Beveridgean Phillips Curve

Pascal Michaillat, Emmanuel Saez

August 2026

This paper develops a Beveridgean model of the Phillips curve. While the New Keynesian Phillips curve is based on monopolistic pricing under price-adjustment costs, the Beveridgean Phillips curve is based on directed-search pricing under price-adjustment costs. Under directed search, prices respond to slack instead of marginal costs. The resulting Phillips curve has three properties that match recent US evidence. First, it delivers the divine coincidence: inflation is on target whenever unemployment is at its efficient level—the full-employment rate of unemployment (FERU), the geometric mean of the unemployment and vacancy rates. Second, it is convex in the unemployment-inflation plane: inflation responds more strongly to unemployment when the labor market is inefficiently tight than when it is inefficiently slack. The convexity comes from the Beveridge curve, which makes unemployment a hyperbolic function of tightness. Third, shifts of the Beveridge curve move the FERU and therefore shift the unemployment-inflation Phillips curve, providing a microfoundation for cost-push shocks. The model also implies that the Fed's dual mandate is internally consistent, that a soft landing from an inefficiently tight labor market is possible, that the NAIRU is not well defined, and that the zero lower bound is not topsy-turvy.

Michaillat: University of California–Santa Cruz. Saez: University of California–Berkeley. We thank Allegra Saggese for helpful comments.

1. Introduction

The post-pandemic spike in inflation in the United States, which coincided with wide labor-market fluctuations, revealed new facts about the Phillips curve. First, the Phillips curve ensures the divine coincidence: inflation is on target whenever the labor market is efficient; inflation rises above target whenever the labor market is inefficiently tight; and inflation falls below target whenever the labor market is inefficiently slack (Benigno and Eggertsson 2023, 2024a; Gitti 2024). Second, inflation is driven by slack rather than marginal costs: labor market tightness—the ratio of job vacancies to unemployed workers—predicts inflation better than the unemployment rate or measures of marginal costs (Barnichon, Oliveira, and Shapiro 2021; Barnichon and Shapiro 2022, 2024). Third, the Phillips curve is not linear: it is much steeper when the labor market is inefficiently tight than when the labor market is inefficiently slack (Babb and Detmeister 2017; Smith, Timmermann, and Wright 2025; Benigno and Eggertsson 2023, 2024a). These empirical findings are summarized by figure 4.

To explain these facts, we propose a Beveridgean model of the Phillips curve. The economy is a matching economy: households purchase and consume services produced by other households, and all trades are mediated by a matching function (Petrongolo and Pissarides 2001)—which generates unemployment, help-wanted ads, and a Beveridge curve. To generate price dynamics, we introduce price competition through directed search (Moen 1997): each household posts help-wanted ads that offer a specific price for services, and job seekers choose which ads to pursue, trading off the price offered against the time required to find a job. To ensure that slack fluctuates, we introduce price rigidity through quadratic price-adjustment costs (Rotemberg 1982). In the model, directed search produces a Phillips curve linking labor-market slack to inflation, and the price-adjustment costs ensure that the Phillips curve is nonvertical.

The Phillips curve guarantees the divine coincidence: inflation is on target if and only if the labor market operates efficiently. The coincidence follows from two ingredients. The first is the efficiency of directed search (Moen 1997): absent price-adjustment costs, price-setters would drive labor market tightness exactly to its efficient level, because at that level the resources lost to unemployment and recruiting are minimized. With price-adjustment costs, price-setters deviate from the efficient tightness only insofar as inflation deviates from the inflation that people find normal. The second ingredient is that the central bank anchors normal inflation at its inflation target—an assumption that is realistic for the modern United States, where long-run inflation expectations have been fixed slightly above 2% for three decades (Hazell et al. 2022, figure 2). Combining the two ingredients, inflation is on target exactly when tightness is efficient—that is, exactly when unemployment is at its efficient level, the full-employment rate of unemployment (FERU).

In the unemployment-inflation plane, the Phillips curve is convex: steep when un-

employment is inefficiently low and flat when unemployment is inefficiently high. The reason is simple. Inflation is roughly linear in labor market tightness, but unemployment is a hyperbolic function of tightness through the Beveridge curve, so inflation is convex in unemployment. We prove this analytically when the matching function is Cobb-Douglas and symmetric, and we verify it numerically in the calibrated model. The convexity also explains why empirical studies detect a kink in the Phillips curve at an unemployment rate of 4.2% (Babb and Detmeister 2017; Smith, Timmermann, and Wright 2025)—the average FERU in the postwar United States—since a convex curve looks kinked in finite samples. Existing models instead generate a kink from constraints on wage movements: wages that are downwardly rigid (Chiarella et al. 2003), or rigid only when the labor market is slack (Benigno and Eggertsson 2023; Gitti 2024).

The model also explains why the unemployment-inflation Phillips curve is unstable while its tightness formulation is not. Whenever the Beveridge curve shifts, the FERU moves, so the point of divine coincidence travels in the unemployment-inflation plane and the Phillips curve shifts. Beveridge shocks therefore microfound the cost-push shocks of the New Keynesian literature. The US Beveridge curve shifted dramatically outward when the pandemic hit, pushing the FERU from 4.0% to 6.7% in one quarter; the FERU then slowly came back down. The efficient tightness, by contrast, is pinned down by two stable sufficient statistics—the Beveridge elasticity and the recruiting cost—so it barely moves when the Beveridge curve shifts. This is why the tightness-inflation Phillips curve performed so much better than the unemployment-inflation Phillips curve during the pandemic (Barnichon, Oliveira, and Shapiro 2021; Barnichon and Shapiro 2022, 2024).

The divine coincidence has several implications for monetary policy. The Fed's full-employment and price-stability mandates coincide: by maintaining the economy at full employment, the Fed automatically keeps inflation on target. The interest rate that achieves both mandates is the full-employment rate of interest (FERI). A soft landing from an inefficiently tight labor market is also possible. Because unemployment is a hyperbolic function of tightness, cooling the labor market from its 2022 peak (tightness of 2) to efficiency (tightness of 1) raises unemployment only from 3.5% to 4.4%—almost exactly what materialized by 2025. The convexity of the Phillips curve further suggests a state-contingent choice of targets: the Fed should target the FERU when the economy is inefficiently slack and the inflation target when the economy is inefficiently tight. Finally, the NAIRU is not well defined in the model—inflation is stable at any unemployment rate—and the zero lower bound is not the topsy-turvy world of the New Keynesian model (Eggertsson and Krugman 2012): at the zero lower bound the model admits a unique solution that behaves like a normal slump, even if the bound is permanent.

2. Slack and full employment in the United States

Before developing the model, we present how we measure slack, full employment, and the unemployment gap in the United States. The framework builds on Michaillat and Saez (2021a, 2024b).

2.1. Unemployment, vacancies, and tightness

We measure labor market slack from three series covering 1929–2024: the unemployment rate, the vacancy rate, and their ratio, labor market tightness.

The unemployment rate u is the number of job seekers as a share of the labor force. Between 1948 and 2024, it is computed from the Current Population Survey (BLS 2025f,c). Between 1929 and 1947, it comes from the historical series constructed by Petrosky-Nadeau and Zhang (2021). The unemployment rate is sharply countercyclical. Over 1929–2024, it averages 6.4%; over the more stable 1948–2019 period, it averages 5.7%.

The vacancy rate v is the number of job openings as a share of the labor force. Between 2001 and 2024, it is computed from the Job Openings and Labor Turnover Survey (BLS 2025d,c). Between 1951 and 2000, it comes from the vacancy proxy constructed by Barnichon (2010) from the Conference Board's help-wanted advertising index. Between 1929 and 1950, it comes from the vacancy proxy constructed by Petrosky-Nadeau and Zhang (2021) from MetLife's help-wanted index. The vacancy rate is sharply procyclical. Over 1929–2024, it averages 3.2%.

Labor market tightness is the ratio of the two rates: $\theta = v/u$. It is the number of vacant jobs per job seeker. Tightness is a natural measure of slack in any matching model: it determines how quickly job seekers find jobs and how quickly employers fill vacancies.

Because unemployment is countercyclical and vacancies are procyclical, the two rates are negatively related: they trace out a Beveridge curve (Beveridge 1944). In the United States, the Beveridge curve is close to a rectangular hyperbola, $uv = A$. When the log vacancy rate is regressed on the log unemployment rate, after removing a slow-moving trend from both series to absorb the decennial shifts of the curve, the estimated elasticity over 1948–2019 is exactly −1.00, with an R^2 of 0.87 (figure 1). The estimates of the Beveridge elasticity on each stable branch of the curve tell the same story: they remain between 0.84 and 1.02 throughout the postwar period (Michaillat and Saez 2021a, figure 6).

2.2. Efficient tightness and the FERU

Slack is inefficient when it is excessive, but some slack is unavoidable—and even desirable. Reducing unemployment requires posting more vacancies, and vacancies absorb labor: recruiting takes work—designing and advertising job vacancies, screening and interviewing candidates, negotiating contracts. In the United States, servicing one vacancy requires

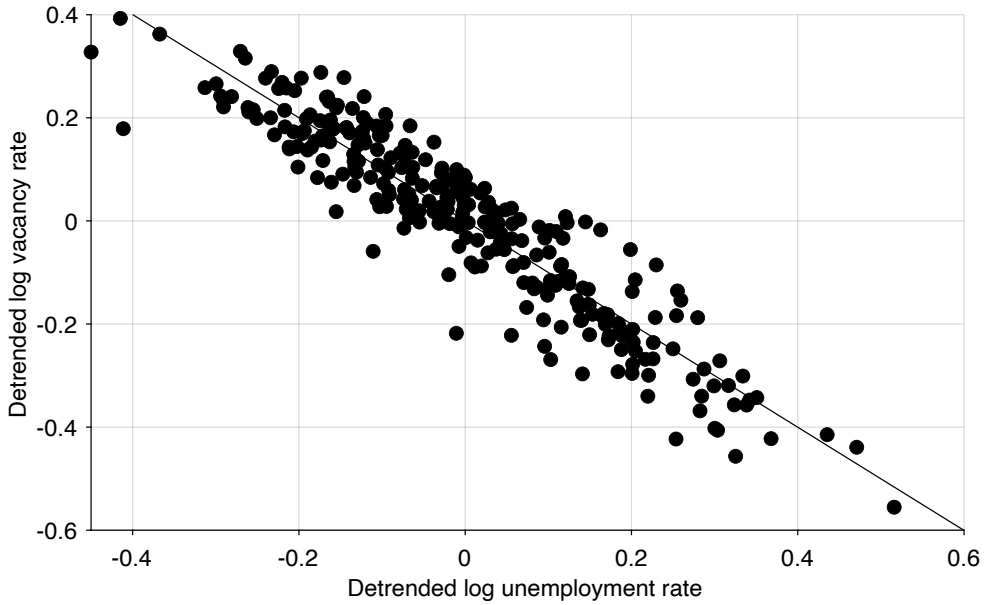


FIGURE 1. The US Beveridge curve is a rectangular hyperbola

The unemployment rate comes from the sources listed in section 2. The figure plots the logarithm of the vacancy rate against the logarithm of the unemployment rate, both detrended using a HP filter with smoothing parameter 10,000. The data are quarterly, 1948–2019. The regression line has a slope of −1.00, with $R^2 = 0.87$, indicating that the US Beveridge curve is isoelastic with elasticity $\beta = 1$: a rectangular hyperbola.

about one full-time worker: estimates from the National Employer Survey imply 0.92 workers per vacancy (Michaillat and Saez 2021a, p. 11), and estimates from a survey by Bersin & Associates imply 1.16 workers per vacancy (Michaillat and Saez 2024b, p. 332). Hence the number of vacancies effectively measures the number of recruiters, and the recruiting cost is $\kappa = 1$ worker per vacancy.

The number of producers in the economy—workers who are neither unemployed nor recruiting—is therefore $(1 - u - \kappa v)h$, where h is the labor force. Efficiency requires choosing the point on the Beveridge curve $v(u)$ that maximizes the number of producers, which means minimizing $u + \kappa v(u)$. The solution takes a simple sufficient-statistic form (Michaillat and Saez 2021a):

$$\theta^* = \frac{1}{\kappa\beta}, \tag{1}$$

where $\beta = -d\ln(v)/d\ln(u)$ is the elasticity of the Beveridge curve. Section 4 microfounds this condition in the model. With $\kappa = 1$ and $\beta = 1$ in US data, the efficient tightness is

$$\theta^* = 1.$$

The US labor market is inefficiently slack whenever tightness is below 1, inefficiently tight

whenever tightness is above 1, and efficient—at full employment—when there are exactly as many vacant jobs as job seekers.[1]

The unemployment rate that prevails at the efficient tightness is the full-employment rate of unemployment (FERU), denoted u^*. Because the Beveridge curve is a rectangular hyperbola, $uv = A$, and full employment requires $v^* = u^*$, the FERU satisfies $(u^*)^2 = A = uv$, so the FERU is the geometric mean of the unemployment and vacancy rates (Michaillat and Saez 2024b):

$$u^* = \sqrt{uv}. \tag{2}$$

The FERU is strictly positive: full employment is not zero unemployment. Zero unemployment is infeasible—the Beveridge curve bounds unemployment away from zero—and undesirable—pushing unemployment toward zero would divert a vast amount of labor toward recruiting.

2.3. A century of slack

Figure 2 displays US labor market tightness over 1929–2024, along with the full-employment level of 1. Over the period, tightness averages 0.73 and generally remains below 1: the US labor market is typically inefficiently slack, especially in recessions. Tightness plunged to 0.03 during the Great Depression and climbed to 6.8 at the end of World War 2. The postwar period is quieter: tightness peaked at 1.60 in 1953:Q1, during the Korean War, and bottomed at 0.16 in 2009:Q3, during the Great Recession.

The labor market is not always inefficiently slack, however. Before 2018, it had been inefficiently tight only during major wars: World War 2, the Korean War, and the Vietnam War. Keynes (1936, p. 322) doubted that a labor market could reach full employment in peacetime, and for eighty years he was right. But since 2018, the labor market has been inefficiently tight twice in peacetime: just before the coronavirus pandemic (2018:Q3–2020:Q1) and in its aftermath (2021:Q3–2024:Q4). In 2022:Q2, tightness reached 1.98, a value last seen in 1945. The post-pandemic period is therefore a rarity: the only episode of peacetime overheating on record.[2]

[1] This criterion echoes Beveridge (1944, p. 18), who posited that full employment means "having always more vacant jobs than unemployed men, not slightly fewer jobs." The criterion is also known in policy circles: asked in 2022 which tightness the Fed might target, its chair responded, "I think when we got to one-to-one in the, you know, in the late teens, we thought that was a pretty good number" (Powell 2022, pp. 12–13). Our criterion is more symmetric than Beveridge's, however: more vacancies than job seekers signals inefficient tightness, just as more job seekers than vacancies signals inefficient slack. The logic is also different: we do not interpret vacancies as unmet labor demand but as a measure of the labor devoted to recruiting.

[2] People at the time were aware of the overheating. Processing about 20,000 news articles per month from major US newspapers, Caldara, Iacoviello, and Yu (2026, figure 2) construct an index of shortages in the United States from 1900 to 2024. The episodes of labor shortage that they detect exactly overlap with the episodes of inefficient tightness in figure 2: World War 2, the Korean War, the Vietnam War, and the post-pandemic period.

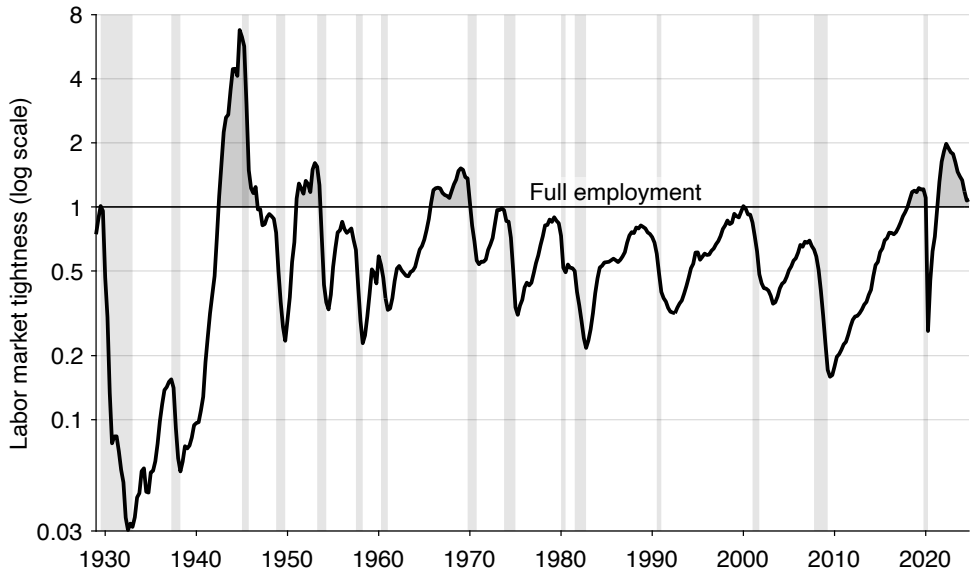
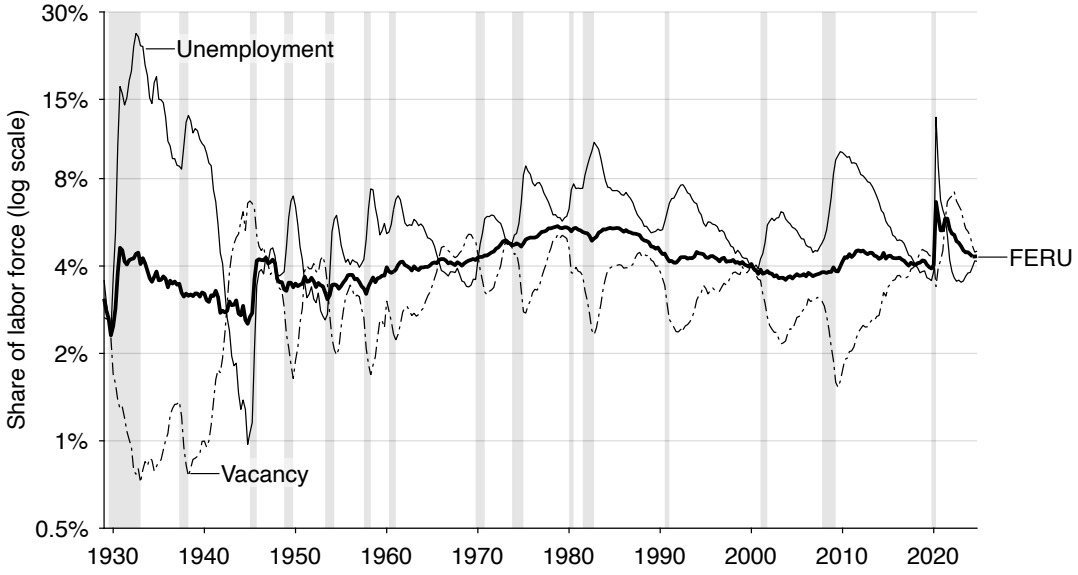


FIGURE 2. Labor market tightness in the United States, 1929–2024

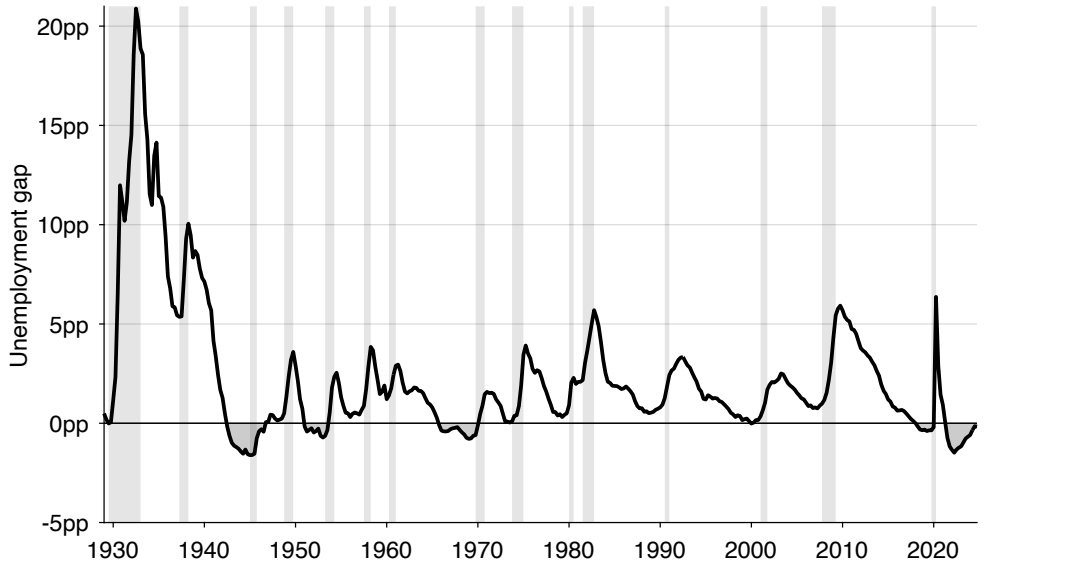
Labor market tightness is $\theta = v/u$, where the unemployment and vacancy rates come from the sources listed in section 2. Shaded areas indicate recessions dated by the NBER (2023). The labor market is at full employment when tightness equals 1, inefficiently slack when tightness is below 1, and inefficiently tight when tightness is above 1.

Figure 3A displays the FERU computed from (2). Over 1929–2024, the FERU averages 4.1%, and it is quite stable: it remains between 2.3% and 6.7% over almost a century. During the postwar period, 1948–2019, the FERU averages 4.2% and remains between 3.1% and 5.5%. The Beveridge curve shifts in and out over the decades, but the shifts are usually too small to move the FERU much. The exceptions are the three major catastrophes of the century, which shifted the Beveridge curve sharply outward and raised the FERU abruptly: the Great Depression (FERU up 2.3pp between 1929:Q4 and 1930:Q4), the demobilization after World War 2 (up 1.7pp between 1944:Q4 and 1945:Q4), and above all the coronavirus pandemic. When the US economy shut down in the spring of 2020, the FERU jumped from 4.0% in 2020:Q1 to 6.7% in 2020:Q2—an increase of 2.7pp in one quarter. The FERU then slowly came back down, falling below 5% in 2022 and below 4.5% in 2023.

Figure 3B displays the unemployment gap, $u - u^*$: the distance of the labor market from full employment. Over 1929–2024, the gap averages +2.3pp, confirming that the US labor market is generally inefficiently slack. The gap reached +20.9pp in 1932:Q3, at the depth of the Great Depression, and −1.6pp in 1945:Q1, at the peak of the war effort. In the postwar period, the gap peaked at +5.9pp in 2009:Q4, during the Great Recession. During the pandemic, the labor market experienced sharp departures from full employment in both directions. The gap first spiked to +6.4pp in 2020:Q2, its highest level since 1945. The



A. FERU



B. Unemployment gap

FIGURE 3. FERU and unemployment gap in the United States, 1929–2024

The unemployment and vacancy rates come from the sources listed in section 2. The FERU is $u^* = \sqrt{uv}$, so on a logarithmic scale it is the midpoint between the unemployment and vacancy rates. The unemployment gap is $u - u^*$. Shaded areas indicate recessions dated by the NBER (2023).

labor market then recovered quickly, reached full employment in the middle of 2021, and overshot: the gap fell to −1.5pp in 2022:Q2, its lowest level since 1945, before shrinking back to −0.2pp in 2024:Q4.

3. Three facts about the Phillips curve

With the measurement framework in place, we turn to three facts about the Phillips curve in the United States.

3.1. First fact: divine coincidence

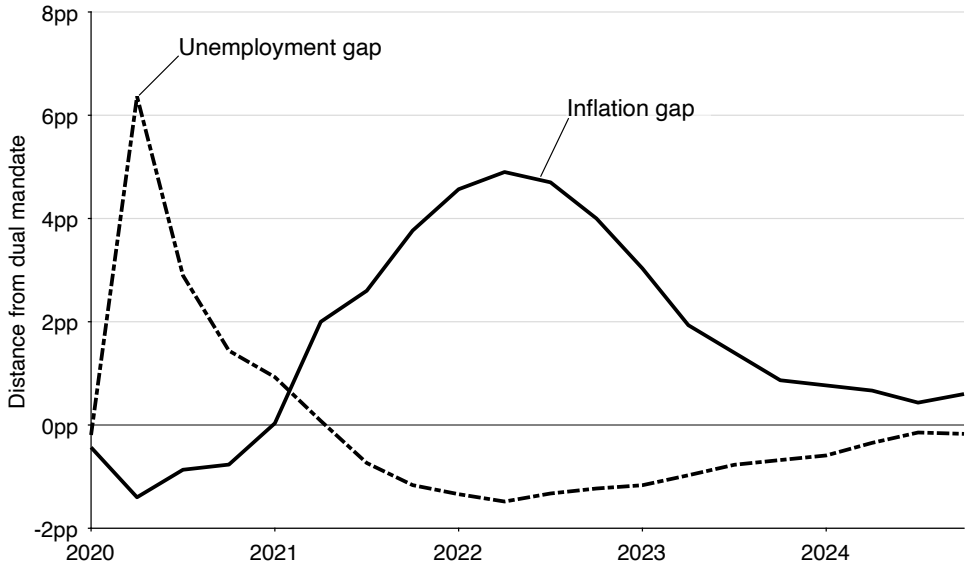
The first fact is that inflation is on target whenever the labor market is at full employment—the divine coincidence.

Figure 4A illustrates the parallel trajectories followed by the inflation gap and the unemployment gap during the pandemic and after (2020–2024). Inflation is measured as the annual change in the price index for personal consumption expenditures (BLS 2024); the inflation gap is inflation minus the 2% target (Board of Governors of the Federal Reserve System 2026). The unemployment gap comes from figure 3B.

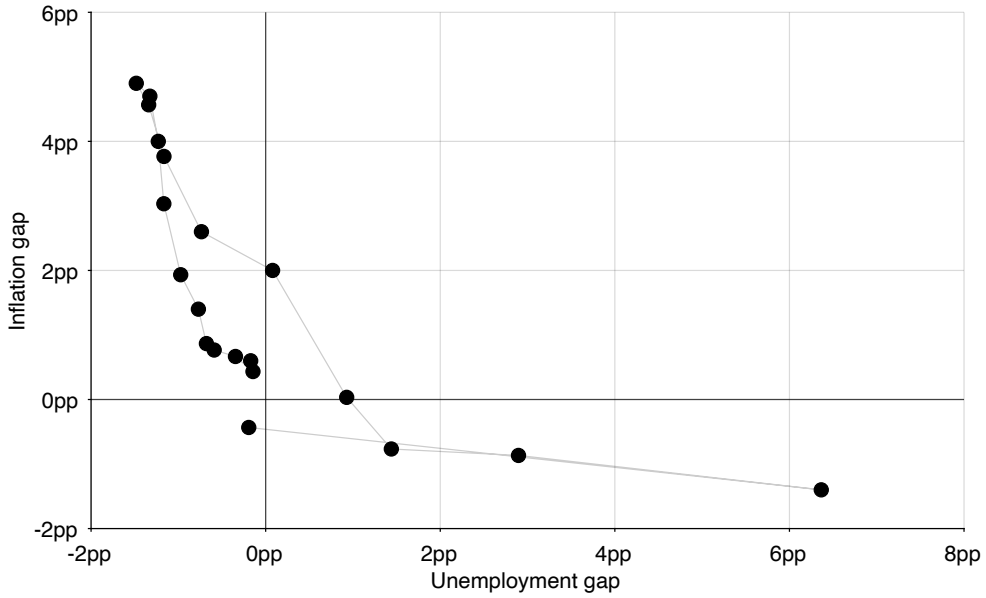
The two gaps evolve in tandem. Both are close to 0 at the beginning of 2020: the unemployment gap is −0.2pp and the inflation gap is −0.4pp—the economy is close to the point of divine coincidence. The following quarter, the unemployment gap spikes to +6.4pp while the inflation gap bottoms at −1.4pp. Both gaps then recover, returning to 0 in the first half of 2021. The inflation gap keeps rising and the unemployment gap keeps falling until 2022:Q2, when the inflation gap peaks at +4.9pp and the unemployment gap bottoms at −1.5pp. After that, both gaps normalize: by 2024:Q3, the inflation gap is down to +0.4pp and the unemployment gap up to −0.1pp. The economy is back at full employment and price stability.

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. The coincidence occurred at the beginning of 2020, in the first half of 2021, and at the end of 2024. The two gaps are also negatively related: inflation is above target whenever unemployment is below the FERU, and below target whenever unemployment is above the FERU.

This basic evidence is supported by other work. Using aggregate US data over 2008–2022, Benigno and Eggertsson (2023, figure 4) find that when labor market tightness is 1, inflation is on target at 2%. Using metropolitan-level US data over 2001–2022, Gitti (2024, figure 1) finds the same pattern, in noisier form: when local tightness is 1, local inflation is around 2%. Before the pandemic, Carney (2017) argued that the divine coincidence held in the United States. The idea has a long policy pedigree: in the 1960s, the Council of Economic Advisers believed that at full employment the economy was not subject to



A. Time series



B. Scatter plot

FIGURE 4. Phillips curve in the United States, 2020–2024

The unemployment gap comes from figure 3B. 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 BLS (2024). The inflation gap is the inflation rate minus the inflation target of 2%. These are the inflation measure and inflation target used by the Federal Open Market Committee (Board of Governors of the Federal Reserve System 2026).

inflationary pressures (Bernanke 2022, p. 18); and over the years the Fed indicated through its FOMC statements that it saw its two mandates as equivalent (Todd 2012, pp. 119–121, p. 130).

3.2. Second fact: slack, not marginal costs, drives inflation

The second fact is that inflation is driven by labor market slack, and that slack is best measured by labor market tightness.

Barnichon, Oliveira, and Shapiro (2021) and Barnichon and Shapiro (2022, 2024) show that a Phillips curve formulated with labor market tightness performs better than the traditional Phillips curve formulated with the unemployment rate: incorporating tightness improves inflation forecasts relative to unemployment alone, and relative to standard marginal-cost proxies. The superiority of the tightness formulation is especially visible in the pandemic period, when the Beveridge curve was unstable. Relatedly, Ball, Leigh, and Mishra (2022) argue that the post-pandemic spike in inflation was so severe because labor market tightness was so elevated.

The tight connection between slack and inflation is visible in figure 4: between 2020 and 2024, the inflation gap tracks the unemployment gap—and therefore labor market tightness—closely at every phase of the cycle: the collapse in 2020, the overheating of 2021–2022, and the normalization of 2023–2024.

3.3. Third fact: the Phillips curve is convex

The third fact is that the Phillips curve is much steeper when the labor market is inefficiently tight than when it is inefficiently slack.

The convexity appears clearly in the scatter plot of figure 4B. In early 2020, the economy was inefficiently slack, and most of the adjustment occurred through unemployment: the unemployment gap spiked to +6.4pp while the inflation gap fell only to −1.4pp. Once the economy became inefficiently tight after 2021, most of the adjustment occurred through inflation: the inflation gap spiked to +4.9pp while the unemployment gap fell only to −1.5pp. The Phillips curve is thus flat to the right of the point of divine coincidence and steep to the left of it.

More systematic evidence points to the same pattern. 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 it is above 4.2%. Similar nonlinearity is found in metropolitan US data (Babb and Detmeister 2017; Smith, Timmermann, and Wright 2025; Gitti 2024) and in international data (Smith, Timmermann, and Wright 2025; Benigno and Eggertsson 2024b). The kink is detected at an unemployment rate of 4.2%—exactly the average FERU

in the postwar United States (figure 3A). These studies describe the pattern as a kink, whereas the Phillips curve developed in this paper is smooth but strictly convex; the two descriptions are hard to distinguish in finite samples, because a convex Phillips curve looks kinked.[3]

4. Beveridgean model of the Phillips curve

We now develop the model.[4] In a New Keynesian model, price dynamics are driven by marginal costs (see for instance Bilbiie and Kanzig 2024, equation (1)). Here, by contrast, price dynamics are driven by the amount of slack, in the tradition of Phillips (1958) and of the early matching models of Hansen (1970) and Mortensen (1970). When the economy is inefficiently slack, buyers of services are pushed to lower the prices they pay. When the economy is inefficiently tight, they are pushed to raise them.

The economy is a matching economy composed of a measure one of large households, indexed by $j \in [0, 1]$. Households purchase and consume services produced by workers from other households; they cannot consume the services of their own workers.[5] All trades are mediated by a matching function, which produces slack. Because the model has a single market that blends the product and labor markets, slack can equivalently be interpreted as unemployment or idle capacity; we adopt the labor-market interpretation throughout.

4.1. Matching market for services

In each household j, $h_j > 0$ workers are part of the labor force; the aggregate labor force is $h = \int_0^1 h_j \, dj$. Each employed worker produces $k > 0$ services per unit time, so the economy's productive capacity is kh services per unit time.

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 produces k services per unit time. She remains so until they separate, which occurs at rate $\lambda > 0$.

To recruit workers, each household j posts V_j help-wanted ads. The ads specify the price P_j that the household plans to pay for the services purchased. Each help-wanted ad

[3] Benigno and Eggertsson (2023, figure 4) also find that inflation is more responsive to tightness when tightness is above 1 than when it is below 1. The model developed here does not produce this property: inflation is essentially linear in tightness (section 7).

[4] A textbook development of the underlying slackish business cycle model appears in Michaillat (2026). The Phillips-curve chapter of that book was motivated by an earlier version of the present paper (Michaillat and Saez 2024a).

[5] This is in the spirit of Diamond (1982), who developed a model of the product market in which self-employed workers must trade because they cannot consume their own production. Modeling the US economy as a service economy is an accurate simplification: in 2024, 86% of US workers were employed in service-providing industries (BLS 2025a,b).

is serviced by $\kappa > 0$ recruiters per unit time; the recruiters are employees of household j devoted to recruiting other workers.

Job search is directed: workers choose which household to apply to. In that way, each household j constitutes a submarket. Each job seeker decides which submarket to join based on the price P_j offered for services and the submarket tightness $\theta_j = V_j/U_j$, where U_j is the number of unemployed workers vying for jobs in household j. The tightness matters because it determines how quickly workers find a job. Job seekers therefore trade off how much they are paid on the job with how long it takes to find the job.

In each submarket j, a matching function $m(U_j, V_j)$ determines the flow of new matches from the number of job seekers and help-wanted ads. The matching function has constant returns to scale and is increasing in both arguments (Petrongolo and Pissarides 2001). Submarket tightness determines the job-finding rate, $f(\theta_j) = m(1, \theta_j)$, which is increasing in θ_j, and the ad-filling rate, $q(\theta_j) = f(\theta_j)/\theta_j$, which is decreasing in θ_j. We denote by $\eta(\theta) \in (0, 1)$ the elasticity of the matching function with respect to unemployment, so the elasticity of f is $1 - \eta(\theta)$ and the elasticity of q is $-\eta(\theta)$.[6]

4.2. Balanced flows: unemployment, vacancies, and matching wedge

The unemployment rate in submarket j is $u_j = U_j/h_j$, where h_j is the labor force attached to submarket j. We assume that labor flows in each submarket are balanced: the number of workers who find a job with household j at any point in time, $f(\theta_j)U_j$, equals the number who leave their jobs, $\lambda(1 - u_j)h_j$. Balanced flows imply that the unemployment rate is a function of submarket tightness, $u_j = u(\theta_j)$, with

$$u(\theta) = \frac{\lambda}{\lambda + f(\theta)}. \tag{3}$$

The unemployment rate is strictly decreasing in tightness: when the market is tighter, workers find jobs faster, so fewer workers are unemployed at any point in time. The ad rate in submarket j is $v_j = V_j/h_j$; with balanced flows it is given by $v(\theta) = \theta u(\theta)$. As tightness moves, the pair $[u(\theta), v(\theta)]$ traces out the Beveridge curve of the model.

Balanced flows also determine how much labor is devoted to recruiting. Household j's l_j employees produce $y_j = kl_j$ services, but the V_j help-wanted ads absorb κ recruiters each, so a total of $k\kappa V_j$ services are devoted to recruiting instead of consumption. Household j's consumption is therefore $c_j = k\left[l_j - \kappa V_j\right]$, which can be written

$$c_j = kh_j\left[1 - u_j - \kappa v_j\right]. \tag{4}$$

[6] Throughout the paper, ϵ_y^x denotes the elasticity of x with respect to y: $\epsilon_y^x = d\ln(x)/d\ln(y)$. With a Cobb-Douglas matching function $m(U, V) = \mu U^\sigma V^{1-\sigma}$, which we use in the quantitative work, the matching elasticity is constant: $\eta(\theta) = \sigma$.

To maintain l_j employees, the household must fill λl_j ads per unit time, which requires $V_j = \lambda l_j / q(\theta_j)$ ads. Hence output and consumption are related by a matching wedge:

$$y_j = \left[1 + \tau(\theta_j)\right] c_j, \qquad \tau(\theta) = \frac{\kappa\lambda}{q(\theta) - \kappa\lambda}. \tag{5}$$

The matching wedge $\tau(\theta)$ is positive and strictly increasing in tightness: in a tighter market, ads take longer to fill, so more labor must be devoted to recruiting for each unit of consumption. The wedge diverges when tightness reaches the upper bound $\bar{\theta}$, defined by $q(\bar{\theta}) = \kappa\lambda$; at that point all employed labor is devoted to recruiting. Combining (4) and (5), consumption satisfies

$$c_j = \frac{1 - u(\theta_j)}{1 + \tau(\theta_j)} k h_j, \qquad \frac{1 - u(\theta_j)}{1 + \tau(\theta_j)} = 1 - u(\theta_j) - \kappa v(\theta_j). \tag{6}$$

Two elasticities will be used repeatedly. First, from (3), the elasticity of the unemployment rate with respect to tightness is

$$\epsilon_\theta^u = -\frac{f(\theta)}{\lambda + f(\theta)} \left[1 - \eta(\theta)\right] = -\left[1 - \eta(\theta)\right] \left[1 - u(\theta)\right], \tag{7}$$

which also implies $\epsilon_\theta^{1-u} = \left[1 - \eta(\theta)\right] u(\theta)$. Second, from (5), $1 + \tau = q/(q - \kappa\lambda)$, so the elasticity of $1 + \tau(\theta)$ with respect to tightness is

$$\epsilon_\theta^{1+\tau} = \epsilon_\theta^q \left[1 - (1 + \tau)\right] = \eta(\theta)\tau(\theta). \tag{8}$$

4.3. Efficient allocation

Consider a social planner who chooses tightness to maximize consumption of services—the only source of welfare in this economy.[7] From (6), aggregate consumption is $c(\theta) = kh\left[1 - u(\theta)\right] / \left[1 + \tau(\theta)\right]$, so

$$\epsilon_\theta^c = \epsilon_\theta^{1-u} - \epsilon_\theta^{1+\tau} = \left[1 - \eta(\theta)\right] u(\theta) - \eta(\theta)\tau(\theta).$$

The elasticity is positive at low tightness and negative at high tightness, so consumption is maximized at the efficient tightness θ^*, where the two terms exactly offset:

$$\frac{\tau(\theta^*)}{u(\theta^*)} = \frac{1 - \eta(\theta^*)}{\eta(\theta^*)}. \tag{9}$$

[7] Unemployed workers produce nothing in the model. This is consistent with the evidence: the value of job seekers' home production, net of the psychological cost of idleness, appears negligible (Michaillat and Saez 2021a, p. 11).

The efficiency condition balances the two costs of slack: at θ^*, the labor lost to unemployment and the labor lost to recruiting are optimally traded off. When the economy is inefficiently slack ($\theta < \theta^*$), the wedge-to-unemployment ratio τ/u is below its efficient level $(1-\eta)/\eta$; when the economy is inefficiently tight ($\theta > \theta^*$), the ratio is above it.

The efficiency condition (9) is the model counterpart of the sufficient-statistic formula (1). Indeed, using (3) and (5), condition (9) can be rearranged as $\kappa\theta\left[\lambda + \eta(\theta) f(\theta)\right] = \left[1 - \eta(\theta)\right] f(\theta)$, which is exactly $\theta^* \beta(\theta^*)\kappa = 1$, where

$$\beta(\theta) = \frac{1}{1-\eta(\theta)}\left[\eta(\theta) + \frac{u(\theta)}{1-u(\theta)}\right] \tag{10}$$

is the elasticity of the model's Beveridge curve. In the model, the Beveridge curve is not exactly a rectangular hyperbola—its elasticity varies slightly with tightness—so the efficient tightness is not exactly 1.[8] The FERU of the model is the unemployment rate at the efficient tightness: $u^* = u(\theta^*)$.

4.4. Directed search and the price-tightness tradeoff

Households observe the prices and tightnesses in all submarkets and send their job seekers to the submarkets that suit them best. This arbitrage links prices and tightnesses across submarkets—and this link is the basis for the Phillips curve.

When workers are unemployed, they receive no income; when they are employed, they produce k services per unit time, each paid a price P_j. In submarket j, a fraction $1 - u(\theta_j)$ of the attached workers are employed at any point in time, so the expected income per worker in submarket j is $kP_j\left[1 - u(\theta_j)\right]$. Households send their workers to the submarkets offering the highest expected income per worker. In their role as employers, households are aware of this arbitrage, so they offer prices that allow them to compete with other households. Arbitrage therefore imposes that $P_j\left[1 - u(\theta_j)\right]$ is the same across all submarkets:

$$P_j\left[1 - u(\theta_j)\right] = P\left[1 - u(\theta)\right], \tag{11}$$

where P is the price level and θ is the aggregate tightness—the ratio of the aggregate number of help-wanted ads to the aggregate number of job seekers. If a household offered a lower expected income than the rest, it would simply attract no applicants.

The arbitrage condition (11) determines the tightness faced by a household as a func-

[8] With a penalized Cobb-Douglas matching function, the model's Beveridge curve is an exact rectangular hyperbola and the efficient tightness is exactly 1, at the cost of somewhat more cumbersome algebra (Michaillat and Saez 2024a). In the calibrated model with a standard Cobb-Douglas matching function, $\beta(\theta^*) = 0.98$ and $\theta^* = 1.03$: the difference is immaterial.

tion of the relative price it posts, $p_j = P_j/P$:

$$\theta_j(p_j) = u^{-1}\left(1 - \frac{1-u(\theta)}{p_j}\right). \tag{12}$$

Since u^{-1} is strictly decreasing, submarket tightness is decreasing in the relative price: a household that posts a generous price attracts many job seekers, so it faces a slack submarket and recruits easily; a household that posts a stingy price attracts few job seekers, so it faces a tight submarket and recruits slowly. In that way, households have monopsony power, and they face a downward-sloping demand-like curve in the price-tightness plane. Taking logs in (11) and differentiating, $1 + \epsilon_\theta^{1-u} \cdot \epsilon_p^\theta = 0$, so the elasticity of the curve is

$$\epsilon_p^\theta = \frac{-1}{\left[1-\eta(\theta_j)\right] u(\theta_j)}, \tag{13}$$

where we used $\epsilon_\theta^{1-u} = [1-\eta]\, u$ from section 4.2.

The tradeoff faced by household j is then clear. All the services it purchases are paid the price P_j, so its expenditure is $P_j y_j = \left[1+\tau(\theta_j)\right] P_j c_j$. The effective price of consumed services is not P_j but $\left[1+\tau(\theta_j)\right] P_j$: it includes the cost of the services consumed and the cost of the services devoted to recruiting. A low posted price P_j economizes on the price per service but imposes a high tightness θ_j and thus a high matching wedge $\tau(\theta_j)$.

4.5. Efficiency without price-adjustment costs

As a benchmark, consider a household that is free to set any price. It minimizes the effective price of services, $\left[1+\tau(\theta_j)\right] P_j$, subject to the arbitrage condition (11). Using the arbitrage condition, the effective price can be written

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

Hence, by choosing its price, the household effectively chooses submarket tightness θ_j to maximize $\left[1-u(\theta_j)\right] / \left[1+\tau(\theta_j)\right]$—which, by (6), is the same as maximizing $1-u_j-\kappa v_j$. This is exactly the social planner's problem of section 4.3, so the household selects the efficient tightness: $\theta_j = \theta^*$, where θ^* satisfies (9). Here we have just recovered the efficiency result of Moen (1997): directed search produces efficient outcomes, because pricing households internalize the planner's tradeoff between unemployment and recruiting.

4.6. Price rigidity

Generally, however, the labor market is not efficient: as figure 3B shows, the US unemployment gap is large and countercyclical. To generate inefficient fluctuations, prices must be somewhat rigid.

Because search is directed, prices are chosen optimally by buyers—they are not set by norms—so we introduce rigidity through a cost of changing prices. As in Rotemberg (1982), household j incurs a quadratic cost when inflation in its submarket, $\pi_j(t) = \dot{P}_j(t)/P_j(t)$, departs from normal inflation π^n—the inflation that people expect and consider fair. The flow disutility is

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

where the parameter $\varpi > 0$ governs the price-adjustment cost.

What does this cost capture? When prices grow less than normal, the workers of household j feel shortchanged: the price P_j directly determines their pay, $P_j k$, and workers' morale—and performance—dips when their pay does not grow as expected (Bewley 1999, 2005). The quadratic cost reflects the cost of restoring morale, or the cost of diminished performance.[9] Conversely, when prices grow more than normal, household j is unhappy to pay more for the same services: higher-than-normal inflation upsets customers, who feel unfairly treated (Shiller 1997). The quadratic cost captures this displeasure.[10]

5. Households' problem

We now present and solve the optimization problem of households. Households have two decisions to make: how much to save, and how to set the prices of the services they purchase. From the optimal saving decision we will derive the aggregate demand curve; from the optimal pricing decision, the Phillips curve.

5.1. Utility function

Household j cares about its consumption of services and its relative wealth, and it incurs the price-adjustment cost. The household's nominal wealth $B_j(t)$ is held in government bonds; its relative real wealth is $\Big[B_j(t) - B(t)\Big]/P(t)$, where $B(t) = \int_0^1 B_i(t)\,di$ is average

[9] 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.... 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."

[10] Eyster, Madarasz, and Michaillat (2021) provide a complete model of why price increases anger customers, and how sellers internalize such anger.

nominal wealth. Household j maximizes the discounted sum of flow utilities,

$$\int_0^\infty \exp(-\delta t)\left[c_j(t)^{1-\alpha} + \frac{1}{a}\cdot\frac{B_j(t)-B(t)}{P(t)} - \frac{\varpi}{2a}\left(\pi_j(t)-\pi^n\right)^2\right]dt, \tag{14}$$

where $\delta > 0$ is the time discount rate, $\alpha \in (0,1)$ governs the concavity of utility over consumption, and $a > 0$ indicates the taste for consumption relative to the other components of utility.

The unusual ingredient is the presence of relative wealth in the utility function, as in the business cycle models by Michaillat and Saez (2021b, 2022). People save not only to smooth consumption but also because they enjoy the social status conferred by wealth—an observation running from Keynes (1919), Fisher (1930), and Lewis (1955) to modern neuroscientific and survey evidence (Camerer, Loewenstein, and Prelec 2005; Coy 2021).[11] Beyond realism, wealth in the utility function makes aggregate demand nondegenerate: it allows the model to accommodate interest-rate pegs, permanent zero-lower-bound episodes, and other policies that would produce anomalies in the New Keynesian model (Michaillat and Saez 2021b).

5.2. Budget constraint

Household j's nominal budget constraint is

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

where $i(t)$ is the nominal interest rate on bonds, y_{ij} is the quantity of services sold by household j's workers in submarket i, y_j is the quantity of services purchased by household j, and $T(t)$ is a lump-sum tax that finances the interest payments on government bonds.

The two flow terms simplify. Since a fraction $1-u(\theta_i)$ of the workers sent to submarket i are employed, and the arbitrage condition (11) equalizes expected income across submarkets, the household's labor income is $\int_0^1 P_i y_{ij}\,di = P\left[1-u(\theta)\right]kh_j$. And given the matching wedge (5) and the price-tightness link (12), the household's expenditure is $P_j y_j = P_j\left[1+\tau\left(\theta_j\left(P_j/P\right)\right)\right]c_j$. Translating the budget constraint into real terms, with real wealth $b_j = B_j/P$ and real interest rate $r(t) = i(t) - \pi(t)$, where $\pi = \dot{P}/P$ is aggregate inflation:

$$\dot{b}_j(t) = r(t)b_j(t) + \left[1-u(\theta(t))\right]kh_j - p_j(t)\left[1+\tau\left(\theta_j\left(p_j(t)\right)\right)\right]c_j(t) - \frac{T(t)}{P(t)}. \tag{15}$$

Besides the budget constraint, household j faces the law of motion of its relative price

[11] In a survey of 2,000 older Americans analyzed by Coy (2021), more than half of respondents do not plan to spend down their wealth in retirement, and about 65% agree that "saving as much as I can makes me happy and fulfilled."

$p_j = P_j/P$:

$$\dot{p}_j(t) = \left[\pi_j(t) - \pi(t)\right] p_j(t). \tag{16}$$

5.3. First-order conditions

Household j chooses paths for consumption $c_j(t)$ and inflation $\pi_j(t)$ to maximize (14) subject to (15) and (16). We solve this optimal-control problem with the Hamiltonian

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

where the control variables are consumption c_j and inflation π_j; the state variables are real wealth b_j and relative price p_j; and the costate variables are ψ_j^b and ψ_j^p.

The first-order condition with respect to consumption, $\partial\mathcal{H}/\partial c_j = 0$, gives

$$(1-\alpha)c_j^{-\alpha} = \psi_j^b\left[1 + \tau(\theta_j)\right] p_j. \tag{17}$$

The first-order condition with respect to inflation, $\partial\mathcal{H}/\partial\pi_j = 0$, gives

$$\frac{\varpi}{a}\left[\pi_j - \pi^n\right] = \psi_j^p p_j. \tag{18}$$

The first-order condition with respect to real wealth, $\partial\mathcal{H}/\partial b_j = \delta\psi_j^b - \dot{\psi}_j^b$, gives

$$\dot{\psi}_j^b = \left[\delta - r\right]\psi_j^b - \frac{1}{a}. \tag{19}$$

The final first-order condition, with respect to the relative price, is $\partial\mathcal{H}/\partial p_j = \delta\psi_j^p - \dot{\psi}_j^p$. Expanding the derivative of the expenditure term, the key object is

$$p_j\frac{\tau'(\theta_j)}{1+\tau(\theta_j)}\theta_j'(p_j) = \epsilon_\theta^{1+\tau}\cdot\epsilon_p^\theta = \eta(\theta_j)\tau(\theta_j)\cdot\frac{-1}{\left[1-\eta(\theta_j)\right]u(\theta_j)} = -\frac{\eta(\theta_j)}{1-\eta(\theta_j)}\cdot\frac{\tau(\theta_j)}{u(\theta_j)},$$

where we combined the elasticities (8) and (13). Noting also that $\left[1+\tau(\theta_j)\right]c_j = y_j$, the condition simplifies to

$$\dot{\psi}_j^p = \left[\delta + \pi - \pi_j\right]\psi_j^p + \psi_j^b y_j\left[1 - \frac{\eta(\theta_j)}{1-\eta(\theta_j)}\cdot\frac{\tau(\theta_j)}{u(\theta_j)}\right]. \tag{20}$$

The term in square brackets is central to what follows: by the efficiency condition (9),

it is zero exactly when submarket tightness is efficient, positive when the submarket is inefficiently slack, and negative when it is inefficiently tight.

5.4. Symmetric solution

We focus on a symmetric solution, in which all households behave the same and all submarkets are the same: all submarket tightnesses equal the aggregate tightness θ, all relative prices equal 1, and all submarket inflation rates equal aggregate inflation π. Dropping the index j, the first-order conditions become

$$(1-\alpha)c(t)^{-\alpha} = \psi^b(t)\left[1+\tau(\theta(t))\right] \tag{21}$$

$$\frac{\varpi}{a}\left[\pi(t)-\pi^n\right] = \psi^p(t) \tag{22}$$

$$\dot{\psi}^b(t) = \left[\delta - r(t)\right]\psi^b(t) - \frac{1}{a} \tag{23}$$

$$\dot{\psi}^p(t) = \delta\psi^p(t) + \psi^b(t)y(t)\left[1 - \frac{\eta(\theta(t))}{1-\eta(\theta(t))}\cdot\frac{\tau(\theta(t))}{u(\theta(t))}\right]. \tag{24}$$

6. Solution of the model in normal times

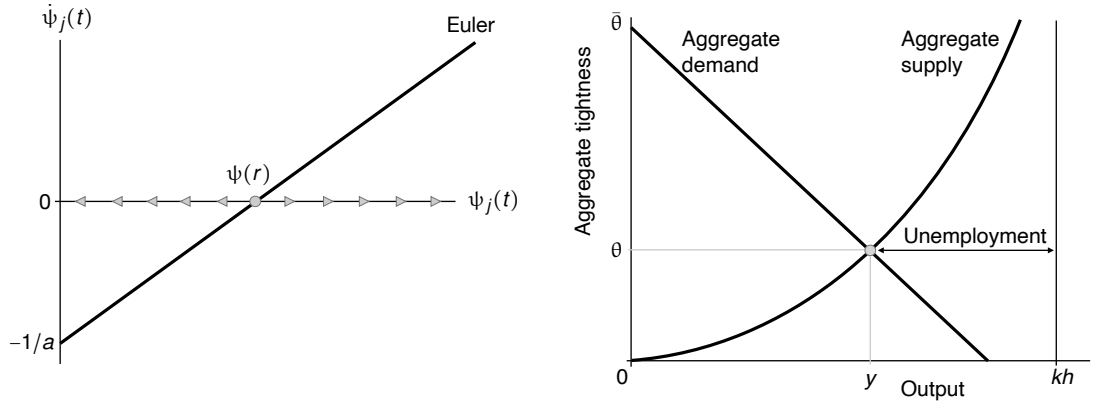
We now solve the model in normal times, when the nominal interest rate is positive. Section 10 covers the zero lower bound.

6.1. Monetary policy

Monetary policy has two dimensions in the model: setting interest rates and communicating about future inflation.

First, the central bank maintains a fixed real interest rate $r < \delta$, by setting the nominal interest rate to $i(t) = r+\pi(t)$ at every point in time. There is no Taylor rule: the central bank adjusts the nominal rate one-for-one with inflation to keep fixed the price that matters for aggregate demand—the real interest rate. The Taylor rule appears in the New Keynesian model mostly for convenience, to ensure the uniqueness of the solution (Cochrane 2011); here the solution is unique without it, so we dispense with it.

Second, through communication and management of expectations, the central bank anchors normal inflation at its inflation target: $\pi^n = \pi^*$, with $\pi^* = 2\%$ in the United States (Board of Governors of the Federal Reserve System 2026). This assumption is central to the divine coincidence, and it is realistic: since 1995, the long-run inflation forecast elicited in the Survey of Professional Forecasters has been exceedingly stable, ever so slightly above 2% (Hazell et al. 2022, figure 2). Section 9.5 explores what happens when the anchoring fails.



A. Phase line of the Euler equation B. Aggregate demand and supply

FIGURE 5. Solution of the model in normal times

The Euler equation is given by (23); its equilibrium point is $\psi(r) = 1/[(\delta - r)a]$. The aggregate supply curve is given by (27). The aggregate demand curve is given by (26). The intersection of the two curves determines aggregate tightness and output; unemployment is read off the distance between output and capacity kh.

6.2. Aggregate demand

The aggregate demand is determined by the two first-order conditions describing consumption and saving: (21) and (23). Because the central bank fixes the real interest rate, this block is isolated from inflation: inflation matters for aggregate demand only through the real rate, which is fixed.

The differential equation (23) is the Euler equation of the model. It is an autonomous linear differential equation in ψ^b; since $\delta - r > 0$, it is a source, with equilibrium point

$$\psi^b = \frac{1}{(\delta - r)a}. \tag{25}$$

The phase line of the equation is plotted in figure 5A. The costate variable ψ^b is not predetermined, so it can jump at any time. If it jumped above the equilibrium point, it would diverge upward, so consumption would fall to zero, violating feasibility; if it jumped below, consumption would diverge, eventually violating the resource constraint. Hence the costate jumps at time 0 to the equilibrium point (25) and stays there.[12] Plugging (25) into (21) gives consumption, and multiplying by $1 + \tau(\theta)$ gives the aggregate demand:

$$y^d(\theta, r) = \frac{[(1-\alpha)(\delta - r)a]^{1/\alpha}}{[1 + \tau(\theta)]^{1/\alpha - 1}}. \tag{26}$$

The aggregate demand gives the quantity of services demanded by households when they

[12]The equilibrium value (25) has a simple interpretation: households keep consumption constant only if the returns on wealth—financial returns $r\psi^b$ plus hedonic returns $1/a$—exactly offset time discounting $\delta\psi^b$.

save optimally. It is decreasing in tightness—because a tighter economy imposes a larger matching wedge, making consumption effectively more expensive—and decreasing in the real interest rate.

6.3. Aggregate supply and equilibrium tightness

By symmetry, the aggregate supply is the quantity of services provided when flows are balanced:

$$y^s(\theta) = [1 - u(\theta)]\, kh. \tag{27}$$

The aggregate supply is increasing in tightness—in a tighter economy, fewer workers are unemployed—and it is sharply convex in a tightness-output plane: flat when tightness is low, steep when tightness is high (figure 5B).

The tightness that solves the model equalizes aggregate demand and supply:

$$y^d(\theta, r) = y^s(\theta). \tag{28}$$

Since demand is decreasing in tightness and supply is increasing, the solution is unique. It is illustrated in figure 5B. Equation (28) implicitly defines tightness as a function of the real interest rate, $\theta(r)$; from tightness, all quantities follow: output $y = y^s(\theta)$, unemployment $u(\theta)$, consumption $c(\theta)$.

6.4. Phillips curve

Inflation is determined by the two remaining first-order conditions, which describe optimal pricing: (22) and (24).

Equation (22) says that inflation is a linear function of the costate variable on the relative price: $\pi(t) = \pi^* + a\psi^p(t)/\varpi$. The costate variable follows the differential equation (24), which is linear and a source since $\delta > 0$. If the costate did not jump to the equilibrium point of the equation, it would explode, dragging inflation to $\pm\infty$ and generating infinitely negative utility. So the costate jumps to the equilibrium point:

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

where we used $\psi^b = 1/[(\delta - r)a]$. Accordingly, inflation jumps at time 0 to

$$\pi = \pi^* + \frac{y(\theta)}{\varpi\delta(\delta - r)}\left[\frac{\eta(\theta)}{1-\eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)} - 1\right]. \tag{29}$$

Equation (29) is the Phillips curve of the model. It links inflation to aggregate tightness,

which is itself determined by aggregate demand and supply. The curve has three main properties.

First, the divine coincidence holds. The term in square brackets measures the productive inefficiency of the economy: by the efficiency condition (9), it is zero exactly when tightness is efficient. So inflation is on target ($\pi = \pi^*$) if and only if the economy is at full employment ($\theta = \theta^*$, $u = u^*$). When the economy is inefficiently tight, the term is positive and inflation is above target; when the economy is inefficiently slack, the term is negative and inflation is below target. Thanks to the coincidence, the central bank faces no tradeoff between inflation and unemployment: by bringing unemployment to the FERU, it mechanically brings inflation to target. The coincidence arises because people take the inflation target as normal inflation; section 9.5 shows how the coincidence breaks when they do not.

Second, the slope of the Phillips curve is governed by the price-adjustment cost ϖ. When prices are very costly to change, ϖ is large and inflation responds little to departures from efficiency. When prices are easy to change, inflation responds forcefully. In the limit $\varpi \to \infty$, inflation is fixed at π^*; in the limit $\varpi \to 0$, the Phillips curve is vertical at θ^* and the economy is always efficient, as in section 4.5.

Third, inflation is strictly increasing in tightness at full employment and above. Output $y(\theta)$, the matching elasticity $\eta(\theta)$, and the matching wedge $\tau(\theta)$ are increasing in tightness, while the unemployment rate $u(\theta)$ is decreasing, so above full employment the bracket is positive and growing while output is positive and growing: inflation grows with tightness. Around full employment the bracket is zero, so the slope of the Phillips curve is determined by the derivative of the bracket, which is strictly positive. Below full employment, two forces oppose: output falls with slack while the (negative) bracket shrinks in magnitude. Near full employment the bracket dominates and inflation still increases with tightness; at very low tightness the output term can dominate and inflation may locally decrease with tightness. In the calibrated model, this non-monotonicity only arises at tightness below 0.08—far outside the observed range—so in practice inflation is increasing in tightness.

What is the intuition for the Phillips curve? When inflation is above normal, a buyer could reduce its price-adjustment cost by lowering its rate of inflation. Since pricing is optimal, there cannot exist any profitable deviation, so lowering inflation must be costly for another reason: with lower inflation, the buyer's posted price would drop relative to other buyers' prices, and this price reduction must hurt. That is only possible if the price is already below the effective-price-minimizing level—which means, since submarket tightness is inversely related to the posted price, that tightness is above the efficient level (section 4.5). Hence, when inflation is above normal, the economy must be inefficiently tight. Symmetrically, when inflation is below normal, the economy must be inefficiently

slack. Price dynamics are thus driven by slack: when workers are scarce, buyers must post higher prices to attract them; when workers are abundant, buyers can lower the prices they pay.[13]

7. Quantitative Phillips curve

To assess the shape of the Phillips curve, we calibrate the model to the US economy and plot the curve—first in the tightness-inflation plane, then in the unemployment-inflation plane.

7.1. Calibration

The calibration is summarized in table 1. As much as possible, we use US evidence from the 1948–2019 period, during which the labor market operated in a consistent fashion; we exclude the disruptions caused by the Great Depression, World War 2, and the coronavirus pandemic.

Several parameters are normalized or directly measured. Labor productivity, labor force, and the taste for consumption are normalized to $k = h = a = 1$ at the baseline. The recruiting cost is set to $\kappa = 1$ worker per vacancy, as measured by the recruiting surveys discussed in section 2.2. The job-separation rate is set to $\lambda = 3.6\%$ per month, the 2001–2019 average of the total-separation rate measured in JOLTS (BLS 2025e). The real interest rate is set to $r = 0.5\%$ per year, the US natural rate estimated by Williams (2017, table 1). The discount rate is set to $\delta = 91\%$ per year, the truncated mean of the annual discount rates estimated in the 54 experimental studies surveyed by Frederick, Loewenstein, and O'Donoghue (2002, table 1) and Andersen et al. (2014, table 3).[14]

The remaining parameters are calibrated to match key US facts. The matching function is Cobb-Douglas, $m(U, V) = \mu U^{\sigma} V^{1-\sigma}$, so the matching elasticity is $\eta(\theta) = \sigma$. We set σ to generate a Beveridge elasticity of 1, as measured in figure 1. Inverting (10) at the average unemployment rate $u = 5.7\%$ gives $\sigma = 0.47$. We then set the matching efficacy μ so that the unemployment rate is at its 1948–2019 average of 5.7% when tightness is at its 1948–2019 average of 0.65; equation (3) gives $\mu = 0.75$ per month. The concavity of consumption utility is set to $\alpha = 0.35$; this is the value that generates a labor share of 65% in the isomorphic production version of the model (Autor et al. 2020; Michaillat and Saez 2024a). Under this calibration, the efficient tightness is $\theta^* = 1.03$ and the FERU is $u^* = u(\theta^*) = 4.5\%$.

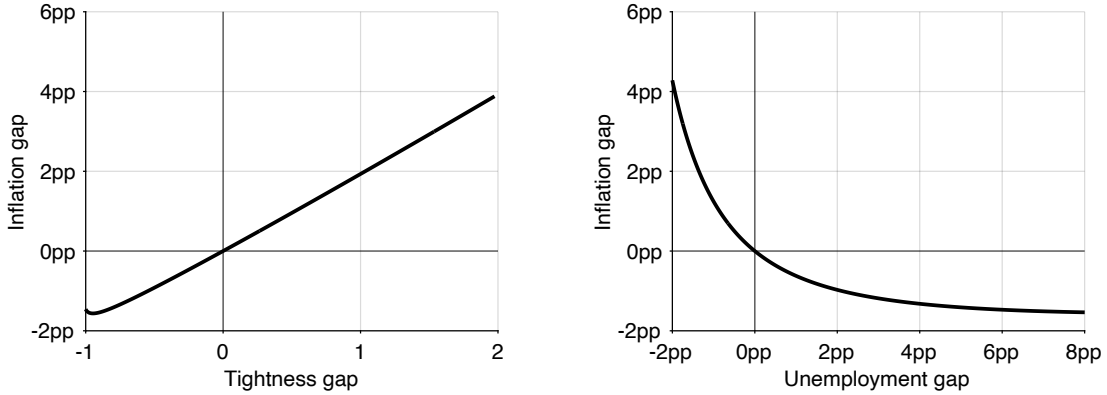
[13] We have assumed that buyers set prices. The results are exactly the same if sellers set the prices of their services instead, as Michaillat and Saez (2024a) show: workers raise their prices when they are in high demand and lower them when demand slumps.

[14] With wealth in the utility function, the discount rate is not tied to the average real interest rate, so it can be calibrated from experimental evidence—which consistently points to double-digit discount rates (Cohen et al. 2020). Such impatience may not be surprising considering that annual interest rates on US credit cards are typically around 20% (Rossman 2026), with penalty rates near 30% (Irby 2026).

TABLE 1. Calibrated values of the parameters and variables of the model

Value	Description	Source or target
	A. Normalized parameters	
$k = 1$	Labor productivity	Normalization
$h = 1$	Labor force	Normalization
$a = 1$	Taste for consumption	Normalization (baseline)
	B. Parameters calibrated from direct measurement	
$\kappa = 1$	Recruiting cost	Michaillat and Saez (2021a, 2024b)
$\lambda = 3.6\%$	Monthly job-separation rate	BLS (2025e)
$\delta = 91\%$	Annual time discount rate	Frederick, Loewenstein, and O'Donoghue (2002); Andersen et al. (2014)
$r = 0.5\%$	Annual real interest rate	Williams (2017)
	C. Parameters calibrated to match target	
$\alpha = 0.35$	Concavity of consumption utility	Labor share = 65% (Autor et al. 2020)
$\sigma = 0.47$	Matching elasticity	Beveridge elasticity = 1 (figure 1)
$\mu = 0.75$	Monthly matching efficacy	Unemployment rate = 5.7%
$\varpi = 58$	Price-adjustment cost	Repricing cost = 1.22% of revenue (Zbaracki et al. 2004)
	D. Variables calibrated from historical evidence	
$\theta = 0.65$	Labor market tightness	Average over 1948–2019
$u = 5.7\%$	Unemployment rate	Average over 1948–2019
	E. Implied variables	
$v = 3.7\%$	Vacancy rate	$v = \theta u$
$f(\theta) = 0.60$	Monthly job-finding rate	$f(\theta) = \mu\theta^{1-\sigma}$
$q(\theta) = 0.92$	Monthly ad-filling rate	$q(\theta) = \mu\theta^{-\sigma}$
$\tau(\theta) = 4.1\%$	Matching wedge	Equation (5)
$\theta^* = 1.03$	Efficient tightness	Equation (9)
$u^* = 4.5\%$	FERU	$u^* = u(\theta^*)$

The model is monthly. The discount rate and real interest rate are annual; they are converted to monthly values in the simulations.



A. Tightness-inflation curve B. Unemployment-inflation curve

FIGURE 6. Phillips curves in the calibrated model

The Phillips curve is given by equation (29). Tightness and unemployment are related by equation (3). The tightness gap is $\theta - \theta^*$; the unemployment gap is $u - u^*$; the inflation gap is $\pi - \pi^*$. The model is calibrated as in table 1.

The final parameter is the price-adjustment cost ϖ, which determines the slope of the Phillips curve. We calibrate it from microevidence collected by Zbaracki et al. (2004). Using time-and-motion methods, they study the pricing process of a large industrial firm and find that the physical, managerial, and customer costs of changing prices amount to 1.22% of the firm's revenue in a given year. That year, the firm changed the price of 25% of its products, and most prices changed by about 4%. Based on this evidence, we set ϖ so that if inflation is 4% on 25% of the services purchased by a household, the cost amounts to 1.22% of the household's purchases in utility terms:

$$\frac{\varpi}{2a} \times 0.04^2 \times 0.25 = 0.0122 \times y.$$

At the average tightness of 0.65, $y = 0.943$ and $a = 1$, so $\varpi = 58$.

7.2. Inflation and tightness

Figure 6A plots the Phillips curve (29) in the tightness-inflation plane, in gap form: inflation gap $\pi - \pi^*$ against tightness gap $\theta - \theta^*$. The curve goes through the point of divine coincidence: the inflation gap is zero when the tightness gap is zero. Inflation is increasing in tightness over the entire relevant range (any tightness above 0.08), and the curve is broadly linear in tightness over a wide range (tightness between 0.08 and 3): the nonmonotonicity discussed in section 6.4 is visible only at extremely low tightness and is irrelevant in practice.

7.3. Inflation and unemployment: convexity

The picture changes once the Phillips curve is represented in the traditional unemployment-inflation plane (figure 6B). The unemployment rate is a strictly decreasing function of tightness, given by (3), so tightness is a strictly decreasing function of the unemployment rate, $\theta(u)$, and the figure plots $\pi(\theta(u))$.

As expected, inflation is decreasing in unemployment. But inflation is now a strongly convex function of unemployment. The curve crosses the point of divine coincidence—inflation of $\pi^* = 2\%$ at unemployment of $u^* = 4.5\%$—and it is much steeper when unemployment is inefficiently low than when it is inefficiently high. When the unemployment gap rises to +8pp, the inflation gap falls by just 1.5pp: inflation falls to 0.5%. But when the unemployment gap drops to −2pp, the inflation gap rises to +4.2pp: inflation climbs to 6.2%.[15]

The convexity is not specific to the calibration: it can be established analytically under the common assumption that the matching function is Cobb-Douglas and symmetric.

PROPOSITION 1. *Assume that the matching function is Cobb-Douglas and symmetric:* $m(u, v) = \mu\sqrt{uv}$. *Let* $\bar{u} = \kappa\lambda^2/(\mu^2 + \kappa\lambda^2)$ *be the unemployment rate at the upper tightness bound* $\bar{\theta}$. *Then the Phillips curve* $\pi(u)$ *given by* (29) *is strictly decreasing in u around the FERU* u^* *and strictly convex in u on* $(\bar{u}, 1)$.

The proof is in appendix A. The logic is simple. With a symmetric Cobb-Douglas matching function, the ratio of matching wedge to unemployment admits the closed form $\tau(u)/u = (1 - \bar{u})/(u - \bar{u}) - 1/u$, which is strictly decreasing and strictly convex in u—a hyperbola in u. The Phillips curve inherits this shape: inflation explodes as unemployment approaches its feasible minimum $\bar{u}$ and flattens out as unemployment grows.

The convexity means that the model produces relatively large fluctuations in inflation and small fluctuations in unemployment when the economy is inefficiently tight—and the reverse when the economy is inefficiently slack. This is the pattern visible in the post-pandemic data (figure 4B): the empirical Phillips curve drawn by the US economy in 2020–2024 is quite close to the Phillips curve simulated from the model (figure 6B). For the same fluctuations in the unemployment gap, the inflation gap varies by comparable amounts in the data and in the simulations. The convexity also explains why econometric studies detect a kink at the postwar FERU of 4.2% (section 3.3): a convex curve estimated with two linear branches yields a steep branch below the FERU and a flat branch above it.

[15] The shape of the simulated Phillips curve in figure 6B is uncannily similar to the shape of the original curve estimated by Phillips (1958, figure 1) from UK data for 1861–1913.

7.4. Nonlinearity of the model

The Phillips curve is almost linear in tightness yet strongly convex in unemployment. How can the same curve have such different curvature in the two representations? The reconciliation is the Beveridge mapping $u(\theta)$: unemployment is a hyperbolic function of tightness, so a nearly linear $\pi(\theta)$ generates a convex $\pi(u)$.

The near-linearity in tightness does not mean that the model is linear. The aggregate supply curve (27) is sharply nonlinear in the tightness-output plane: aggregate demand shocks have small effects on tightness in a slack economy—they affect mostly output—but large effects on tightness in a tight economy. Since inflation is linear in tightness, inflation behaves in the same way: it responds little to demand shocks when the economy is slack and much more when the economy is tight. Equivalently, in the unemployment representation: the supply curve is linear in unemployment, $y^s = (1-u)kh$, so unemployment responds similarly to demand shocks in booms and slumps, but inflation responds sharply to unemployment when the economy is tight—because the Phillips curve is convex—and barely when the economy is slack. Whether the nonlinearity is assigned to the supply curve or to the Phillips curve is a matter of representation; the economics is the same.

8. Response to business-cycle shocks

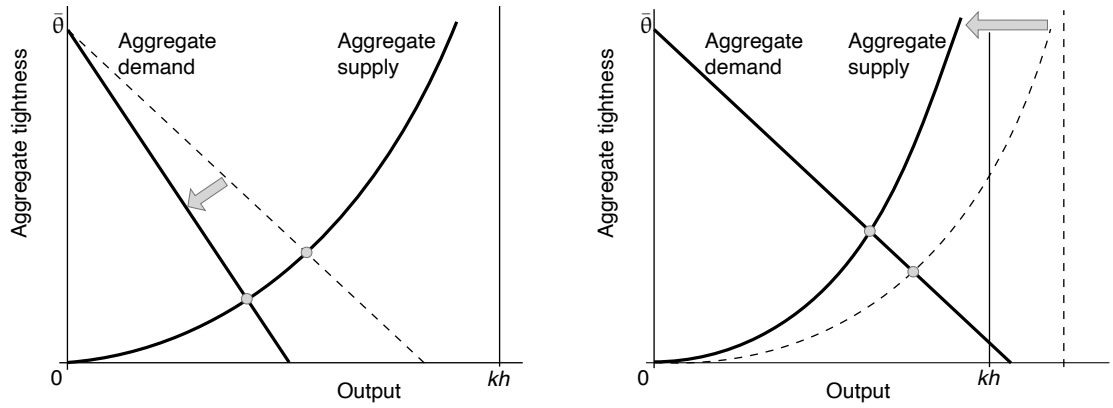
We now use the model to study business-cycle shocks. Because all jump variables move instantly to their equilibrium values, the response to unexpected permanent shocks is obtained by comparative statics: the economy jumps from the old intersection of aggregate demand and supply to the new one, and inflation jumps along the Phillips curve.

8.1. Aggregate demand shocks

The most natural negative aggregate demand shock is a decrease in the taste for consumption a: households become thriftier, desiring to save more and consume less.[16] The shock rotates the aggregate demand curve inward (figure 7A). Tightness and output fall; unemployment rises. Since the efficient tightness (9) does not depend on a, the FERU is unchanged, so the unemployment gap widens one-for-one with unemployment. Along the Phillips curve, the drop in tightness lowers inflation: the economy slides down the curve, producing the standard negative comovement of unemployment and inflation.

The effects of demand shocks are sharply state-dependent. When the economy is slack, the aggregate supply curve is flat, so demand shocks move mostly output and unemployment, and barely inflation. When the economy is tight, the supply curve is steep, so demand shocks move mostly tightness and inflation, and barely unemployment. In

[16] A decrease in the time discount rate δ operates in the same way.



A. Negative aggregate demand shock B. Negative aggregate supply shock

FIGURE 7. Response of the model to aggregate demand and supply shocks

The aggregate supply curve is given by equation (27). The aggregate demand curve is given by equation (26). A negative aggregate demand shock is a decrease in the taste for consumption a; it rotates the aggregate demand curve inward. A negative aggregate supply shock is a decrease in productive capacity kh; it shifts the aggregate supply curve inward.

the unemployment-inflation plane, this is just the convexity of the Phillips curve at work: demand shocks generate large unemployment movements and small inflation movements in slumps, and the reverse in booms. Both phases of the pandemic episode fit this pattern: the collapse of demand in 2020, hitting a slack economy, moved unemployment enormously and inflation little; the surge of demand in 2021–2022, hitting a tight economy, moved inflation enormously and unemployment little (figure 4B).

8.2. Aggregate supply shocks

A negative aggregate supply shock is a decrease in productive capacity kh—caused by a decrease in labor productivity k or in the labor force h. The shock shifts the aggregate supply curve inward (figure 7B). Output falls but tightness rises: people want to consume a similar amount of services while fewer services are offered, so the economy gets tighter and unemployment falls. Along the Phillips curve, the rise in tightness raises inflation. Supply shocks are therefore stagflationary in the model: they lower output while raising inflation. And because the efficient tightness does not depend on k or h, the drop in unemployment translates into a negative unemployment gap: the economy becomes inefficiently tight, which is why inflation rises.

8.3. Beveridge shocks: microfounded cost-push shocks

The third type of shock is specific to the Beveridgean model: shocks that shift the Beveridge curve. In the model, the position of the Beveridge curve is determined by the

job-separation rate λ and the matching efficacy μ. An increase in λ or a decrease in μ shifts the Beveridge curve outward: at any tightness, unemployment is higher.

A Beveridge shock differs from demand and supply shocks in one respect: it moves the FERU. The efficiency condition (9) depends on λ and μ through $u(\theta)$ and $\tau(\theta)$, so an outward shift of the Beveridge curve raises the efficient unemployment rate $u^* = u(\theta^*)$—as it did dramatically during the pandemic, when the FERU jumped from 4.0% to 6.7% in one quarter (figure 3A). Because the point of divine coincidence sits at (u^*, π^*), the unemployment-inflation Phillips curve shifts outward with the Beveridge curve: the same inflation now obtains at higher unemployment.

Beveridge shocks therefore provide a microfoundation for the cost-push shocks of the New Keynesian literature. In that literature, cost-push shocks are reduced-form disturbances appended to the Phillips curve to generate a policy tradeoff and explain the curve's instability. Here, the Phillips curve shifts whenever the Beveridge curve shifts, and shifts of the Beveridge curve are observable in the data—so the Phillips-curve shifts are measurable in real time through the FERU formula $u^* = \sqrt{uv}$.

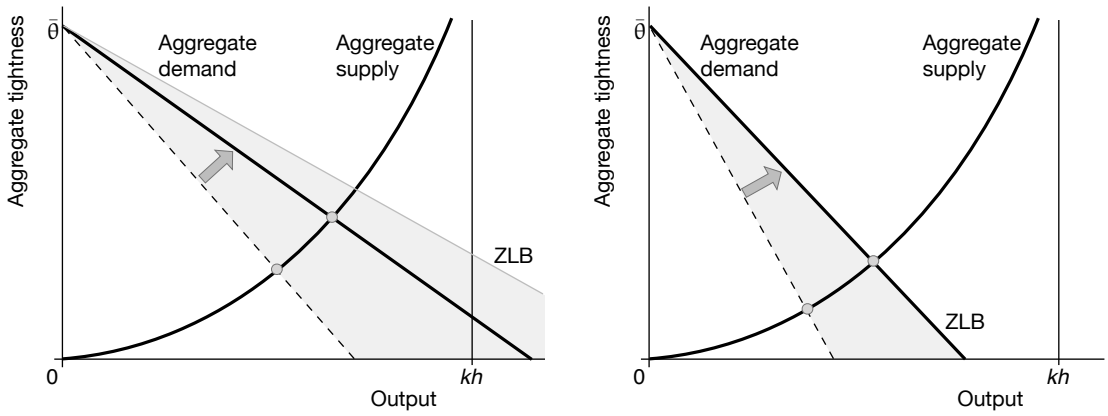
The divine coincidence survives Beveridge shocks: the coincidence point moves, but inflation remains on target if and only if unemployment is at the (new) FERU. No shock in the model can generate inflation at full employment. The tightness formulation of the Phillips curve is also far more stable than the unemployment formulation. The efficient tightness is pinned down by the Beveridge elasticity and the recruiting cost, $\theta^* = 1/(\kappa\beta)$, and both statistics are stable in US data—so θ^* barely moves when the Beveridge curve shifts, while u^* moves a lot. This explains fact 2: why tightness predicts inflation so much better than unemployment, especially in the pandemic period when the Beveridge curve was unstable (Barnichon, Oliveira, and Shapiro 2021; Barnichon and Shapiro 2022, 2024).

9. Implications for monetary policy

This section turns to the implications of the Beveridgean Phillips curve for monetary policy.

9.1. Satisfying the dual mandate: the FERI

In the United States, the Federal Reserve has a dual mandate: maintaining full employment and price stability. In the model, both mandates are achieved at once. By selecting the real interest rate r, monetary policy moves the aggregate demand curve along the aggregate supply curve (figure 8A). To satisfy the dual mandate, policy must bring the aggregate demand curve to the point where tightness is efficient: there, the Beveridge curve guarantees full employment ($u = u^*$) and the Phillips curve guarantees price stability ($\pi = \pi^*$).



A. Interest rate cuts B. Aggregate demand at the ZLB

FIGURE 8. Monetary policy in the model

The aggregate demand and supply curves are constructed in figure 5B. In panel A, 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 ZLB curve is the most outward position that the aggregate demand curve can reach, obtained when $i = 0$. In panel B, aggregate demand is so depressed that even at the zero lower bound it intersects aggregate supply below the efficient tightness.

We call the real interest rate that brings the economy to full employment the full-employment rate of interest (FERI), denoted r^*. Because of the divine coincidence, the FERI also delivers price stability. It solves $y^d(\theta^*, r^*) = y^s(\theta^*)$, which from (26) and (27) gives

$$r^* = \delta - \frac{1}{a} \cdot \frac{1 + \tau(\theta^*)}{(1-\alpha)c(\theta^*)^{-\alpha}}, \tag{30}$$

where $c(\theta^*)$ is the efficient level of consumption. The corresponding nominal interest rate is $i^* = r^* + \pi^*$. The Fed can achieve its dual mandate whenever $i^* \geq 0$. If aggregate demand is so depressed—the taste for consumption a so low—that $r^* < -\pi^*$, the dual-mandate policy would violate the zero lower bound; the central bank then sets $i = 0$, and the analysis of section 10 applies (figure 8B).

9.2. Soft landing

After the pandemic, inflation rose well above target—the inflation gap peaked at +4.9pp in 2022—while the labor market became inefficiently tight (figure 4). Once inflation was high, the policy debate centered on a mechanical question: could the Federal Reserve achieve a soft landing—bring inflation back down without a material rise in unemployment?

In the model, inflation is governed by tightness through the Phillips curve, so cooling inflation requires reducing tightness. Monetary policy operates through aggregate demand: raising the real interest rate depresses demand, which reduces tightness along

the aggregate supply curve. Since unemployment moves inversely with tightness, a strict soft landing is impossible: unemployment always rises when tightness falls. But the rise can be modest when the labor market starts from a very tight position, because the aggregate supply curve is convex: when tightness is high, a given reduction in tightness raises unemployment only a little.

The Beveridge curve makes the required rise in unemployment easy to quantify. By definition, $\theta = v(u)/u$ along the Beveridge curve, so differentiating in elasticity form gives $1 = -\beta\epsilon^u_\theta - \epsilon^u_\theta$, and therefore

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

In the United States, $\beta = 1$, so the elasticity of unemployment with respect to tightness is $-1/2$. If the Fed brings tightness down from 2—its value at the peak of the post-pandemic boom (2022:Q2)—to the efficient level of 1, tightness falls by 50%, so unemployment rises by 25%: from 3.5% at the peak of the boom to $3.5\% \times 1.25 = 4.4\%$. The effect on unemployment is not zero, but it is far from a recession. And this is almost exactly what happened: by July 2025, tightness had come down to 0.99 and the unemployment rate stood at 4.3% (Michaillat and Saez 2026). A soft landing from an inefficiently tight labor market is therefore possible, because the labor market starts on the steep portion of the aggregate supply curve.

9.3. Choice of target variables

Since the two mandates coincide, targeting the inflation rate or targeting the FERU is equivalent in the model. In practice, however, the Fed should target the variable that is most volatile—to observe most clearly departures from the dual mandate.

The slope of the Phillips curve determines which variable is most volatile. If the Phillips curve is steep, demand shocks mostly move inflation, so inflation is the more informative target. If the Phillips curve is flat, demand shocks mostly move unemployment, so unemployment is the more informative target. Because the Phillips curve is convex—flat when unemployment is inefficiently high, steep when it is inefficiently low—the implication is state-contingent: the Fed should target the FERU when the economy is inefficiently slack, and target 2% inflation when the economy is inefficiently tight. Since inflation and tightness vary commensurately (the tightness-inflation curve is broadly linear), an equivalent prescription when tightness is easier to measure than inflation is to target the FERU in slumps and the efficient tightness $\theta^* = 1$ in booms.

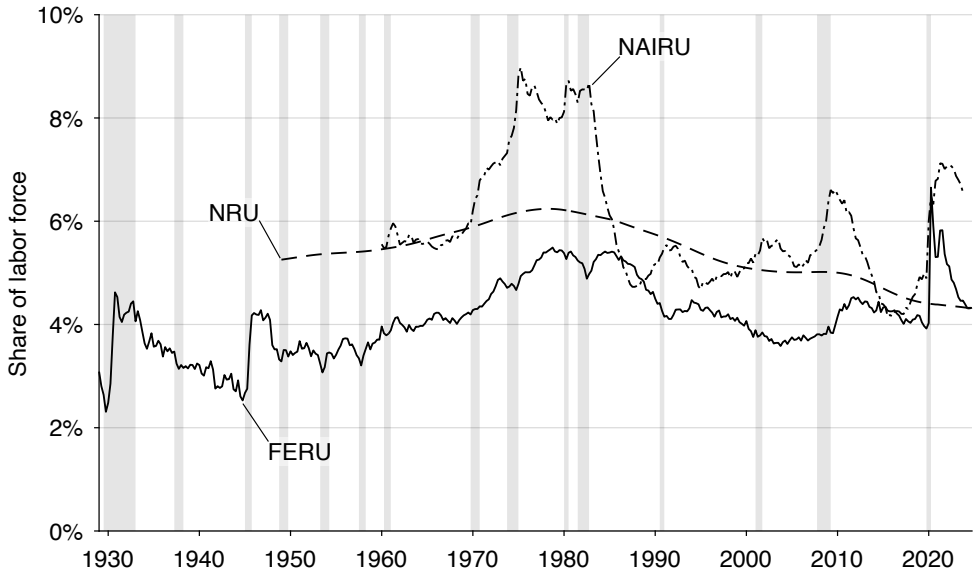


FIGURE 9. FERU, NRU, and NAIRU in the United States, 1929–2024

The FERU comes from figure 3A. The NRU is constructed by the US Congressional Budget Office (2024) for 1949–2024. The NAIRU is constructed by Crump et al. (2024, figure 2) for 1960–2023. Shaded areas indicate recessions dated by the NBER (2023).

9.4. The FERU is not the NRU or the NAIRU

US policymakers have long measured full employment with the natural rate of unemployment (NRU) or the nonaccelerating-inflation rate of unemployment (NAIRU).[17] Neither concept measures full employment in the model.

The NRU is a slow-moving trend of the unemployment rate; the widely used series produced by the US Congressional Budget Office (2024) is anchored on the assumption that the labor market was at full employment in 2005 and adjusted for the demographic composition of the labor force. Without a theory of efficiency, there is no reason for the trend of unemployment to be efficient. Indeed, between 1949 and 2019 the NRU is always above the FERU—by 1.2pp on average (figure 9). The values of the FERU for the 1950s and 1960s (between 3% and 4%) are actually in line with the full-employment estimates used by the Council of Economic Advisers at the time—3% until 1956 and 4% in the 1960s (Duboff 1977, pp. 8–10)—while the NRU is much above them.

The NAIRU is the unemployment rate at which inflation is stable; it is estimated from accelerationist Phillips curves. In the model, the NAIRU is simply not well defined:

[17] See Backhouse, Forder, and Laskaridis (2023) for a history of the two concepts. The NAIRU was introduced—as the noninflationary rate of unemployment—by Modigliani and Papademos (1975), and developed by Baily (1976) and Baily and Tobin (1977). The Council of Economic Advisers still defines full employment as "the lowest unemployment rate consistent with stable inflation" (US Council of Economic Advisers 2024, p. 24).

inflation is stable at any unemployment rate. If unemployment is at the FERU, inflation is stable at the target; if unemployment is above the FERU, inflation is stable below the target; if unemployment is below the FERU, inflation is stable above the target. The Phillips curve is not accelerationist. This may explain why NAIRU estimates are notoriously imprecise (Staiger, Stock, and Watson 1997; Laubach 2001), and why the NAIRU produced by state-of-the-art techniques (Crump et al. 2024) behaves so erratically—reaching 9.0% in 1975 and 7.0% in 2022 (figure 9)—for what is supposed to be a measure of full employment.

9.5. Ineffective monetary communication

The divine coincidence rests on the assumption that the central bank anchors normal inflation at its target: $\pi^n = \pi^*$. If monetary communication is ineffective, so that $\pi^n \neq \pi^*$, the derivations are unchanged except that π^* is replaced by π^n in the Phillips curve (29). It is as if a shifter $\pi^n - \pi^*$ were added to the curve:

$$\pi = (\pi^n - \pi^*) + \pi^* + \frac{y(\theta)}{\varpi\delta(\delta - r)}\left[\frac{\eta(\theta)}{1-\eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)} - 1\right]. \tag{31}$$

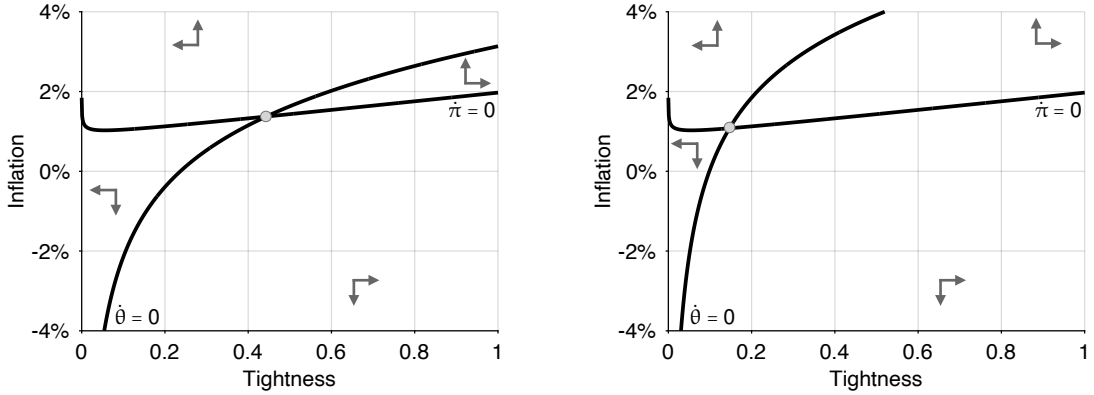
At full employment, inflation is now π^n instead of π^*: the divine coincidence breaks down, and the two mandates cannot be achieved simultaneously. If people believe that inflation will normally run above target ($\pi^n > \pi^*$), then inflation is too high when the economy is at full employment; equivalently, when inflation is on target, the labor market is inefficiently slack. Hence, in this model, the coincidence is not so much "divine" as the product of careful and effective messaging by the central bank. The good news is that such messaging appears to have succeeded in the United States since the mid-1990s—long-run inflation expectations have remained glued to 2% (Hazell et al. 2022)—which is why the coincidence is visible in modern US data.

10. Solution of the model at the zero lower bound

In normal times, the central bank sets the real interest rate, which separates the dynamics of tightness from those of inflation: tightness is determined by aggregate demand and supply, and inflation follows from the Phillips curve. At the zero lower bound (ZLB), this separation disappears. The nominal rate is stuck at 0, so the real rate is just negative inflation, $r = -\pi$, and tightness and inflation are jointly determined by a two-dimensional dynamical system.

10.1. When does the ZLB bind?

The economy reaches the ZLB after a large drop in aggregate demand—for instance a large increase in the marginal utility of wealth $1/a$. The ZLB binds when, even at a zero nominal



A. Source at the zero lower bound B. Source with lower aggregate demand

FIGURE 10. Phase diagrams at the zero lower bound in the calibrated model

The $\dot{\theta} = 0$ nullcline is given by equation (35). The $\dot{\pi} = 0$ nullcline is given by equation (36). The model is calibrated as in table 1. In panel A, the taste for consumption is $a = 1.66$. In panel B, the taste for consumption is lower, set to $a = 1.62$.

rate, aggregate demand falls short of the efficient level of output: $y^d(\theta^*, -\pi^*) < y^s(\theta^*)$ (figure 8B). This occurs when the marginal utility of wealth satisfies

$$\frac{1}{a} > \frac{(\delta + \pi^*)(1-\alpha)(c^*)^{-\alpha}}{1+\tau^*}, \tag{32}$$

where c^* and τ^* are the efficient values of consumption and the matching wedge.

10.2. Tightness-inflation dynamics

The dynamical system describing the model at the ZLB is derived in appendix B; here we report the main results. The system comprises two differential equations:

$$\frac{\dot{\theta}}{\theta}[\alpha[1-\eta(\theta)]u(\theta) + (1-\alpha)\eta(\theta)\tau(\theta)] = \frac{[1+\tau(\theta)]c(\theta)^{\alpha}}{a(1-\alpha)} - \delta - \pi \tag{33}$$

$$\dot{\pi} = \delta\left[\pi - \pi^*\right] + \frac{a(1-\alpha)}{\omega}c(\theta)^{1-\alpha}\left[1 - \frac{\eta(\theta)}{1-\eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)}\right]. \tag{34}$$

The $\dot{\theta}$ equation arises from the consumption-saving condition (23); the $\dot{\pi}$ equation arises from the pricing condition (24).

Figure 10 displays the phase diagram of the system in the calibrated model. The $\dot{\theta} = 0$ nullcline combines the aggregate demand and supply curves: as inflation rises, the real rate falls, so aggregate demand moves up along aggregate supply, raising tightness. The

nullcline is

$$\pi = \frac{1}{a(1-\alpha)}\left[1+\tau(\theta)\right]c(\theta)^{\alpha} - \delta. \tag{35}$$

The $\dot{\pi} = 0$ nullcline is based on the Phillips curve, describing inflation as an increasing function of tightness over the relevant range:

$$\pi = \pi^{*} + \frac{a(1-\alpha)}{\delta\varpi}c(\theta)^{1-\alpha}\left[\frac{\eta(\theta)}{1-\eta(\theta)}\cdot\frac{\tau(\theta)}{u(\theta)} - 1\right]. \tag{36}$$

This nullcline remains below the inflation target on $(0, \theta^{*})$, and it is flatter—closer to $\pi = \pi^{*}$—when prices are more rigid (higher ϖ).

The system always admits an equilibrium (θ^{z}, π^{z}) with $\theta^{z} \in (0, \theta^{*})$, and when price rigidity ϖ is large enough the equilibrium is unique (appendix B). Around the equilibrium, the system is a source: in all four quadrants of the phase diagram, trajectories move away from the equilibrium (figure 10). The solution of the model at the ZLB is therefore unique, although monetary policy is completely passive: tightness and inflation must jump to the equilibrium and stay there; had they jumped anywhere else, they would have diverged.

10.3. The ZLB is not topsy-turvy

The uniqueness result means that the model can be used to study ZLB episodes of any duration—even permanent ones. The unique ZLB solution is a normal slump: tightness is below its efficient level, unemployment is above the FERU, and inflation is below target. Comparative statics remain normal as well. A negative aggregate demand shock—a further decrease in the taste for consumption a—shifts the $\dot{\theta} = 0$ nullcline inward (from figure 10A to figure 10B): tightness falls, unemployment rises, and inflation falls. Demand shocks operate at the ZLB just as they do in normal times.

This behavior is entirely different from the New Keynesian model, in which the ZLB 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): when the ZLB lasts long, output and inflation collapse to implausible levels, the model admits no solution with a permanent ZLB, and policies acquire extraordinarily large effects (Michaillat and Saez 2021b; Del Negro, Giannoni, and Patterson 2023). None of these anomalies arise here. The difference comes from wealth in the utility function, which keeps aggregate demand well defined even under an interest-rate peg—and the ZLB is nothing but a peg at zero.

11. Conclusion

This paper develops a Beveridgean model of the Phillips curve. The economy is a matching economy in which households purchase services from other households, and all trades are mediated by a matching function—generating unemployment, vacancies, and a Beveridge curve. Directed search generates price competition: households post prices for the services they buy, and job seekers arbitrage across submarkets. Price-adjustment costs generate price rigidity. Combining the two—Moen-Rotemberg pricing—produces a nonvertical Phillips curve anchored at the efficient allocation.

The model explains the three facts that the post-pandemic economy revealed about the US Phillips curve. Inflation is on target exactly when unemployment is at the FERU, because directed search is efficient and the Fed anchors normal inflation at its target. Slack drives inflation, and the tightness formulation of the Phillips curve is more stable than the unemployment formulation, because the efficient tightness is pinned down by two stable statistics while the FERU moves with the Beveridge curve. And the Phillips curve is convex in the unemployment-inflation plane: steep when the labor market is inefficiently tight, flat when it is inefficiently slack. The convexity comes from the hyperbolic shape of the Beveridge curve. It explains why the drop in inflation was so muted when unemployment spiked in 2020, and why inflation flared up so quickly when the labor market overheated in 2021–2023.

The Fed's two mandates coincide, and both are achieved by setting the real interest rate to the FERI. The FERU is measurable in real time as $u^* = \sqrt{uv}$, unlike the NRU and NAIRU—the latter of which is not even well defined under a Beveridgean Phillips curve. A soft landing from an inefficiently tight labor market is feasible: the model's prediction for the post-pandemic cooling—unemployment rising from 3.5% to 4.4% as tightness returns to 1—matches what transpired by 2025. And the zero lower bound poses no theoretical difficulties: the model admits a unique solution there, even if the ZLB is permanent.

Beyond the Phillips curve, Moen-Rotemberg pricing could be useful in other matching models. Matching models of the labor market, product market, or entire economy require some price rigidity to generate realistic fluctuations in tightness (Shimer 2005; Michaillat and Saez 2015). Existing mechanisms rely on random search: fixed price norms (Shimer 2004; Hall 2005); credible bargaining (Hall and Milgrom 2008; Petrosky-Nadeau and Zhang 2021); staggered Nash bargaining (Gertler and Trigari 2009); and rigid price norms (Blanchard and Gali 2010; Michaillat 2012, 2024; Michaillat and Saez 2015, 2019). Moen-Rotemberg pricing offers an alternative built on directed search (Wright et al. 2021). The degree of rigidity is disciplined by measurable price-adjustment costs, and the resulting Phillips curve is automatically anchored at the efficient allocation.

References

Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutstrom. 2014. "Discounting Behavior: A Reconsideration." *European Economic Review* 71: 15–33.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics* 135 (2): 645–709. https://doi.org/10.1093/qje/qjaa004.

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.

Backhouse, Roger E., James Forder, and Christina Laskaridis. 2023. "The Natural Rate of Unemployment and the NAIRU." *European Economic Review* 159: 104563. https://doi.org/10.1016/j.euroecorev.2023.104563.

Baily, Martin N. 1976. "Contract Theory and the Moderation of Inflation by Recession and by Controls." *Brookings Papers on Economic Activity* 7 (3): 585–633. https://doi.org/10.2307/2534370.

Baily, Martin N. and James Tobin. 1977. "Macroeconomic Effects of Selective Public Employment and Wage Subsidies." *Brookings Papers on Economic Activity* 8 (2): 511–544. https://doi.org/10.2307/2534409.

Ball, Laurence M., Daniel Leigh, and Prachi Mishra. 2022. "Understanding US Inflation During the COVID Era." *Brookings Papers on Economic Activity* 53 (2): 1–54. https://doi.org/10.1353/eca.2022.a901276.

Barnichon, Regis. 2010. "Building a Composite Help-Wanted Index." *Economics Letters* 109 (3): 175–178. https://doi.org/10.1016/j.econlet.2010.08.029.

Barnichon, Regis, Luiz E. Oliveira, and Adam H. Shapiro. 2021. "Is the American Rescue Plan Taking Us Back to the 60s?" FRBSF Economic Letter 2021-27. https://perma.cc/7P6C-HSMG.

Barnichon, Regis and Adam H. Shapiro. 2022. "What's the Best Measure of Economic Slack?" FRBSF Economic Letter 2022-04. https://perma.cc/88SA-ABFN.

Barnichon, Regis and Adam Hale Shapiro. 2024. "Phillips Meets Beveridge." *Journal of Monetary Economics* 148: 103660. https://doi.org/10.1016/j.jmoneco.2024.103660.

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.

Benigno, Pierpaolo and Gauti B. Eggertsson. 2024a. "Revisiting the Phillips and Beveridge Curves: Insights from the 2020s Inflation Surge." Jackson Hole Economic Policy Symposium Paper. https://perma.cc/BQ5E-6MFK.

Benigno, Pierpaolo and Gauti B. Eggertsson. 2024b. "The Slanted-L Phillips Curve." *AEA Papers and Proceedings* 114: 84–89. https://doi.org/10.1257/pandp.20241051.

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

Beveridge, William H. 1944. *Full Employment in a Free Society*. London: George Allen & Unwin.

Bewley, Truman. 2005. "Fairness, Reciprocity, and Wage Rigidity." In *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*, edited by Herbert Gintis, Samuel Bowles, Robert Boyd, Ernst Fehr, Christopher Ansell, and Ken Binmore, chap. 11. Cambridge, MA: MIT Press.

Bewley, Truman F. 1999. *Why Wages Don't Fall During a Recession*. Cambridge, MA: Harvard University Press.

Bilbiie, Florin O. and Diego R. Kanzig. 2024. "Greed? Profits, Inflation, and Aggregate Demand." NBER Working Paper 31618. https://doi.org/10.3386/w31618.

Blanchard, Olivier J. and Jordi Gali. 2010. "Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment." *American Economic Journal: Macroeconomics* 2 (2): 1–30. https://doi.org/10.1257/mac.2.2.1.

BLS. 2024. "Personal Consumption Expenditures: Chain-type Price Index." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/PCEPI.

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.

BLS. 2025c. "Civilian Labor Force Level." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/CLF16OV.

BLS. 2025d. "Job Openings: Total Nonfarm." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/JTSJOL.

BLS. 2025e. "Total Separations: Total Nonfarm." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/JTSTSR.

BLS. 2025f. "Unemployment Level." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/UNEMPLOY.

Board of Governors of the Federal Reserve System. 2026. "Statement on Longer-Run Goals and Monetary Policy Strategy." https://perma.cc/7BBR-9UFG.

Caldara, Dario, Matteo Iacoviello, and David Yu. 2026. "Measuring Shortages since 1900." https://perma.cc/Q3K8-Z7QM.

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

Carney, Mark. 2017. "Lambda." Speech given at the London School of Economics. https://perma.cc/7JYQ-NEYB.

Chiarella, Carl, Peter Flaschel, Gang Gong, and Willi Semmler. 2003. "Nonlinear Phillips curves, complex dynamics and monetary policy in a Keynesian macro model." *Chaos, Solitons & Fractals* 18 (3): 613–634. https://doi.org/10.1016/S0960-0779(02)00680-X.

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

Cohen, Jonathan D., Keith M. Ericson, David Laibson, and John M. White. 2020. "Measuring Time Preferences." *Journal of Economic Literature* 58 (2): 299–347. https://doi.org/10.1257/jel.20191074.

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.

Crump, Richard K., Stefano Eusepi, Marc Giannoni, and Aysegul Sahin. 2024. "The Unemployment-Inflation Trade-off Revisited: The Phillips Curve in COVID Times." *Journal of Monetary Economics* 145: 103580. https://doi.org/10.1016/j.jmoneco.2024.103580.

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.

Duboff, Richard B. 1977. "Full Employment: The History of a Receding Target." *Politics & Society* 7 (1): 1–25. https://doi.org/10.1177/003232927700700101.

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.

Eyster, Erik, Kristof Madarasz, and Pascal Michaillat. 2021. "Pricing under Fairness Concerns." *Journal of the European Economic Association* 19 (3): 1853–1898. https://doi.org/10.1093/jeea/jvaa041.

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

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature* 40 (2): 351–401.

Gertler, Mark and Antonella Trigari. 2009. "Unemployment Fluctuations with Staggered Nash Wage Bargaining." *Journal of Political Economy* 117 (1): 38–86. https://doi.org/10.1086/597302.

Gitti, Giulia. 2024. "Nonlinearities in the Regional Phillips Curve with Labor Market Tightness." https://perma.cc/F8QP-FAE6.

Hall, Robert E. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness." *American Economic Review* 95 (1): 50–65. https://doi.org/10.1257/0002828053828482.

Hall, Robert E. and Paul R. Milgrom. 2008. "The Limited Influence of Unemployment on the Wage Bargain." *American Economic Review* 98 (4): 1653–1674. https://doi.org/10.1257/aer.98.4.1653.

Hansen, Bent. 1970. "Excess Demand, Unemployment, Vacancies, and Wages." *Quarterly Journal of Economics* 84 (1): 1–23. https://doi.org/10.2307/1879397.

Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jon Steinsson. 2022. "The Slope of the Phillips Curve: Evidence from U.S. States." *Quarterly Journal of Economics* 137 (3): 1299–1344. https://doi.org/10.1093/qje/qjac010.

Hirsch, Morris W., Stephen Smale, and Robert L. Devaney. 2013. *Differential Equations, Dynamical Systems, and an Introduction to Chaos*. 3rd ed. Boston: Academic Press.

Irby, LaToya. 2026. "What Is a Penalty APR?" Experian. https://perma.cc/8JTK-EVZG.

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

Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money*. New York: Macmillan.

Laubach, Thomas. 2001. "Measuring the NAIRU: Evidence from Seven Economies." *Review of Economics and Statistics* 83 (2): 218–231. https://doi.org/10.1162/00346530151143761.

Lewis, W. Arthur. 1955. *The Theory of Economic Growth*. George Allen & Unwin.

Michaillat, Pascal. 2012. "Do Matching Frictions Explain Unemployment? Not in Bad Times." *American Economic Review* 102 (4): 1721–1750. https://doi.org/10.1257/aer.102.4.

1721.

Michaillat, Pascal. 2024. "Modeling Migration-Induced Unemployment." arXiv:2303.13319. https://doi.org/10.48550/arXiv.2303.13319.

Michaillat, Pascal. 2026. "A Theory of Economic Slack." https://pascalmichaillat.org/18/. Draft manuscript.

Michaillat, Pascal and Emmanuel Saez. 2015. "Aggregate Demand, Idle Time, and Unemployment." *Quarterly Journal of Economics* 130 (2): 507–569. http://doi.org/10.1093/qje/qjv006.

Michaillat, Pascal and Emmanuel Saez. 2019. "Optimal Public Expenditure with Inefficient Unemployment." *Review of Economic Studies* 86 (3): 1301–1331. https://doi.org/10.1093/restud/rdy030.

Michaillat, Pascal and Emmanuel Saez. 2021a. "Beveridgean Unemployment Gap." *Journal of Public Economics Plus* 2: 100009. https://doi.org/10.1016/j.pubecp.2021.100009. Data available at https://github.com/pmichaillat/unemployment-gap.

Michaillat, Pascal and Emmanuel Saez. 2021b. "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.

Michaillat, Pascal and Emmanuel Saez. 2022. "An Economical Business-Cycle Model." *Oxford Economic Papers* 74 (2): 382–411. https://doi.org/10.1093/oep/gpab021.

Michaillat, Pascal and Emmanuel Saez. 2024a. "Beveridgean Phillips Curve." arXiv:2401.12475v2. https://doi.org/10.48550/arXiv.2401.12475.

Michaillat, Pascal and Emmanuel Saez. 2024b. "$u^* = \sqrt{uv}$: The Full-Employment Rate of Unemployment in the United States." *Brookings Papers on Economic Activity* 55 (2). https://doi.org/10.1353/eca.2024.a964373.

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

Modigliani, Franco and Lucas Papademos. 1975. "Targets for Monetary Policy in the Coming Year." *Brookings Papers on Economic Activity* 6 (1): 141–166.

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

Mortensen, Dale T. 1970. "Job Search, the Duration of Unemployment, and the Phillips Curve." *American Economic Review* 60 (5): 847–862. https://www.jstor.org/stable/1818285.

NBER. 2023. "US Business Cycle Expansions and Contractions." https://perma.cc/B4LH-BVT8.

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

Petrongolo, Barbara and Christopher A. Pissarides. 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature* 39 (2): 390–431. https://doi.org/10.1257/jel.39.2.390.

Petrosky-Nadeau, Nicolas and Lu Zhang. 2021. "Unemployment Crises." *Journal of Monetary Economics* 117: 335–353. https://doi.org/10.1016/j.jmoneco.2020.01.009.

Phillips, A. W. 1958. "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957." *Economica* 25 (100): 283–299. https://doi.org/10.1111/j.1468-0335.1958.tb00003.x.

Powell, Jerome H. 2022. "May 2022 Federal Open Market Committee Press Conference." Press Conference Transcript, Board of Governors of the Federal Reserve System. https://perma.cc/7DRL-WFQX.

Rossman, Ted. 2026. "Current Credit Card Interest Rates." Bankrate. https://perma.cc/24DA-7EUL.

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

Shiller, Robert J. 1997. "Why Do People Dislike Inflation?" In *Reducing Inflation: Motivation and Strategy*, edited by Christina D. Romer and David H. Romer, chap. 1. Chicago: University of Chicago Press. http://www.nber.org/chapters/c8881.

Shimer, Robert. 2004. "The Consequences of Rigid Wages in Search Models." *Journal of the European Economic Association* 2 (2/3): 469–479. https://doi.org/10.1162/154247604323068159.

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (1): 25–49. https://doi.org/10.1257/0002828053828572.

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.

Staiger, Douglas O., James H. Stock, and Mark W. Watson. 1997. "The NAIRU, Unemployment and Monetary Policy." *Journal of Economic Perspectives* 11 (1): 33–49. https://doi.org/10.1257/jep.11.1.33.

Todd, Tim. 2012. *The Balance of Power: The Political Fight for an Independent Central Bank, 1790–present*. 2nd ed. Public Affairs Department of the Federal Reserve Bank of Kansas City.

US Congressional Budget Office. 2024. "Noncyclical Rate of Unemployment." FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/NROU.

US Council of Economic Advisers. 2024. "Economic Report of the President." https://perma.cc/2JGM-QG4Z.

Williams, John C. 2017. "Comment on 'Safety, Liquidity, and the Natural Rate of Interest'." *Brookings Papers on Economic Activity* 48 (1): 295–303.

Wright, Randall, Philipp Kircher, Benoit Julien, and Veronica Guerrieri. 2021. "Directed Search and Competitive Search Equilibrium: A Guided Tour." *Journal of Economic Literature* 59 (1): 90–148. https://doi.org/10.1257/jel.20191505.

Zbaracki, Mark J., Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. 2004. "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets." *Review of Economics and Statistics* 86 (2): 514–533.

Appendix A. Convexity of the Phillips curve

This appendix proves proposition 1: the Phillips curve $\pi(u)$ is strictly convex in u under the assumption that the matching function is Cobb-Douglas and symmetric,

$$m(u, v) = \mu\sqrt{uv}.$$

It is difficult to obtain more general analytical results; for other matching functions, numerical analysis is required, as in section 7.

A.1. Symmetric Cobb-Douglas matching function

With a symmetric Cobb-Douglas matching function, the trading rates are

$$f(\theta) = \mu\sqrt{\theta} \quad \text{and} \quad q(\theta) = \frac{\mu}{\sqrt{\theta}}, \tag{A1}$$

and the matching elasticity is constant at $\eta(\theta) = 1/2$, so

$$\frac{\eta(\theta)}{1-\eta(\theta)} = 1. \tag{A2}$$

The unemployment rate is given by (3), so

$$u(\theta) = \frac{\lambda}{\lambda + \mu\sqrt{\theta}}.$$

Inverting this strictly decreasing relationship, we write tightness as a strictly decreasing function of the unemployment rate:

$$\theta(u) = \left[\frac{\lambda(1-u)}{\mu u}\right]^2. \tag{A3}$$

The highest possible value for tightness is $\theta \to \bar{\theta}$, when $\tau(\theta) \to \infty$. The upper bound $\bar{\theta}$ is defined such that $q(\bar{\theta}) = \kappa\lambda$, so here

$$\bar{\theta} = \left(\frac{\mu}{\kappa\lambda}\right)^2.$$

The lowest possible unemployment rate is achieved when $\theta \to \bar{\theta}$. We denote this minimum unemployment rate by $\bar{u} = u(\bar{\theta})$; it is given by

$$\bar{u} = \frac{\kappa\lambda^2}{\mu^2 + \kappa\lambda^2}. \tag{A4}$$

A.2. Preliminary result

LEMMA A1. *The ratio between the matching wedge and unemployment rate admits the closed form*

$$\frac{\tau(u)}{u} = \frac{1-\bar{u}}{u-\bar{u}} - \frac{1}{u}. \tag{A5}$$

The ratio is strictly decreasing and strictly convex in u on $(\bar{u}, 1)$.

PROOF. The proof has three steps: derive the closed form (A5), then sign the first and second derivatives.

First, combining (A1) and (A3) allows us to express the ad-filling rate as a function of the unemployment rate:

$$q(u) = \frac{\mu^2}{\lambda} \cdot \frac{u}{1-u}.$$

Using expression (5), we can express the matching wedge as a function of the unemployment rate too:

$$\tau(u) = \frac{\kappa\lambda}{\frac{\mu^2}{\lambda}\frac{u}{1-u} - \kappa\lambda} = \frac{\kappa\lambda^2(1-u)}{(\mu^2+\kappa\lambda^2)u - \kappa\lambda^2} = \frac{\bar{u}(1-u)}{u-\bar{u}},$$

where the last equality uses the definition (A4) of $\bar{u}$. Then, the ratio between matching wedge and unemployment rate is a simple function of the unemployment rate:

$$\frac{\tau(u)}{u} = \frac{\bar{u}(1-u)}{u(u-\bar{u})} = \frac{\bar{u}-u+u-\bar{u}u}{u(u-\bar{u})} = \frac{1-\bar{u}}{u-\bar{u}} - \frac{1}{u},$$

which establishes (A5).

Second, differentiating (A5) once with respect to the unemployment rate, we get

$$\frac{d\tau/u}{du} = \frac{1}{u^2} - \frac{1-\bar{u}}{(u-\bar{u})^2}. \tag{A6}$$

The sign of the derivative is the sign of

$$(u-\bar{u})^2 - (1-\bar{u})u^2 = \bar{u}\left[\bar{u} - 2u + u^2\right] = \bar{u}\left[(\bar{u}-u) + u(u-1)\right].$$

Since $0 < \bar{u} < u < 1$, we have $u - 1 < 0$ and $\bar{u} - u < 0$, so the first derivative is strictly negative: τ/u is strictly decreasing in u.

Third, differentiating (A6) with respect to the unemployment rate, we get the second derivative of the ratio:

$$\frac{d^2\tau/u}{du^2} = \frac{2(1-\bar{u})}{(u-\bar{u})^3} - \frac{2}{u^3}.$$

The sign of the second derivative is the sign of

$$\begin{aligned}(1-\bar{u})u^3-(u-\bar{u})^3 &= \bar{u}\left[-u^3+\bar{u}^2+3u^2-3u\bar{u}\right]\\ &= \bar{u}\left[u^2(1-u)+(\bar{u}-u)^2+u(u-\bar{u})\right].\end{aligned}$$

Since $0<\bar{u}<u<1$, all the terms in the square bracket are positive: $u^2(1-u)>0$, $(\bar{u}-u)^2>0$, and $u(u-\bar{u})>0$. Hence the second derivative is strictly positive: $\tau(u)/u$ is strictly convex in u. □

A.3. Proof of proposition 1

With lemma A1 in hand, we can now prove the proposition.

PROOF OF PROPOSITION 1. The Phillips curve $\pi(\theta)$ is given by (29), where here (A2) holds. We compose it with the function $\theta(u)$ given by (A3) to obtain a Phillips curve linking inflation to unemployment:

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

The ratio $\tau(u)/u$ is given by (A5) and output is given by $y(u)=(1-u)kh$. Given the linear expression for output, we can rewrite the Phillips curve as

$$\pi(u)=\pi^*+\frac{kh}{\varpi\delta(\delta-r)}\cdot(1-u)\cdot\left[\frac{\tau(u)}{u}-1\right]. \tag{A7}$$

We differentiate (A7) twice. The first derivative is

$$\pi'(u)=\frac{kh}{\varpi\delta(\delta-r)}\cdot\left[(1-u)\cdot\frac{d\tau/u}{du}+1-\frac{\tau(u)}{u}\right].$$

At efficiency, by (9) and (A2), $\tau(u^*)/u^*=1$, so the sign of $\pi'(u^*)$ is the sign of $(1-u^*)d(\tau/u)/du$. By lemma A1, $d(\tau/u)/du<0$, so $\pi'(u^*)<0$: in the vicinity of the FERU, the Phillips curve is strictly decreasing in the unemployment rate.

Differentiating again, we obtain

$$\pi''(u)=\frac{kh}{\varpi\delta(\delta-r)}\cdot\left[(1-u)\cdot\frac{d^2\tau/u}{du^2}-2\cdot\frac{d\tau/u}{du}\right].$$

By lemma A1, $d(\tau/u)/du<0$ and $d^2(\tau/u)/du^2>0$, so both terms in the square brackets are strictly positive, and $\pi''(u)>0$ on $(\bar{u},1)$: the Phillips curve is strictly convex in the unemployment rate. □

Appendix B. Dynamical system at the zero lower bound

This appendix derives and studies the dynamical system describing the behavior of tightness and inflation at the zero lower bound.

B.1. Costate variables, inflation, and tightness

The first step is to express the costate variables ψ^b and ψ^p as functions of inflation π and tightness θ. Using these relationships, we transform the differential equations for the costate variables into differential equations for inflation and tightness.

For the costate variable ψ^p, the relationship is directly given by (22):

$$\psi^p = \frac{\varpi}{a}\left[\pi - \pi^*\right], \qquad \dot{\psi}^p = \frac{\varpi}{a}\dot{\pi}. \tag{A8}$$

For the costate variable ψ^b, the relationship arises from (21). In the model, aggregate consumption is directly determined by tightness, from (6):

$$c(\theta) = \frac{1-u(\theta)}{1+\tau(\theta)} \cdot kh. \tag{A9}$$

Using the function $c(\theta)$, condition (21) relates ψ^b to tightness:

$$\psi^b(\theta) = (1-\alpha) \cdot \frac{c(\theta)^{-\alpha}}{1+\tau(\theta)}. \tag{A10}$$

Notice also, which will be useful below, that because $y(\theta) = c(\theta)\left[1+\tau(\theta)\right]$,

$$\psi^b(\theta) y(\theta) = (1-\alpha) c(\theta)^{1-\alpha}. \tag{A11}$$

By the chain rule, $\dot{\psi}^b = (\psi^b/\theta) \cdot \epsilon_\theta^{\psi^b} \cdot \dot{\theta}$, so we need the elasticity of ψ^b with respect to tightness. From (8), $\epsilon_\theta^{1+\tau} = \eta(\theta)\tau(\theta)$; from (7), $\epsilon_\theta^{1-u} = \left[1-\eta(\theta)\right]u(\theta)$. Hence, from (A9), the elasticity of consumption is

$$\epsilon_\theta^c = \left[1-\eta(\theta)\right]u(\theta) - \eta(\theta)\tau(\theta),$$

and from (A10), the elasticity of the costate variable is

$$\epsilon_\theta^{\psi^b} = -\left[\alpha\left[1-\eta(\theta)\right]u(\theta) + (1-\alpha)\eta(\theta)\tau(\theta)\right].$$

The time derivative of the costate variable is therefore related to that of tightness by

$$\frac{\dot{\psi}^b}{\psi^b} = -\left[\alpha\left[1-\eta(\theta)\right]u(\theta) + (1-\alpha)\eta(\theta)\tau(\theta)\right] \cdot \frac{\dot{\theta}}{\theta}. \tag{A12}$$

B.2. Differential equations for inflation and tightness

We first transform the differential equation for the costate variable ψ^p, given by (24), into a differential equation for inflation. We replace ψ^p and $\dot{\psi}^p$ by the expressions (A8), and $\psi^b y$ by the expression (A11). We obtain equation (34):

$$\dot{\pi} = \delta\left[\pi - \pi^*\right] + \frac{(1-\alpha)a}{\varpi} \cdot c(\theta)^{1-\alpha}\left[1 - \frac{\eta(\theta)}{1-\eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)}\right].$$

The nullcline $\dot{\pi} = 0$ is given by (36). This nullcline is close to, but not the same as, the Phillips curve in normal times (29): in normal times the costate variable ψ^b is fixed by the real interest rate, whereas at the zero lower bound it is jointly determined by tightness and inflation.

Next, we transform the differential equation for the costate variable ψ^b, given by (23), into a differential equation for tightness. We divide the equation by ψ^b and set the real interest rate to $r = -\pi$:

$$\frac{\dot{\psi}^b}{\psi^b} = (\delta + \pi) - \frac{1}{a\psi^b}.$$

Replacing ψ^b and $\dot{\psi}^b$ by the expressions (A10) and (A12) yields equation (33):

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

The nullcline $\dot{\theta} = 0$ is given by (35).

B.3. Equilibrium of the dynamical system

We now determine the equilibrium of the dynamical system composed of (33) and (34). At the equilibrium, $\dot{\pi} = \dot{\theta} = 0$, so the equilibrium lies at the intersection of the two nullclines. Let us introduce auxiliary functions to simplify the analysis:

$$P(\theta) = \pi^* + \frac{(1-\alpha)a}{\delta\varpi} \cdot c(\theta)^{1-\alpha}\left[\frac{\eta(\theta)}{1-\eta(\theta)} \cdot \frac{\tau(\theta)}{u(\theta)} - 1\right]$$

$$Q(\theta) = \frac{1}{a(1-\alpha)} \cdot \left[1+\tau(\theta)\right]c(\theta)^{\alpha} - \delta$$

$$R(\theta) = \frac{\theta}{\alpha\left[1-\eta(\theta)\right]u(\theta) + (1-\alpha)\eta(\theta)\tau(\theta)}.$$

The nullcline (36) is $\pi = P(\theta)$ and the nullcline (35) is $\pi = Q(\theta)$. Let $S(\theta) = Q(\theta) - P(\theta)$: the tightness that solves the model satisfies $S(\theta) = 0$.

At $\theta = 0$, we have $u(0) = 1$ so $c(0) = 0$ and $S(0) = -(\delta + \pi^*) < 0$. At $\theta = \theta^*$, we have

$\tau/u = [1-\eta]/\eta$ by (9), so

$$S(\theta^*) = \frac{1}{a(1-\alpha)} \cdot \left[1+\tau(\theta^*)\right] c(\theta^*)^{\alpha} - (\delta + \pi^*).$$

The economy is at the zero lower bound only when aggregate demand is depressed enough, so that the marginal utility of wealth satisfies (32)—and (32) is equivalent to $S(\theta^*) > 0$. We have established that $S(0) < 0$ and $S(\theta^*) > 0$, and S is continuous on $[0, \theta^*]$; by the intermediate value theorem, the equation $S(\theta) = 0$ admits a solution on $(0, \theta^*)$. The dynamical system therefore admits an equilibrium on $(0, \theta^*)$, which we denote (θ^z, π^z).

Is the equilibrium unique? It would be if S were strictly increasing on $(0, \theta^*)$. The function Q is strictly increasing on $(0, \theta^*)$, as $\tau(\theta)$ and $c(\theta)$ are increasing on the interval. But the function P is not monotone (indeed $P(0) = P(\theta^*) = \pi^*$), so we cannot conclude in general. However, when prices are rigid enough—when ϖ is large enough—the fluctuations in P become arbitrarily small: at the limit $\varpi \to \infty$, $P(\theta) \to \pi^*$ for any θ. So for ϖ large enough, S never returns to 0 once it has crossed 0 at θ^z, and the equilibrium is unique. The equilibrium is unique under the calibration of figure 10. In that case, since $S(0) < 0$ and $S(\theta^*) > 0$, S crosses 0 from below, so $S'(\theta^z) > 0$. In the rest of the appendix we assume that price rigidity ϖ is sufficient to guarantee uniqueness.

B.4. Classification of the dynamical system

Finally, we determine whether the dynamical system is a source, sink, or saddle around the zero-lower-bound equilibrium, by computing the Jacobian of the system, which describes the local dynamics (Hirsch, Smale, and Devaney 2013, pp. 61–64).

With the auxiliary functions, the differential equation (33) is $\dot{\theta} = R(\theta)\left[Q(\theta) - \pi\right]$ and the differential equation (34) is $\dot{\pi} = \delta\left[\pi - P(\theta)\right]$. The Jacobian of the system $[\pi(t), \theta(t)]$ is

$$\mathbf{J} = \begin{bmatrix} \delta & -\delta P'(\theta^z) \\ -R(\theta^z) & R(\theta^z)Q'(\theta^z) \end{bmatrix},$$

where we simplified the Jacobian using $Q(\theta^z) - \pi^z = 0$, which holds by definition of the equilibrium. The trace and determinant of the Jacobian are

$$\begin{aligned} \operatorname{tr}(\mathbf{J}) &= \delta + R(\theta^z)Q'(\theta^z) \\ \det(\mathbf{J}) &= \delta R(\theta^z)\left[Q'(\theta^z) - P'(\theta^z)\right] = \delta R(\theta^z)S'(\theta^z). \end{aligned}$$

We know that $\delta > 0$, $R(\theta^z) > 0$, and $Q'(\theta^z) > 0$, so the trace is strictly positive. We also know that $S'(\theta^z) > 0$, so the determinant is strictly positive. Thus, the dynamical system describing inflation and tightness is a source around the zero-lower-bound equilibrium.