A Theory of Economic Slack

Pascal Michaillat

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

15 Slackish Phillips curve 3
15.1 Overview of the model 3
15.2 Slackish markets for services 4
15.3 Equilibrium aggregate supply 5
15.4 Directed search and price-tightness tradeoff 5
15.5 Saving and pricing by households 10
15.6 Monetary policy 13
15.7 Equilibrium aggregate demand 14
15.8 Phillips curve 14
15.9 Solution of the model in normal times 18
15.10 Implications for the conduct of monetary policy 18
15.11 Ineffective monetary communication 19
15.12 Empirical evidence on the Phillips curve 20
15.13 The economy after the coronavirus pandemic 23
15.14 Summary 25

Bibliography 27

CHAPTER 15.

Slackish Phillips curve

So far in the book we have assumed that inflation was exogenously determined by price norms. However, the post-pandemic spike in inflation in the United States, which coincided with a very hot economy revealed that inflation might arise when the economy is inefficiently tight. This situation had not occurred in the United States in the modern, post-Volcker era of monetary policy since the US economy had not been inefficiently tight since the 1970s.

In this chapter, we insert a Phillips curve linking inflation to unemployment in the slackish business cycle model of chapter 14. The Phillips curve emerges from directed search: workers search for buyers offering the highest prices for their services, knowing of course that they will face more competition for high-paying jobs and less for low-paying jobs. The Phillips curve ensures divine coincidence, so inflation is on target when the economy is at full employment. The Phillips curve is smooth, but once it is written in the unemployment-inflation plane it is strictly convex in unemployment: inflation responds more strongly to unemployment when the labor market is inefficiently tight than when it is inefficiently slack.

15.1. Overview of the model

The model follows the structure of the slackish business cycle model of chapter 14. Households purchase and consume services produced by other households. All the trades on the market for services are mediated by a matching function, which produces slack.

However, here the market is divided into many submarkets in which each household recruits workers to provide services for them. In each submarket, a household posts help-wanted ads that offer a specific price per service. Job search is directed: job seekers choose which submarket to join, depending on the price per service offered and of course the number of job seekers and help-wanted ads in that market. From a job seeker's perspective, the most attractive submarkets offer high prices, have many help-wanted ads, and have few other job seekers competing for the available positions.

Because search is directed, households must decide not only how many help-wanted ads they post but also what price they offer to pay for services rendered. A high price is expensive but allows households to recruit workers faster; a low price is cheap but imposes a long recruiting time.

This price-setting decision generates a Phillips curve. As usual in monetary models, to ensure that the Phillips curve is nondegenerate, we assume that households face price-adjustment costs—reflecting the type of fairness constraints discussed in chapter 6.

15.2. Slackish markets for services

Services are sold through long-term worker-household relationships. Once a worker has matched with a household, she becomes a full-time employee of the household, where she provides k services per unit time. She remains so until they separate, which occurs at rate λ.

To recruit workers, each household $j \in [0, 1]$ posts V_j help-wanted ads. These ads specify the price P_j that the household plans to pay for the services purchased. Each help-wanted ad is serviced by κ recruiters per unit time. The recruiters are employees of household j devoted to recruiting other workers.

Workers of other households only work for or apply to one household. In that way, each household j constitutes a submarket. Each worker decides which submarket to join, based on the price P_j offered for services and the submarket tightness θ_j. The submarket tightness is the ratio of the number of help-wanted ads and job seekers: $\theta_j = V_j/U_j$, where U_j is the number of unemployed workers from other households vying for jobs in household j. The tightness matters because it determines how quickly workers find a job. So job seekers trade off how much they are paid on the job with how long it takes to find a job.

In each submarket j, a matching function $m(U_j, V_j)$ determines the flow of new matches based on the number of job seekers and help-wanted ads. The matching function m satisfies the assumptions from chapter 4. Tightness θ_j determines the job-finding and ad-filling rates, as usual.

15.3. Equilibrium aggregate supply

In each submarket j, there are l_{ij} employed workers from household i, and a total $l_j = \int_0^1 l_{ij}(t)\,di$ employed workers from all households. There are also U_{ij} unemployed workers from household i, and a total $U_j = \int_0^1 U_{ij}(t)\,di$ unemployed workers from all households. The total number of workers from household i who are attached to submarket j is $h_{ij} = l_{ij} + U_{ij}$, and the labor force attached to submarket j is $h_j = \int_0^1 h_{ij}(t)\,di$.

Because it takes time to find a job in any submarket j, not all h_j workers are employed. The unemployment rate in submarket j is

$$u_j = \frac{U_j}{h_j}.$$

We assume that labor flows in each individual submarket are balanced: the number of workers who find a job with household j at any point in time equals the number who leave their jobs. This assumption implies that the unemployment rate in submarket j is a function of submarket tightness: $u_j = u(\theta_j)$, where the function $u(\theta)$ is given by (9.8).

The equilibrium submarket supply is the amount of services provided at any point in time given the submarket tightness, and given that submarket flows are balanced:

$$y_j^s(\theta_j) = \left[1 - u(\theta_j)\right] kh_j. \tag{15.1}$$

The equilibrium aggregate supply is the amount of services provided at any point in time across the entire economy, given that all submarket flows are balanced: $y^s = \int_0^1 y_j^s(\theta_j)\,dj$.

15.4. Directed search and price-tightness tradeoff

Household j observes the prices and tightnesses in all submarkets, and pick the one that's best for them. Indeed, different household j may offer different prices P_j for services rendered. They may also post a different number of help-wanted ads and attract a different number of workers, generating a different tightness in submarket j—making it easier or harder to find a job there. In this section, we analyze how directed search links tightnesses and prices in all submarkets. This link will be the basis for the Phillips curve.

15.4.1. Arbitrage condition

When workers are unemployed, they do not receive any income. When they are employed, they produce k services per unit time, each paid a price P_j. Furthermore, in the submarket, the unemployment rate is $u(\theta_j)$, so a fraction $u(\theta_j)$ of the workers sent to submarket j are in unemployment, and a fraction $1 - u(\theta_j)$ are in employment. Accordingly, household i's expected income when they send h_{ij} workers to submarket j is $kP_j[1 - u(\theta_j)]h_{ij}$. The expected income per worker is $kP_j[1 - u(\theta_j)]$.

Household i does not prefer one employer or the other, so they pick their employer based solely on the expected income in the associated submarket j. They arbitrage across submarkets and send their workers to the submarket offering the highest expected income per worker. In their role as employer, households are aware of this arbitrage, so all households offer to pay a price for services that allows them to compete with other households. Arbitrage across submarkets therefore imposes that

$$P_j[1 - u(\theta_j)]$$

is the same across all submarkets j. Sellers direct their search toward the most attractive buyers, which induces competition across all buyers. Through competition, buyers set prices so sellers are indifferent across all buyers. If a buyer set a lower price than the rest, so the expected income on their submarket is lower than the rest, they would simply not attract any applicants.

By arbitrage, there is a price level p such that for all j,

$$P_j \cdot [1 - u(\theta_j)] = P \cdot [1 - u(\theta)], \tag{15.2}$$

where the aggregate tightness is the ratio of the aggregate number of help-wanted ads to the aggregate number of job seekers, given by

$$\theta = \frac{\sum_j V_j}{\sum_j U_j}.$$

15.4.2. Effect of price on tightness

The price posted by household j determines the tightness θ_j in their submarket, and therefore the ease with which the household is able to recruit workers. For instance, if the household offers to pay a high price for services, it attracts many workers—who want to work for a generous employer—so it can tap into a large pool of job seekers. This means that the household faces a low tightness, and that it can easily recruit workers.

This matters to household j because when it is harder to recruit—when help-wanted ads take longer to fill—they must devote more services to recruiting, and cannot consume as much. Indeed, the l_j workers in household j's employ produce $y_j = kl_j$ services. But the V_j help-wanted ad require each κ recruiters, so a total of $k\kappa V_j$ services are allocated to recruiting instead of consumption. Household j's consumption is therefore only

$$c_j = y_j - k\kappa V_j = k\left[l_j - \kappa V_j\right] = kh_j\left[1 - u_j - \kappa v_j\right], \tag{15.3}$$

where $u_j = U_j/h_j = 1 - l_j/h_j$ is the unemployment rate in submarket j and $v_j = V_j/h_j$ is the ad rate in submarket j.

Because we assume that flows are balanced in submarket j, we can express the gap between output and consumption of services using a matching wedge:

$$y_j = \left[1 + \tau(\theta_j)\right] c_j, \tag{15.4}$$

where the matching wedge $\tau(\theta_j)$ is given by (9.13).

Then, the arbitrage condition (15.2) tells us that when the price posted in submarket j is P_j, so the relative price in submarket j is $p_j = P_j/P$, then tightness in submarket j is

$$\theta_j(p_j) = u^{-1}\left(1 - \frac{1 - u(\theta)}{p_j}\right). \tag{15.5}$$

The function u^{-1} is strictly decreasing, so the submarket tightness $\theta_j(p_j)$ is decreasing in the relative price p_j. This means that a high relative price leads to low tightness and thus a low matching wedge $\tau(\theta_j)$. In fact, using again the results from appendix B, we see that the elasticity of the function $\theta_j(p_j)$ with respect to p_j is:

$$\epsilon_p^\theta = \frac{1}{\epsilon_\theta^u} \cdot \frac{-[1 - u(\theta)]/p_j}{1 - [1 - u(\theta)]/p_j} \cdot (-1) = \frac{-1}{[1 - \eta(\theta_j)]u(\theta_j)}, \tag{15.6}$$

where the elasticity ϵ_θ^u of $u(\theta)$ with respect to θ comes from (14.13) and is evaluated at θ_j, and we used $[1 - u(\theta)]/p_j = 1 - u(\theta_j)$.

In that way, employers have monopsonistic power in this economy. By offering higher prices for services, they attract more workers to their submarket, which allows them to recruit more easily. Conversely, if they choose to pay lower prices, fewer job seekers will join the submarket, and it will be difficult to recruit.

Household j therefore faces a tradeoff between the price per service, P_j, and the submarket tightness θ_j, which itself determines the amount of services that are devoted to matching, $\tau(\theta_j)$. All workers in household j are paid a price P_j per service. The expenditure by household j on workers therefore is

$$P_j y_j = \left[1 + \tau(\theta_j)\right] P_j c_j.$$

The effective price of services is not just P_j but $[1 + \tau(\theta_j)]P_j$. The price involves the price per service as well as the matching cost. Household j would prefer to pay a low price P_j and face a low tightness θ_j. But that is not possible, as (15.5) shows—hence the tradeoff.

The aim of the household is to minimize the effective price for services, $[1+\tau(\theta_j)]P_j$, which includes the cost of services consumed as well as the cost of services required for matching per unit of consumption. A low price P_j economizes on the price per service, but it imposes a high tightness θ_j and thus a high wedge $\tau(\theta_j)$.

15.4.3. Efficiency without price-adjustment costs

For reference, we describe what would happen in a submarket without any price-adjustment cost. Household j is free to set any price P_j she wants to minimize the effective price of services, $[1+\tau(\theta_j)]P_j$.

That is, the household chooses P_j to minimize $[1+\tau(\theta_j)]P_j$ subject to the arbitrage condition (15.2). Because of the arbitrage condition, the effective service price can be written

$$[1+\tau(\theta_j)]P_j = P\cdot[1-u(\theta)]\cdot\frac{1+\tau(\theta_j)}{1-u(\theta_j)}.$$

Hence, by choosing a price P_j, household j sets submarket tightness θ_j to minimize $[1+\tau(\theta_j)]/[1-u(\theta_j)]$, which is equivalent to maximizing $[1-u(\theta_j)]/[1+\tau(\theta_j)]$.

This in turn is equivalent to choosing tightness θ_j to maximize $1-u_j-\kappa v_j$. Indeed, notice by combining (15.1) and (15.4) that consumption by household j is

$$c_j = \frac{1-u(\theta_j)}{1+\tau(\theta_j)}kh_j.$$

But consumption is also given by (15.3), which shows that

$$\frac{1-u(\theta_j)}{1+\tau(\theta_j)} = 1-u_j-\kappa v_j.$$

This is equivalent to choosing the unemployment rate u_j to minimize $u_j+\kappa v(u_j)$, where the unemployment and ad rates are related by the Beveridge curve (9.5), which prevails here as flows are assumed to be balanced in submarket j. This is exactly the social planner's problem studied in chapter 10. Thus, tightness θ_j and unemployment rate u_j are chosen efficiently, so that condition (10.10) holds. Here we have just recovered the efficiency result of Moen (1997): directed search produces efficient market outcomes.

The key intuition is that directed search lets households internalize the planner's problem when choosing posted prices. In the absence of price-adjustment costs, each household's private pricing problem reproduces the planner's condition, so submarket tightness is efficient.

15.4.4. Introducing price rigidity

Generally, the unemployment rate is not efficient (chapter 13). To introduce inefficiency in the model, prices must be somewhat rigid—to break the efficiency result from Moen (1997).

Given that search is directed instead of random, prices are determined by buyers to minimize the effective cost of services—prices are not determined by norms here. To introduce price rigidity, we therefore assume that changing prices is costly for buyers. As in Rotemberg (1982), household j incurs a quadratic price-adjustment cost when inflation in submarket j departs from normal inflation π^n. The inflation in submarket j is

$$\pi_j(t) = \frac{\dot{P}_j(t)}{P_j(t)}. \tag{15.7}$$

The flow disutility caused by actual inflation deviating from normal inflation is

$$\frac{\varpi}{2} \cdot \left(\pi_j - \pi^n\right)^2.$$

The parameter $\varpi > 0$ governs the price-adjustment cost. This quadratic cost appears in the household j's utility function.

What does this price-adjustment mean? What does it capture in reality? First of all, when prices fall, or increase less than normal, workers in household j feel shortchanged. Indeed, the price p_j directly determines their salary, $p_j k$. And we saw in chapter 7 that workers' morale dips when their wages do not grow as expected, and that their performance may dip as well. The quadratic cost reflects the cost of restoring workers' morale, or the cost caused by diminished performance, when price growth is below the normal growth π^n.[1]

When prices rise, or increase more than normal, household j might be unhappy to pay more for the same services. We saw in chapter 7 that higher-than-normal inflation upsets customers, who feel unfairly treated when they shop and have to pay more for the same good or service. Here we assume that this displeasure with higher prices take the form of a quadratic cost when price growth is above the normal growth π^n.

[1]The typical Keynesian wage-floor idea is that workers do not tolerate a drop in the level of nominal wages. But as Okun (1981, p. 93) notes, the same idea might apply to the growth rate of nominal wages, as we assume here: "Workers are likely to develop some notion of how fast they expect their wages to rise.... Any subsequent slowdown below the norm must be justified and explained to experienced workers. If the employer stresses the norm, a Keynesian wage-floor phenomenon may apply to the rate of increase of wages, and not just to their level."

15.5. Saving and pricing by households

We now present and solve the optimization problem of households. Households have essentially two decisions to make: how much income to save, and how to set prices for the services they purchase. From the optimal saving decision we will derive the aggregate demand curve, and from the optimal pricing decision the Phillips curve. We will then use these curves together with the aggregate supply curve to solve the model.

15.5.1. Household's utility function

Just as in chapter 14, households care about their consumption of services and their social status, derived from their relative real wealth. Here too, wealth is held in government bonds. In addition households incur a cost from changing the prices they pay for services.

Hence, household j maximizes the discounted sum of flow utilities,

$$\int_0^\infty \exp(-\delta t)\left[c_j(t)^{1-\alpha} + \frac{1}{a}\cdot\frac{B_j(t)-B(t)}{P(t)} - \frac{\varpi}{2}\cdot\left(\pi_j(t)-\pi^n\right)^2\right]dt, \tag{15.8}$$

where $\delta > 0$ is the time discount rate, $a > 0$ indicates the taste for consumption relative to saving and building social status, $B_j(t)$ is nominal wealth, held in government bonds, $B(t) = \int_0^1 B_j(t)\,dj$ is average nominal wealth in the economy, and $P(t)$ is the price level, given by (15.2).

15.5.2. Household's budget constraint

Household j's nominal budget constraint is similar to that in chapter 14 except that the income of the household now comes from employment in different households that might pay different prices for services. The budget constraint therefore is

$$\dot{B}_j(t) = i(t)\cdot B_j(t) + \int_0^1 P_i(t)y_{ij}(t)\,di - P_j(t)y_j(t) - T(t),$$

where $T(t)$ is a lump-sum tax levied by the government.

Given that the unemployment rate is $u(\theta_i)$ in each submarket i, and given the arbitrage condition (15.2), the household's income can be rewritten as follows:

$$\begin{aligned}\int_0^1 P_i(t)y_{ij}(t)\,di &= \int_0^1 P_i(t)[1-u(\theta_i(t))]kh_{ij}(t)\,di\\ &= p(t)[1-u(\theta(t))]k\int_0^1 h_{ij}(t)\,di\\ &= p(t)[1-u(\theta(t))]kh_j.\end{aligned}$$

Given the matching wedge (15.3) and tightness-price link (15.5), the household's expenditure can be written as a function of consumption and price:

$$P_j(t)y_j(t) = P_j(t)\left[1+\tau\left(\theta_j\left(\frac{P_j(t)}{P(t)}\right)\right)\right]c_j(t).$$

Accordingly, the nominal budget constraint can be written as

$$\dot{B}_j(t) = i(t)\cdot B_j(t) + p(t)[1-u(\theta(t))]kh_j - P_j(t)\left[1+\tau\left(\theta_j\left(\frac{P_j(t)}{P(t)}\right)\right)\right]c_j(t) - T(t).$$

Then, just as in chapter 14, we translate the nominal budget constraint into a real budget constraint:

$$\dot{b}_j(t) = r(t)\cdot b_j(t) + [1-u(\theta(t))]kh_j - p_j(t)\left[1+\tau\left(\theta_j\left(p_j(t)\right)\right)\right]c_j(t) - \frac{T(t)}{P(t)}, \tag{15.9}$$

where $b_j(t) = B_j(t)/P(t)$ is real bond holdings. The real interest rate is defined by $r(t) = i(t) - \pi(t)$, where the inflation rate is $\pi(t) = \dot{P}(t)/P(t)$.

15.5.3. Law of motion for household's price

Besides the budget constraint, household j faces another constraint: the law of motion for the price it pays for services. The law of motion follows from (15.7), and says that any price change requires some inflation, which itself is costly. The law of motion for the nominal price $P_j(t)$ is

$$\dot{P}_j(t) = \pi_j(t)P_j(t).$$

Hence, using a result of the sort of (14.5), we find that the law of motion for the relative price $p_j(t) = P_j(t)/P(t)$ is

$$\dot{p}_j(t) = \left[\pi_j(t) - \pi(t)\right]p_j(t). \tag{15.10}$$

15.5.4. Household's problem

We now solve household j's maximization problem, which is to maximize (15.8) subject to the laws of motion (15.9) and (15.10). We do so using a Hamiltonian (result G.23):

$$\begin{aligned}\mathcal{H}(t, c_j(t), \pi_j(t), b_j(t), p_j(t)) &= c_j(t)^{1-\alpha} + \frac{1}{a}\left[b_j(t) - b(t)\right] - \frac{\varpi}{2}\left[\pi_j(t) - \pi^n\right]^2 \\ &+ \psi_j^b(t)\left\{r(t)b_j(t) + [1-u(\theta(t))]kh_j - \left[1+\tau\left(\theta_j\left(p_j(t)\right)\right)\right]p_j(t)c_j(t) - \frac{T(t)}{P(t)}\right\} \\ &+ \psi_j^p(t)\left[\pi_j(t) - \pi(t)\right]p_j(t).\end{aligned}$$

The control variables are consumption $c_j(t)$ and inflation $\pi_j(t)$. The state variables are real bond holdings $b_j(t)$ and relative price $p_j(t)$. The costate variables are $\psi_j^b(t)$, which applies to the law of motion of real bond holdings (15.9), and $\psi_j^p(t)$, which applies to the law of motion of the relative price (15.10).

We now study the first-order conditions that characterize the household's optimal behavior. We begin with the first-order condition with respect to consumption: $d\mathcal{H}/dc_j = 0$. It gives

$$(1-\alpha)c_j(t)^{-\alpha} = \psi_j^b(t)\left[1+\tau(\theta_j(t))\right]p_j(t). \tag{15.11}$$

Next we turn to the first-order condition with respect to inflation: $d\mathcal{H}/d\pi_j = 0$. It yields

$$\varpi\left[\pi_j(t) - \pi^n\right] = \psi_j^p(t)\,p_j(t). \tag{15.12}$$

The next first-order condition is that with respect to real bond holdings: $d\mathcal{H}/db_j = \delta\psi_j^b(t) - \dot{\psi}_j^b(t)$. It gives the same condition as in chapter 14:

$$\dot{\psi}_j^b(t) = [\delta - r(t)]\,\psi_j^b(t) - \frac{1}{a}. \tag{15.13}$$

The final first-order condition is that with respect to the relative price: $d\mathcal{H}/dp_j = \delta\psi_j^p(t) - \dot{\psi}_j^p(t)$. This condition becomes

$$\begin{aligned}\dot{\psi}_j^p(t) - \delta\psi_j^p(t) &= \psi_j^p(t)[\pi(t) - \pi_j(t)] \\ &\quad + \psi_j^b(t)c_j(t)[1+\tau(\theta_j(p_j(t)))]\left[1 + p_j(t)\frac{\tau'(\theta_j(t))}{1+\tau(\theta_j(p_j(t)))}\theta_j'(p_j(t))\right].\end{aligned}$$

We can rewrite the following term:

$$p_j(t)\frac{\tau'(\theta_j(t))}{1+\tau(\theta_j(p_j(t)))}\theta_j'(p_j(t)) = \frac{\theta_j(t)}{1+\tau(\theta_j(p_j(t)))}\cdot\tau'(\theta_j(t))\cdot\frac{p_j(t)}{\theta_j(t)}\cdot\theta_j'(p_j(t)) = \epsilon_\theta^{1+\tau}\cdot\epsilon_p^\theta,$$

where the elasticity $\epsilon_\theta^{1+\tau}$ is given by (7.12) and the elasticity ϵ_p^θ is given by (15.6), both evaluated at $\theta_j(t)$. Hence, the term boils down to

$$p_j(t)\frac{\tau'(\theta_j(t))}{1+\tau(\theta_j(p_j(t)))}\theta_j'(p_j(t)) = -\frac{\eta(\theta_j(t))}{1-\eta(\theta_j(t))}\cdot\frac{\tau(\theta_j(t))}{u(\theta_j(t))}.$$

Noting lastly that $[1+\tau(\theta_j(t))]c_j(t) = y_j(t)$, we simplify the first-order condition to

$$\dot{\psi}_j^p(t) = [\delta + \pi(t) - \pi_j(t)]\psi_j^p(t) + \psi_j^b(t)y_j(t)\left[1 - \frac{\eta(\theta_j(t))}{1-\eta(\theta_j(t))}\cdot\frac{\tau(\theta_j(t))}{u(\theta_j(t))}\right]. \tag{15.14}$$

15.5.5. Symmetric solution

We focus on a symmetric solution to the model, in which all households behave the same, and all submarkets are therefore the same. In this symmetric situation, all submarket tightnesses θ_j are the same and equal the aggregate tightness θ. All submarket prices P_j are the same and equal the price level P, so all relative prices p_j equal 1, and submarket inflation π_j is the same as aggregate inflation π.

We can drop the index j on all individual variables.

The first-order conditions are necessary for an interior solution to the household's problem. So the symmetric solution must satisfy them. Under symmetry, the conditions (15.11), (15.12), (15.13), and (15.14) simplify as follows:

$$(1-\alpha)c(t)^{-\alpha} = \psi^b(t)\,[1+\tau(\theta(t))] \tag{15.15}$$

$$\varpi\left[\pi(t) - \pi^n\right] = \psi^p(t) \tag{15.16}$$

$$\dot{\psi}^b(t) = [\delta - r(t)]\,\psi^b(t) - \frac{1}{a} \tag{15.17}$$

$$\dot{\psi}^p(t) = \delta\psi^p(t) + \psi^b(t)y(t)\left[1 - \frac{\eta(\theta(t))}{1-\eta(\theta(t))}\cdot\frac{\tau(\theta(t))}{u(\theta(t))}\right]. \tag{15.18}$$

15.6. Monetary policy

Before solving the model, we must specify monetary policy. Monetary policy has two dimensions in the model: setting interest rates and communicating around the inflation that can be expected in the future.

First, the central bank maintains a fixed real interest rate r. This is achieved by setting the nominal interest rate to $i(t) = r + \pi(t)$ at any point. There is no Taylor rule here: the central bank adjusts the nominal interest rate one-for-one with inflation to keep the real interest rate fixed. The central bank focuses on the real interest rate because this is the relevant price in the model: it is the relevant interest rate for aggregate demand.

Second, we assume that through effective communication and management of expectations, the Federal Reserve anchors people's inflation expectations to the inflation target, which is 2% in practice (FOMC 2026). Hence, we assume that normal inflation π^n is just the Federal Reserve's inflation target π^*. That is, the Federal Reserve manages to convince people that prices will grow at their target inflation rate. This assumption plays a central role in producing the divine coincidence.

15.7. Equilibrium aggregate demand

We now derive the aggregate demand in the model. Just as in chapter 14, the aggregate demand is determined by two first-order conditions: (15.15) and (15.17).

These equations are the same equations as in chapter 14, so the aggregate-demand block is not affected by directed search or endogenous inflation. The reason is that by setting a real interest rate, monetary policy isolates aggregate demand from inflation. Inflation matters for aggregate demand only insofar as it affects the real interest rate. With the typical Taylor rule assumed in New Keynesian model, inflation and real rate are linked. But that assumption appears there for convenience—to ensure determinacy of the equilibrium (Cochrane 2011). Here the equilibrium is determinate without it, so we simplify the analysis by making the reasonable assumption that the central bank fixes directly the price that matters: the real interest rate.

As in chapter 14, to respect the transversality condition and no-Ponzi condition, the costate variable ψ^b cannot diverge to ∞, and because the differential equation (15.17) is a source, the costate variable must jump to its equilibrium value at time 0, which is again given by equation (14.9): $\psi^b = 1/[(\delta - r)a]$.

Then, again as in chapter 14, the household's consumption is a function of the real interest rate and aggregate tightness. The function is obtained by plugging the value of the costate variable in (15.15). This gives again equation (14.10), from which we derive the same aggregate demand as in equation (14.11):

$$y^d(\theta, r) = \frac{[(1-\alpha)(\delta - r)a]^{1/\alpha}}{[1+\tau(\theta)]^{1/\alpha - 1}}. \tag{15.19}$$

15.8. Phillips curve

Finally, we derive the Phillips curve in the model. The Phillips curve comes from the two first-order conditions that we have not used so far: (15.16) and (15.18).

Equation (15.16) simply says that inflation is a linear function of the costate variable associated with the relative price: $\pi(t) = \pi^* + \psi^p(t)/\varpi$.

The costate variable $\psi^p(t)$ is given by (15.18), which is a linear first-order differential equation. The equation is a source since $\delta > 0$. So if the costate variable does not jump to the equilibrium point of the differential equation, it exhibits explosive dynamics: going to ∞ as $t \to \infty$. Such dynamics imply that inflation would also diverge to $\pm\infty$, which cannot be optimal as that generates infinitely negative utility. So this costate variable must jump to the equilibrium point, just like the costate variable associated with real wealth:

$$\psi^p = \frac{y(\theta)}{\delta(\delta - r)a} \cdot \left[\frac{\eta(\theta)}{1-\eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)} - 1\right].$$

Accordingly, inflation jumps at time 0 to the following value:

$$\pi = \pi^{*} + \frac{y(\theta)}{\varpi\delta(\delta - r)a} \cdot \left[\frac{\eta(\theta)}{1 - \eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)} - 1\right]. \tag{15.20}$$

Equation (15.20) is the Phillips curve of the model. It links inflation to aggregate tightness, which is itself determined from the aggregate demand and supply curves. It has a few striking properties, which we review now.

The first property of the Phillips curve is that the divine coincidence holds: inflation is on target ($\pi = \pi^{*}$) whenever the economy operates efficiently—whenever the economy is at full employment ($u = u^{*}$). This can be seen because, when the economy operates efficiently, the term in square brackets in the Phillips curve is 0, as we established in equation (10.11).

In fact the term in square brackets in (15.20) measures the productive inefficiency of the economy. When the economy is inefficiently tight, the term is positive, so inflation is above target. When the economy is inefficiently slack, the term is negative, so inflation is below target. Furthermore, the response of inflation to departures from efficiency is governed by the price-adjustment cost ϖ. When prices are very costly to change, ϖ is large and inflation does not respond much to changes in tightness. By contrast, when prices are easy to change, ϖ is small and inflation responds much more forcefully to changes in tightness.

Thanks to the divine coincidence, the central bank does not face an inflation-unemployment tradeoff here: by bringing the unemployment rate to the FERU, it mechanically brings inflation to its target. Of course, the divine coincidence arises because we assumed that people take the inflation target as normal inflation. If that's not the case, the divine coincidence breaks down, as we discuss below. But when it's the case, the government automatically achieves its price mandate when it achieves its employment mandate.

We saw that inflation is above target when tightness is inefficiently high and below target when tightness is inefficiently low. But, more generally, inflation is strictly increasing in tightness around full employment and above. This can be seen because output $y(\theta)$, matching elasticity $\eta(\theta)$, matching wedge $\tau(\theta)$ are increasing in tightness, while the unemployment rate $u(\theta)$ is strictly decreasing in tightness. So above full employment ($\theta > \theta^{*}$), the term in square brackets is positive and growing with tightness and output is positive and growing with tightness so inflation is growing with tightness. Around full employment, the term in square brackets is 0, so the slope of the Phillips curve is determined by the derivative of the term in square brackets, which is strictly positive. Once again, inflation is increasing in tightness.

The logic actually continues to hold even below full employment ($\theta < \theta^{*}$) as long as output is large enough. Below full employment there are two opposing forces: output is

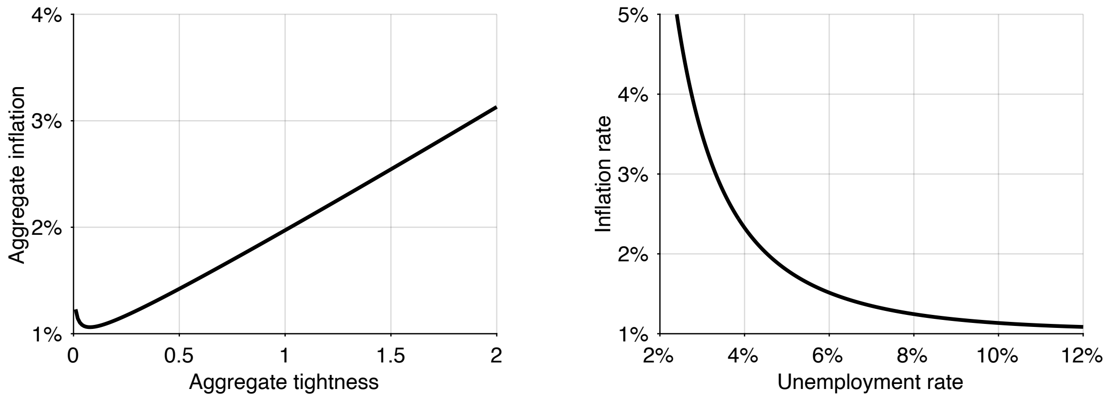

A. Tightness-inflation curve

B. Unemployment-inflation curve

FIGURE 15.1. Phillips curves in the calibrated model

The Phillips curve is given by equation (15.20). Tightness and unemployment are related by equation (9.8). The model is calibrated as in table 11.2, with in addition labor productivity set to $k = 1$ and price-adjustment costs set to $\varpi = 100$.

increasing with tightness; the term in square brackets is increasing with tightness, but that term is negative and becoming less negative. Hence the product of the two terms is negative and could be increasing or decreasing. When tightness is close to full employment and the term in square brackets is close to 0, the change in the term dominates and inflation is increasing in tightness. When tightness is much lower and the term is square brackets is very negative, the negative change in output dominates and inflation might decrease with tightness.

This non-monotonicity can be seen in figure 15.1A, which plots the Phillips curve under the calibration of chapter 11 (see table 11.2). In this calibration, however, inflation is increasing in tightness over the relevant range of tightness: for any tightness above 0.08. This means that generally, as expected, inflation rises as the economy becomes tighter.

The final numerical property is that the Phillips curve appears almost linear in tightness under this calibration. This is in contrast with the aggregate supply curve, which is sharply nonlinear under the same calibration (the shape of the aggregate supply curve appears in figure 9.1B). So while output increases in a very convex fashion with tightness, inflation increases almost linearly with tightness.

Once we represent the Phillips curve in a more traditional unemployment-inflation plane, the curvature of the Phillips curve changes drastically (figure 15.1B). The unemployment rate is a strictly decreasing function of tightness, given by (9.8), so tightness is a strictly decreasing function of the unemployment rate, $\theta(u)$. The figure plots $\pi(\theta(u))$, where $\pi(\theta)$ is given by (15.20). As expected, the inflation rate is strictly decreasing in the unemployment rate, but interestingly, inflation is now a strictly convex function of unemployment. In the calibration, the point of divine coincidence occurs when inflation is $\pi^* = 2\%$ and unemployment is $u^* = 4.5\%$. The Phillips curve goes through it, and it is much steeper when the unemployment rate is inefficiently low than when it is inefficiently high. When the unemployment rate goes up from u^* to 12%, inflation falls by less than 1pp. But when the unemployment rate drops from u^* to 2%, inflation increases by more than 3pp.

Although inflation varies smoothly with unemployment, the model produces relatively large fluctuations in inflation when the economy is inefficiently tight and relatively small fluctuations in inflation when the economy is inefficiently slack. Conversely, the model produces relatively small fluctuations in unemployment when the economy is inefficiently tight and relatively large fluctuations in unemployment when the economy is inefficiently slack. Here we obtain the result numerically using the calibration in table 11.2. In appendix F, we establish analytically that the Phillips curve $\pi(u)$ is strictly convex under the common assumption that the matching function is Cobb-Douglas and symmetric ($m(u, v) = \mu\sqrt{uv}$).

The near-linearity in the tightness-inflation representation is consistent with strong curvature in the unemployment-inflation representation because unemployment is a nonlinear transform of tightness through the Beveridge mapping $u(\theta)$, so a nearly linear $\pi(\theta)$ can generate a convex $\pi(u)$.

We have seen that price dynamics are driven by the amount of slack in the economy. When the economy is inefficiently slack, sellers are pushed to reduce their prices below what is normal. Conversely, when the economy is inefficiently tight, sellers are pushed to raise their prices above what is normal. What is the intuition for this behavior?

When inflation is above normal, a buyer can reduce its price-adjustment cost by lowering its rate of inflation. Since pricing is optimal, however, there cannot exist any profitable deviation from the current situation. This means that the buyer must incur a commensurate cost when it lowers its rate of inflation. With lower inflation, the price offered by the buyer drops relative to the prices of other buyers. The absence of profitable deviation imposes that the price reduction is costly, so the price must already be below the effective-price-minimizing level. And since submarket tightness is inversely related to the buyer's price, submarket tightness must be above the effective-price-minimizing tightness, which is just the efficient tightness (section 15.4.3). In other words, when inflation is above normal, tightness must be inefficiently high—otherwise profitable deviations would exist for price-setting households.

The same logic holds when inflation is below normal. Then a buyer can reduce its price-adjustment cost by raising its rate of inflation. With higher inflation, the price offered by the buyer rises relative to the prices of other sellers. The absence of profitable deviation imposes that the price increase is costly, so the price must already be above the effective-price-minimizing level. And since submarket tightness is inversely related to the buyer's price, submarket tightness must be below the effective-price-minimizing level, which is the efficient tightness. So when inflation is below normal, tightness must be inefficiently low.

15.9. Solution of the model in normal times

We now solve the slackish business cycle model. We concentrate on normal times, when the nominal interest rate is positive.

By symmetry, the aggregate supply is simply:

$$y^s(\theta) = [1 - u(\theta)]\, kh. \tag{15.21}$$

The aggregate demand is given by (15.19). The tightness that solves the model equalizes aggregate supply and demand:

$$y^d(\theta, r) = y^s(\theta). \tag{15.22}$$

Then the inflation rate is given by the Phillips curve (15.20) and the unemployment rate by (9.8), and so on.

15.10. Implications for the conduct of monetary policy

This section now reviews a few implications of the slackish Phillips curve for the conduct of monetary policy.

15.10.1. Satisfying the dual mandate

Both mandates are achieved by moving the Euler curve along the Phillips curve to arrive at the point where $u = u^*$ and $\pi = \pi^*$. This can be done for instance through monetary policy, which affects the nominal interest rate i and therefore the location of the Euler curve (15.19). The efficient nominal interest rate i^* ensures that inflation is on target ($\pi = \pi^*$) and unemployment is efficient ($u = u^*$). From the Euler curve (15.19), we obtain an expression for the efficient nominal interest rate i^*:

$$1 - u^* = \frac{\delta - i^* + \pi^*}{\sigma \cdot l}$$

so that the efficient nominal interest rate is

$$i^* = \pi^* + \delta - \sigma \cdot (1 - u^*) \cdot l. \tag{15.23}$$

When the nominal interest rate is set to i^*, the model admits a steady-state solution in which the divine coincidence prevails: $(\pi, u) = (\pi^*, u^*)$ satisfies both (15.21) and (15.19) when i is given by (15.23).

Such divine equilibrium exists only when $i^* \geq 0$. If $i^* < 0$, then the divine equilibrium is not a solution of the model since it would violate the zero lower bound constraint that $i \geq 0$. In that case the central bank would resort to setting $i = 0$.

At the zero lower bound, monetary policy can no longer pin down the real interest rate independently of inflation, so the normal-times decomposition of the model no longer applies directly. We therefore focus in this chapter on normal times and defer the full ZLB analysis to future work.

15.10.2. Choice of target variables

In the model, since the full-employment and price-stability mandates of the Fed coincide, aiming for the inflation target or the efficient unemployment rate is completely equivalent. In practice, however, the Fed should target the variable that is the most volatile—to observe more clearly departures from the dual mandate.

The slope of the Phillips curve in turn determines which of inflation and unemployment is the most volatile. If the Phillips curve is steep, aggregate-demand shocks will mostly generate movements in inflation, and targeting inflation will be easier. By contrast, if the Phillips curve is flat, aggregate-demand shocks will mostly generate movements in unemployment, and targeting unemployment or tightness will be easier.

In the model, the Phillips curve appears convex: flat when unemployment is inefficiently high, and steep when unemployment is inefficiently low. The implication is that the Fed should target the efficient unemployment rate u^* when the economy is inefficiently slack; and it should target an inflation rate of 2% when the economy is inefficiently hot.

15.11. Ineffective monetary communication

If monetary communication is ineffective, the divine coincidence breaks down. What we mean by ineffective monetary communication is that the inflation that people consider normal, π^n, is not aligned with the inflation target, π^*.

All the derivations are the same with ineffective monetary communication, except that π^* is replaced by π^n in the Phillips curve (15.20). Equivalently, it is as if a shifting term $\pi^n - \pi^*$ was added to the Phillips curve:

$$\pi = (\pi^n - \pi^*) + \pi^* + \frac{y(\theta)}{\varpi\delta(\delta - r)a} \cdot \left[\frac{\eta(\theta)}{1 - \eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)} - 1\right]. \tag{15.24}$$

Because of the shifter $\pi^n - \pi^* \neq 0$, when the economy is at full employment (and the

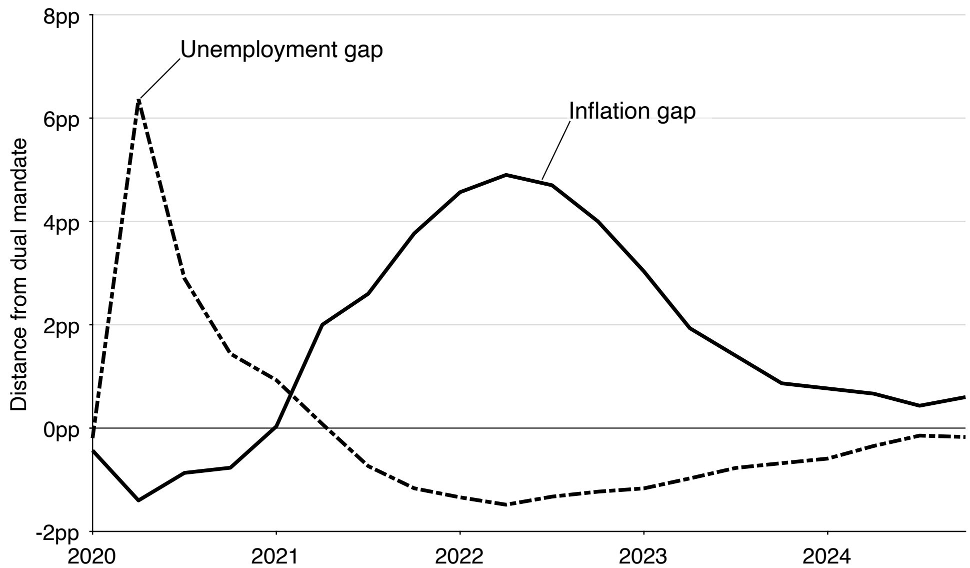

FIGURE 15.2. Phillips curve in the United States, 2020–2024

The unemployment gap comes from figure 13.3. The inflation rate is the quarterly average of the percentage change from a year ago of the monthly seasonally adjusted personal consumption expenditure price index, which is produced by the BEA (2026). The inflation gap is the inflation rate minus the inflation target of 2%.

term in square brackets is 0), inflation is at π^n instead of π^*. So the divine coincidence breaks down: being at full employment does not guarantee price stability.

The two mandates of the central bank are not aligned here, and cannot be achieved at the same time. For instance if people believe that inflation will normally be above target ($\pi^n > \pi^*$), then inflation is too high when the economy is at full employment. Equivalently, when inflation is on target, tightness is too low, and the unemployment rate is inefficiently high.

Hence, in this model, the coincidence is not so much "divine" as the result of careful and effective messaging by the central bank.

15.12. Empirical evidence on the Phillips curve

In this section, we review recent US evidence supporting the slackish Phillips curve. In the modern US economy, the divine coincidence appears to hold, the inflation gap moves negatively with the unemployment gap, and the Phillips curve in the unemployment-inflation plane appears much steeper below the FERU than above it.

15.12.1. Divine coincidence

Figure 15.2 illustrates the parallel trajectories followed by the inflation gap and unemployment gap in the United States during the pandemic and after (2020–2024). Inflation is measured as annual change in the price index for personal consumption expenditures, which is itself produced by the BEA (2026). The inflation gap is the inflation rate minus an inflation target of 2%. These are the inflation measure and inflation target used by the Federal Open Market Committee (FOMC 2026). The unemployment gap is computed in chapter 13 (figure 13.3).[2]

The main observation is that unemployment gap and inflation gap evolve in tandem between 2020 and 2024. Both gaps are close to 0 at the beginning of 2020: the unemployment gap is just −0.2pp while the inflation gap is −0.4pp. So the economy is close to full employment and price stability—close to the point of divine coincidence.

Then, the following quarter, the unemployment gap rises very sharply and the inflation gap drops very sharply. In 2020:Q2, the unemployment gap reaches +6.4pp while the inflation gap bottoms at −1.4pp.

After that both gaps recover. They first return to 0 in 2021:Q1 for inflation and the following quarter for unemployment. Then, the inflation gap keeps rising and the unemployment gap keeps falling. In 2022:Q2, the inflation gap peaks at +4.9pp while the unemployment gap bottoms at −1.5pp.

After reaching these extreme values, both gaps normalize until the end of 2024. The inflation gap shrinks to reach +0.4pp in 2024:Q3. The unemployment gap shrinks to reach −0.1pp in 2024:Q3. So the economy is essentially back at full employment and stable prices at the end of 2024.

This parallel pattern suggests that the divine coincidence prevails in the modern US economy. The inflation gap is 0 whenever the unemployment gap is 0—or the inflation rate is on target whenever the unemployment rate is at the FERU—or prices are stable whenever the economy is at full employment. This coincidence occurred at the beginning of 2020, the beginning of 2021, and the end of 2024.

Moreover, inflation gap and unemployment gap are negatively related: inflation rises above target whenever the unemployment rate is below the FERU; and inflation falls below target whenever the unemployment rate is above the FERU.

The basic evidence of divine coincidence presented in figure 15.2 is buttressed by other work. Benigno and Eggertsson (2023) use aggregate data for inflation and labor

[2] The unemployment gap was computed by focusing on efficiency of the labor market, not the entire economy. This restriction is visible in two ways. The unemployment rate was measured from household data (CPS), and therefore does not include the sort of idle labor in firms described in chapter 2, which would ideally be included to compute aggregate tightness. The vacancy rate was measured from firm data (JOLTS), and therefore does not include the sort of buying effort by households described in chapter 2, which would ideally be included to compute aggregate tightness. Nevertheless, in absence of better data, we use that labor-market unemployment gap in lieu of an aggregate unemployment gap for our analysis.

market tightness in the United States. They find clear evidence of divine coincidence in the 2008–2022 period. When the labor market is efficient, which corresponds to a tightness of 1 (figure 13.1), inflation is on target at 2% (Benigno and Eggertsson 2023, figure 4).

In the 1960s, the Council of Economic Advisors believed in the divine coincidence since they thought that at full employment the economy was not subject to inflationary pressures (Bernanke 2022, p. 18). And over the years the Fed through its FOMC statements often indicated that it believed in the divine coincidence, because it believed that price stability necessarily guaranteed full employment, so that its two mandates were equivalent (Thornton 2012, pp. 119–121, p. 130).

15.12.2. Convexity

Finally, it seems not only that the divine coincidence holds in the United States, but also that the unemployment-inflation Phillips curve appears convex.

The convexity of the Phillips curve can be seen in figure 15.2. In early 2020, the drop in inflation gap (down to –1.4pp) is much more subdued than the coinciding spike in unemployment gap (up to +6.4pp). So the economy was inefficiently slack and most of the movement occurred with the unemployment gap. By contrast, when the economy became inefficiently tight after 2021, most of the movement occurred with the inflation gap. In 2022, the spike in inflation gap (up to +4.9pp) was much larger than the drop in unemployment gap (down to –1.5pp). This indicates that the unemployment-inflation Phillips curve is much flatter when the economy is inefficiently slack than when it is inefficiently tight.

More generally, Babb and Detmeister (2017, table 4) and Smith, Timmermann, and Wright (2025, table 5) find that inflation is much more responsive to unemployment along the US Phillips curve when the unemployment rate is below 4.2% than when the unemployment rate is above 4.2%. In the postwar United States, the efficient unemployment rate is just 4.2% (figure 13.2), so the kink is detected at the efficient unemployment rate. It is true that these studies describe the pattern as a kink, whereas the slackish Phillips curve is smooth, but a convex Phillips curve can look kinked in finite samples.

Benigno and Eggertsson (2023, figure 4) even find that inflation is more responsive to labor market tightness along the US Phillips curve when tightness is inefficiently high (above 1) than when tightness is inefficiently low (below 1). In the numerical illustration of figure 15.1A, the slackish Phillips curve does not display this property: it is essentially linear in tightness. One way to generate the nonlinearity observed by Benigno and Eggertsson is to introduce an asymmetric price-adjustment cost, as Michaillat and Saez (2024) do: when wage cuts are more costly than price hikes, the tightness-inflation Phillips curve is steeper when the economy is inefficiently tight than when it is inefficiently slack.

15.13. The economy after the coronavirus pandemic

In this final section, we use the slackish business cycle model and its Phillips curve to revisit economic fluctuations after the coronavirus pandemic.

15.13.1. Post-pandemic shocks

We consider next an unusual business cycle shock: a shock that shifts the Beveridge curve and then the aggregate supply curve. Such shock involves a change in either the job-separation rate λ or the matching efficacy μ. Both an increase in job separations and a decrease in matching efficacy shift the Beveridge curve outward, which raises the efficient unemployment rate u^*.

A negative supply shock leads to an inward shift of the aggregate supply curve (figure 14.2B). In response to the negative shock, tightness is higher, but the unemployment rate is higher, and inflation is higher. But the key is that the unemployment gap is lower (it has become negative) and inflation is higher. Indeed the efficient unemployment rate has increased more than actual unemployment, so the unemployment rate is now inefficiently low. Such excessive tightness leads to higher inflation.

If the Euler curve remains the same after the outward shift of the Phillips curve—for example because the central bank is not aware of it—then we obtain a burst of inflation after the adverse shock to the Phillips curve. Since the US Beveridge curve has shifted dramatically outward in the aftermath of the coronavirus pandemic (figure 2.7), the flare-up in inflation in 2021–2023 might partly result from this dramatic shift.

The model predicts that inflation rises above target whenever the labor market is inefficiently tight. The model therefore provides an explanation for the flare-up in inflation in 2021–2023. The US labor market was inefficiently tight after the coronavirus pandemic—in fact tighter than at any point since the end of World War 2 (figure 13.1). That excessive tightness must have fueled the flare-up in inflation, all the more so because the Phillips curve $\pi(u)$ is convex: a given tightening of the labor market moves inflation more when unemployment is already low than when unemployment is already high.

15.13.2. Soft landing

After the pandemic, inflation was too high: the inflation gap peaked at 5pp in 2022 (figure 15.2). The Fed wanted to achieve a soft landing: to reduce inflation without raising unemployment.

However, in our model, a strict soft landing was not desirable, since unemployment was too low: the unemployment gap bottomed at -1.5pp in 2022. Moreover, in the model, a soft landing is not possible.

Monetary policy operates through the aggregate demand curve. To reduce inflation, tightness must fall, so monetary policy must raise the real interest rate (figure 14.2A). From this, we can see that strict soft landing is not possible—the unemployment rate always increases after a decrease in tightness.

Even though a strict soft landing isn't possible, unemployment might not increase by much, especially when it is low. This is because the aggregate supply curve is convex in slackish models, as we established in chapter 9. When tightness is high and unemployment is low, the aggregate supply curve is steep so as tightness is reduced, and tightness falls, the unemployment rate doesn't increase by a lot. Thus, it is true that if the central bank cools the economy at a spot where there is little unemployment, a reduction in tightness generates a small increase in unemployment. On the other hand, if we try to cool the economy in the flatter region of the aggregate supply curve, a further reduction in tightness increases the unemployment rate a lot. Thus, a "loose" soft landing is possible with our convex aggregate supply curve: the increase in unemployment might be small.

In our model, we can compute the elasticity of the unemployment rate to tightness along the aggregate supply curve—and relate it to the elasticity of the Beveridge curve, which we have measured (chapter 2). By definition, $\theta = v(u)/u$, where $v(u)$ is the Beveridge curve—which is isomorphic to the aggregate supply curve. Implicitly differentiating in elasticity, we get

$$1 = -\beta \cdot \epsilon^u_\theta - \epsilon^u_\theta,$$

where $\beta = -\epsilon^v_u$ is the elasticity of the Beveridge curve. Hence, the elasticity of the unemployment rate to tightness is directly determined by the Beveridge elasticity:

$$\epsilon^u_\theta = \frac{-1}{1+\beta}.$$

In the United States, $\beta = 1$, so $\epsilon^u_\theta = -1/2$.

Therefore, if the Fed wants to bring labor market tightness down from 2, its value at the peak of the post-pandemic boom (2022:Q2) to 1 (the pre-pandemic, efficient level), the unemployment rate would increase by 50%/2 = 25%, since labor market tightness falls by 50%. Given that the unemployment rate was 3.5% at the peak of the post-pandemic boom (2022:Q3), it would increase to 3.5% × 1.25 = 4.4%. Thus, the effect on unemployment wouldn't be extreme, but it would not be zero either. In fact, the increase predicted by the Beveridge curve matches almost exactly the increase in unemployment observed after the pandemic: in July 2025, labor market tightness reached 0.99, and the unemployment rate was at 4.3% (Michaillat and Saez 2026).

15.14. Summary

This chapter develops a slackish model of the Phillips curve that links inflation to economic slack. Price dynamics are driven by the amount of slack in the economy. When the economy is inefficiently slack, sellers reduce their prices to attract customers. Conversely, when the economy is inefficiently tight, sellers raise their prices. Without price rigidity, the model would be efficient, as in Moen (1997)'s model. But with price rigidity, the model generates a Phillips curve that describes inflation as a function of the unemployment gap.

The slackish Phillips curve ensures divine coincidence: inflation is on target if and only if the unemployment rate is efficient. In the unemployment-inflation plane, the Phillips curve is strictly convex in unemployment (see also appendix F). Inflation is therefore more sensitive to the unemployment gap when the labor market is inefficiently tight than when it is inefficiently slack.

Recent US evidence supports the model's predictions. The divine coincidence holds: inflation rose above target in 2021 as the unemployment rate fell below the FERU, and inflation has declined since its 2022 peak as unemployment has moved back toward the FERU. The unemployment-inflation Phillips curve also appears steeper when the labor market is inefficiently tight than when it is inefficiently slack.

Bibliography

Babb, Nathan R. and Alan K. Detmeister. 2017. "Nonlinearities in the Phillips Curve for the United States: Evidence Using Metropolitan Data." Finance and Economics Discussion Series 2017-070. https://doi.org/10.17016/FEDS.2017.070.

BEA. 2026. "Personal Consumption Expenditures: Chain-type Price Index." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/PCEPI.

Benigno, Pierpaolo and Gauti B. Eggertsson. 2023. "It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve." NBER Working Paper 31197. https://doi.org/10.3386/w31197.

Bernanke, Ben S. 2022. *21st Century Monetary Policy*. New York: W. W. Norton.

Cochrane, John H. 2011. "Determinacy and Identification with Taylor Rules." *Journal of Political Economy* 119 (3): 565–615. https://doi.org/10.1086/660817.

FOMC. 2026. "Statement on Longer-Run Goals and Monetary Policy Strategy." https://perma.cc/7BBR-9UFG.

Michaillat, Pascal and Emmanuel Saez. 2024. "Beveridgean Phillips Curve." arXiv:2401.12475. https://doi.org/10.48550/arXiv.2401.12475.

Michaillat, Pascal and Emmanuel Saez. 2026. "Automated Business Cycle Dashboard." https://pascalmichaillat.org/dashboard/.

Moen, Espen R. 1997. "Competitive Search Equilibrium." *Journal of Political Economy* 105 (2): 385–411. https://doi.org/10.1086/262077.

Okun, Arthur M. 1981. *Prices and Quantities: A Macroeconomic Analysis*. Washington, DC: Brookings Institution.

Rotemberg, Julio J. 1982. "Monopolistic Price Adjustment and Aggregate Output." *Review of Economic Studies* 49 (4): 517–531. https://doi.org/10.2307/2297284.

Smith, Simon, Allan Timmermann, and Jonathan H. Wright. 2025. "Breaks in the Phillips Curve: Evidence from Panel Data." *Journal of Applied Econometrics* 40 (2): 131–148. https://doi.org/10.1002/jae.3102.

Thornton, Daniel L. 2012. "The Dual Mandate: Has the Fed Changed Its Objective?" *Federal Reserve Bank of St. Louis Review* 94 (2): 117–134. https://doi.org/10.20955/r.94.117-134.