DOI: 10.1111/jmcb.12898

JORDI GALÍ

Insider–Outsider Labor Markets, Hysteresis, and Monetary Policy

I develop a version of the New Keynesian model with insider–outsider labor markets and hysteresis that can account for the high persistence of European unemployment. I study the implications of that environment for the design of monetary policy. The optimal policy calls for strong emphasis on (un)employment stabilization which a standard interest rate rule fails to deliver, with the gap between the two increasing in the degree of hysteresis. Two simple targeting rules are shown to approximate well the optimal policy. The properties of the model and effects of different policies are analyzed through the lens of the labor wedge and its components.

JEL codes: E24, E31, E32
Keywords: monetary policy tradeoffs, New Keynesian model, unemployment fluctuations, wage stickiness, wage Phillips curve.

Much discussion on the European unemployment problem has tended to focus on its high level, relative to the United States and other advanced economies. But a look at the path of the European unemployment rate over the past four decades points to another defining characteristic of that variable: its high persistence. The latter property has been emphasized by many authors, going back to Blanchard and Summer's influential *hysteresis* paper.[1]

I am grateful for comments and suggestions to Sanjay Chugh, Juanjo Dolado, Yuriy Gorodnichenko, Nils Gottfries, Andy Levin, Silvana Tenreyro, Antonella Trigari, and two anonymous referees, as well as seminar and conference participants at the NBER Summer Institute, Oxford-NY Fed workshop, Banque de France, EUI-ADEMU, IAE-MacFinRobods Workshop, European Central Bank (ECB)-Fed International Research Forum on Monetary Policy, EABCN-University of Zurich Conference, Colegio Carlo Alberto, Riksbank, Norges Bank, IIES, Einaudi Institute and National Bank of Ukraine. Christian Höynck, Cristina Manea, and Alain Schlaepfer provided excellent research assistance. I acknowledge Financial support from the CERCA Programme/Generalitat de Catalunya and the Severo Ochoa Programme for Centres of Excellence.

Jordi Galí *is at CREI, Universitat Pompeu Fabra, and Barcelona School of Economics* (*E-mail*: jgali@crei.cat).

Received January 12, 2018; and accepted in revised form July 21, 2020.

1. Blanchard and Summers (1986). See Ball (2009, 2014) and Blanchard (2018) for an empirical analysis of hysteresis across a number of OECD countries.

Journal of Money, Credit and Banking, Supplement to Vol. 54, No. S1 (February 2022)

Can the standard New Keynesian model, the workhorse framework of modern macroeconomics, account for the high persistence of European unemployment? The analysis below suggests that the answer is a negative one. In particular, I show that simulations of a (realistically calibrated) version of that model tend to generate fluctuations in the unemployment rate that are too short-lived relative to the data.

Motivated by the previous observation, I develop a variant of the New Keynesian model whose equilibrium properties can be more easily reconciled with the evidence on unemployment persistence. The modified model, inspired by the seminal work of Blanchard and Summers (1986), Gottfries and Horn (1987), and Lindbeck and Snower (1988), has two key distinctive features: (i) insider–outsider labor markets, and (ii) hysteresis. The first feature leads unions to give a disproportionate weight to a subset of the labor force– the *insiders*– when setting wages.[2] The second feature implies that the measure of insiders evolves endogenously over time as a function of employment. I show how a calibrated version of the modified model can generate a degree of unemployment persistence comparable to that observed in the data, in response to a variety of shocks, and under a "realistic" monetary policy rule that responds to inflation and output growth. Under such a rule, large negative deviations of employment from its efficient level do not bring about significant deflationary pressures, since unions largely ignore the interests of outsiders (the nonemployed) when setting wages. As a result, such deviations do not elicit a stabilizing response from the central bank, thus rendering them highly persistent, even in the face of transitory shocks.

Having made a case for insider–outsider labor markets and hysteresis as a potential explanation for the high persistence of European unemployment, I turn to the implications of that environment for the design of monetary policy. I derive the optimal policy with commitment, characterize its equilibrium implications, and compare them to those implied by the baseline simple interest rate rule. I show that the gap between the economy's response to different shocks under the two policies is increasing in the degree of hysteresis. Furthermore, the welfare gains induced by a switch from the simple rule to the optimal policy are shown to be large, the more so the higher the degree of hysteresis.

Motivated by the behavior of the insider–outsider economy under the optimal policy, I propose two simple targeting rules that get around the complexity of the optimal policy while approximating the latter's outcomes. The first rule, which I label *n-targeting,* seeks to fully stabilize employment (or, equivalently, wage inflation). The second, labeled *u-targeting*, keeps the unemployment rate constant. Both rules do an excellent job at approximating the outcomes of the fully optimal policy, including its welfare consequences.

Finally, I study the model's implications regarding the relation between hysteresis and labor wedge volatility, and how that relation is affected by the policy rule in

2. Wage setting by insiders is only one of the several mechanisms that have been proposed in the literature as a source of hysteresis in the unemployment rate. Other mechanisms include the lower employability of long-term unemployed, or the impact of unemployment on labor market institutions. See Blanchard (2018) for a discussion.

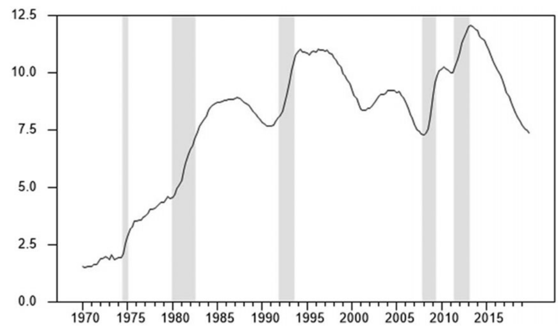

Fig. 1. Unemployment Rate in the Euro Area

place. The analysis yields a number of interesting insights. In particular, it shows that the adoption of the optimal policy (or the simple targeting rules that approximate it) reduce dramatically the volatility of the labor wedge by stabilizing one of its two components, namely, the wage markup. That reduction is larger the higher is the degree of hysteresis in the economy.

The analysis and findings below, though based on a highly stylized model, convey a message that is likely to remain valid in more general settings: monetary policies that focus excessively on inflation stabilization may fail to keep the economy close to its optimal level of activity in environments in which inflation is not responsive to such gaps, even if the latter are large and persistent. The economy with insider–outsider labor markets and strong hysteresis analyzed below provides a clear example of such an environment.

The paper is organized as follows. Section 1 presents evidence of high unemployment persistence in the euro area. Section 2 develops the New Keynesian model with insider–outsider labor markets. Section 3 analyzes the ability of that model to generate unemployment persistence, and contrasts it with the standard New Keynesian model. Section 4 derives the optimal monetary policy in the presence of insider–outsider labor markets, and characterizes the implied equilibrium. Section 5 analyzes the properties of two alternative simple targeting rules. Section 6 studies the connection between hysteresis and the labor wedge. Section 7 concludes.

1. EVIDENCE

The high persistence of European unemployment is apparent in Figure 1, which displays the unemployment rate for the euro area over the sample period 1970Q1–2019Q4, together with Centre for Economic Policy Research (CEPR)-dated recessions (as shaded areas).[3] The unemployment rate can be seen to wander about a

3. Source: ECB's Area Wide Model quarterly data set, originally constructed by Fagan et al. (2001) and subsequently updated by ECB. I am using update 18 of that data set, extended to include observations for 2018 and 2019 drawn from the ECB Statistical Data Warehouse.

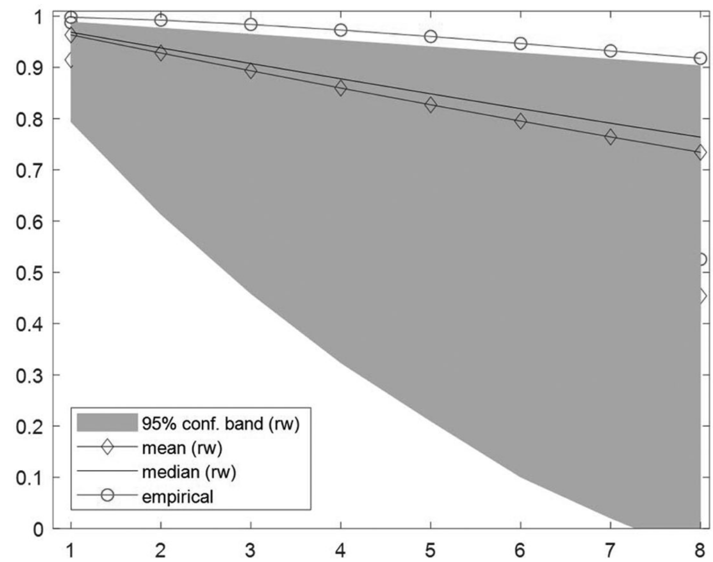

Fig. 2a. Euro Area Unemployment Rate: Autocorrelogram (1970Q1-2019Q4)

(seemingly) upward trend, showing variations that are smooth and highly persistent, and with no clear tendency to gravitate toward some constant long-run equilibrium value.[4] Each recession episode seems to pull the unemployment rate toward a new plateau, around which it appears to stabilize. The unemployment rate eventually declines as the economy recovers, or increases further if a new recession hits (as in 1980 or 2012). While the upward trend seems to have flattened out since the 1990s, fluctuations in the unemployment rate have remained highly persistent through the end of the sample.

The previous visual assessment is confirmed by the estimated autocorrelogram for the euro area unemployment rate over the 1970Q1–2019Q4 period, which is displayed in Figure 2a (line with circles). The estimated autocorrelations decay very slowly, a trademark of highly persistent time series. As a benchmark for comparison, the figure also shows the median and mean estimates (as well as 95% confidence bands) of the distribution of the estimated autocorrelogram for a random walk (without drift), based on 5,000 simulated time series with the same number of observations as our sample (200 observations). Note that the estimated autocorrelogram for the euro area unemployment rate lies outside the confidence interval, and well above the median and mean autocorrelations associated with the random walk, pointing to greater persistence than the latter process.

4. The latter observation is in stark contrast to the U.S. unemployment rate, which has fluctuated around a value close to 5% throughout the postwar period. See Galí (2015b) for a comparison of the persistence of the unemployment rate in the United States and the euro area.

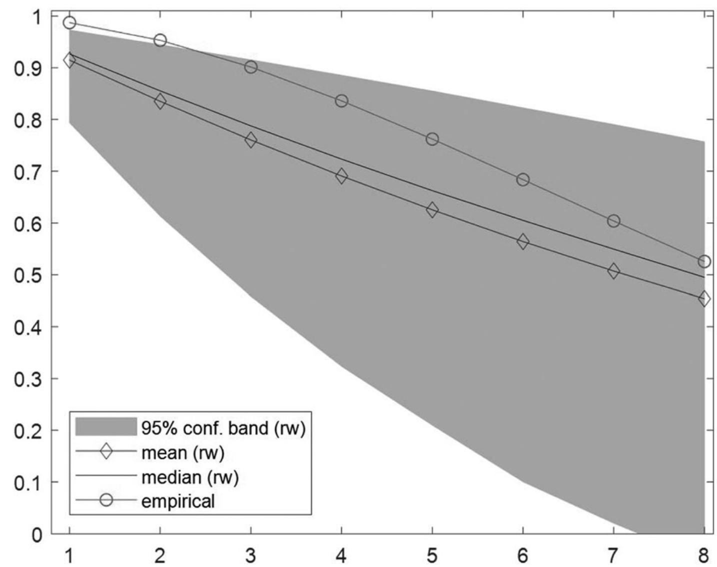

Fig. 2b. Euro Area Unemployment Rate: Autocorrelogram (1999Q1-2019Q4)

TABLE 1
ADF Unit Root Tests

Sample Period	*1 Lag*	*4 Lags*
1970Q1–2019Q4	−2.42 (−2.87)	−2.28 (−2.87)
1999Q1–2019Q4	−2.36 (−2.89)	−1.70 (−2.89)

Note: *t*-Statistics of augmented Dickey-Fuller tests (with intercept) for the null of a unit root in the unemployment rate. Shown in brackets are the critical values for the null of a unit root, with a 5 % significance (adjusted for sample size).

Measures of persistence of the unemployment rate over the full sample may be distorted by the large upward adjustment experienced by that variable in the 1970s and 1980s. When I restrict the sample to the monetary union period (1999Q1–2019 Q4; henceforth, the European Monetary Union (EMU) period) the estimated autocorrelogram comes down quite a bit, but keeps pointing to very high persistence in the unemployment rate, with most of the estimates lying well inside the confidence interval associated with the random walk (see Figure 2b).[5]

The outcome of unit root tests applied to the euro area unemployment rate tends to accord with the previous evidence. In particular, and as reported in Table 1, an augmented Dickey–Fuller (ADF) test (with 1 and 4 lags) does not reject the null of

5. The statistics for the random walk shown in Figure 2b are based on 5,000 simulations of 84 observations each.

unit root in the unemployment rate at a 5% significance level for the full sample period. An identical outcome obtains when I start the sample period in 1999Q1.

The evidence above makes it clear that the unemployment rate in the euro area displays very high persistence. Here I do not take a stance as to whether it has or does not have a unit root. Yet, it is clear that the observed persistence is not far from that of a random walk, even during the more recent EMU period. As I argue below, the standard New Keynesian model has a hard time accounting for that persistence. Next I propose a variation on that model that has the potential to generate highly persistent unemployment fluctuations.

2. A NEW KEYNESIAN MODEL WITH INSIDER–OUTSIDER LABOR MARKETS

In the present section, I modify an otherwise standard New Keynesian framework by embedding in it a model of wage setting along the lines of insider–outsider models of the labor market. With the exception of the assumptions on wage setting, the environment is similar to that described in Galí (2015a, chapter 7), in which the household block of the New Keynesian model is reformulated in order to generate a meaningful concept of unemployment.

2.1 Households

I assume a large number of identical households. Each household has a continuum of members distributed uniformly over the unit square. Each member is indexed by a pair $(j, s) \in [0, 1] \times [0, 1]$. The first index, $j \in [0, 1]$, represents the type of labor service ("occupation") that she is specialized in. The second index, $s \in [0, 1]$, determines her disutility from work. The latter is given by χs^{φ} if she is employed and zero otherwise, where $\chi > 0$ and $\varphi > 0$ are exogenous parameters. Employed individuals work a constant number of hours. Employment for each occupation, $\mathcal{N}_t(j) \in [0, 1]$, is demand-determined and taken as given by each household, which allocates it to the members with the lowest work disutility among those specialized in the given occupation, that is, $s \in [0, \mathcal{N}_t(j)]$. Full risk sharing within the household is assumed. Given the separability of preferences, this implies the same level of consumption for all household members, independently of their occupation or employment status.

The household's period utility is given by the integral of its members' utilities:

$$
\begin{aligned}
U(C_t, \{\mathcal{N}_t(j)\}; Z_t) &\equiv \left(\log C_t - \chi \int_0^1 \int_0^{\mathcal{N}_t(j)} s^{\varphi} ds dj \right) Z_t \\
&= \left(\log C_t - \chi \int_0^1 \frac{\mathcal{N}_t(j)^{1+\varphi}}{1+\varphi} dj \right) Z_t ,
\end{aligned}
$$

where $C_t \equiv (\int_0^1 C_t(i)^{1-\frac{1}{\epsilon_{p,t}}} di)^{\frac{\epsilon_{p,t}}{\epsilon_{p,t}-1}}$ is a consumption index, with $C_t(i)$ being the quantity consumed of good i, for all $i \in [0, 1]$. Parameter $\epsilon_{p,t}$ denotes the elasticity of substitution, which is (possibly) time-varying. The exogenous preference shifter $z_t \equiv \log Z_t$ is assumed to follow an $AR(1)$ process:

$$z_t = \rho_z z_{t-1} + \varepsilon_t^z,$$

where $\rho_z \in [0, 1]$ and ε_t^z is a white noise process.

Each household seeks to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t U(C_t, \{\mathcal{N}_t(j)\}; Z_t)$$

subject to a sequence of flow budget constraints given by

$$\int_0^1 P_t(i)C_t(i)di + Q_t B_t \leq B_{t-1} + \int_0^1 W_t(j)\mathcal{N}_t(j)dj + D_t, \tag{1}$$

where $P_t(i)$ is the price of good i, $W_t(j)$ is the nominal wage for occupation j, B_t represents purchases of a nominally riskless one-period discount bond paying one unit of account ("money"), Q_t is the price of that bond, and D_t denotes dividends from the ownership of firms.[6] $\beta \in [0, 1]$ is the household's discount factor.

Independently of the nature of wage setting, the household's problem above gives rise to two types of optimality conditions: a set of optimal demand schedules for each consumption good and a standard intertemporal optimality condition (or Euler equation). Those take the familiar form (using lower case letters to denote logs):

$$c_t(i) = -\epsilon_{p,t}(p_t(i) - p_t) + c_t$$

for all $i \in [0, 1]$, and

$$c_t = E_t\{c_{t+1}\} - (i_t - E_t\{\pi_{t+1}^p\} - \rho) + (1 - \rho_z)z_t,$$

where $\pi_t^p \equiv p_t - p_{t-1}$ denotes price inflation, and $\rho \equiv -\log \beta$ is the discount rate.[7]

Following Galí (2011, 2015a), I define $L_t(j)$ as the marginal participant for occupation j, determined by condition:

$$\frac{1}{C_t} \frac{W_t(j)}{P_t} = \chi L_t(j)^{\varphi}.$$

6. The above sequence of period budget constraints is supplemented with a solvency condition that prevents the household from engaging in Ponzi schemes.

7. See Woodford (2003) or Galí (2015b) for a derivation of these and other equilibrium conditions unrelated to the labor market.

Taking logs and aggregating over all occupations one can derive the following equation for aggregate participation:

$$\omega_t = c_t + \varphi l_t + \xi, \tag{2}$$

where $\omega_t \equiv w_t - p_t$ is the average (log) real wage, $w_t \equiv \int_0^1 w_t(j)dj$ is the average (log) nominal wage, $l_t \equiv \int_0^1 l_t(j)dj$ can be interpreted as the (log) labor force (or participation), and $\xi \equiv \log \chi$.

The unemployment rate can be (naturally) defined as

$$u_t \equiv l_t - n_t, \tag{3}$$

where $n_t \equiv \int_0^1 n_t(j)dj$ is (log) aggregate employment, which is demand determined. The unemployment rate thus defined satisfies the relation:

$$\mu_t^w = \varphi u_t, \tag{4}$$

where μ_t^w is defined by

$$\mu_t^w \equiv \omega_t - (c_t + \varphi n_t + \xi), \tag{5}$$

that is, μ_t^w is the gap between the average (log) real wage and the average (log) marginal rate of substitution between employment and consumption, $mrs_t \equiv c_t + \varphi n_t + \xi$. Henceforth, I refer to μ_t^w as the *wage markup* and assume $\mu_t^w \geq 0$.[8] Note also that, under the above assumptions, relation (4) holds independently of how wages are set.

2.2 FIRMS

I assume the existence of a continuum of differentiated goods $i \in [0, 1]$, each produced by a monopolistic competitor, with a production function:

$$Y_t(i) = A_t N_t(i)^{1-\alpha}, \tag{6}$$

where $Y_t(i)$ denotes the output of good i, A_t is an exogenous technology parameter common to all firms, and $N_t(i)$ is a Constant Elasticity of Substitution (CES) function of the quantities of the different types of labor services employed by firm i, whose elasticity of substitution is given by ϵ_w. Cost minimization by firms gives rise to the labor demand schedule (13) introduced below. Technology is assumed to follow an $AR(1)$ process in logs, that is,

$$a_t = \rho_a a_{t-1} + \varepsilon_t^a,$$

8. A negative wage markup would imply that at least some individuals are working for a wage below their relevant marginal rate of substitution, which would amount to forced labor.

where $a_t \equiv \log A_t$ and $\rho_a \in [0, 1]$.

Price setting is staggered *à la* Calvo, with a constant fraction θ_p of firms that keep prices unchanged in any given period. Aggregation of price-setting decisions, gives rise to an inflation equation of the form (around a zero-inflation steady state)

$$\pi_t^p = \beta E_t\{\pi_{t+1}^p\} - \lambda_p(\mu_t^p - x_t), \tag{7}$$

where

$$\mu_t^p \equiv \log(1-\alpha) + a_t - \alpha n_t - \omega_t - \log(1-\tau) \tag{8}$$

is the average price markup, $\lambda_p \equiv \frac{(1-\theta_p)(1-\beta\theta_p)}{\theta_p}\frac{1-\alpha}{1-\alpha+\alpha\epsilon_p} > 0$, τ is a constant wage subsidy, and $x_t \equiv \log \frac{\epsilon_{p,t}}{\epsilon_{p,t}-1}$ is the desired or *natural* price markup.[9] The latter is assumed to follow an $AR(1)$ process with mean $\mu^p \equiv \log \frac{\epsilon_p}{\epsilon_p - 1}$ and autoregressive coefficient $\rho_x \in [0, 1)$.

Note that one can rewrite the price markup gap as follows:

$$\mu_t^p - x_t = -\frac{\alpha}{1-\alpha}\widetilde{y}_t - \widetilde{\omega}_t, \tag{9}$$

where $\widetilde{y}_t \equiv y_t - y_t^n$ and $\widetilde{\omega}_t \equiv \omega_t - \omega_t^n$ denote, respectively, the *output* and *wage gaps*, defined as the log deviation between the each of those variables and their natural (i.e., flexible price and wage) counterparts, y_t^n and ω_t^n (to be determined below). By combining (7) and (9), we obtain a version of the New Keynesian Phillips curve:

$$\pi_t^p = \beta E_t\{\pi_{t+1}^p\} + \varkappa_p\widetilde{y}_t + \lambda_p\widetilde{\omega}_t, \tag{10}$$

where $\varkappa_p \equiv \frac{\alpha\lambda_p}{1-\alpha}$.

Goods market equilibrium requires that $c_t = y_t$ for all t, which combined with the household's Euler equation implies:

$$y_t = E_t\{y_{t+1}\} - (i_t - E_t\{\pi_{t+1}^p\} - \rho) + (1-\rho_z)z_t. \tag{11}$$

Given equilibrium output, employment is given by

$$(1-\alpha)n_t = y_t - a_t. \tag{12}$$

9. See chapter 3 in Galí (2015a) for a derivation of (7). Note that the average (log) marginal cost ψ_t is given by

$$\psi_t = w_t + \log(1-\tau) - \log(1-\alpha) - a_t + \alpha n_t.$$

The expression for the average price markup follows from combining the previous expression with the definition $\mu_t^p \equiv p_t - \psi_t$.

2.3 Wage Setting

Next I turn to a description of wage setting. First, I describe the features of wage setting that are common to both versions of the New Keynesian model analyzed below. In both cases, I adopt the Calvo model of staggered nominal wage setting, which assumes that a constant fraction $1 - \theta_w$ of occupations (or the unions representing them), drawn randomly from the set of existing occupations, are allowed to reset their nominal wage in any given period. When setting the new wage $w_t^*(j)$, a union representing occupation j takes into account current and (expected) future demand for the labor services of its members, as given by:

$$n_{t+k|t}(j) = -\epsilon_w(w_t^*(j) - w_{t+k}) + n_{t+k} \tag{13}$$

for $k = 1, 2, 3, \ldots$where $n_{t+k|t}(j)$ denotes period $t + k$ (log) employment for occupation j whose wage has been reset for the last time in period t, and n_{t+k} is (log) aggregate employment in period $t + k$. Note that $\epsilon_w > 1$ is the wage elasticity of labor demand. Equation (13) can be derived from firms' cost minimization.[10] The wage set by the union for a specific occupation is assumed to be enforceable, that is, it cannot be underbid by the unemployed.

As a result, the evolution of the average (log) nominal wage is described by the difference equation:

$$w_t = \theta_w w_{t-1} + (1 - \theta_w)w_t^*, \tag{14}$$

where $w_t^* \equiv (1 - \theta_w)^{-1} \int_{j \in \mathcal{J}_t} w_t^*(j)dj$, where $\mathcal{J}_t \subset [0, 1]$ represents the subset of occupations resetting their wage in period t. Thus, w_t^* is the average newly set wage in period t, expressed in logs.[11]

The previous features are common to the two models of wage setting considered below. Next I describe the features that are specific to each model.

Wage setting in the standard new Keynesian model. In the standard New Keynesian model (e.g., Erceg et al. 2000) it is assumed that, when resetting the wage, each union seeks to maximize the utility of the representative household, to which all union members (employed or unemployed) belong.[12] This gives rise to a (log-linearized) wage-setting rule of the form:

$$w_t^* = \mu^w + (1 - \beta\theta_w)\sum_{k=0}^{\infty}(\beta\theta_w)^k E_t\{\underline{w}_{t+k|t}\}, \tag{15}$$

10. See, for example, chapter 6 in Galí (2015a).

11. The previous equation, like others used in the present analysis, is log-linear approximations of the exact equilibrium conditions in a neighborhood of a zero-inflation steady state. See Galí (2015a) for detailed derivations.

12. See, for example, Galí (2015a, chapter 6) for a discussion of the union's problem and a derivation of the optimal wage-setting rule.

where $\underline{w}_{t+k|t} \equiv p_{t+k} + c_{t+k} + \varphi n_{t+k|t} + \xi$ is the "target wage" in $t+k$ for a union that is resetting its wage in period t, and where $\mu^w \equiv \log \frac{\epsilon_w}{\epsilon_w - 1}$ is the desired or *natural wage markup*, which is assumed to be constant. It is easy to show that the latter is the wage markup that any union (acting independently) would choose if wages were fully flexible (i.e., adjusted every period), given a labor demand schedule with a constant wage elasticity $\epsilon_w > 1$.

Combining (14) and (15) allows one to derive the wage inflation equation:

$$\pi_t^w = \beta E_t\{\pi_{t+1}^w\} - \lambda_w(\mu_t^w - \mu^w), \tag{16}$$

where $\pi_t^w \equiv w_t - w_{t-1}$ denotes wage inflation, $\mu_t^w = \omega_t - (c_t + \varphi n_t + \xi)$ is the average wage markup introduced above, and $\lambda_w \equiv \frac{(1-\theta_w)(1-\beta\theta_w)}{\theta_w(1+\epsilon_w\varphi)} > 0$.

One can combine (16) with the relation between the age markup and the unemployment rate (4) derived above to obtain the following New Keynesian wage Phillips curve, linking wage inflation and unemployment:

$$\pi_t^w = \beta E_t\{\pi_{t+1}^w\} - \lambda_w\varphi(u_t - u), \tag{17}$$

where $u \equiv \frac{\mu^w}{\varphi}$ is the *natural* rate of unemployment, that is, the unemployment rate that would obtain under flexible wages (and, hence, a constant wage markup μ^w).

It is easy to see that the previous model of wage setting guarantees the tendency of the unemployment rate to gravitate toward its natural rate. Thus, equation (17) implies that in the face of a current or anticipated high (low) unemployment rate (relative to the natural rate u), wages will tend to decrease (increase), thus lowering (raising) marginal cost, inflation, and the interest rate (through a policy rule like the one introduced below) and, as a result, boosting (dampening) output and reducing (increasing) the unemployment rate. Greater wage flexibility (i.e., a low θ_w) would be reflected in a larger λ_w value and, *ceteris paribus*, a faster convergence of the unemployment rate to its natural counterpart.

The implied stationarity of the unemployment rate becomes apparent by noting that the wage-setting rule (15) can be equivalently rewritten as

$$(1 - \beta\theta_w)\sum_{k=0}^{\infty}(\beta\theta_w)^k E_t\left\{\mu_{t+k|t}^w\right\} = \mu^w, \tag{18}$$

where $\mu_{t+k|t}^w \equiv w_t^* - \underline{w}_{t+k|t}$ is the prevailing markup k periods after the wage is set and conditional on the latter remaining in place. Thus, when reoptimizing, unions choose a wage such that a weighted average of the wage markups expected to prevail over the life of the newly set wage equals the desired or frictionless wage markup μ^w. Since all wage-setting unions behave in a similar way, the economy's average wage markup μ_t^w will fluctuate about μ^w. Accordingly, and given (4), the unemployment rate will display mean-reverting fluctuations about the constant natural rate u.

For future reference, I derive the natural level of employment in the standard New Keynesian model by combining (5) and (8) after imposing $\mu_t^w = \mu^w$ and $\mu_t^p = x_t$ for all t, the wage and price markups prevailing in the equilibrium with flexible wages and prices. This yields

$$n_t^n = \frac{\log(1-\alpha) - \xi - (x_t + \mu^w) - \log(1-\tau)}{1+\varphi}$$

with the corresponding natural levels of output and wages given by $y_t^n = a_t + (1 - \alpha)n_t^n$ and $\omega_t^n = \log(1-\alpha) + a_t - \alpha n_t^n - x_t - \log(1-\tau)$.

An insider–outsider model of wage setting. Insider–outsider models of the labor market, as originally developed in Blanchard and Summers (1986), Gottfries and Horn (1987), and Lindbeck and Snower (1988), emphasize the segmentation of the labor force between insiders and outsiders and the dominant role of the former in wage determination. In the words of Blanchard and Summers (1986):

> "...there is a fundamental asymmetry in the wage-setting process between insiders who are employed and outsiders who want jobs. Outsiders are disenfranchised and *wages are set with a view to ensuring the jobs of insiders*. Shocks that lead to reduced employment change the number of insiders and thereby change the subsequent equilibrium wage rate, giving rise to hysteresis..."

Here, I adapt the Blanchard and Summers (1986) version of the insider–outsider model in order to make it consistent with the Calvo wage-setting formalism so that it can be readily embedded in the standard New Keynesian model. thus, I assume that a union resetting the wage for occupation j in period t chooses a wage, $w_t^*(j)$, such that the following condition is satisfied:

$$(1-\beta\theta_w)\sum_{k=0}^{\infty}(\beta\theta_w)^k E_t\left\{n_{t+k|t}(j)\right\} = n_t^*(j) \tag{19}$$

with $n_{t+k|t}(j)$ given by (13), for $k = 0, 1, 2 \ldots$. In words, the wage is set so that a weighted average of expected employment in occupation j conditional on the wage remaining effective equals some employment *target* $n_t^*(j)$. The latter target can be interpreted as representing the measure of *insiders* in occupation j.[13]

Thus, and in contrast with the standard New Keynesian model in which, when setting a nominal wage, unions target a weighted average of expected wage markups, in the insider–outsider model they target an identical weighted average of expected employment levels.

I follow Blanchard and Summers (1986) and assume that the measure of insiders (and, hence, the employment target) in any given occupation j evolves over time according to the difference equation:

$$n_t^*(j) = \gamma n_{t-1}(j) + (1-\gamma)n^*, \tag{20}$$

13. A possible justification for this type of behavior may involve some deviation from perfect consumption risk sharing within households, with each individual's consumption being related to her individual wage income. A formal treatment is beyond the scope of the present model.

where n^* is the union's *long-run target* for (log) employment, which is assumed to be common across occupations. Parameter $\gamma \in [0, 1]$ determines the extent to which actual employment in a given occupation changes the measure of insiders. This is the property referred to in the literature as *hysteresis*.

Beyond the particular specification chosen, the motivation behind that assumption is the notion that the concerns of employed workers are given a disproportionate weight in the bargaining of wages. This may be the case for a variety of reasons: they are more likely to participate or remain close to the bargaining process, they are the ones with the ability to strike and hence are an important source of the union's bargaining power, they are more likely to pay their union fees, and so on. On the other hand, those who are unemployed are, to some extent, disenfranchised from the wage-setting process.

When wages are fully flexible ($\theta_w = 0$), all unions reset wages every period, with (19) implying $n_t^n = n_t^n(j) = n_t^*$ for all $j \in [0, 1]$ and all t. Under the assumption that initial employment is at its target level (i.e., $n_{-1} = n^*$), it follows from (20) that $n_t^n = n^*$ for all t, that is, the *natural* level of employment in the insider–outsider model is constant and equal to the target employment n^* at all times. Accordingly, the corresponding natural levels of output and wage are given $y_t^n = a_t + (1-\alpha)n^*$ and $\omega_t^n = \log(1-\alpha) + a_t - \alpha n^* - x_t - \log(1-\tau)$ for all t.

Next I derive an equation describing the evolution of wage inflation in the insider–outsider model when wages are sticky. Substituting (13) into (19) yields the wage-setting rule:

$$w_t^*(j) = -\frac{1}{\epsilon_w} n_t^*(j) + (1-\beta\theta_w) \sum_{k=0}^{\infty} (\beta\theta_w)^k E_t \left\{ w_{t+k} + \frac{1}{\epsilon_w} n_{t+k} \right\}. \tag{21}$$

Averaging (21) over all the unions resetting their wage in period t (i.e., over all $j \in \mathcal{J}_t$) and letting $n_t^* \equiv (1-\theta_w)^{-1} \int_{j \in \mathcal{J}_t} n_t^*(j) dj$ is the average (log) employment target for those unions we have:

$$w_t^* = -\frac{1}{\epsilon_w} n_t^* + (1-\beta\theta_w) \sum_{k=0}^{\infty} (\beta\theta_w)^k E_t \left\{ w_{t+k} + \frac{1}{\epsilon_w} n_{t+k} \right\}, \tag{22}$$

$$n_t^* = \gamma n_{t-1} + (1-\gamma) n^*. \tag{23}$$

Combining (22) with (14) yields (after some algebra) the following wage inflation equation for the insider–outsider economy:

$$\pi_t^w = \beta E_t\{\pi_{t+1}^w\} + (1-\gamma)\lambda_n(1-\beta\theta_w)\widehat{n}_t + \gamma\lambda_n \Delta n_t, \tag{24}$$

where $\widehat{n}_t \equiv n_t - n^*$ and $\lambda_n \equiv \frac{1-\theta_w}{\theta_w \epsilon_w}$, which is decreasing in the degree of wage rigidities (θ_w) and the wage elasticity of labor demand (ϵ_w). Note that both the (log) employment change Δn_t and the deviation of employment from steady state, $\widehat{n}_t$, are the drivers of fluctuations in wage inflation. The larger is the degree of hysteresis γ, the smaller is the weight of $\widehat{n}_t$ and the larger that of Δn_t in the determination of wage inflation. To the extent that the responsiveness of wage inflation to the degree of labor market tightness $\widehat{n}_t$ acts as a stabilizing device, one may conjecture from the previous finding that, *ceteris paribus*, a large value of γ will tend to weaken that stabilizing mechanism, leading to more persistent fluctuations in employment (and unemployment). This conjecture is confirmed by the simulations reported below.

Two limiting cases of the model are worth pointing out. When $\gamma = 0$, then we have

$$\pi_t^w = \beta E_t\{\pi_{t+1}^w\} + \lambda_n(1 - \beta\theta_w)\widehat{n}_t \tag{25}$$

with only the employment gap $\widehat{n}_t$ emerging as a driving variable. At the opposite extreme, we have the case of $\gamma = 1$. In that case, singled out in Blanchard and Summers (1986) as one of special interest, the set of insiders corresponds to the workers employed at the end of the previous period, with no weight attached to the unemployed in the wage-setting decision. Accordingly, equation (24) collapses to

$$\pi_t^w = \beta E_t\{\pi_{t+1}^w\} + \lambda_n \Delta n_t \tag{26}$$

with the *change* in employment being the only driving force of wage inflation. Accordingly, the *level* of employment itself does not generate the pressure on wage inflation that would act as a self-correcting mechanism. As shown below, stabilization of (un)employment in that extreme case will require a more direct policy response, one not intermediated by inflation.

Wage inflation equation (24), like its counterpart in the standard New Keynesian model (17), abstracts from many features that are likely to be relevant in actual economies, including indexation and nonlinearities (e.g., resulting from downward nominal wage rigidities). The simplicity of both formulations, however, facilitates the optimal monetary policy analysis below, while focusing on the qualitative differences between the two models. Despite that simplicity, evidence favorable to the presence of a hysteresis effect in European wage inflation dynamics is reported in a companion paper (Galí 2015b). In particular, I show how equation (26), which assumes full hysteresis, can account reasonably well for the evolution of wage inflation in the euro area over the 1999–2014 period. By contrast, wage inflation equation (17), associated with the standard New Keynesian model, is shown to imply fluctuations in wage inflation of an order of magnitude larger than those observed during that period.

2.4 *Efficient Allocation, Steady State, and Equilibrium Dynamics*

Efficient allocation. The *efficient* allocation, that is, the one that maximizes households' utility given the economy's resource constraints, is easy to characterize. Employment is identical across firms and occupations, and all goods are consumed in identical quantities. The efficiency condition equating the marginal rate of substitution and the marginal product of labor implies a constant optimal level of employment, given by:

$$n_t^e \equiv \frac{\log(1-\alpha) - \xi}{1+\varphi} \equiv n^e.$$

The efficient level of output is thus given by

$$y_t^e \equiv a_t + (1-\alpha)n^e,$$

that is, it responds one for one to changes in technology, while being invariant to preference and/or markup shocks.

The efficient allocation provides a useful benchmark in some of the analyses below.

Steady state. The (zero inflation) steady state of the decentralized economy is not invariant to the assumed wage-setting environment. Thus in the standard model steady-state employment is given by

$$n \equiv \frac{\log(1-\alpha) - (\xi + \mu^w + \mu^p) - \log(1-\tau)}{1+\varphi},$$

while the steady-state unemployment rate satisfies:

$$u = \frac{\mu^w}{\varphi}.$$

Note that steady-state efficiency can be attained by setting a wage subsidy $\tau = 1 - \exp\{-(\mu^w + \mu^p)\} > 0$.

In the model with insider–outsider labor markets with $\gamma \in [0, 1)$ the (zero inflation) steady state is given by the long-run employment target n^*, which is assumed to be common across unions. Thus, $n = n^*$ in that case. The steady-state unemployment rate is given by

$$\begin{aligned} u &= \frac{\mu^w}{\varphi} \\ &= \frac{1}{\varphi}[(\log(1-\alpha) - \alpha n^* - \mu^p - \log(1-\tau)) - (y + \varphi n^* + \xi)] \\ &= \frac{1}{\varphi}[(\log(1-\alpha) - \log(1-\tau) - \mu^p - \xi - (1+\varphi)n^*]. \end{aligned}$$

On the other hand, when $\gamma = 1$ there is no well-defined (zero inflation) steady state: any (constant) level of employment is in principle consistent with a steady state. The welfare analysis below is restricted to the case in which a well-defined steady state exists (i.e., $\gamma \in [0, 1)$), with the employment target assumed to correspond to the efficient level of employment, that is, $n = n^* = n^e$. In that special case, we have

$$u = -\frac{\log(1-\tau) + \mu^p}{\varphi}, \tag{27}$$

which, in turn, requires $\tau \geq 1 - \exp\{-\mu^p\}$ in order to rule out "forced labor" (i.e., negative unemployment) in the steady state.

Equilibrium dynamics. Equations (2), (3), (10), (11), and (12), the identity

$$\omega_t \equiv \omega_{t-1} + \pi_t^w - \pi_t^p \tag{28}$$

and wage inflation equation (17) (standard model) or (24) (insider–outsider model), together with the corresponding expressions for y_t^n and ω_t^n define the nonpolicy block of the model. In order to close the model one must supplement the previous equilibrium conditions with a description of a monetary policy rule that (directly or indirectly) determines the nominal interest rate i_t. For the baseline simulations below, I assume an interest rate rule of the form:

$$i_t = \phi_i i_{t-1} + (1-\phi_i)[\rho + \phi_\pi \pi_t^p + \phi_y \Delta y_t]. \tag{29}$$

For values of ϕ_i close to unity (as assumed in the simulations below), the previous rule is similar to the one proposed in Orphanides (2006) and Smets (2010) as a good approximation to ECB policy. Below, I used their proposed calibration and interpret (29) as a stylized but "realistic" simple interest rate rule.

3. UNEMPLOYMENT PERSISTENCE IN THE NEW KEYNESIAN MODEL

Can the New Keynesian model account for the observed persistence of euro area unemployment? In the present section, I try to provide an answer to that question by simulating a calibrated version of the New Keynesian model under the two wage-setting regimes considered (standard and insider–outsider), and use the generated time series to determine the persistence of unemployment, which are then compared to the corresponding properties in the data.

3.1 Calibration

Table 2 lists the baseline settings for the model parameters used in the simulations. The relevant period is taken to be a quarter. Parameters ϵ_p is set to 3.8. That value is associated with a steady-state price markup of 35%, and is consistent with the evidence used in the calibration of the ECB's New Area-Wide Model (NAWM) of Christoffel et al. (2008). Given that setting, a value of $1/4$ for parameter α is roughly

TABLE 2
CALIBRATION

φ	Curvature of labor disutility	3.4
β	Discount factor	0.99
α	Decreasing returns to labor	0.25
ϵ_w	Elasticity of substitution (labor)	4.3
ϵ_p	Elasticity of substitution (goods)	3.8
θ_p	Calvo index of price rigidities	0.75
θ_w	Calvo index of wage rigidities	0.75
ϕ_i	Lagged interest rate coefficient	0.9
ϕ_π	Inflation coefficient	1.5
ϕ_y	Output growth coefficient	0.5
ρ_a	Technology shock: Autoregressive coefficient	0.89
ρ_x	Markup shock: Autoregressive coefficient	0.39
ρ_z	Demand shocks: Autoregressive coefficient	0.91

consistent with the observed average labor income share in the euro area.[14] Parameter ϵ_w is set to 4.3, again based on the evidence discussed in Christoffel et al. (2008). Given that setting for ϵ_w, and using the approach developed in Galí (2011), a value of φ equal to 3.4 can be shown to be consistent with a steady-state unemployment rate of 7.8%, the average unemployment rate in the euro area over the 1970–2019 period.[15] As to the discount factor, I set $\beta = 0.99$, as is common practice in the business cycle literature. I set the Calvo wage and price stickiness parameters, θ_p and θ_w, to 0.75, which implies an average duration of individual wages and prices of four quarters. That setting is roughly consistent with the bulk of the microevidence for the euro area (see, e.g., Álvarez et al. (2006) and ECB (2009)). As to the interest rate rule coefficients, I assume $\phi_\pi = 1.5$, $\phi_y = 0.5$, and $\phi_i = 0.9$. That calibration is close to the one proposed in Orphanides (2006) and Smets (2010) as a good approximation to ECB policy. Finally, I calibrate the persistence and volatility of the exogenous driving forces based on the estimates of the NAWM in Christoffel et al. (2008). Specifically, I set $\rho_a = 0.89$, $\rho_x = 0.39$, and $\rho_z = 0.91$, which correspond, respectively, to their estimates of the autoregressive coefficient of transitory technology shocks, (domestic) price markup shocks, and (domestic) risk premium shocks.[16]

14. Note that in the steady state the following relation holds:

$$\frac{WN}{PY} = (1-\alpha)\left(1 - \frac{1}{\epsilon_p}\right).$$

15. Galí (2015a) shows that the φ, ϵ_w and the steady-state unemployment rate u are related according to equation:

$$\varphi u = \log \frac{\epsilon_w}{\epsilon_w - 1}.$$

Interestingly, the resulting setting for φ is nearly identical to the calibrated value in the NAWM of Christoffel et al. (2008).

16. The analysis below focuses on the persistence of unemployment and is always conditional on a specific shock, so I do not need to calibrate the variance of the driving forces. It should also be clear that

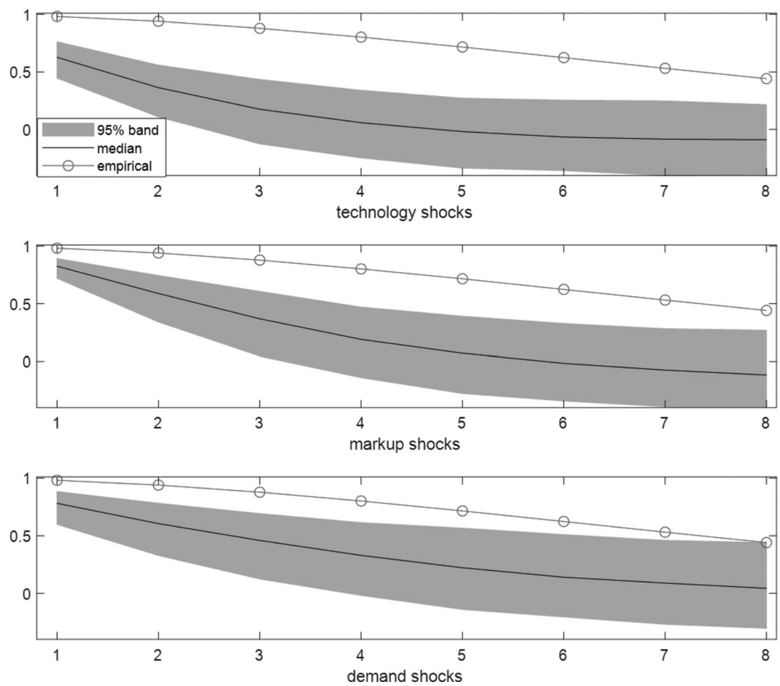

Fig. 3a. Unemployment Persistence–Standard Model versus Data

3.2 Unemployment Persistence in the Standard New Keynesian Model

I simulate the *standard* New Keynesian model under the above baseline calibration to evaluate its ability to generate the degree of unemployment persistence observed in euro area data. In what follows, I take the estimated autocorrelogram for the unemployment rate over the EMU period (1999Q1–2019Q4) as a rough "target" to match. More specifically, I generate 500 draws of 84 observations each, conditional on each of the three exogenous shocks being the driver of fluctuations (one at a time). For each draw, I estimate the autocorrelation of the simulated unemployment rate, for up to 8 lags. Figure 3a shows, for each type of shock, the *median* as well as the (centered) 95% band of the distribution of estimated unemployment autocorrelations (across 500 simulations) of the standard model under the baseline calibration. In addition, each plot also displays, as a benchmark, the corresponding empirical autocorrelogram for the EMU period.

The findings reported in Figure 3a suggest that the standard New Keynesian model has clear difficulties to match the persistence of European unemployment,

I restrict the analysis to the three assumed shocks for simplicity, with no claim that they span the range of relevant shocks in the euro area.

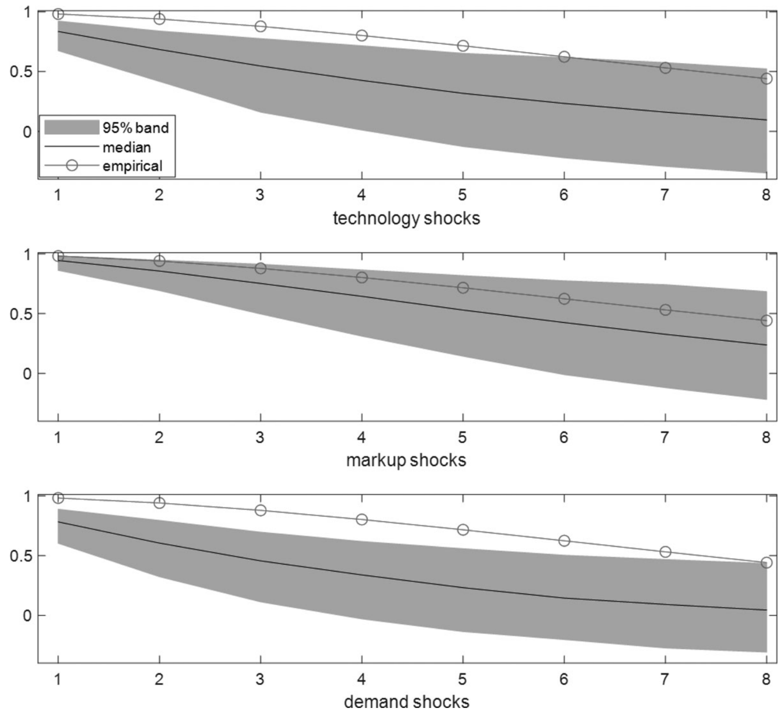

Fig. 3b. Unemployment Persistence–Standard Model versus Data

independently of the nature of the shock driving those fluctuations: the autocorrelations generated by the model, while positive and high at short lags, decline much faster than in the data. Furthermore, the empirical autocorrelation lies well outside the 95% confidence interval for the three shocks.

Figure 3b displays similar statistics based on model simulations that assume extreme persistence of the three driving forces. Specifically, I set $\rho_a = \rho_x = \rho_z = 0.99$, instead of the settings based on the estimates in Christoffel et al. (2008). Despite the (counterfactual) extreme persistence of the shocks, the median autocorrelation generated by the model remains well below its empirical counterpart for all shocks. Only in the case of markup shocks does the confidence band for the model autocorrelations contains the empirical autocorrelogram. The previous findings suggests that the stabilizing mechanisms embedded in the standard model and discussed above, are quite effective at (partly) insulating unemployment from the exogenous disturbances hitting the economy.

The adjustment of wages in response to the unemployment gap, captured in equation (17), is a key element in the stabilization of unemployment. In other words, the degree of unemployment persistence is not independent of the degree of wage rigidities. This is illustrated by the unemployment autocorrelations generated by the

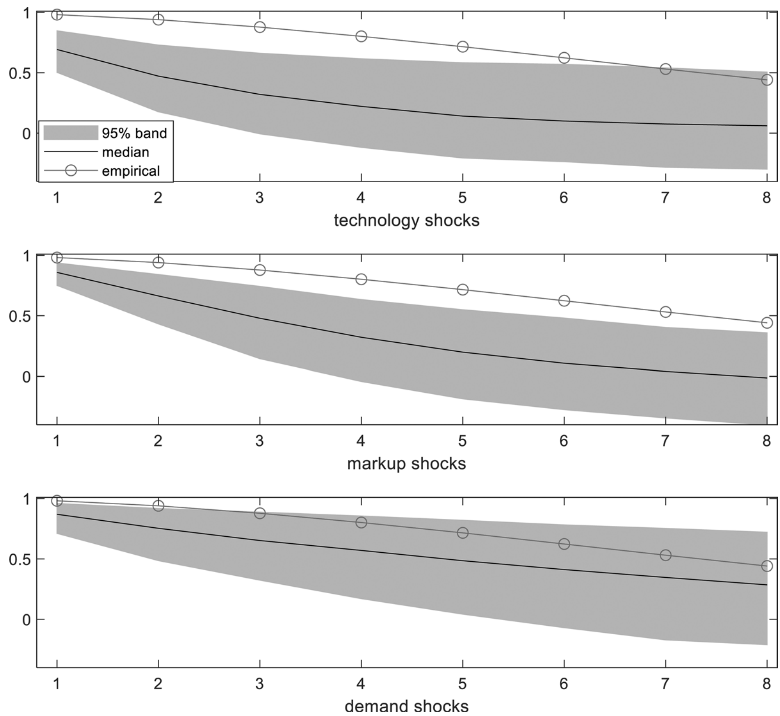

Fig. 3c. Extreme Wage Stickiness: Unemployment Persistence–Standard Model versus Data

model when I modify the baseline calibration by increasing the degree of wage stickiness. In particular, I assume $\theta_w = 0.95$, which implies an average duration of an individual wage of 5 years (!). Figure 3c displays the median autocorrelogram (with the corresponding confidence bands) across simulations under the modified θ_w value, together with the empirical autocorrelogram. Note that the model's unemployment autocorrelations are now uniformly higher at all lags, and for all shocks, but they still remain below their empirical counterparts. Note that the latter fall within the 95% confidence interval in the case of demand shocks. But not even the previous finding can be seen as a success for the standard model, since the required level of wage stickiness needed to match the observed unemployment persistence is highly unrealistic, being clearly at odds with the microevidence.

From the previous exercise, I conclude that a calibrated version of the standard New Keynesian model, under a "realistic" policy rule, cannot account for the high persistence of European unemployment, at least under plausible calibrations of the persistence of shocks and the degree of wage stickiness. A reasonable conjecture is that the failure of the standard model may lie in its treatment of the labor market itself, which may be at odds with the European reality. Next I study whether the insider–outsider version of the New Keynesian model does a better job at accounting for the observed unemployment persistence.[17]

3.3 Unemployment Persistence in the New Keynesian Model with Insider–Outsider Labor Markets and Hysteresis

Next I repeat the exercise described in the previous subsection using a version of the New Keynesian model with insider–outsider labor markets, as described above. Again, I simulate the model 500 times, conditional on each shock and obtain a set of artificial time series with 84 observations for each draw. All parameters are set to their baseline values. The insider–outsider model contains an additional parameter measuring the degree of hysteresis (γ).

After experimenting with alternative settings, a number of findings emerge that I summarize next. First, under $\gamma = 0$, that is, in the absence of hysteresis (and, hence, with a constant employment target), the persistence of unemployment is very similar (though not identical) to that in the New Keynesian model. This is true even though, as discussed above, their wage-setting rules are different (one targets a constant level of employment, the other targets a constant wage markup). Second, and irrespective of the shock considered, the model's implied autocorrelation of unemployment increases with γ. Third, given the assumed shock persistence and degree of wage stickiness, a γ value close to one is needed in order to approximate the unemployment persistence observed in the data. The latter finding is illustrated in Figure 4, which displays the model-based autocorrelations (median and 95% confidence bands) under the assumption that $\gamma = 0.99$, together with the empirical autocorrelogram. While the empirical autocorrelations still remain above the corresponding medians across model simulations at most lags, the gap between the two is noticeably smaller than that observed in Figure 3a for the standard model. As shown in Figure 4, this is especially true when demand shocks are the source of fluctuations, in which case the model-based and empirical autocorrelograms lie almost on top of each other.

Figure 5 illustrates graphically the role of hysteresis as a source of unemployment persistence, by showing the dynamic responses of the unemployment rate to each of the three shocks considered, in both the standard and the insider–outsider models, and in the latter case under three alternative settings of the hysteresis parameter: $\gamma = 0$, $\gamma = 0.9$, and $\gamma = 1$. The size of the shock is normalized to 1% in all cases, while their sign is chosen so that the unemployment rate increases on impact.[18] Several results

17. In Galí (2015b), I discuss possible sources of unemployment rate nonstationarity in the New Keynesian model. In addition to the hysteresis model proposed below, I point to nonstationarity in the desired wage markup and/or in the inflation target as possible addition sources of a unit root in the unemployment rate. As argued in that paper, however, some of the implications of those alternative hypothesis are hard to reconcile with the observed joint behavior of wage inflation and the unemployment rate.

18. This leads to a choice of a positive technology and markup shocks, and a negative demand shock. The fact that, under a simple interest rate rule like the one assumed here, a positive technology shock raises the unemployment rate in the New Keynesian model is well known from the literature (see, e.g., Galí, Smets, and Wouters 2012). It is also consistent with the empirical evidence (e.g., Barnichon 2010).

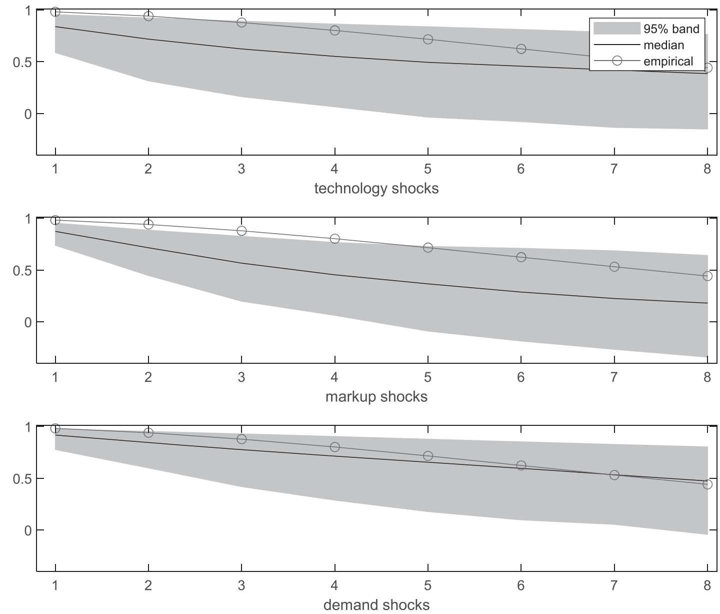

Fig. 4. Baseline Calibration: Unemployment Persistence: Insider-Outsider Model versus Data

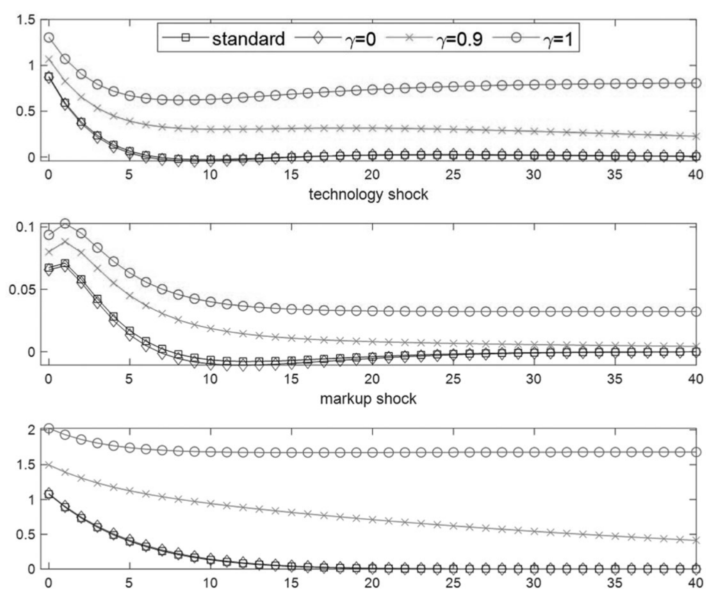

Fig. 5. Hysteresis and Unemployment Persistence

captured in Figure 5 are worth mentioning. First, the unemployment response in the insider–outsider model when $\gamma = 0$ is very similar (though not identical) to that generated by the standard model. Second, there is a strong positive relation between the size of γ and the observed persistence of the unemployment response. Finally, under full hysteresis ($\gamma = 1$), and under the assumed policy rule, the unemployment rate (as well as employment and output, not shown) changes permanently in response to any of the shocks, that is, it displays a unit root. Accordingly, any shock will generally have a permanent effect on the level of those variables, even when the shock itself is transitory. Such large and persistent effects on unemployment (and other variables) in the presence of strong hysteresis are due to the lack of significant deflationary pressures despite the persistent low activity, since the employment gap has little influence on wage inflation when γ is high, and none when it is equal to one (see equation (24)). As a result, such deviations do not elicit a suitable response from the central bank, unless the latter seeks to prevent those deviations to begin with (as in the optimal policy analyzed below).

Next I turn to the analysis of the implications of the insider–outsider model for the design of monetary policy.

4. OPTIMAL MONETARY POLICY WITH INSIDER–OUTSIDER LABOR MARKETS

In the present section, I derive the optimal monetary policy in the New Keynesian model with insider–outsider labor markets developed above. In doing so, I examine the role played by the degree of hysteresis (as measured by parameter γ) in shaping the response of unemployment to different shocks, with a focus on the differential response under the optimal policy relative to a simple interest rate rule.

4.1 The Optimal Monetary Policy Problem

In the analysis below, I assume that unions' long-term employment goal corresponds to the efficient level of employment. Formally,

$$n^* = n^e = \frac{\log(1-\alpha) - \xi}{1+\varphi}.$$

The previous assumption simplifies the analysis while allowing me to focus on the role of hysteresis without the (well-understood) complications arising from an inefficient steady state.[19] In addition, I restrict myself to model calibrations consistent

19. See, for example, chapter 5 in Galí (2015a). That assumption plays a role similar to the presence of an "optimal" employment subsidy in standard analyses of the optimal monetary policy in the New Keynesian model.

with stationarity of employment and the output gap, which rules out the extreme case of full hysteresis ($\gamma = 1$). Under the previous assumptions, one can approximate (up to second order) the representative household's welfare losses in a neighborhood of the steady state, expressed as a fraction of steady-state consumption, by the function:

$$\frac{1}{2}E_0 \sum_{t=0}^{\infty} \beta^t \left((1+\varphi)(1-\alpha)\widehat{n}_t^2 + \frac{\epsilon_p}{\lambda_p}(\pi_t^p)^2 + \frac{\epsilon_w(1-\alpha)}{\lambda_w}(\pi_t^w)^2 \right). \tag{30}$$

The loss function (30) is equivalent to that used in the standard New Keynesian model. The reason is that while the Calvo wage dynamics equation (14) (which holds in both the standard and the insider–outsider model) is used in the derivation of the loss function, this is not the case for the wage setting equation (15), so the latter's replacement by (21) has no impact on the form of the welfare loss function.[20]

The monetary authority seeks to minimize (30) subject to:

$$\pi_t^p = \beta E_t\{\pi_{t+1}^p\} + \lambda_p \alpha \widehat{n}_t + \lambda_p \widetilde{\omega}_t, \tag{31}$$

$$\pi_t^w = \beta E_t\{\pi_{t+1}^w\} + \lambda_n(1-\gamma)(1-\beta\theta_w)\widehat{n}_t + \lambda_n \gamma \Delta n_t], \tag{32}$$

$$\widetilde{\omega}_{t-1} \equiv \widetilde{\omega}_t - \pi_t^w + \pi_t^p + \Delta a_t - \Delta x_t \tag{33}$$

for $t = 0, 1, 2,$ ..together with some initial conditions for $\widetilde{\omega}_{-1}$ and $\widehat{n}_{-1}$. Note that (31) and (32) correspond to equations (10) and (24) derived above, while (33) uses the definition of the wage gap and the expression for the natural wage for the insider–outsider model derived above.

Let $\{\zeta_{1,t}\}$, $\{\zeta_{2,t}\}$, and $\{\zeta_{3,t}\}$ denote the sequence of Lagrange multipliers associated with the previous constraints, respectively. The optimality conditions for the optimal policy problem are thus given by

$$\begin{aligned}(1+\varphi)(1-\alpha)\widehat{n}_t \; &+ \lambda_p \alpha \zeta_{1,t} + \lambda_n(1-(1-\gamma)\beta\theta_w)\zeta_{2,t} \\ &- \lambda_n \gamma \beta E_t\{\zeta_{2,t+1}\} = 0,\end{aligned} \tag{34}$$

$$\frac{\epsilon_p}{\lambda_p}\pi_t^p - \Delta\zeta_{1,t} + \zeta_{3,t} = 0, \tag{35}$$

$$\frac{\epsilon_w(1-\alpha)}{\lambda_w}\pi_t^w - \Delta\zeta_{2,t} - \zeta_{3,t} = 0, \tag{36}$$

20. Note that the loss function is often expressed in terms of the output gap instead of the output gap (see, e.g., Galí 2015a, chapter 6). The relation between the two is given by $\widetilde{y}_t = (1-\alpha)\widehat{n}_t$, given the constancy of the natural level of employment.

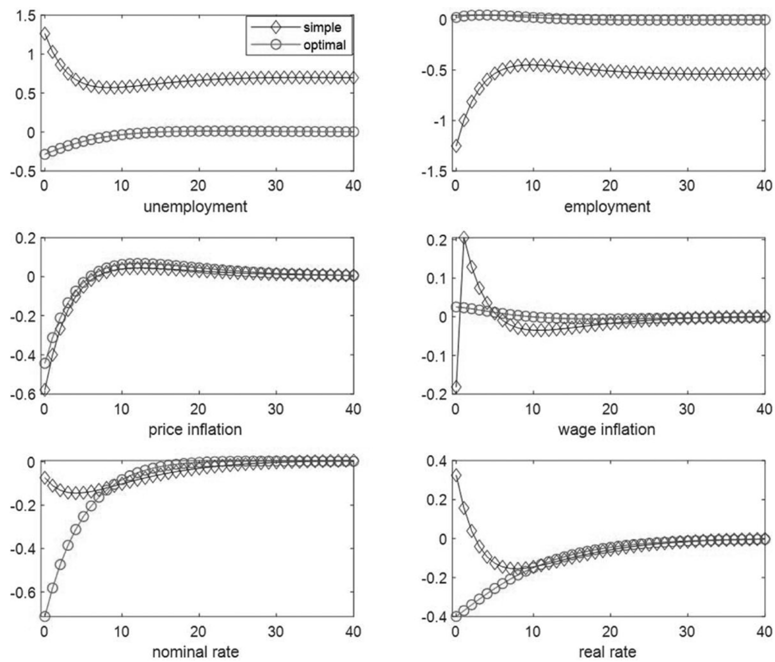

Fig. 6a. The Dynamic Effects of Technology Shocks: Optimal Policy versus Simple Rule

$$\lambda_p \zeta_{1,t} + \zeta_{3,t} - \beta E_t\{\zeta_{3,t+1}\} = 0 \tag{37}$$

for $t = 0, 1, 2, \ldots$which, together with the constraints (31), (32), and (33) given $\zeta_{1,-1} = \zeta_{2,-1} = 0$ and an initial condition for $\widetilde{\omega}_{-1}$ and $\widehat{n}_{-1}$, characterize the solution to the optimal policy problem.

Next, I describe the economy's response to shocks under the optimal policy, and compare it to that under the simple rule.

4.2 Dynamic Responses to Shocks and Welfare: Optimal Policy versus Simple Rule

Figures 6a–6c display the response of several macro variables to shocks to technology (Figure 6a), desired markups (Figure 6b), and demand (Figure 6c), under two policies: the optimal policy (represented by the lines with circles) and the simple rule (lines with diamonds). The six variables considered are: the unemployment rate, (log) employment, price inflation, wage inflation and the nominal and real interest rate. The degree of hysteresis is set at $\gamma = 0.99$, a value in the range needed to generate realistic unemployment persistence. The remaining parameters (including the coefficients in the simple policy rule) are kept at their baseline settings, as in the simulations of Section 3. As in the latter, the size of the shock is normalized to 1% in all cases, while the sign is chosen so that under the simple rule unemployment increases on impact.

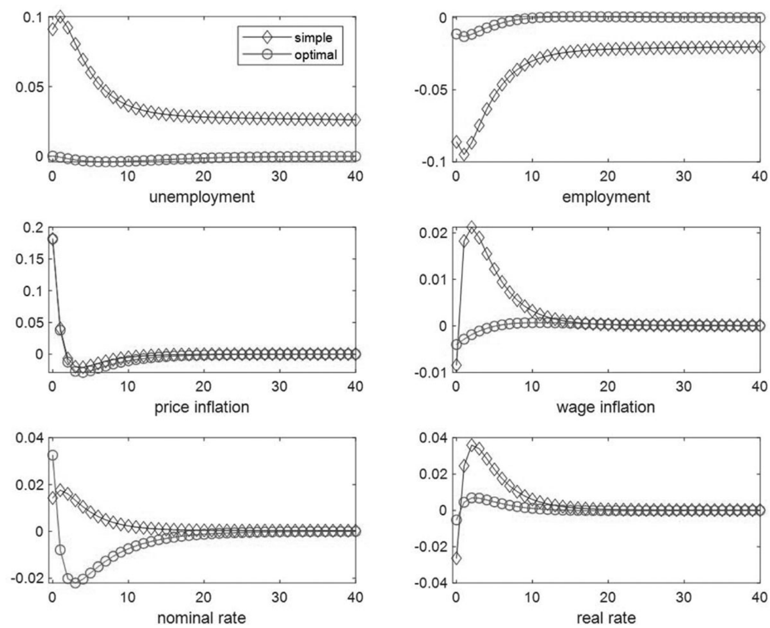

Fig. 6b. The Dynamics Effects of Markup Shocks: Optimal Policy versus Simple Rule

Several findings are worth stressing. First, under the optimal policy the unemployment rate remains highly stable, in comparison to the responses under the simple rule, which display very high persistence. This is true independently of the shock impinging on the economy. It takes an extreme form in the case of demand shocks, in response to which the optimal policy *fully* stabilizes the unemployment rate.[21]

The large discrepancies between the optimal and simple rules observed for unemployment are also visible for the remaining variables, irrespective of the shock. The gap is more persistent in the case of the unemployment rate and (log) employment, but shorter lived (albeit sizable) for the other variables.

Due to the high degree of hysteresis assumed, the large and persistent negative deviations of employment from its (constant) efficient level do not bring about significant deflationary pressures, since the employment gap has little influence on wage inflation when γ is high (see equation (24)). As a result, such deviations do not elicit a suitable response from the central bank under the simple rule. On the other

21. It is interesting to note that the difference in the sign of the unemployment response to a positive technology shock. As noted above, it is well known that under a simple rule like the one assumed here unemployment rises in response to a positive technology shock. Under the optimal policy, however, unemployment declines (albeit by a small amount). This raises wage inflation which, together with the slight price deflation, contributes to the increase in real wages needed to approximate the flexible price and wage equilibrium.

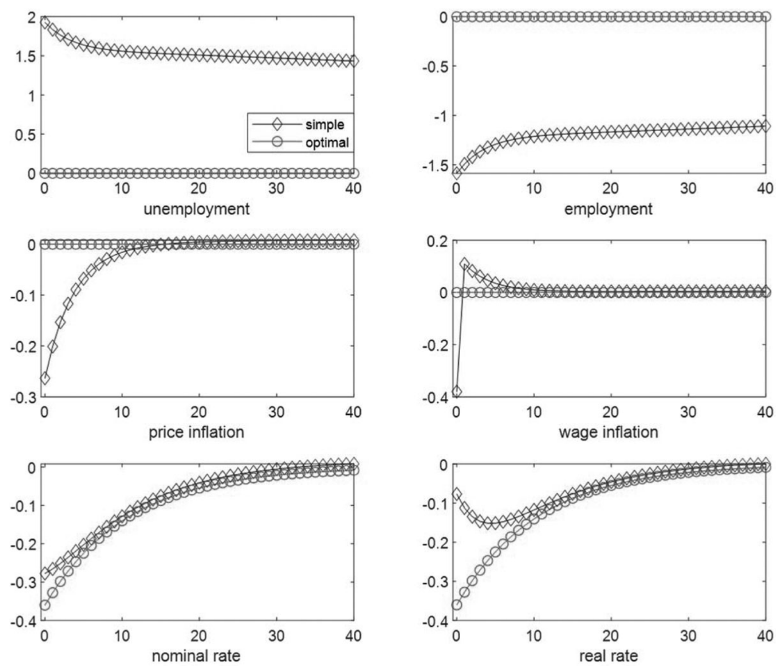

Fig. 6c. The Dynamic Effects of Demand Shocks: Optimal Policy versus Simple Rule

hand, the optimal policy seeks to prevent those deviations to begin with, through a more aggressive and persistent reduction in interest rates, either on impact or over time.

The nontrivial gap between the responses of employment and price and wage inflation under the two policies suggests that the adoption of the optimal policy may bring about considerable welfare gains relative to the simple rule. Next I show that this is the case, especially when, as in the previous analysis, the degree hysteresis is high.

Figure 7 displays the welfare losses under both the simple rule and the optimal policy, conditional on each shock, and as a function of the degree of hysteresis γ. For each shock, welfare losses, expressed as a fraction of steady-state consumption, are normalized to be equal to one under the simple rule and in the absence of hysteresis ($\gamma = 0$). The plots show that, under the simple rule (line with diamonds), welfare losses are relatively insensitive to γ for low values of the latter, but when γ reaches a certain level (which varies across shocks), welfare losses rise very fast. Thus, an economy with strong hysteresis in which the central bank follows the assumed simple rule may experience welfare losses that can be a large multiple of those of an economy with low or no hysteresis (but otherwise identical). This is particularly true when fluctuations are driven by demand shocks, as implied by the findings shown in Figure 7.

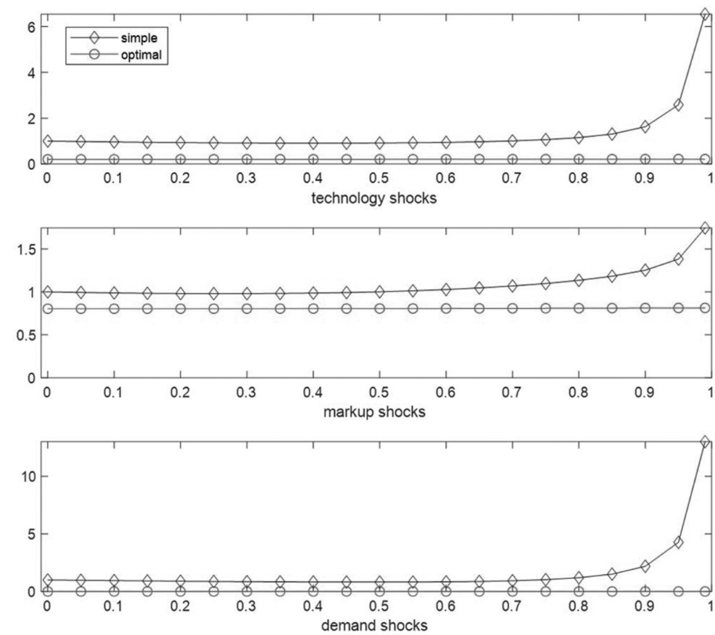

Fig. 7. Hysteresis and Welfare Losses: Optimal Policy versus Simple Rule

By contrast, under the optimal policy (line with circles) welfare losses are largely insensitive to the degree of hysteresis. This is literally so in the case of demand shocks, for which welfare losses under the optimal policy are zero independently of γ. In the case of technology and markup shocks welfare losses under the optimal policy are only a fraction of the welfare losses under the simple rule and vary little with γ, with the changes being visually unperceptible in the figure.

The findings reported in Figure 7 make clear that the adoption of the optimal policy would imply a substantial reduction of welfare losses relative to the simple rule. This reduction is, of course, extreme in the case of demand shocks since in that case welfare losses are brought down to zero under the optimal policy. For technology and markup shocks, the gains from adopting the optimal policy increase monotonically with the degree of hysteresis (in absolute as well as in percent terms). To put it differently, the costs of following the simple rule as opposed to the optimal policy are larger in economies that feature strong hysteresis, independently of the source of fluctuations.

Unfortunately, the large gains that the adoption of the optimal policy can bring about in an economy with a high degree of hysteresis cannot hide the fact that, in practice, it is hard to imagine a central bank adopting a policy of such complexity (as reflected in the four difference equations that describe it), with the communication challenges that it would generate. This motivates the next subsection, where I evaluate the potential of two simple targeting rules to approximate the optimal policy.

5. SIMPLE TARGETING RULES FOR THE INSIDER–OUTSIDER NEW KEYNESIAN MODEL

A look at the dynamic responses under the optimal policy in Figures 6a–6c point to a common feature across shocks: the high stability of both unemployment and employment, which hardly respond to any of the shocks. In the case of demand shocks that stability is complete, since the optimal policy fully insulates those three variables (as well as price inflation) from the shock, thus preserving the efficient allocation. Note also that wage inflation remains nearly unchanged in response to the three shocks, which follows from the high stability of employment, and given the close relation between wage inflation and employment implied by the wage Phillips curve (24).

By contrast, the simple rule analyzed in the previous section (and meant to approximate the historical behavior of the ECB) focuses instead on the (partial) stabilization of price inflation and output growth, while failing to provide the *anchor* for employment or unemployment that would eliminate (or at least reduce) the large size and high persistence of the deviations of those variables from their optimal levels in response to shocks under that rule. The option of increasing the size of the coefficient on output growth, or to replace it with (detrended) output does not seem desirable since it may overstabilize that variable in the face of shocks that change its efficient level, possibly permanently (e.g., technology shocks), so I choose not to pursue that avenue.[22]

The previous discussion suggests that a rule that seeks to stabilize employment (or, equivalently, wage inflation, given (24)) or unemployment may approximate well the optimal policy rule, independently of the degree of hysteresis. To assess the merits of those rules, I compute the welfare losses associated with the following simple targeting rules. The first rule, which I label *n-targeting*, requires that $n_t = n_t^e$ for all t (or, equivalently, $\pi_t^w = 0$ for all t). The second rule, labeled *u-targeting*, calls for $u_t = u$ for all t, where u is the unemployment rate consistent with efficient employment, given by (27). Figures 8a and 8b display the dynamic responses of the same six macro variables shown in Figure 6 to (respectively) a technology and a markup shock, under *n-targeting* (line with diamonds), *u-targeting* (line with squares), and the optimal policy (line with circles). The responses to a demand shock under the two targeting rules are identical to those implied by the optimal policy (see Figure 6c) and are thus not displayed. As Figures 8a and 8b make clear, the gap between the response of the different variables under the two simple targeting policies and that under the optimal policy is generally very small, once the scale of the graphs is taken into account. In the

22. Of course, adding the level of the output gap as an argument would help attain the desired objective, though the unobservability of that variable disqualifies that rule as a “simple” rule, since it could not be implemented in practice. In fact, as is well known from textbook analyses of the New Keynesian model (e.g., Woodford (2003) or Galí (2015)) a policy that fully stabilizes the output gap approximates very well the optimal policy, for reasonable calibrations. See Coibion et al. (2018) for a discussion of the pitfalls of output gap measures based on existing estimates of potential output.

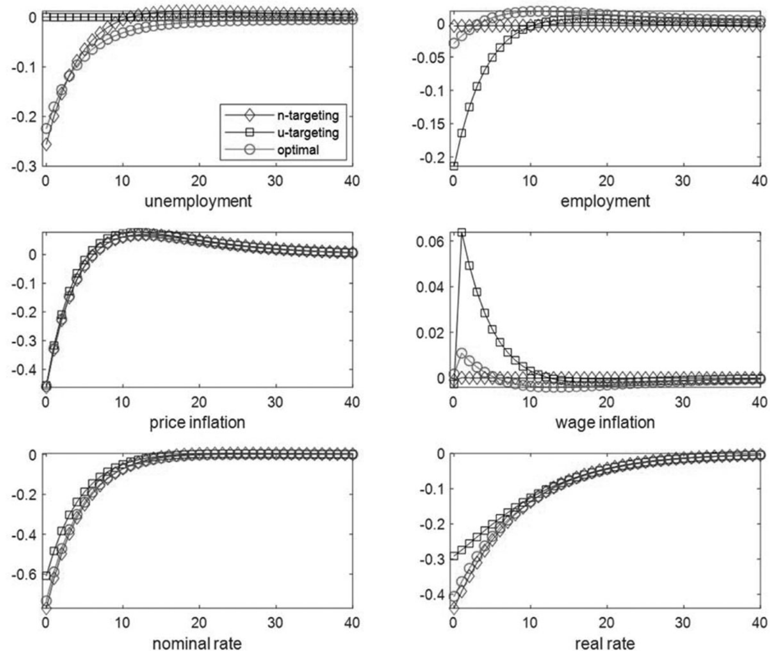

Fig. 8a. The Dynamic Effects of Technology Shocks: Optimal Policy versus Simple Targeting Rules

case of the *n-targeting* rule, that gap is close to negligible for both shocks. In the case of the *u-targeting*, some small but noticeable differences arise relative to the optimal policy in the short-run responses of employment and unemployment to a technology shock, but the size of the gap is much smaller than that implied by the simple rule, as shown in Figure 6.

The previous findings are complemented by Figure 9 which displays the welfare losses associated with the two simple targeting rules and the optimal policy, conditional on each shock, and as a function of the degree of hysteresis γ, with all other parameters at their baseline values. The comparable welfare losses under the simple rule, much larger in size, are displayed in Figure 7 discussed above.

Given the similarity of the impulse responses, it is not surprising that the welfare losses under the two targeting rules (relative to the simple rule) have the same order of magnitude and are quantitatively similar to those associated with the optimal policy (though obviously larger or equal than the latter, by construction). In the case of demand shocks, the welfare losses are zero in the three cases, since the two targeting policies, like the optimal policy, succeed in replicating the efficient allocation. In the case of technology shocks, we see that the welfare losses under the *n-targeting* rule remain uniformly below those associated with the *u-targeting* rule, for any γ value. This not the case, however, when I condition on markup shocks, in which case

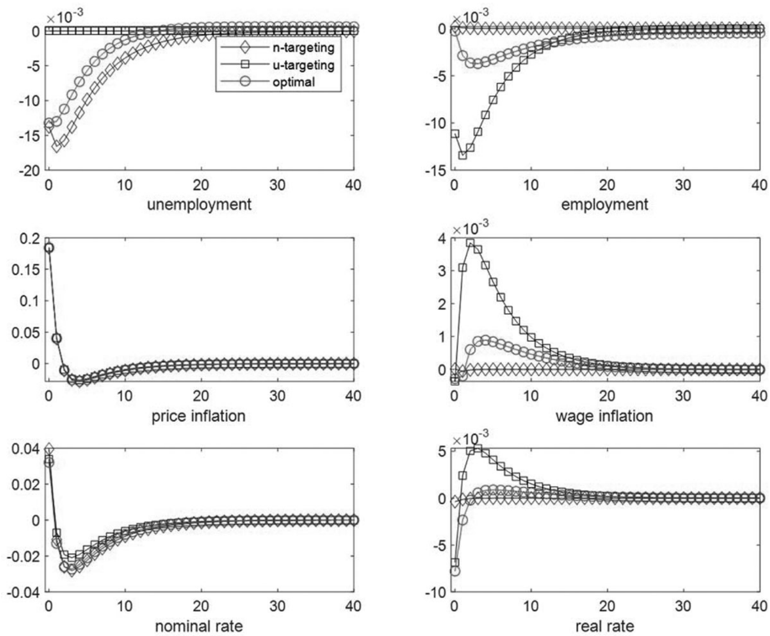

Fig. 8b. The Dynamic Effects of Markup Shocks: Optimal Policy versus Simple Targeting Rules

there is an intermediate region of γ values for which the *u-targeting* rule appears to dominate the *n-targeting* rule, though the differences among the three rules are quantitatively tiny.

The previous analysis points to the benefits from pursuing a policy rule that targets employment or unemployment, especially in the presence of strong hysteresis. The next section I analyze the implications of hysteresis on the behavior of the labor wedge, and provide some insights on the channels through which the optimal policy and the two simple targeting rules improves on the outcomes from the simple rule.

6. HYSTERESIS AND THE LABOR WEDGE

In this section, I use the insider–outsider version of the New Keynesian model developed above to study the relation between hysteresis and the "labor wedge." The latter variable, denoted by μ_t, is defined as

$$\mu_t \equiv mpn_t - mrs_t,$$

where mpn_t is the (log) marginal product of labor and mrs_t is the (log) marginal rate of substitution between labor and consumption. That gap, often referred to in the

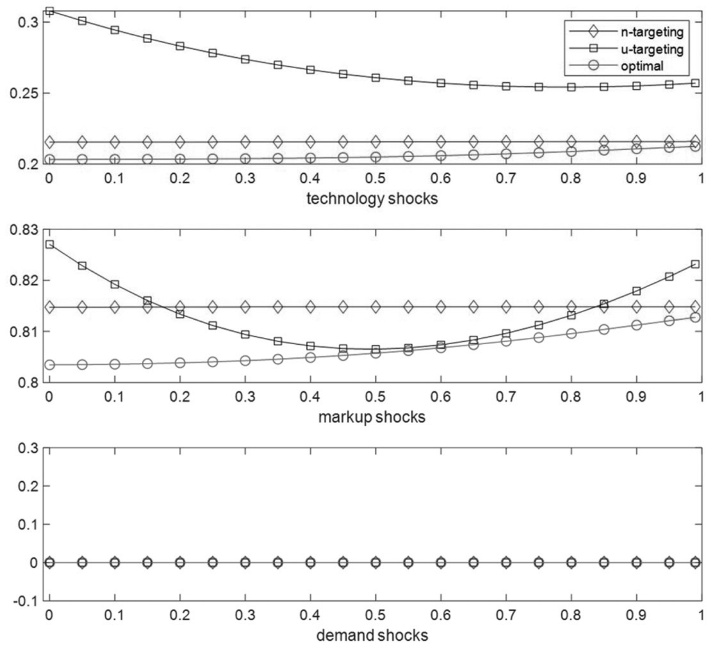

Fig. 9. Hysteresis and Welfare Losses: Optimal Policy versus Simple Targeting Rules

literature as the "labor wedge," can be interpreted as an index of the extent to which labor is underutilized relative to the efficient allocation. Several authors have constructed measures of the labor wedge and provided evidence of a strong countercyclical behavior of that variable, confirming the traditional view of recessions as periods in which the utilization of labor is *inefficiently* low relative to booms (see, e.g., Hall 1997, Chari et al. 2007, Galí et al. 2007, Shimer 2009, Galí 2011, Karabarbounis 2014, among others). Analyses of the labor wedges implied by alternative theoretical models and their comparison with their empirical counterpart may help guide modeling efforts and shed some light on the channels through which alternative policies affect the aggregate efficiency of allocations and welfare. This is the main motivation behind the present section.[23]

The labor wedge can be viewed as the sum of two variables that are central to the New Keynesian framework: the price markup $\mu_t^p \equiv mpn_t - \omega_t$ and the wage markup $\mu_t^w \equiv \omega_t - mrs_t$. Those variables measure the contribution of the goods and labor markets, respectively, to the aggregate labor wedge. Empirical analyses of the labor wedge and its two components in advanced economies, including the euro area, point

23. See, for example, Shimer (2009) and Chugh et al. (2018) for an illustration of such a use of the labor wedge.

TABLE 3

HYSTERESIS AND LABOR WEDGE VOLATILITY

	Labor wedge			*Wage markup*			*Price markup*		
	$\gamma = 0$	$\gamma = 0.9$	$\gamma = 0.99$	$\gamma = 0$	$\gamma = 0.9$	$\gamma = 0.99$	$\gamma = 0$	$\gamma = 0.9$	$\gamma = 0.99$
Technology									
Simple	1.0	1.81	6.42	0.69	1.59	6.35	0.31	0.30	0.31
Optimal	0.09	0.09	0.07	0.31	0.26	0.23	0.24	0.24	0.25
n-Targeting	0.0	0.0	0.0	0.24	0.24	0.24	0.24	0.24	0.24
u-Targeting	0.27	0.26	0.26	0.0	0.0	0.0	0.27	0.26	0.26
Markup									
Simple	1.0	1.42	3.11	0.74	1.18	2.98	0.28	0.26	0.26
Optimal	0.18	0.13	0.10	0.09	0.12	0.16	0.23	0.22	0.21
n-Targeting	0.0	0.0	0.0	0.21	0.21	0.21	0.21	0.21	0.21
u-Targeting	0.23	0.22	0.22	0.0	0.0	0.0	0.23	0.22	0.22
Demand									
Simple	1.0	2.72	10.43	0.92	2.67	10.41	0.07	0.07	0.09
Optimal	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
n-Targeting	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	00.0
u-Targeting	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

NOTE: For each shock and policy rule, the table reports the model-implied standard deviation of the labor wedge, the wage markup and the price markup for three different calibrations of γ, and with the standard deviation of the labor wedge under the simple rule with no hysteresis normalized to unity for each shock.

to fluctuations in the wage markup μ_t^w as the main source of the high volatility and countercyclical behavior of the labor wedge, with the price markup μ_t^p displaying instead much less volatility and a largely acyclical pattern.[24]

Note that in the New Keynesian model developed above, the efficient allocation corresponds to $\mu_t = 0$ for all t and is associated with a constant level of employment $n^e = \frac{\log(1-\alpha)-\xi}{1+\varphi}$. Fluctuations in the labor wedge, on the other hand, satisfy the simple relation

$$\begin{aligned} \mu_t &= mpn_t - mrs_t \\ &= -(1+\varphi)\widehat{n}_t, \end{aligned}$$

where $\widehat{n}_t \equiv n_t - n^e$. Thus, if we take (log) employment as a reference variable, it follows that the model implies, by construction, a strongly countercyclical pattern for the labor wedge, in accordance with the empirical evidence. Note also that the first term in the welfare loss function (30) is proportional to the volatility of the labor wedge, thus motivating the interest of an analysis of the factors underlying that volatility. In particular, in the context of the present paper, one may want to study the relation between hysteresis and labor wedge volatility, how that relation is affected by the monetary policy rule in place, and the role played by each of the components of the labor wedge in shaping that relation.

Table 3 reports some statistics that shed some light on the previous questions. For each shock and policy rule, the table reports the model-implied standard deviation

24. See Galí et al. (2007) for U.S. evidence, Galí (2011) for euro area evidence, and Karabarbounis (2014) for 15 OECD countries evidence.

of the labor wedge, the wage markup, and the price markup for three different calibrations of γ (0, 0.9, and 0.99). For each shock, the standard deviation of the labor wedge under the simple rule with no hysteresis ($\gamma = 0$) is normalized to unity. The remaining parameters are set at their baseline values.

Several results are worth pointing out. First, under the simple rule and for all shocks, there is a strong positive relation between the degree of hysteresis and the volatility of the labor wedge. Second, under the simple rule (meant to approximate actual historical policy), fluctuations in the wage markup account for the bulk of the volatility in the labor wedge, in a way consistent with the evidence mentioned above. Third, the adoption of the optimal policy or either of the targeting rules considered leads to a considerable reduction in the volatility of both the labor wedge and its wage markup component, with those volatilities becoming largely independent of the degree of hysteresis. By construction, in the case of the *n-targeting* rule the volatility of the labor wedge is zero, whereas under *u-targeting* the volatility of the wage markup is zero. Note also that, conditional on demand shocks, the volatility of the three wedges is zero under the optimal policy and the two targeting rules, for all real variables (other than the real interest rate) are fully isolated from the shock. Fourth, the volatility of the price markup is generally small, and largely unrelated to the degree of hysteresis. It is, nevertheless, slightly lower under the optimal policy and/or either targeting rule, relative to the simple rule.

The previous findings shed some light on the channels through which hysteresis ends up affecting an important component of welfare losses—namely, inefficient employment fluctuations, as reflected in the labor wedge, largely associated with wage markup fluctuations. The switch from the simple rule to the optimal policy (or the targeting policies that approximate it) reduces the volatility of the labor wedge and the wage markup considerably, especially when in the presence of strong hysteresis. The previous findings partly explain why a policy that succeeds in stabilizing employment (like the *n-targeting* rule) performs so well from a welfare viewpoint. The following section is motivated by that observation. Given the central role played by wage markup fluctuations in the overall fluctuations of the labor wedge, the findings above also explain why a policy that seeks to stabilizes the unemployment rate (and, hence, the wage markup) like the *u-targeting* rule does almost as well as the optimal policy.

7. CONCLUDING REMARKS

The high persistence of European unemployment constitutes a challenge for conventional macro models, including the standard New Keynesian model. In the present paper, I have proposed a modified version of that model that can generate highly persistent unemployment. The main modification consists of combining insider–outsider labor markets and hysteresis, as in Blanchard and Summers (1986), with the

Calvo-type price and wage-setting structure characteristic of the New Keynesian model. In the modified model, the degree of hysteresis needs to be substantial in order to generate European levels of persistence. Under full hysteresis, unemployment and other real variables may experience permanent deviations from their efficient levels, even in response to shocks that are transitory. Such deviations, even if large, do not necessarily generate inflationary pressures (of either sign) and hence may not elicit a suitable response from an inflation-focused central bank.

The presence of hysteresis effects has important implications for the conduct of monetary policy. Specifically, I have shown that the optimal monetary policy calls for a more aggressive stabilization of employment and/or unemployment than implied by a conventional interest rate rule. The welfare gains from shifting to the optimal policy are considerable, and increasing in the degree of hysteresis. Furthermore, the analysis above suggests that the outcome of the optimal policy can be approximated well by a simple targeting rule that seeks to stabilize employment or unemployment. Finally, an exploration of the connection between hysteresis and the labor wedge, points to a strong degree of hysteresis, combined with a (realistic) simple interest rate rule, as a potential source of a highly volatile labor wedge (through its wage markup component), the adoption of an optimal policy (or the targeting rules that approximate it) makes it possible, however, to nearly shut down that effect of hysteresis.

The findings in the paper, while based on a highly stylized model, should raise a warning flag on monetary policy strategies that put too much weight on inflation stabilization, for in environments where inflation is unresponsive to inefficient levels of employment, those strategies run the risk of chronifying those inefficiencies.

LITERATURE CITED

Álvarez, Luis J., Emmanuel Dhyne, Marco M. Hoeberichts, Claudia Kwapil, Hervé Le Bihan, Patrick Lünnemann, Fernando Martins, Roberto Sabbatini, Harald Stahl, Philip Vermeulen, and Jouko Vilmunen. (2006) "Sticky Prices in the Euro Area: A Summary of New Micro Evidence." *Journal of the European Economic Association*, 4, 575–84.

Ball, Laurence. (2009) "Hysteresis in Unemployment," In *Understanding Inflation and the Implications for Monetary Policy*, edited by Jeff Fuhrer, Yolanda K. Kodrzycki, Jane Sneddon Little, and Giovanni P. Olivei, pp. 361–82. Cambridge, MA: MIT Press.

Ball, Laurence. (2014) "Long-term Damage from the Great Recession in OECD Countries." *European Journal of Economics and Economic Policies: Intevention*, 11, 149–60.

Barnichon, Régis. (2010) "Productivity and Unemployment over the Business Cycle." *Journal of Monetary Economics*, 57, 1013–25

Blanchard, Olivier. (2018) "Should We Reject the Natural Rate Hypothesis." *Journal of Economic Perspectives*, 32, 97–120.

Blanchard, Olivier, and Lawrence Summers. (1986) "Hysteresis and the European Unemployment Problem. " In *NBER Macroeconomics Annual*, edited by Stanley Fischer, pp. 15–78. Cambridge. MA: MIT Press.

Chari, Varadarajan, Patrick Kehoe, and Ellen McGrattan. (2007) "Business Cycle Accounting." *Econometrica*, 75, 781–836.

Chugh, Sanjay K., Wolfgang Lechthaler, and Christian Merkl. (2018) “Optimal Fiscal Policy with Labor Selection.” *Journal of Economic Dynamics & Control*, 94, 142–89.

Christoffel, Kai, Gunter Coenen, and Anders Warner. (2008) “The New Area-Wide Model of the Euro Area: A Micro-Founded Open Economy Model for Forecasting and Policy Analysis.” ECB Working Paper No. 944.

Coibion, Olivier, Yuriy Gorodnichenko, and Mauricio Ulate. (2018) “The Cyclical Sensitivity in Measures of Potential Output.” *Brookings Papers on Economic Activity*, 343–411.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin. (2000) “Optimal Monetary Policy with Staggered Wage and Price Contracts,” *Journal of Monetary Economics* 46, 281–314.

European Central Bank. (2009) “Wage Dynamics in Europe: Final Report of the Wage Dynamics Network (WDN).” Directorate General Research.

Fagan, Gabriel, Jerome Henry, and Ricardo Mestre. (2001) “An Area-Wide Model (AWM) for the Euro Area.” ECB Working Paper Series No. 42.

Hall, Robert. (1997) “Macroeconomic Fluctuations and the Allocation of Time.” *Journal of Labor Economics*, 15, 223–50.

Galí, Jordi. (2011) *Unemployment Fluctuations and Stabilization Policies: A New Keynesian Perspective*. Cambridge, MA: MIT Press.

Galí, Jordi. (2015a) *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework*, 2nd ed. Princeton: Princeton University Press.

Galí, Jordi. (2015b) “Hysteresis and the European Unemployment Problem Revisited.” In *Inflation and Unemployment in Europe*, pp. 53–79. Proceedings of the ECB Forum on Central Banking, European Central Bank, Frankfurt am Main.

Galí, Jordi, Mark Gertler, and David López-Salido. (2007) “Markups, Gaps and the Welfare Costs of Business Fluctuations.” *Review of Economics and Statistics*, 89, 44–59.

Galí, Jordi, Frank Smets, and Raf Wouters. (2012) “Unemployment in an Estimated New Keynesian Model.” In *NBER Macroeconomics Annual 2011*, edited by D. Acemoglu and M. Woodford Chicago: University of Chicago Press. pp. 329–60. Chicago:IL.

Gottfries, Nils, and Henrik Horn. (1987) “Wage Formation and the Persistence of Unemployment.” *Economic Journal*, 97, 877–84.

Karabarbounis, Loukas. (2014) “The Labor Wedge: MRS vs. MPN.” *Review of Economic Dynamics*, 17, 206–23.

Lindbeck, Assar, and Dennis J. Snower. (1988) *The Insider-Outsider Theory of Employment and Unemployment*. Cambridge, MA: MIT Press.

Orphanides, Athanasios. (2006) “Review of the ECB’s Strategy and Alternative Approaches.” Contribution to *The ECB and Its Watchers*, Center for Financial Studies, Frankfurt.

Shimer, Robert. (2009) “Convergence in Macroeconomics: The Labor Wedge.” *American Economic Journal: Macroeconomics*, 1, 267–97.

Smets, Frank. (2010) “Comment on Chapters 6 and 7.” In *The Euro: The First Decade*, edited by M. Buti, S. Deroose, V. Gaspar, and J. N. Martins, pp. 259–79. Cambridge: Cambridge University Press.

Woodford, Michael. (2003) *Interest and Prices. Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press. pp. 259–79.