either

[4]The calibration is summarized in table 11.2. Most of the parameters have similar interpretations in the two models, so it seems appropriate to use the calibration. The sole relevant parameter with a different interpretation is α: it determines the shape of the production function in the labor market model and the shape of the utility function in this business cycle model. Both functions are concave, and we keep the same value of α in the two models for transparency and simplicity.

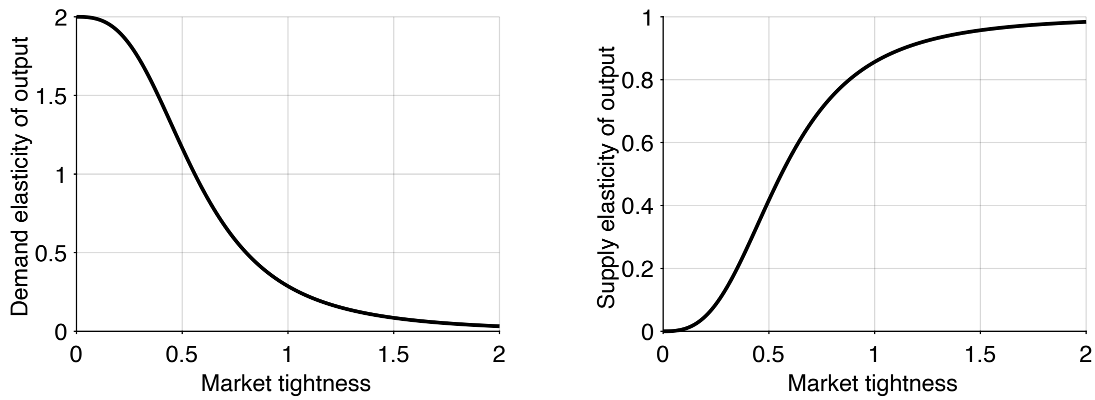

A. Semielasticity of unemployment to the marginal utility of wealth

B. Semielasticity of unemployment to aggregate capacity

FIGURE 14.3. Effects of aggregate shocks on unemployment in the slackish business cycle model

The semielasticity of unemployment to the marginal utility of wealth is given by (14.15). The semielasticity of unemployment to aggregate capacity is given by (14.16). The model is calibrated as in table 11.2.

labor productivity, a, or labor force participation, h.

After a decrease in labor force participation or productivity, the aggregate capacity shrinks, so the aggregate supply shifts inward (figure 14.2B). As a result, aggregate tightness increases but output decreases. Note that unlike under aggregate demand shocks, tightness and output move in opposite directions for aggregate supply shocks—the usual pattern that we have been observing since chapter 5. The unemployment rate $u(\theta)$ decreases because tightness increases. So here, people want to consume a similar amount of services, but fewer services are offered in the economy, so the economy becomes tighter: there is less slack, and labor services are sold more easily.

Although shocks to productivity and labor force have similar effects on output and tightness, they have a different effect on the employment level. After a reduction in labor force, the number of employed workers l drops: fewer workers want to work, so fewer workers work. This can be seen formally: $l = y/a$ and y decreases, so l decreases. After a reduction in productivity, on the other hand, the number of employed workers rises: workers are less productive so more workers are required to provide demanded services. Formally, $l = (1 - u)h$ and u decreases, so l actually increases here. In this model, the fact that output and employment move in opposite directions is specific to productivity shocks.

14.8.4. Quantitative effects of aggregate supply shocks

To quantify the effect of aggregate supply shocks on the unemployment rate, we derive the semielasticity of unemployment with respect to aggregate capacity, ε^{u}_{ah}. The semielasticity is defined as follows:

$$\varepsilon^u_{ah} = ah \cdot \frac{du}{dah}.$$

It is related to the elasticity of unemployment with respect to aggregate capacity by $\varepsilon^u_{ah} = u \cdot \epsilon^u_{ah}$.

Using the analysis of supply shocks in slackish markets, it is straightforward to obtain the elasticity of tightness with respect to aggregate capacity. Formula (7.15) applies here too, so

$$\epsilon^\theta_{ah} = \frac{-\epsilon^s_{ah}}{\epsilon^s_\theta - \epsilon^d_\theta},$$

where ϵ^s_{ah} is the elasticity of aggregate supply with respect to aggregate capacity, $\epsilon^s_{ah} = 1$, ϵ^d_θ is the elasticity of aggregate demand with respect to tightness, given by (7.13), and ϵ^s_θ is the elasticity of aggregate supply with respect to tightness, given by (9.14). (We have also set $\epsilon^d_{ah} = 0$ in (7.15) since aggregate demand does not depend on aggregate capacity.) Combining these elements, we obtain

$$\epsilon^\theta_{ah} = \frac{-1}{[1 - \eta(\theta)]\, u(\theta) + \frac{1-\alpha}{\alpha} \cdot \eta(\theta)\tau(\theta)}.$$

Finally, using the elasticity of unemployment with respect to tightness given by (14.14), we find that the semielasticity of the unemployment rate with respect to aggregate capacity is

$$\varepsilon^u_{ah} = u \cdot \epsilon^u_\theta \cdot \epsilon^\theta_{ah} = \frac{1 - u(\theta)}{1 + \frac{1-\alpha}{\alpha} \cdot \frac{\eta(\theta)}{1-\eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)}}. \tag{14.16}$$

The semielasticity ε^u_{ah} gives the percentage-point increase in the unemployment rate when the aggregate capacity rises by 1%. We see that it is positive, confirming that the unemployment rate rises when aggregate capacity rises (a positive aggregate supply shock). The effect is sharply dependent on tightness. As with semielasticity (14.15), we see that when the economy is operating efficiently, the semielasticity is just $(1 - u^*)\alpha \approx \alpha$. We also see that the semielasticity is likely decreasing in θ around efficiency.

Figure 14.3B validates these results in the calibrated slackish business cycle model. The semielasticity is around α when the economy is efficient (tightness around 1). The semielasticity is for the most part decreasing in tightness. Why is that? We know that output responds strongly to supply shocks in tight markets and weakly in slack markets. But the response of output, $y = (1 - u)ah$, is the mechanical response to capacity ah, dampened by the response of the unemployment rate u. So, it follows that there is little dampening in tight markets but significant dampening in slack markets. This then means that the unemployment rate responds little in tight markets but significantly in slack

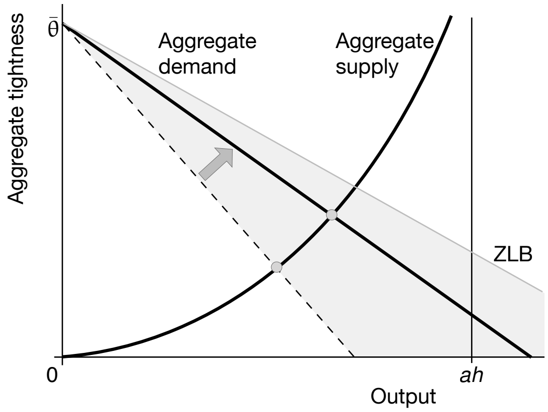

A. Decrease in nominal interest rate

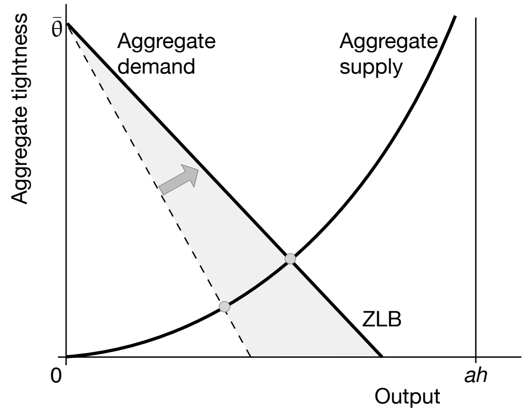

B. Nominal interest rate at the zero lower bound

FIGURE 14.4. Conventional monetary policy in the slackish business cycle model

The aggregate demand and supply curves are constructed in figure 14.1B. The gray cone indicates all the positions that the aggregate demand curve can reach when the nominal interest rate is reduced from its current level. The zero-lower-bound curve shows the position of the aggregate demand curve at the zero lower bound; it is the most outward position that the aggregate demand curve can reach; it is obtained from equation (14.12) when $i = 0$.

markets—just as the figure shows.

14.9. Monetary policy

We now show how conventional monetary policy steers aggregate demand by modulating the nominal interest rate. In doing so, it affects aggregate tightness and thus the unemployment rate.

14.9.1. How does conventional monetary policy operate?

We start by describing how the model responds to conventional monetary policy. We consider an unexpected permanent decrease in the nominal interest rate i. As inflation π remains constant, the real interest rate $r = i - \pi$ falls. When the real rate falls, wealth has lower returns, so households desire to save less and consume more, which boosts aggregate demand (equation (14.12)). Conventional monetary policy therefore operates just like an aggregate demand shock.

Graphically, when the nominal interest rate drops, the aggregate demand curve rotates outward, so output and tightness increase (figure 14.4A). Since tightness rises, the unemployment rate decreases. And since the efficient tightness and unemployment rate are unaffected by monetary policy, the unemployment gap moves with the unemployment rate: it shrinks whenever the nominal interest rate is reduced.

14.9.2. The monetary multiplier

Next, we quantify the effect of conventional monetary policy on the unemployment rate. We compute and study the monetary multiplier,

$$\frac{du}{di}.$$

The monetary multiplier gives the percentage-point increase in the unemployment rate when the nominal interest rate increases by 1pp.

We follow the usual steps to compute the monetary multiplier. First, as formula (7.10) applies here too, we get the elasticity of tightness with respect to the interest rate:

$$\epsilon_i^\theta = \frac{\epsilon_i^d}{\epsilon_\theta^s - \epsilon_\theta^d},$$

where ϵ_i^d is the elasticity of aggregate demand with respect to the interest rate, given by

$$\epsilon_i^d = -\frac{1}{\alpha} \cdot \frac{i}{\delta + \pi - i},$$

and ϵ_θ^d and ϵ_θ^s are the elasticities of aggregate demand and supply with respect to tightness, given by (7.13) and (9.14). Combining these elements, we find that the elasticity is

$$\epsilon_i^\theta = -\frac{i}{\delta + \pi - i} \cdot \frac{1}{\alpha\left[1 - \eta(\theta)\right] u(\theta) + (1-\alpha)\eta(\theta)\tau(\theta)}.$$

Next, introducing the elasticity of unemployment with respect to tightness, given by (14.14), we find that the monetary multiplier is

$$\frac{du}{di} = \frac{u}{i} \cdot \epsilon_\theta^u \cdot \epsilon_i^\theta = \frac{1 - u(\theta)}{\delta + \pi - i} \cdot \frac{1}{\alpha + (1-\alpha) \cdot \frac{\eta(\theta)}{1-\eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)}}.$$

We see that the monetary multiplier takes the same expression as the semielasticity of unemployment with respect to the marginal utility of wealth, given by (14.15), up to the $1/(\delta - r)$ factor. This means that the properties of the multiplier can be immediately derived from those of the semielasticity. First, the monetary multiplier is positive, confirming that the unemployment rate rises when the nominal interest rate increases. Second, the monetary multiplier is also sharply slack-dependent. When the economy is operating efficiently, the multiplier is $(1 - u^*)/(\delta - r)$, and around efficiency, the multiplier is likely to decrease with tightness. Numerically, up to a factor $1/(\delta - r)$, the monetary multiplier behaves as described in figure 14.3A.

TABLE 14.2. Timeline of the zero-lower-bound episode and forward guidance

Marginal utility of wealth $\mathcal{W}'(0)$	Nominal interest rate i	Equilibrium costate variable ψ
A. Initial normal times: $t < 0$		
$\mathcal{W}'_n$	$i_n > 0$	$\psi_n = \mathcal{W}'_n/(\delta + \pi - i_n)$
B. ZLB episode: $0 < t < t_z$		
$\mathcal{W}'_z > \mathcal{W}'_n$	0	$\psi_z = \mathcal{W}'_z/(\delta + \pi) > \psi_n$
C. Forward guidance: $t_z < t < t_z + \Delta$		
$\mathcal{W}'_n$	0	$\psi_f = \mathcal{W}'_n/(\delta + \pi) < \psi_n$
D. Terminal normal times: $t_z + \Delta < t$		
$\mathcal{W}'_n$	$i_n > 0$	ψ_n

The parameter $t_z > 0$ gives the duration of the zero-lower-bound episode. The parameter $\Delta > 0$ gives the duration of forward guidance.

14.10. Forward guidance at the zero lower bound

In this section, we move away from conventional monetary policy and instead study the effects of forward guidance around zero-lower-bound episodes. At the zero lower bound, the nominal interest rate is stuck at zero, so conventional monetary policy cannot operate. This is why the central bank must turn to forward guidance—to circumvent the zero-lower-bound constraint.

Forward guidance is a type of monetary policy that consists in announcing ahead of time future interest rates. The goal is to make promises that can stimulate the economy today, although the current interest rate is stuck at zero. Forward guidance became popular around the time of the Great Recession because policymakers had no choice—since the zero lower bound became binding in many countries.

14.10.1. Timing of forward guidance

We consider a two-stage scenario where a forward-guidance episode follows a zero-lower-bound episode, as in Cochrane (2017). The complete scenario is presented in table 14.2.

Temporary zero-lower-bound episodes and forward guidance require a dynamic analysis. We therefore come back to the Euler equation (14.9). This is a differential equation that governs the dynamics of the costate variable $\psi(t)$, which indicates how tight the finances of the representative household are, and therefore how much they want to consume.

Importantly, at any point in time, the costate variable determines the model's aggregate tightness $\theta(t)$. Indeed, (14.1), (14.3), and (14.7) tell us that

$$\psi(t) = \frac{1-\alpha}{(ah)^{\alpha}} \cdot \frac{1}{[1+\tau(\theta(t))]^{1-\alpha}\,[1-u(\theta(t))]^{\alpha}}. \tag{14.17}$$

Since $\tau(\theta)$ is strictly increasing in θ while $u(\theta)$ is strictly decreasing in θ, the equation tells us that the costate variable $\psi(\theta)$ is a strictly decreasing function of tightness θ, which means that aggregate tightness is a strictly decreasing function of the costate variable at any point in time. So, once we determine the dynamics of the costate variable, we can immediately trace the dynamics of aggregate tightness, and then of the unemployment rate and all the other variables in the model. Whenever the costate variable is high, the economy is slack, and conversely, whenever the costate variable is low, the economy is tight.

Before time 0, the economy is in normal times: the marginal utility of wealth is at a low, normal level: $\mathcal{W}'(0) = \mathcal{W}'_n$. The nominal interest rate is strictly positive, at $i_n > 0$, to keep the equilibrium costate variable ψ at a normal level $\psi_n = \mathcal{W}'_n/(\delta + \pi - i_n)$. This keeps the equilibrium tightness at the desired level θ_n, which is related to ψ_n via equation (14.17).

Between times 0 and t_z, there is a zero-lower-bound episode. The marginal utility of wealth is elevated, depressing aggregate demand: $\mathcal{W}'(0) = \mathcal{W}'_z > \mathcal{W}'_n$. To compensate, the nominal interest rate is lowered to $i = 0$. However, this is insufficient to keep tightness at the desired level: the equilibrium costate variable ψ is elevated, taking a value $\psi_z = \mathcal{W}'_z/(\delta + \pi) > \mathcal{W}'_n/(\delta + \pi - i_n) = \psi_n$. Accordingly, equilibrium tightness is depressed during the zero-lower-bound episode: $\theta_z < \theta_n$.

To alleviate the situation, the central bank makes a forward-guidance promise at time 0: that it will maintain the nominal interest rate at 0 for a duration Δ once the ZLB is over.

After time t_z, the marginal utility of wealth returns at its normal, lower level: $\mathcal{W}'(0) = \mathcal{W}'_n < \mathcal{W}'_z$. Between times t_z and $t_z + \Delta$, the central bank fulfills its forward-guidance promise and keeps the nominal interest rate at 0. As a result, the equilibrium costate variable ψ is low, taking a value $\psi_f = \mathcal{W}'_n/(\delta + \pi) < \mathcal{W}'_n/(\delta + \pi - i_n) = \psi_n$. After time $t_z + \Delta$, monetary policy returns to normal, $i = i_n$, the equilibrium costate variable is back at its desirable level, ψ_n, and the equilibrium tightness is back at its desirable level, θ_n.

14.10.2. How does the zero lower bound affect the economy?

To fix ideas, we start with a zero-lower-bound episode without forward guidance. We analyze the episode using the phase lines in figure 14.5A. We go backward in time. After time t_z, monetary policy maintains the economy at its desired location by setting the appropriate nominal interest rate. The costate variable $\psi(t)$ is maintained at its normal value $\psi_n = \mathcal{W}'_n/(\delta + \pi - i_n)$.

We look next at the zero-lower-bound episode, between times 0 and t_z. Since trajectories are continuous, the economy is at the same point at the end of the zero-lower-bound episode and after it. However, during the zero-lower-bound episode, the point $\psi = \psi_n$ is not an equilibrium: it is below the equilibrium point, so $\dot{\psi} < 0$ at that point. Hence, the costate variable decreases over time between $t = 0$ and $t = t_z$ to reach that point. From this

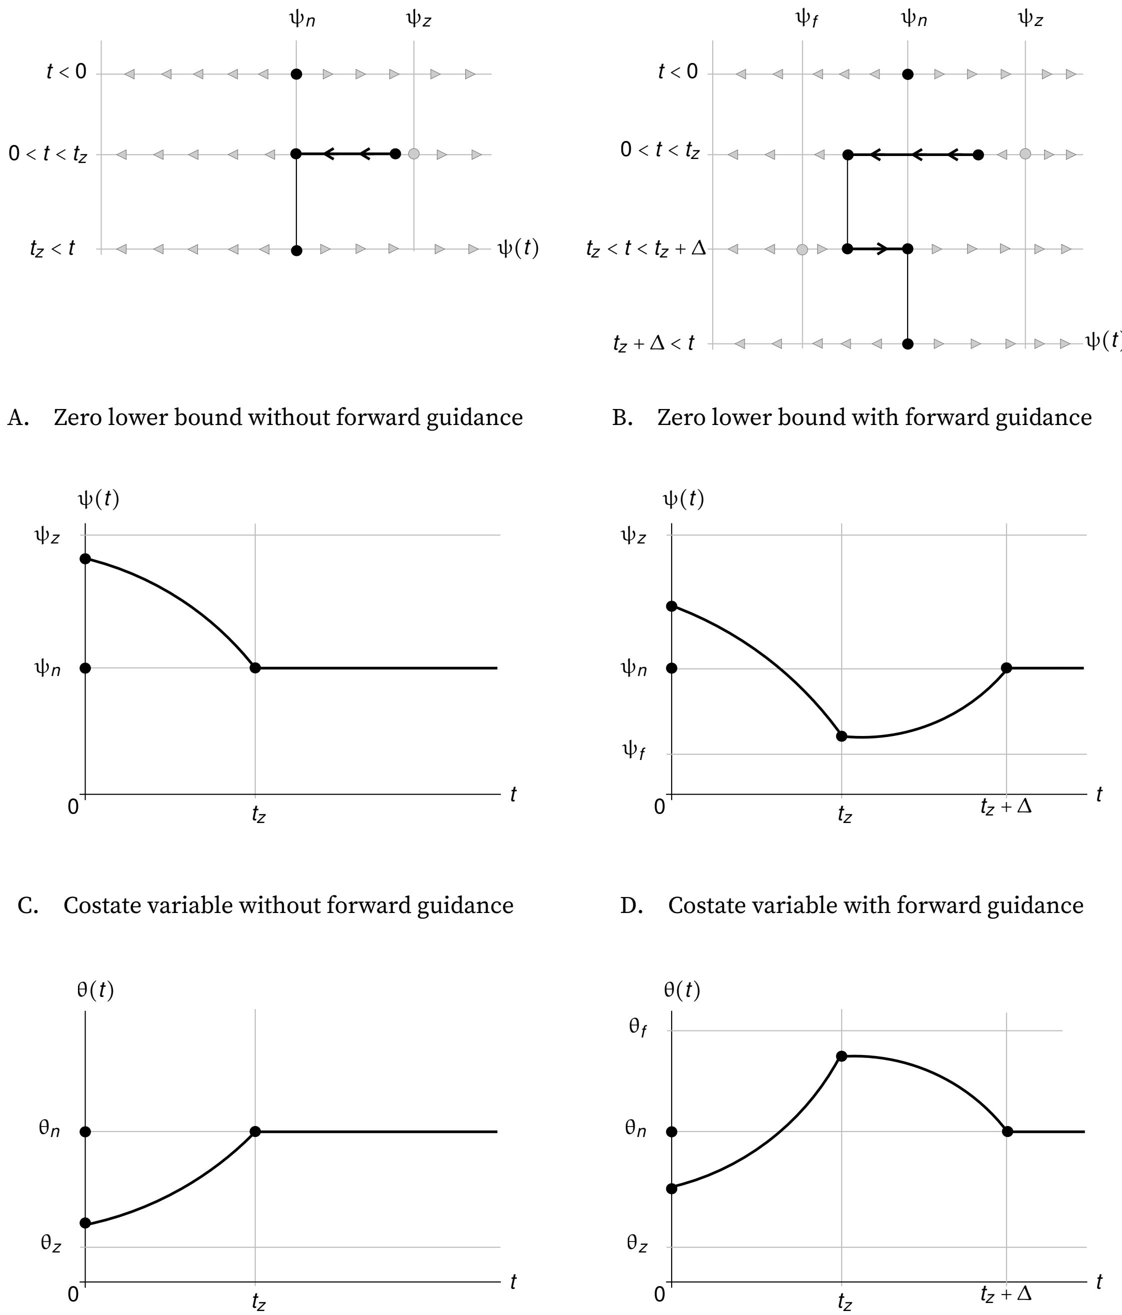

A. Zero lower bound without forward guidance

B. Zero lower bound with forward guidance

C. Costate variable without forward guidance

D. Costate variable with forward guidance

E. Aggregate tightness without forward guidance

F. Aggregate tightness with forward guidance

FIGURE 14.5. Zero-lower-bound episodes without and with forward guidance

The timeline of the analysis is presented in table 14.2. Panels A and B display the phase lines of the Euler equation (14.9) at the different stages of the analysis. Panels C and D describe how the costate variable $\psi(t)$ is moving over time—as determined from the phase lines. Panels E and F describe how aggregate tightness $\theta(t)$ is moving over time—as inferred from (14.17).

we infer that at time 0, when the zero-lower-bound episode starts, the costate variable $\psi(0)$ must jump somewhere between the zero-lower-bound equilibrium, ψ_z, and the normal equilibrium, ψ_n.

Note that the costate variable can only jump at time 0, because this is the only time when some new information is revealed. Without new information, the optimality of the behavior of households requires that the costate variable evolves continuously over time.

Overall, we see that the costate variable is above normal during the zero-lower-bound episode: it initially jumps to its highest point and then slowly decreases to its normal value (figure 14.5C). Since aggregate tightness is negatively related to the costate variable, as (14.17) shows, we infer that tightness jumps down at time 0, when the marginal utility of wealth jumps up and the zero lower bound becomes binding. After that, the aggregate tightness slowly recovers to reach its normal level at time t_z (figure 14.5E).

14.10.3. How does forward guidance operate?

Now, how can forward guidance improve things at the zero lower bound? Conventional monetary policy improves things by reducing the nominal interest rate to 0 at time 0, as soon as the negative aggregate demand shock hits. But we assume that the increase in the marginal utility of wealth is so large that the reduction in interest rate is insufficient to maintain the economy at the desired level: $\psi_n < \psi_z$. To supplement conventional monetary policy, the central bank therefore introduces forward guidance.

We analyze the zero-lower-bound episode with forward guidance using the phase lines in figure 14.5B. We go backward in time again. After time $t_z + \Delta$, monetary policy maintains the economy at its desired location by setting the appropriate nominal interest rate. The costate variable is maintained at $\psi = \psi_n$.

Between times t_z and $t_z + \Delta$, the economy is in forward guidance. Since trajectories are continuous, the economy is at the same point at the end of forward guidance and after it. However, during the forward-guidance episode, the point $\psi = \psi_n$ is not an equilibrium: it is above the equilibrium point, $\psi = \psi_f$, so $\dot{\psi} > 0$ at that point. Hence, the costate variable increases over time between $t = t_z$ and $t = t_z + \Delta$ to reach that point. This means that at the beginning of forward guidance, the costate variable must be below the normal equilibrium ψ_n. We also see that when the forward guidance lasts longer (higher Δ), the costate variable is pushed to lower levels at the beginning of forward guidance.

We then turn to the zero-lower-bound episode, between times 0 and t_z. At time t_z, the costate variable is lower than the normal level and lower than the zero-lower-bound equilibrium point, $\psi = \psi_z$. This means that $\dot{\psi} < 0$ at that point. Hence, the costate variable decreases over time between $t = 0$ and $t = t_z$ to reach that point. This implies that at time 0, when the zero-lower-bound episode starts, the costate variable $\psi(0)$ must jump somewhere between the zero-lower-bound equilibrium, $\psi = \psi_z$, and the forward-guidance equilibrium, $\psi = \psi_f$. We also see that when the zero lower bound lasts longer (higher t_z), the costate variable is pushed to higher levels at time 0.

In sum, we see that the costate variable is below normal during the forward-guidance episode. During the zero lower bound, the costate is higher than during forward guidance, but it may be above or below normal, depending on the durations of the zero-lower-bound and forward-guidance episodes. If the zero-lower-bound episode is long enough, the costate variable initially jumps to its highest point, then slowly decreases, and finally increases back to its normal level (figure 14.5D). Since aggregate tightness is negatively related to the costate variable, tightness jumps down at time 0, when the zero lower bound becomes binding. After that, the aggregate tightness slowly increases until time t_z, before decreasing back down to reach its normal level at time $t_z + \Delta$ (figure 14.5F).

The boom engineered during forward guidance therefore improves the situation at the zero lower bound. Instead of reaching the normal equilibrium at time t_z, the economy reaches a point with above-normal tightness at that time, so at any time before t_z, tightness is higher than without forward guidance, so the economy is less slack. Forward guidance is more powerful if it lasts longer and if the zero lower bound is shorter.

14.10.4. Closed-form expressions for the costate variable

So far we have used the phase lines for a qualitative analysis of the zero lower bound and forward guidance. But, because the Euler equation (14.9) is an autonomous linear first-order differential equation, it is not difficult to obtain closed-form expressions for the costate variable. We need to find an expression on the zero-lower-bound episode and a separate expression on the forward-guidance episode. We go backward in time again. From these expressions, we could derive the dynamics of the aggregate tightness $\theta(t)$ using equation (14.17).

During forward guidance, the costate variable's dynamics are governed by

$$\dot{\psi}(t) - (\delta + \pi)\psi(t) = -\mathcal{W}_n',$$

with boundary condition $\psi(t_z + \Delta) = \psi_n$. Applying result F.28, we find that during forward guidance, when $t_z < t < t_z + \Delta$, the costate variable is given by

$$\psi(t) = \psi_f + \left(\psi_n - \psi_f\right) \exp(-(\delta + \pi)(t_z + \Delta - t)).$$

This expression actually tells us where the costate must be at the end of the zero-lower-bound episode, and at the start of forward guidance. When $t = t_z$, the costate variable equals

$$\psi(t_z) = \psi_f \cdot [1 - \exp(-(\delta + \pi)\Delta)] + \psi_n \cdot \exp(-(\delta + \pi)\Delta). \tag{14.18}$$

So the costate variable is between the forward-guidance and normal equilibrium points, ψ_f and ψ_n, and it is closer and closer to the forward-guidance equilibrium point, so lower and lower, as the forward guidance lasts longer (higher Δ).

During the zero lower bound, the costate variable's dynamics are governed by

$$\dot{\psi}(t) - (\delta + \pi)\psi(t) = -\mathcal{W}'_z,$$

with boundary condition that $\psi(t_z)$ is given by (14.18). Thus, we find that during the zero lower bound, when $0 < t < t_z$, the costate variable is given by

$$\begin{aligned}\psi(t) = {} & \psi_z \cdot [1 - \exp(-(\delta + \pi)(t_z - t))] \\ & + \left(\psi_f [1 - \exp(-(\delta + \pi)\Delta)] + \psi_n \exp(-(\delta + \pi)\Delta)\right) \exp(-(\delta + \pi)(t_z - t)).\end{aligned} \tag{14.19}$$

Here we see again that the costate variable is higher when the zero lower bound lasts longer (higher t_z) but lower when forward guidance lasts longer (higher Δ).

14.10.5. Forward-guidance puzzle

One possible reason why forward guidance became popular at the time of the Great Recession is that the New Keynesian model predicts that the policy might be extremely powerful.

Michaillat and Saez (2021, proposition 4) isolates two results that describe the potency of forward guidance in the New Keynesian model. Just as described in table 14.2, let's consider a zero-lower-bound episode of duration t_z followed by a forward-guidance episode of duration Δ. The first result is that there exists a threshold Δ^* such that a forward guidance longer than Δ^* transforms a zero-lower-bound episode of any duration into a boom. In addition, when forward guidance is longer than Δ^*, a long-enough forward guidance or a long-enough zero lower bound generates an arbitrarily large boom.

However, researchers quickly realized that in the data, future monetary policy has much more subdued effects on output and inflation (Del Negro, Giannoni, and Patterson 2023). The effects of future monetary policy on current output and inflation are implausibly strong in the New Keynesian model. The gap between the New Keynesian model's prediction and the empirical evidence is known as the forward-guidance puzzle.

The slackish model does not suffer from the forward-guidance puzzle, as forward guidance only has limited effects in the model—especially if the zero-lower-bound episode lasts a long time. We can use the closed-form expressions above to establish these results.

First, there exists a threshold t_z^* such that a zero-lower-bound episode longer than t_z^* prompts a slump—a costate variable above normal and thus tightness below normal—irrespective of the duration of forward guidance. That is, for $t_z > t_z^*$, and for any Δ,

$\psi(0) > \psi_n$, which implies $\theta(0) < \theta_n$ through (14.17).

This can be seen using the expression (14.19) for the costate variable during the zero-lower-bound episode. At time 0, the costate variable jumps to

$$\psi(0) = \psi_z \cdot [1 - \exp(-(\delta + \pi)t_z)] + \psi(t_z)\exp(-(\delta + \pi)t_z),$$

where $\psi(t_z)$ is given by (14.18). The value $\psi(t_z)$ is a weighted average of ψ_f and ψ_n, so the lowest value that it can take, irrespective of the duration of forward guidance, is ψ_f (which is reached if the forward-guidance episode lasts an infinitely long time). Since $\psi(0)$ is itself a weighted average of $\psi(t_z)$ and ψ_n, a lower bound on $\psi(0)$ is

$$\psi_z \cdot [1 - \exp(-(\delta + \pi)t_z)] + \psi_f \exp(-(\delta + \pi)t_z). \tag{14.20}$$

That lower bound is a weighted average of ψ_f and ψ_z, and the weight on $\psi_z > \psi_f$ is larger when t_z is larger. Thus, when t_z is large enough, $\psi(0)$ is close enough to ψ_z and therefore above ψ_n. In fact the threshold t_z^* is such that

$$\psi_z \cdot \left[1 - \exp\left(-(\delta + \pi)t_z^*\right)\right] + \psi_f \exp\left(-(\delta + \pi)t_z^*\right) = \psi_n,$$

guaranteeing that any zero-lower-bound episode longer than t_z^* prompts an initial slump. Reshuffling the equation, we can extract a closed-form expression for the zero-lower-bound duration:

$$t_z^* = \frac{1}{\delta + \pi} \cdot \ln\left(\frac{\psi_z - \psi_f}{\psi_z - \psi_n}\right).$$

Any zero-lower-bound episode longer than t_z^* generates a slump, irrespective of the length of the ensuing forward guidance.

Furthermore, the slump approaches the zero-lower-bound equilibrium as the zero-lower-bound duration approaches infinity, irrespective of forward guidance. We can see this by looking at the lower bound on the costate variable given by (14.20). As $t_z \to \infty$, the lower bound converges to ψ_z, which means that $\psi(0)$ itself converges to ψ_z—as $\psi(0)$ is sandwiched between the lower bound and ψ_z. So when the zero lower bound lasts a very long time, there is nothing that forward guidance can do: it is entirely ineffective, and the economy is stuck in the vicinity of the zero-lower-bound equilibrium.

Finally, arbitrarily large booms are impossible in the slackish model. The costate variable always remains between ψ_f and ψ_z, as (14.19) shows, so the costate variable is bounded below by ψ_f, implying that tightness is bounded above by θ_f.

14.11. Fiscal policy

Monetary policy plays a central role for stabilization in the United States and most other developed economies, but it is not always sufficient. In particular, when the nominal interest rate is constrained at the zero lower bound and aggregate demand remains too weak, it is natural to ask whether fiscal policy can provide additional support. Such support is especially important if the zero lower bound is expected to last a long time—because then the effectiveness of forward guidance diminishes greatly.

Indeed, the fiscal mechanism developed in chapter 12 in the context of the labor market carries over directly to the present business cycle environment. If the government decides to purchase an amount g of services to stimulate the economy, aggregate demand is boosted and becomes $y^d(\theta, i) + g$, where $y^d(\theta, i)$ is the private demand given by (14.12). As public spending stimulates aggregate demand, it also stimulates aggregate tightness, given by the modified equation:

$$y^d(\theta, i) + g = y^s(\theta).$$

As tightness rises, the unemployment rate, still given by (14.2), falls.

Hence, in this business cycle model, public spending operates just like public employment in the labor market model of part III. The effect of public spending can be illustrated in the tightness-output diagram of figure 14.3A and looks just like the effect of public employment in figure 12.1A. After an increase in public spending, the aggregate demand curve shifts outward and the economy moves up along the aggregate supply curve.

In fact the multiplier computed in chapter 12 continues to apply in this setting. It remains the case, as in (12.2), that the public-spending multiplier is given by

$$\frac{dy}{dg} = \frac{1}{1 - \epsilon^d_\theta / \epsilon^s_\theta},$$

where ϵ^d_θ is the elasticity of private aggregate demand with respect to aggregate tightness, still given by (7.13), and ϵ^s_θ is the elasticity of aggregate supply with respect to aggregate tightness, still given by (9.14). Hence, the properties of the public-spending multiplier are determined by the positive elasticity ratio

$$-\frac{\epsilon^d_\theta}{\epsilon^s_\theta} = \frac{1-\alpha}{\alpha} \cdot \frac{\eta(\theta)}{1-\eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)}.$$

In chapter 9, we showed that the elasticity ratio determines the effects of demand and supply shocks in the slackish market model. Here, it determines all the properties of the public-spending multiplier.

Just like the public-employment multiplier, the public-spending multiplier is strictly positive, strictly less than 1, and strictly decreasing with tightness. Hence, purchasing public services increases output and reduces unemployment, but it diminishes purchases of private services by households. There is some crowding out of private services by public services: when the government purchases public services, it makes it harder for households to purchase private services because the government and households compete to hire the same workers. Furthermore, when tightness is lower, the multiplier is higher: public spending crowds out private spending less and reduces unemployment more.

14.12. Summary

This chapter develops a slackish model of a service-based economy in which households both produce and consume services. Because all transactions must pass through a matching market, households are unable to sell all of their productive capacity, giving rise to slack. The assumption that households derive utility not only from consumption but also from relative wealth ensures that aggregate demand is well defined. The model isolates the interaction between consumption–saving decisions (aggregate demand) and productive capacity (aggregate supply) via the aggregate tightness, which in turn determines output, consumption, and unemployment.

The model delivers transparent comparative statics. Negative demand shocks reduce tightness and output, and raise unemployment. Negative supply shocks have distinct signatures: declines in labor force participation and declines in productivity both reduce output and raise tightness, but they have opposite implications for employment. Finally, the chapter shows that the model continues to behave perfectly normally at the zero lower bound—although of course conventional monetary policy becomes ineffective then.

At the zero lower bound, forward guidance can stimulate the economy, especially when the zero-lower-bound episode is short. The key mechanism behind the model dynamics at the zero lower bound and during forward guidance is the Euler equation for the costate variable, which maps one-for-one into dynamics for aggregate tightness. This mapping allows temporary zero-lower-bound episodes and forward guidance to be analyzed using transition paths for tightness rather than with comparative statics. The model therefore links intertemporal household behavior directly to slack dynamics.

Fiscal policy also provides a direct stabilization tool in the model. Public spending stimulates aggregate demand, raising tightness and reducing unemployment. The public-spending multiplier—which measures the increase in output for each service purchased by the government—is positive, below one, and larger when the economy is slack.

Bibliography

Anderson, Cameron, John A. D. Hildreth, and Laura Howland. 2015. "Is the Desire for Status a Fundamental Human Motive? A Review of the Empirical Literature." *Psychological Bulletin* 141 (3): 574–601. https://doi.org/10.1037/a0038781.

BLS. 2025a. "All Employees, Goods-Producing." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/USGOOD.

BLS. 2025b. "All Employees, Service-Providing." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/SRVPRD.

Camerer, Colin, George Loewenstein, and Drazen Prelec. 2005. "Neuroeconomics: How Neuroscience Can Inform Economics." *Journal of Economic Literature* 43 (1): 9–64. https://doi.org/10.1257/0022051053737843.

Cheng, Joey T. and Jessica L. Tracy. 2013. "The Impact of Wealth on Prestige and Dominance Rank Relationships." *Psychological Inquiry* 24 (2): 102–108. https://psycnet.apa.org/doi/10.1080/1047840X.2013.792576.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" In *Handbook of Macroeconomics*, vol. 1, edited by John B. Taylor and Michael Woodford, chap. 2. Amsterdam: Elsevier. https://doi.org/10.1016/S1574-0048(99)01005-8.

Cochrane, John H. 2017. "The New-Keynesian Liquidity Trap." *Journal of Monetary Economics* 92: 47–63. https://doi.org/10.1016/j.jmoneco.2017.09.003.

Cole, Harold L., George J. Mailath, and Andrew Postlewaite. 1995. "Incorporating Concern for Relative Wealth Into Economic Models." *Quarterly Review* 19 (3): 12–21. https://doi.org/10.21034/qr.1932.

Coy, Peter. 2021. "Here's Why So Many People Intend to Die With Money in the Bank." Bloomberg Businessweek. https://www.bloomberg.com/news/articles/2021-03-25/here-s-why-so-many-people-intend-to-die-with-money-in-the-bank.

Del Negro, Marco, Marc P. Giannoni, and Christina Patterson. 2023. "The Forward Guidance Puzzle." *Journal of Political Economy Macroeconomics* 1 (1): 43–79. https://doi.org/10.1086/724214.

Diamond, Peter A. 1982. "Aggregate Demand Management in Search Equilibrium." *Journal of Political Economy* 90 (5): 881–894. https://doi.org/10.1086/261099.

Eggertsson, Gauti B. and Paul Krugman. 2012. "Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach." *Quarterly Journal of Economics* 127 (3): 1469–1513. https://doi.org/10.1093/qje/qjs023.

Fisher, Irving. 1930. *The Theory of Interest, as determined by Impatience to Spend Income and Opportunity to Invest it*. New York: Macmillan.

Fiske, Susan T. 2010. "Interpersonal Stratification: Status, Power, and Subordination." In *Handbook of Social Psychology*, vol. 2, edited by Susan T. Fiske, Daniel T. Gilbert, and Gardner Lindzey, chap. 26. Hoboken, NJ: Wiley.

Frank, Robert H. 1985. *Choosing the Right Pond: Human Behavior and the Quest for Status*. Oxford: Oxford University Press.

Heffetz, Ori and Robert H. Frank. 2011. "Preferences for Status: Evidence and Economic Implications." In *Handbook of Social Economics*, vol. 1A, edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, chap. 3. Amsterdam: Elsevier.

Keynes, John Maynard. 1919. *The Economic Consequences of the Peace*. London: Macmillan.

Mattan, Bradley D., Jennifer T. Kubota, and Jasmin Cloutier. 2017. "How Social Status Shapes Person Perception and Evaluation: A Social Neuroscience Perspective." *Perspectives on Psychological Science* 12 (3): 468–507. https://doi.org/10.1177/1745691616677828.

Michaillat, Pascal and Emmanuel Saez. 2021. "Resolving New Keynesian Anomalies with Wealth in the Utility Function." *Review of Economics and Statistics* 103 (2): 197–215. https://doi.org/10.1162/rest_a_00893.

Ridgeway, Cecilia L. 2014. "Why Status Matters for Inequality." *American Sociological Review* 79 (1): 1–16. https://psycnet.apa.org/doi/10.1177/0003122413515997.

Steedman, Ian. 1981. "Time Preference, the Rate of Interest and Abstinence from Accumulation." *Australian Economic Papers* 20 (37): 219–234. https://doi.org/10.1111/j.1467-8454.1981.tb00353.x.

Stock, James H. and Mark W. Watson. 2010. "Modeling Inflation after the Crisis." In *Proceedings of Jackson Hole Economic Policy Symposium*, pp. 173–220. Kansas City, MO: Federal Reserve Bank of Kansas City. https://perma.cc/T4KA-XMVK.

Stock, James H. and Mark W. Watson. 2020. "Slack and Cyclically Sensitive Inflation." *Journal of Money, Credit and Banking* 52 (S2): 393–428. https://doi.org/10.1111/jmcb.12757.

Weiss, Yoram and Chaim Fershtman. 1998. "Social Status and Economic Performance: A Survey." *European Economic Review* 42: 801–820. https://doi.org/10.1016/S0014-2921(97)00137-2.