

The Slanted-L Phillips Curve

By PIERPAOLO BENIGNO AND GAUTI B. EGGERTSSON *

As Tobin (1972, p. 9) aptly stated at the time, the Phillips curve is “an empirical finding in search of theory, like Pirandello characters in search of an author.”

Since Tobin’s writing, the search has persisted. Before the inflation surge of the 2020s the literature converged on the New Keynesian Phillips curve. Employed by all major policy institutions, it has two central properties. 1) Linearity: It is a log-linear relationship between inflation and some measure of economic activity such as labor market tightness. 2) Flatness: A percentage reduction in, e.g., the unemployment rate, results in only a modest increase in inflation. A widely cited estimate by Hazell et al. (2022), for example, suggests that a 1 percentage point drop in unemployment increases inflation by 0.33 percentage points, provided inflation expectations remain anchored. That estimate is based on data from the period 1978-2018.

The modern incarnation of the Phillips curve was subject to a severe stress test during the inflation surge of 2020s in the United States. It is hard to claim it emerged from it with flying colors. As we show in Benigno and Eggertsson (2023, BE from now on), both Wall Street professional forecasters (Survey of Professional Forecasters) as well as policymakers at the Federal Reserve (Summary of Economic Projections) were caught flatfooted. Both failed to anticipate the surge, which started in mid 2021. Moreover, as inflation escalated, they consistently predicted inflation to revert quickly to the Federal Re-

serve’s inflation target. Yet, contrary to these predictions, the surge accelerated well into 2022 until the Federal Reserve started raising rates.

The inflation surge of the 2020s created the largest inflation spike in the U.S. since the Great Inflation of the 1970s. BE suggest that the economic profession failed to anticipate the surge because it disregarded what was once upon a time considered a conventional wisdom: The Phillips curve is highly non-linear. Ironically, the very curve Phillips (1958) first proposed is, in fact, highly non-linear. Indeed, it is one of the central points of Phillips’s seminal paper. Phillips suggests that with “very few unemployed we should expect employers to bid up wages quite rapidly, each firm and each industry being continually tempted to offer a little above the prevailing wage.” In contrast, when unemployment is high, “workers are reluctant to offer their services at less than the prevailing rate,” so “wages fall only very slowly.”

BE argue that the non-linearity of the Phillips curve was overlooked for a simple reason: Empirical evidence for the non-linearity can only be found in U.S. aggregate data from *before* the Great Inflation of the 1970s. Since the Great Inflation of the 1970s serves as the central reference point for most modern observers analyzing inflation dynamics, and tight labor markets played no role in explaining it, this created a blind spot.

Figure 1, extracted from BE, presents a scatter plot of the inflation rate and labor market tightness in the U.S. on quarterly basis for the period 2009-2023 (for earlier periods, refer to BE). Labor market tightness is defined as the ratio between firms’ vacancy rates (v) and workers’ unemployment rates (u). The labor market is tight when there are more jobs firms are looking to fill than there are workers looking

* Benigno: University of Bern, Schanzeneckstrasse 1, 3001 Bern, Switzerland, pierpaolo.benigno@unibe.ch. Eggertsson: Brown University, Robinson Hall, 64 Waterman Street, Providence, RI 02912, United States, gauti.eggertsson@brown.edu. Pierpaolo Benigno gratefully acknowledges support from SNF grant number 100018_207669 and PRIN number 2020PLZR2P_001 - CUP: I83C22000360008. We thank our discussant, Jon Steinsson, for comments.

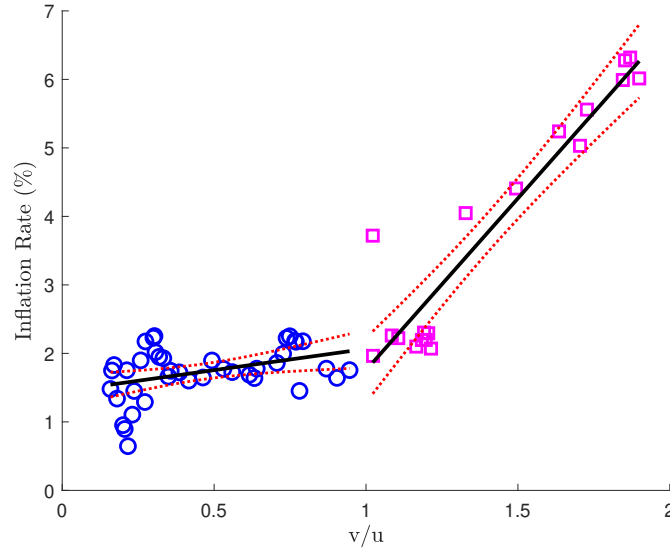


FIGURE 1. UNITED STATES 2009-2023: CPI INFLATION RATE AND VACANCY-TO-UNEMPLOYED RATIO (v/u)

for jobs, i.e. $v/u > 1$. While the exact cut-off point, i.e. 1, is not precisely estimated by BE, Beveridge (1944) argues for it on theoretical grounds. BE use the term *labor shortage* to describe the labor market conditions when $v/u > 1$, a term commonly used in the U.S. during the 2020s inflation surge.¹ Figure 1 suggests a non-linearity when $v/u > 1$, a claim BE establish is statistically significant looking at a longer sample. A key empirical observation is that, outside of 2020s, one needs to look *before* the Great Inflation of the 1970s to find extended periods of labor shortage. BE documents that, aside from the 2020s, there have been four occasions when $v/u > 1$: WWI, WWII, the Korean War and the escalation of Vietnam War spending (along with President Johnson’s tax cuts) in the late 1960s.² Like the 2020s, all these periods were marked by an inflation surge.

This paper presents international evidence that the Phillips curve is non-linear

¹For instance, several stores announced closures during specific hours attributed to “labor shortage” when the labor market was at its tightest. Similarly, many restaurants seated customers at only a third of their capacity due to “labor shortage”.

²Hall and Sargent (2022) discuss evidence of extraordinary monetary and fiscal stimuli, referring to “three wars” by including COVID-19 alongside WWI and WWII.

using the unemployment rate as a proxy for labor market tightness instead of v/u .³ The focus is on the period from the first quarter of 2009 to the third quarter of 2023 which corresponds to the last of the four sub-periods analyzed in BE. Labor shortage becomes prominent towards the end of this period. Our question is whether similar labor shortages were observed in other industrialized countries and, if so, whether they also triggered an inflation surge.

The general conclusion is that for the sample of seven other major industrial countries, the pattern mirrors that of the U.S. As we will see, the results become particularly stark once we focus on unemployment as a measure of slack instead of v/u . What emerges is an slanted-L shaped Phillips curve in unemployment-output space.

At a broad level, the economic mechanism behind the Slanted-L Phillips curve aligns with BE. If the economy operates below full capacity, with idle workers and vacant factories, an increase in nominal spending boosts output (reduces unemployment) with a modest impact on prices. While most factors of production can be in-

³Our reliance on unemployment data is due to lack of comprehensive, comparable data on firm vacancies across countries, a topic reserved for future research.

creased over some period, one way or another, there is one factor fixed over any relevant time horizon: the number of people. Thus, at some point, a firm responding to higher demand will eventually run out of people to hire. This intuitive, and perhaps obvious, observation is what gave rise to the old conventional wisdom that, as a matter of pure logic, the Phillips curve *has to* be nonlinear *at some point*. If firms cannot ramp up production due to a lack of labor, any additional increase in nominal spending results in increased inflation rather than higher output. Alternatively, with output hitting a wall, firms can resort to rationing goods and services instead of raising prices, but we will abstract from this possibility.⁴ Our proposed Phillips curve, in the unemployment and inflation space, is therefore a slanted-L, with the lower leg of the L slightly downward-slanted for reasons we clarify shortly. While the Slanted-L Phillips curve suggests that demand shocks have a much larger inflation impact once the economy enters the vertical part of the slanted-L, it also implies that supply shocks create much larger movements in inflation in that region. The large impact of supply shocks on inflation during labor shortages is discussed in detail and established both empirically and theoretically in BE.

The general perspective proposed in this paper, somewhat surprisingly, reconciles the work of Keynes and Friedman. Keynes's General Theory posits that rigidly downward wages rationalize why an increase in nominal spending increases real output and employment. Yet, Keynes also develops a theory of "demander's" inflation, similar to the neoclassical account of the surge in inflation during World War II, which occurs when "government, investors, and consumers want in real terms... more than... available producible output," noting that "...in peacetime...the size of the cake depends on the amount of work done. But in wartime, the size of the cake is fixed."

⁴Rationing does in fact often become the norm in episodes featuring labor shortages during war times because governments try to contain inflation by price controls.

(Keynes, 1940, p. 4).

The view that the economy is fundamentally asymmetric, as implied by the Slanted-L Phillips Curve, is shared by Friedman's plucking model. In Friedman (1964, 1993), "Output is viewed as bumping along the ceiling of maximum feasible output, except that every now and then it is plucked down by a cyclical contraction." (Friedman, 1964, p.17).⁵ In what follows, Section I describes the evidence, Section II a simple model and an online Appendix details about the data and estimation.

I. International Evidence on the Slanted-L Phillips Curve

Figure 2 shows data on unemployment and inflation in eight advanced economies from the first quarter of 2009 to the third quarter in 2023. The evidence broadly fits our hypothesis. When unemployment declines, inflation gently increases. Once unemployment goes below some critical threshold, however, inflation surges quickly. This threshold, however, differs from country to country.

To formalize the visual impression given by the data, we draw an "L with a slant" without any attempt to add controls or claim identification. At the corner of the L in each country is our measure of the unemployment rate consistent with maximum employment. We label this as u^f .⁶ The slanted right leg of the L is estimated via Ordinary Least Squares regression on the remaining data points.

Figure 3 combines the data from these countries. The thick blue line shows the slanted L, with the slanted leg obtained via the regression:⁷

$$\pi_{i,t} = \underset{(0.2438)}{2.4722} - \underset{(0.0359)}{0.1336} \times u_{i,t}^{dev} + \varepsilon_t$$

where i represents each country (Australia,

⁵See Dupree, Nakamura, and Steinsson (2019) for a recent attempt to resurrect Friedman's plucking model.

⁶It is approximated by calculating the average of the observations within the range from the lowest unemployment level to 0.2 percent above it.

⁷The number of observations is 417 and the $R^2 = 0.032$.

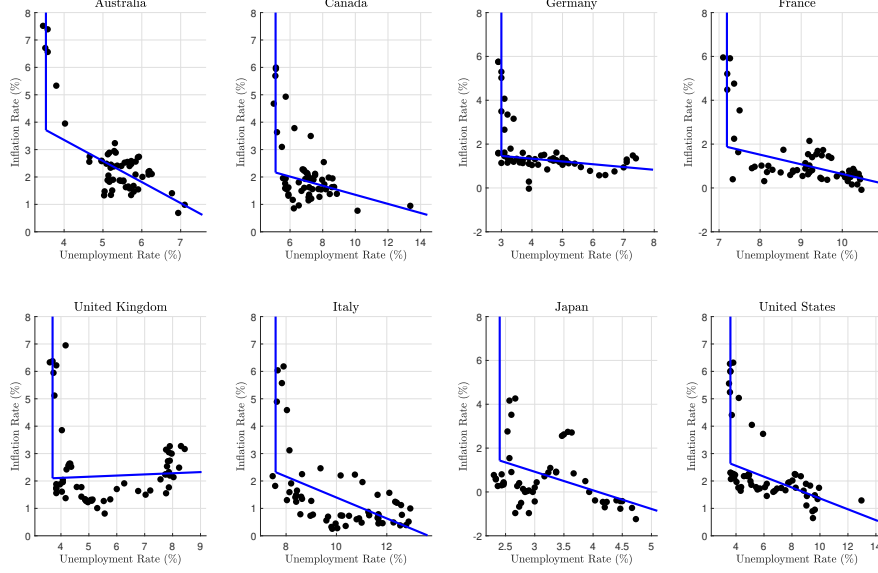


FIGURE 2. INTERNATIONAL EVIDENCE ON INFLATION AND UNEMPLOYMENT TRADE-OFF, 2009-2023.

Canada, Germany, France, United Kingdom, Italy, Japan, United States). Here $u_{i,t}^{dev}$ is the adjusted unemployment rate constructed to be comparable across countries.⁸ The thick black line in Figure 3, however, employs non-linear least squares to estimate the original curve proposed by Phillips:

$$(1) \quad \pi_{i,t} = a + b \left(\frac{1}{u_{i,t}^{dev}} \right)^c$$

where a , b , and c are estimated coefficients. Remarkably, estimating the curve initially proposed by Phillips results an object that strongly resembles the Slanted-L Phillips curve.⁹

II. A Model of the Slanted-L Phillips Curve

A representative household maximizes utility

$$\sum_{t=0}^{\infty} \beta^t U(C_t)$$

⁸For each country, this variable is calculated by subtracting the country-specific unemployment rate at full employment, u^f , and adding the average u^f across all countries.

⁹Phillips originally fitted his curve to unemployment and *wage* inflation instead of price inflation.

where $0 < \beta < 1$ is the rate of time preference, $U(\cdot)$ is a concave function of the consumption good C , subject to:

$$P_t C_t + B_t = (1 + i_{t-1})B_{t-1} + W_t L_t + \Psi_t,$$

where P_t is the price level, B_t one-period risk-free bond that pays interest rate i_t , W_t is the nominal wage, L_t employment, Ψ_t are firms' profits. Each period the household receives an employment endowment \bar{L} so equilibrium employment will be bounded by $0 < L_t \leq \bar{L}$. The household incurs no dis-utility of working.

Firms produce the consumption goods using the technology $Y_t = A_t L_t^\alpha$, where Y_t is output, A_t is a technological factor, and the parameter α is between 0 and 1. Firms maximize profits taking prices and wages as given, yielding optimal labor demand

$$(2) \quad L_t^d = \left(\frac{1}{\alpha A_t} \frac{W_t}{P_t} \right)^{-\frac{1}{1-\alpha}}.$$

If wages are flexible they adjust so that the supply of labor is equal to demand, $L_t^d = \bar{L}$, which we refer to as full employment, i.e., $L^f = \bar{L}$. The unemployment rate at full employment is $u^f = 1 - L^f / \bar{F}$ where \bar{F} is the labor force which is divided be-

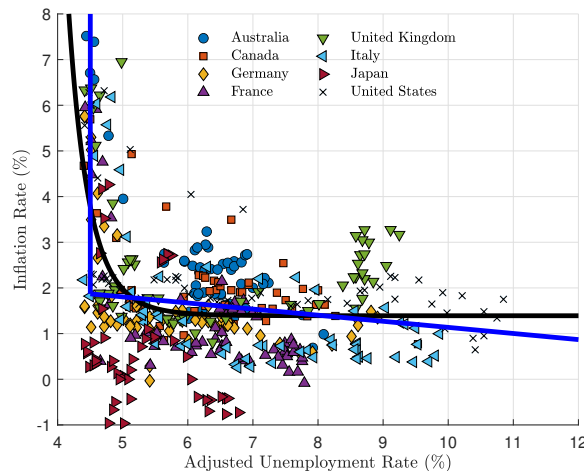


FIGURE 3. INTERNATIONAL EVIDENCE ON INFLATION AND UNEMPLOYMENT TRADE-OFF, 2009-2023 (POOLED DATA).

tween unemployed and employed. For simplicity we assume that at full employment the unemployment rate is zero.¹⁰ Friedman’s notion of “maximum feasible output” is defined as production if all labor is employed, i.e. $Y_t^f = A_t(L^f)^\alpha$, which will be the equilibrium outcome if real wages, w^f , freely adjust:

$$w^f = \alpha(L^f)^{\alpha-1} = \alpha \left(\frac{Y_t^f}{A_t} \right)^{\frac{\alpha-1}{\alpha}}.$$

Consider a macroeconomic policy regime that controls nominal spending, $D_t = P_t Y_t$. At full employment the price level is

$$P_t = \frac{D_t}{A_t(L^f)^\alpha}.$$

Hence, variations in nominal spending have no effect on real output and employment, when wages are flexible. Nominal prices and wages are simply proportional to nominal spending. This environment, in other words, describes the vertical “wall” of the Slanted-L curve representing Friedman’s “ceiling” of a plucking model, or Keynes’ “fixed cake” at war times. Any increase in nominal spending has no effect on output or employment. Instead, it translates directly into inflation.

¹⁰In BE we model frictional unemployment via search and matching so that $u^f > 0$.

We capture the slanted leg of the L supply curve in two steps. First, we assume that workers refuse to accept a job that pays below the prevailing wage, W_t^{norm} , but are willing to accept any work that pays above it. This implies that the equilibrium nominal wage rate is¹¹

$$W_t = \max \{ W_t^{norm}, P_t w^f \}$$

where the first element, W_t^{norm} , captures the wage norm prevailing in the market. If $W_t^{norm} > P_t w^f$ then the equilibrium wage is above the full employment wage so that only part of available workers are employed. In this case labor is rationed, i.e., there is unemployment. If, however, $W_t^{norm} < P_t w^f$, firms bid up wages until all labor is employed.

Second, we assume that the wage norm takes the form:¹²

$$W_t^{norm} = (W_{t-1}(\Pi_t^e)^\gamma)^\lambda (P_t w^f)^{(1-\lambda)}$$

where Π_t^e is expected inflation; and the parameters γ and λ satisfy $0 \leq \gamma \leq 1$ and $0 \leq \lambda \leq 1$.

Keynesian downward nominal wage rigidity is captured by setting $\gamma = 0$ and $\lambda = 1$ so that workers refuse to work if the nomi-

¹¹As in Eggertsson, Mehrotra and Robbins (2019).

¹²In BE we generalize this concept within a search and matching framework and distinguish between new and existing wages.

nal wage is below the last period's wage.¹³ Our more flexible specification is more in line with Phillips idea that we summarized in the introduction. In general, we allow the wage norm to react to market conditions via $P_t w^f$ and inflation expectations.

The equilibrium real wage is then

$$(3) \quad w_t = \max \left\{ \left(w_{t-1} \frac{(\Pi_t^e)^\gamma}{\Pi_t} \right)^\lambda (w^f)^{1-\lambda}, w^f \right\}.$$

Using these ingredients we can characterize an L-shaped Phillips curve in a generic period t , preceded by a period $t-1$ in which wages are at some rate $w_{t-1} = \phi w^f$ for a constant $\phi > 0$.

Denote the natural logarithm of P_t, D_t, Y_t and A_t with lower cases and define $\pi_t = \ln \Pi_t$, $\pi_t^e = \ln \Pi_t^e$, and $v_t = \ln \phi - a_t/\lambda$. Then combining labor demand by the firms, the expression for real wages, the production function, and the definition of unemployment we obtain the L-curve:

$$(4) \quad u_t = u^f$$

if $d_t \geq p_t + y_t^f$ and

$$(5) \quad \pi_t = -\kappa u_t + v_t + \gamma \pi_t^e$$

if $d_t < p_t + y_t^f$. where $\kappa \equiv (1 - \alpha)/\lambda$. This pair of equations provide natural micro-foundations for the L-shape function shown in Figures (2) and (3). Equation (4) is the vertical part of the L while (5) is the slanted leg which becomes more slanted the higher is κ , i.e. the more flexible wages are.

REFERENCES

- Benigno, Pierpaolo and Gauti Eggertsson. 2023. "It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve." NBER Working Paper No. 31197.
- Beveridge William Henry. 1944. *Full employment in a free society*. London, England: Allen and Unwin.
- Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson. 2019. "A Plucking Model of Business Cycles." NBER Working Paper 26352.
- Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins. 2019. "A Model of Secular Stagnation: Theory and Quantitative Evaluation." *American Economic Journal: Macroeconomics*, 11 (1): 1-48.
- Friedman Milton. 1964. "Monetary Studies of the National Bureau." In *The National Bureau Enters Its 45th Year*, New York, NY: National Bureau of Economic Research, 7-25, 44th Annual Report.
- Friedman, Milton. 1993. "The "Plucking Model" of Business Fluctuations Revisited." *Economic Inquiry* 31: 171-177.
- Hall, George J., and Thomas J. Sargent. 2022. "Three World Wars: Fiscal-Monetary consequences." *Proceedings of the National Academy of Sciences* 119.18.
- Hazell, Jonathon, Juan Herrero, Emi Nakamura, Jón Steinsson. 2022 "The Slope of the Phillips Curve: Evidence from U.S. States." *The Quarterly Journal of Economics* 137(3): 1299-1344.
- Keynes, John Maynard. 1936. *The General Theory of Employment, Interest, and Money*. New York.
- Keynes, John Maynard. 1940. *How to Pay for the War*. MacMillan and Co., London.
- Phillips, Alban 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.
- Schmitt-Grohé, Stephanie, Martín Uribe, and Michael Woodford. 2022 *International Macroeconomics: A modern approach*. Princeton University Press.
- Tobin, James. 1972. "Inflation and Unemployment." *The American Economic Review*, 62(1): 1-18.

¹³For a good overview for the evidence on nominal wage rigidities see Schmitt-Grohe, Uribe and Woodford (2022),

ONLINE APPENDIX

Figure 1

Figure 1 shows scatter plots of the annual inflation rate and the labor-market tightness (v/u) for the United States and for the sample 2009 Q1–2023 Q2. Inflation rate (core) is at annual rates and computed using the quarterly CPI core. CPI core quarterly observations are the average of the relevant monthly observations. Data are from U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: All Items Less Food and Energy in U.S. City [CPILFESL], retrieved from FRED, Federal Reserve Bank of St. Louis. The variable v/u is computed as the ratio between the job openings and the unemployment level. Data are monthly. Accordingly, the quarterly series is the average of the relevant monthly observations. Job openings are from U.S. Bureau of Labor Statistics, Job Openings: Total Nonfarm [JTSJOL], retrieved from FRED, Federal Reserve Bank of St. Louis. Unemployment level is from U.S. Bureau of Labor Statistics, Unemployment Level [UNEMPLOY], retrieved from FRED, Federal Reserve Bank of St. Louis. Each plot shows also the fit of the linear regression model with 95% confidence bounds conditional on $v/u < 1$ and $v/u \geq 1$.

Figure 2

Figure 2 presents scatter plots of unemployment rate and core inflation rate at quarterly frequency for the period 2009 Q1 – 2023 Q3, with core inflation represented at annual rates. The ‘L’ function is shown with the vertical line indicating the average of the unemployment rate computed on the observations between its minimum and (minimum + 0.2). The flat segment of the ‘L’ function corresponds to the fitted line derived from linear regression (OLS) between inflation and unemployment for each country. Observations used to draw the vertical line are excluded from this regression analysis. Unemployment Rate: Harmonized Unemployment, Monthly Rates, Total, All Persons, obtained for each country (Australia, Canada, Germany, France, United Kingdom, Italy, Japan, United States) from the Organization for Economic Co-operation and Development via FRED, Federal Reserve Bank of St. Louis. Inflation Rate: Annual percentage change in the Core CPI, corresponding to the series ‘core CPI (standardized) SADJ,’ retrieved for each country from Thomson Reuters Datastream, respectively given by AUCCOR..E, CNCCOR..E, BDC-COR..E, UKCCOR..E, FRCCOR..E, ITCCOR..E, JPCCOR..E, USCCOR..E.

Figure 3

Figure 3 presents scatter plots of adjusted unemployment rate and core inflation rate at quarterly frequency for the period 2009 Q1 – 2023 Q3, with core inflation represented at annual rates. For each country, the adjusted unemployment rate is derived from the time-series data on the unemployment rate of each country by subtracting the unemployment rate at maximum employment, u_f , derived in Figure 2 to draw the vertical line and then adding the average unemployment rate (at maximum employment) across all countries. The ‘L’ function is illustrated, featuring a vertical line representing the average of the (adjusted) unemployment rate computed on the observations between its minimum and (minimum + 0.2). The flat segment of the ‘L’ model corresponds to the fitted line obtained through linear regression (OLS) between inflation and (adjusted) unemployment for all countries. Observations used to draw the vertical line are excluded from this regression analysis. The hyperbolic function corresponds to the non-linear least-squares estimates of

the model

$$\pi_{i,t} = a + b \left(\frac{1}{u_{i,t}^{dev}} \right)^c$$

with the following estimated coefficients $a = 1.3909$, $b = 1.3531e + 09$ and $c = 13.3963$.

Derivation of equation (5).

Consider labor demand

$$L_t^d = \left(\frac{1}{\alpha A_t} \frac{W_t}{P_t} \right)^{\frac{1}{\alpha-1}}$$

and take the log, to obtain:

$$(A1) \quad l_t^d = \frac{1}{\alpha - 1} (w_t - \ln \alpha - a_t),$$

in which $l_t^d = \ln L_t^d$, $w_t = \ln W_t/P_t$ and $a_t = \ln A_t$. Using equation (3) when the wage norm is binding, the (log) real wage is given by

$$w_t = \lambda w_{t-1} + \gamma \lambda \pi_t^e - \lambda \pi_t + (1 - \lambda) w_t^f,$$

in which $\pi_t^e = \ln \Pi_t^e$, $\pi_t = \ln \Pi_t$ and

$$(A2) \quad w_t^f = \ln \alpha + (\alpha - 1) \bar{l},$$

for $\bar{l} = \ln \bar{L}$. Consider the assumption

$$w_{t-1} = \ln \phi + \ln \alpha + (\alpha - 1) \bar{l},$$

it follows that

$$w_t = \lambda \ln \phi + \gamma \lambda \pi_t^e - \lambda \pi_t + w_t^f$$

Therefore we can write (A1) as

$$l_t^d = \frac{1}{\alpha - 1} (\lambda \ln \phi + \gamma \lambda \pi_t^e - \lambda \pi_t + w_t^f - \ln \alpha - a_t),$$

from which, using (A2), it follows

$$l_t^d - \bar{l} = \frac{1}{\alpha - 1} (\lambda \ln \phi + \gamma \lambda \pi_t^e - \lambda \pi_t - a_t).$$

The above equation can be also rewritten as

$$\pi_t = \frac{(1 - \alpha)}{\lambda} (l_t^d - \bar{l}) + \ln \phi + \gamma \pi_t^e - \frac{1}{\lambda} a_t.$$

Note that

$$u_t = 1 - \frac{L_t^d}{\bar{L}}$$

and that for small u , we can use the approximation $u_t \approx -\ln L_t^d/\bar{L}$. We can then write

$$\pi_t = -\kappa u_t + \gamma \pi_t^e + v_t,$$

which is equation (5), having defined

$$\begin{aligned}\kappa &\equiv \frac{(1 - \alpha)}{\lambda}, \\ v_t &\equiv \ln \phi - \frac{1}{\lambda} a_t.\end{aligned}$$