IS THERE A TRADE-OFF BETWEEN INFLATION AND OUTPUT STABILIZATION?

ALEJANDRO JUSTINIANO, GIORGIO E. PRIMICERI, AND ANDREA TAMBALOTTI

ABSTRACT. Not in an estimated DSGE model of the US economy, once we account for the fact that most of the high-frequency volatility in wages appears to be due to noise, rather than to variation in workers' preferences or market power.

1. INTRODUCTION

Between 1954 and 2009, average GDP growth in the U.S. was 3.2 percent per quarter at an annual rate, with a standard deviation of 3.8 percent. The origin of these fluctuations in economic activity is still a matter of debate. Even more controversial is the extent to which these fluctuations reflect the efficient responses of economic agents to changes in fundamentals, as opposed to departures from the efficient allocation of resources. This distinction matters because in the first instance stabilization policy cannot improve welfare, while in the second, policy intervention could be beneficial.

We make two contributions to this debate. First, we show that a sizable fraction of the output fluctuations experienced by the U.S. economy in the postwar period represents changes in its degree of inefficiency, rather than movements of the efficient frontier. Second, we find that a significant reduction in these inefficient output fluctuations would have been compatible with *more stable* inflation, and would have thus resulted in higher welfare. This is because the policy trade-off between real and nominal stabilization is quantitatively negligible in our model economy.

Our results are based on an estimated dynamic stochastic general equilibrium (DSGE) model, which is built around neoclassical growth foundations, but assumes that firms and workers enjoy some monopoly power and cannot set prices freely every period. Relative to

Date: May 2011. This paper is a substantially revised and extended version of an earlier manuscript titled "Potential and Natural Output." We thank Jose Dorich, Gauti Eggertsson, Chris Erceg, Stefano Eusepi, Andrea Ferrero, Jordi Gali, John Leahy, Frank Smets, Lars Svensson, Carl Walsh, Mike Woodford and participants in several conferences and seminars for comments and suggestions, Ging Cee Ng for superb research assistance and especially Vasco Cúrdia for generously sharing his code for the computation of optimal equilibria. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Chicago, New York or the Federal Reserve System.

an environment with perfect competition (and flexible prices), these New-Keynesian features distort the economy's equilibrium away from the efficient allocation. This distortion manifests itself in the form of markups of goods prices over nominal marginal cost and of real wages over the marginal rate of substitution between consumption and leisure. These markups vary over time due to exogenous movements in market competitiveness, and to the stickiness of prices and wages.

We summarize the evolution of these distortions by measuring the distance between actual GDP and the GDP that would emerge in equilibrium if markups were constant at their steady state level. We call this counterfactual object *potential* GDP. This is also the level of output that would be observed in an economy identical to ours, but with flexible prices and wages, and no exogenous variation in the degree of competition. In our model, potential GDP is below its efficient level—the one that would be observed if markets were perfectly competitive—since markups are positive in steady state. However, potential and efficient output share identical log-linear dynamics. Therefore the output gap—the difference between actual and potential GDP—is a useful summary statistic of the movements of the economy away from its efficient frontier. According to our estimates, the output gap is pro-cyclical and often quite large, with a standard deviation of 2.5 percentage points. From this evidence, we conclude that movements in the economy's degree of inefficiency contributed significantly to postwar fluctuations.

Does the existence of inefficient output fluctuations imply that stabilization policy was suboptimal over our sample? In general, the answer is no, since our economy features multiple distortions that result in a complex trade-off between the stability of the output gap and other policy objectives. Most notably, the stabilization of the output gap produces dispersion in the cross-section of prices and wages, which in the aggregate manifests itself as inflation. For example, the exogenous fluctuations in desired markups due to changes in market competitiveness give agents with the chance to adjust their prices a reason to do so, even when the output gap is zero, thus creating a discrepancy between newly set prices and the existing ones. Price dispersion, in turn, produces dispersion in markups and in the supply of goods and labor, which is inefficient because workers and firms are identical and the technology that aggregates their inputs is concave. As a result, producing the potential level of output in the actual economy requires more work than in the counterfactual economy with stable markups. Stated more formally, the full potential allocation, in which markups are constant

across agents and time, is infeasible in our model. Therefore, policymakers must trade-off the stabilization of output around potential with that of price and wage inflation.¹

To evaluate quantitatively the significance of this trade-off in our estimated model, we compute its optimal allocation, focusing in particular on optimal output. This is the counterfactual level of output that would have been observed in the post-war U.S. economy if the nominal interest rate had been set so as to maximize the utility of the model’s representative agent, rather than following the interest rate rule we estimate. The main finding of this exercise, and the central result of the paper, is that optimal and potential output move closely together, and that this near stabilization of the output gap is roughly consistent with the stabilization of wage and price inflation. The surprising implication is that stabilization policy appears to face a negligible trade-off among its three main objectives and that much of the inefficient variation in output uncovered by our estimates could have been eliminated, increasing welfare at the same time.

The key to this finding is our treatment of wages in the estimation of the model. Unlike in most empirical DSGE exercises, exemplified for instance by the seminal work of Smets and Wouters (2007), we match the model’s wage variable with two measures of hourly labor income, allowing for errors in their measurement, along the lines of Boivin and Giannoni (2006a) (see also Gali, Smets, and Wouters (2011)). As a result, most of the high frequency variation that characterizes the individual series on compensation used in estimation is not interpreted as implausibly large variation in the monopoly power of workers. It is instead recognized as measurement error or, more generally, as a mismatch between the data and the model’s wage concept.

This shift in the interpretation of the factors behind the observed high frequency movements in wages—from markup shocks to measurement errors—produces a model that is economically more plausible, empirically more successful, and with radically different normative implications. Indeed, in models estimated with only one wage series, the extreme volatility of wage markup shocks generates an optimal output that is significantly more volatile than actual output. This destabilization of output is the price policymakers must pay to obtain

¹ In environments with flexible wages, no markup shocks and no capital accumulation, the stabilization of the output gap and of aggregate markups are equivalent and produce no price dispersion, thus delivering the efficient allocation. This is what Blanchard and Galí (2007) call the “divine coincidence”. See Woodford (2003) or Galí (2008) for a textbook treatment. In our environment, output gap and markup stabilization are not equivalent, due to the presence of capital accumulation. As a result, the policy tradeoff is not only between output and inflation stabilization, but also involves the composition of demand between consumption and investment.

the desired stabilization of wages, whose dispersion is extremely costly for the representative agent. We present an example of a model of this kind in section 6. In contrast, this extreme tension between output and wage stabilization is virtually absent in the baseline model with multiple wage indicators, where the variation in wage markups is negligible. In fact, in this model, the deviation of the optimal from the potential allocation is altogether minimal, leading to our conclusion that monetary policymakers face a negligible conflict among their objectives.

One important qualification to the absence of a significant policy trade-off is that it depends on the maintained assumption that the low frequency variation in labor supply identified by the model, and reflected in the behavior of hours worked, is mostly due to shifts in households' attitudes towards work in the market. If these secular changes in hours were interpreted instead as coming from a trend in workers' monopoly power, our inference on the movements of the economy's efficient frontier—and therefore the output gap—would change significantly, as also illustrated by Sala, Söderström, and Trigari (2010). However, the properties of the optimal equilibrium, and of optimal output in particular, are barely affected by this alternative interpretation. This suggests that the lack of identification of labor supply shocks has only a minor impact on the normative implications of the model.

This paper is related to a large literature on the estimation of DSGE models (e.g. Rotemberg and Woodford (1997), Lubik and Schorfheide (2004), Boivin and Giannoni (2006b), Smets and Wouters (2007)). Some of these papers explicitly assume that monetary policy reacts to the difference between actual and potential output, but do not focus on the model's predictions regarding the behavior of the output gap. A number of studies, such as Levin, Onatski, Williams, and Williams (2005), Nelson (2005), Andrés, López-Salido, and Nelson (2005) and Edge, Kiley, and Laforge (2008), have tackled this issue more directly, but have found by and large that it is difficult to obtain model-based estimates of the output gap with reasonable cyclical properties. In comparison, our DSGE output gap is more consistent with conventional views of the business cycle, and broadly in line with the estimates of Sala, Soderstrom, and Trigari (2008), who were the first to obtain a cyclical output gap in an estimated DSGE model. In this positive dimension, our work is closely related to Gali, Gertler, and Lopez-Salido (2007), who also measure the extent of inefficient fluctuations, but do so through the lens of the labor wedge, rather than of the output gap.

However, our main contribution is not to find a more reasonable output gap than in previous studies, but rather to interpret this object from a more explicit normative perspective. In this dimension, we make contact with the optimal policy literature in medium-scale DSGE models (e.g. Levin, Onatski, Williams, and Williams (2005), Schmitt-Grohe and Uribe (2004) and Schmitt-Grohé and Uribe (2007)). As these authors, we find that nominal dispersion is key for the normative implications of the model. Unlike them, however, we find virtually no tension between inflation and output gap stabilization, once we recognize that wage markup shocks are likely to be small. The important role of wage markup shocks and labor supply shocks more in general relates our work to Chari, Kehoe, and McGrattan (2009) and Sala, Söderström, and Trigari (2010), although they are not concerned with the characterization of optimal policy.

The rest of the paper is organized as follows. Section 2 provides the details of the theoretical model and section 3 describes the approach to measurement and inference. Sections 4 and 5 present our estimates of potential and optimal output, while sections 6 and 7 analyze the role of labor supply shocks in the results. Section 8 concludes.

2. THE MODEL ECONOMY

This section outlines our baseline model of the U.S. business cycle, which is similar to Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). It is a medium-scale DSGE model with a neoclassical growth core, augmented with several shocks and “frictions”—departures from the simplest assumptions on tastes, technology and market structure—now common in the literature.

The economy is populated by five classes of agents: producers of a final good, intermediate goods producers, households, employment agencies and a government. We now present their optimization problems.

2.1. Final good producers. At every point in time t , perfectly competitive firms produce the final good Y_t by combining a continuum of intermediate goods $\{Y_t(i)\}_i$, $i \in [0, 1]$, according to the technology

$$Y_t = \left[\int_0^1 Y_t(i)^{\frac{1}{1+\Lambda_{p,t}}} di \right]^{1+\Lambda_{p,t}}.$$

Profit maximization and the zero profit condition imply that the price of the final good, P_t , is a CES aggregate of the prices of the intermediate goods, $\{P_t(i)\}_i$

$$P_t = \left[\int_0^1 P_t(i)^{-\frac{1}{\Lambda_{p,t}}} di \right]^{-\Lambda_{p,t}},$$

and that the demand function for intermediate good i is

$$(2.1) \quad Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\frac{1+\Lambda_{p,t}}{\Lambda_{p,t}}} Y_t.$$

The curvature of the aggregator $\Lambda_{p,t}$ determines the degree of substitutability across intermediate goods in the production of the final good and hence the elasticity of demand for each of these intermediates, as shown in (2.1). We model $\Lambda_{p,t}$ as an the exogenous stochastic process

$$\log(1 + \Lambda_{p,t}) \equiv \lambda_{p,t} = (1 - \rho_p)\lambda_p + \rho_p\lambda_{p,t-1} + \varepsilon_{p,t},$$

driven by innovations $\varepsilon_{p,t}$ distributed $i.i.d.N(0, \sigma_p^2)$. We refer to these innovations as *price markup shocks*, since, according to Lerner's formula, $\Lambda_{p,t}$ is the desired net markup of price over marginal cost for intermediate firms. Therefore, $\Lambda_{p,t}$ is also a measure of the (lack of) competitiveness in the intermediate goods market and its exogenous movements are one of the forces driving the fluctuations of the economy away from its efficient frontier.

2.2. Intermediate goods producers. A monopolist produces the intermediate good i using the production function

$$(2.2) \quad Y_t(i) = A_t^{1-\alpha} K_t(i)^\alpha L_t(i)^{1-\alpha} - A_t F,$$

where $K_t(i)$ and $L_t(i)$ denote the amounts of capital and labor employed by firm i . Both of these inputs are homogenous and F is a fixed cost of production, chosen so that profits are zero in steady state (see Rotemberg and Woodford, 1995 or Christiano, Eichenbaum, and Evans, 2005).

A_t represents exogenous labor-augmenting technological progress or, equivalently, a *neutral technology* factor. The level of neutral technology is non-stationary and its growth rate ($z_t \equiv \Delta \log A_t$) follows an AR(1) process

$$z_t = (1 - \rho_z)\gamma + \rho_z z_{t-1} + \varepsilon_{z,t},$$

with $\varepsilon_{z,t}$ $i.i.d.N(0, \sigma_z^2)$.

As in Calvo (1983), every period a fraction ξ_p of intermediate firms cannot optimally choose their price, but reset it according to the indexation rule

$$P_t(i) = P_{t-1}(i)\pi_{t-1}^{\iota_p}\pi^{1-\iota_p},$$

where $\pi_t \equiv \frac{P_t}{P_{t-1}}$ is gross inflation and π is its steady state. This indexation scheme implies no price dispersion in steady state. Therefore, the level of π is inconsequential for welfare, which allows us to abstract from the challenging question of the optimal level of inflation (see Schmitt-Grohe and Uribe (2010) for a recent survey.) In addition, full indexation produces a vertical Phillips curve in the long-run, so that the steady state level of output is independent from that of inflation, regardless of policy. As a consequence, monetary policy cannot bring the economy closer to its efficient frontier on average, even if this shift might be desirable.

The remaining fraction of firms choose their price, $\tilde{P}_t(i)$, by maximizing the present discounted value of future profits

$$E_t \sum_{s=0}^{\infty} \xi_p^s \frac{\beta^s \Lambda_{t+s}}{\Lambda_t} \left\{ \left[\tilde{P}_t(i) \left(\prod_{j=0}^s \pi_{t-1+j}^{\iota_p} \pi^{1-\iota_p} \right) \right] Y_{t+s}(i) - \left[W_t L_t(i) + r_t^k K_t(i) \right] \right\},$$

subject to the demand function 2.1 and the production function 2.2. In this objective, Λ_{t+s} is the marginal utility of nominal income of the representative household that owns the firm, while W_t and r_t^k are the nominal wage and the rental rate of capital.

2.3. Employment agencies. Firms are owned by a continuum of households, indexed by $j \in [0, 1]$. Each household is a monopolistic supplier of specialized labor, $L_t(j)$, as in Erceg, Henderson, and Levin (2000). A large number of competitive “employment agencies” combine these specialized types of labor into a homogenous labor input sold to intermediate firms, according to

$$(2.3) \quad L_t = \left[\int_0^1 L_t(j)^{\frac{1}{1+\Lambda_{w,t}}} dj \right]^{1+\Lambda_{w,t}}.$$

As in the case of the final good production, the elasticity of this aggregator $\Lambda_{w,t}$ corresponds to the desired markup of wages over households’ marginal rate of substitution between consumption and leisure. Unlike for the price markup shock, which is an AR(1) process, we assume instead that $\log(1 + \Lambda_{w,t}) \equiv \lambda_{w,t}$ is *i.i.d.* $N(0, \sigma_w^2)$, for reasons explained in the next subsection. We refer to $\lambda_{w,t}$ as the *wage markup shock*.

Profit maximization by the perfectly competitive employment agencies implies the labor demand function

$$(2.4) \quad L_t(j) = \left(\frac{W_t(j)}{W_t} \right)^{-\frac{1+\Lambda_{w,t}}{\Lambda_{w,t}}} L_t,$$

where $W_t(j)$ is the wage paid by the employment agencies to the supplier of labor of type j , while

$$W_t = \left[\int_0^1 W_t(j)^{-\frac{1}{\Lambda_{w,t}}} dj \right]^{-\Lambda_{w,t}}$$

is the wage paid by intermediate firms for the homogenous labor input sold to them by the agencies.

2.4. Households. Each household j maximizes the utility function

$$(2.5) \quad E_t \left\{ \sum_{s=0}^{\infty} \beta^s b_{t+s} \left[\log(C_{t+s} - hC_{t+s-1}) - \varphi_t \frac{L_{t+s}(j)^{1+\nu}}{1+\nu} \right] \right\},$$

where C_t is consumption and h is the degree of habit formation. The disturbance to the discount factor b_t is an *intertemporal preference shock* and follows the stochastic process

$$\log b_t = \rho_b \log b_{t-1} + \varepsilon_{b,t},$$

with $\varepsilon_{b,t} \sim i.i.d.N(0, \sigma_b^2)$. The disturbance to the disutility of labor φ_t is instead an *intratemporal preference* or *labor supply shock*, as in Hall (1997). This shock enters households' first order conditions for the optimal supply of labor in exactly the same way as the wage markup shock. As a consequence, these two disturbances are not separately identified in this model, when only using data on wages and total hours.² However, the implications of these two shocks for the evolution of potential output differ markedly, as also pointed out by Chari, Kehoe, and McGrattan (2009) and Sala, Söderström, and Trigari (2010).

To disentangle these two disturbances, therefore, we model the labor supply shock as an AR(1) process

$$\log \varphi_t = (1 - \rho_\varphi) \varphi + \rho_\varphi \log \varphi_{t-1} + \varepsilon_{\varphi,t},$$

with $\varepsilon_{\varphi,t} \sim i.i.d.N(0, \sigma_\varphi^2)$. The autocorrelation of the taste shock φ_t , which distinguishes it from the *i.i.d.* wage markup shock $\Lambda_{w,t}$, reflects the prior view that taste shocks are a more plausible device to capture the effects on hours worked of changes in labor force

² Galí (2010) and Galí, Smets, and Wouters (2011) propose a reinterpretation of this same model with an explicit treatment of unemployment, in which the two shocks can be separately identified. For a DSGE model with similar implications, but alternative microfoundations of unemployment, see Christiano, Trabandt, and Walentin (2010).

participation and other low-frequency movements largely unrelated to the business cycle, but are nonetheless evident in our data. We will return to the observational equivalence of these two shocks and to its normative implications in sections 7, where we show that reinterpreting the estimated labor disutility shocks as markup shocks is instead inconsequential for the properties of optimal output.

Since technological progress is non stationary, utility is logarithmic to ensure the existence of a balanced growth path. Moreover, consumption is not indexed by j because the existence of state contingent securities ensures that in equilibrium consumption and asset holdings are the same for all households.

As a result, the household's flow budget constraint is

$$P_t C_t + P_t I_t + T_t + B_t \leq R_{t-1} B_{t-1} + Q_t(j) + \Pi_t + W_t(j) L_t(j) + r_t^k u_t \bar{K}_{t-1} - P_t a(u_t) \bar{K}_{t-1},$$

where I_t is investment, B_t is holdings of government bonds, R_t is the gross nominal interest rate, $Q_t(j)$ is the net cash flow from household's j portfolio of state contingent securities, Π_t is the per-capita profit accruing to households from ownership of the firms and T_t is lump-sum taxes and transfers.

Households own capital and choose the capital utilization rate, u_t , which transforms physical capital into effective capital according to

$$K_t = u_t \bar{K}_{t-1}.$$

Effective capital is then rented to firms at the rate r_t^k . The cost of capital utilization is $a(u_t)$ per unit of physical capital. We parameterize it as $a(u_t) = \rho \frac{u_t^{1+\chi}-1}{1+\chi}$, as in Levin, Onatski, Williams, and Williams (2005), such that in steady state, $u = 1$, $a(1) = 0$ and $\chi \equiv \frac{a''(1)}{a'(1)}$. In the log-linear approximation of the model solution this curvature is the only parameter that matters for the dynamics.

The physical capital accumulation equation is

$$\bar{K}_t = (1 - \delta) \bar{K}_{t-1} + \mu_t \left(1 - S \left(\frac{I_t}{I_{t-1}} \right) \right) I_t,$$

where δ is the depreciation rate. The function S captures the presence of adjustment costs in investment, as in Christiano, Eichenbaum, and Evans (2005), and we parameterize it as $S(I_t/I_{t-1}) = \frac{\zeta}{2} (I_t/I_{t-1} - e^\gamma)^2$. In steady state, $S = S' = 0$ and $S'' = \nu > 0$. This coefficient is also the only one that matters for the log-linear dynamics. The *investment shock* μ_t is a source of exogenous variation in the efficiency with which the final good can be transformed

into physical capital, and thus into tomorrow's capital input. Justiniano, Primiceri, and Tambalotti (2010) show that this variation might stem from technological factors specific to the production of investment goods, as in Greenwood, Hercowitz, and Krusell (1997), but also from disturbances to the process by which these investment goods are turned into productive capital. The investment shock follows the stochastic process

$$\log \mu_t = \rho_\mu \log \mu_{t-1} + \varepsilon_{\mu,t},$$

where $\varepsilon_{\mu,t}$ is *i.i.d.* $N(0, \sigma_\mu^2)$.

As in Erceg, Henderson, and Levin (2000), every period a fraction ξ_w of households cannot freely set its wage, but follows the indexation rule

$$W_t(j) = W_{t-1}(j) (\pi_{t-1} e^{z_{t-1}})^{\iota_w} (\pi e^\gamma)^{1-\iota_w}.$$

The remaining fraction of households chooses instead an optimal wage by maximizing their utility, subject to the labor demand function 2.4.

2.5. Monetary and government policies. When estimating the model and characterizing its positive properties, such as the behavior of potential output, we assume that the short-term nominal interest rate follows a feedback rule, of the type that has been found to provide a good description of actual monetary policy in the United States at least since Taylor (1993). Our specification of this policy rule features interest rate smoothing, a systematic response to deviations of annual inflation from a time varying inflation target, and to deviations of observed annual GDP growth (X_t/X_{t-4}) from its steady state level

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R} \right)^{\rho_R} \left[\left(\frac{\left(\prod_{s=0}^3 \pi_{t-s} \right)^{1/4}}{\pi_t^*} \right)^{\phi_\pi} \left(\frac{(X_t/X_{t-4})^{1/4}}{e^\gamma} \right)^{\phi_X} \right]^{1-\rho_R} e^{\varepsilon_{R,t}},$$

where R is the steady state gross nominal interest rate and $\varepsilon_{R,t}$ is an *i.i.d.* $N(0, \sigma_R^2)$ monetary policy shock.

The inflation target π_t^* evolves exogenously according to the process

$$\log \pi_t^* = (1 - \rho_\pi) \log \pi + \rho_\pi \log \pi_{t-1}^* + \varepsilon_{\pi,t},$$

with $\varepsilon_{\pi,t} \sim i.i.d. N(0, \sigma_\pi^2)$. The primary role of this inflation target is to account for the very low frequency behavior of inflation (Ireland (2007)), which would otherwise imply a high

degree of backward price indexation (Cogley and Sbordone (2008)). We choose to minimize this source of inflation persistence because it affects significantly the relationship between inflation and welfare, but it does not appear to be supported by the micro evidence on firms' pricing behavior.

When characterizing optimal policy, of course, this rule is ignored, and we assume that the central bank sets the interest rate so as to directly maximize the utility of the representative agent.

Fiscal policy is Ricardian. The government finances its budget deficit by issuing short term bonds. Public spending is determined exogenously as a time-varying fraction of output

$$G_t = \left(1 - \frac{1}{g_t}\right) Y_t,$$

where the *government spending shock* g_t follows the stochastic process

$$\log g_t = (1 - \rho_g) \log g + \rho_g \log g_{t-1} + \varepsilon_{g,t},$$

with $\varepsilon_{g,t} \sim i.i.d. N(0, \sigma_g^2)$.

3. MODEL SOLUTION AND ESTIMATION

This section briefly describes the solution and estimation of the model. The first step in the solution consists of rewriting the equilibrium conditions in terms of deviations of the real variables from the non-stationary technology process A_t . Let

$$(3.1) \quad E_t [f(\zeta_{t+1}, \zeta_t, \zeta_{t-1}, e^{\varepsilon_t}, \theta)] = 0,$$

denote the collection of these equilibrium conditions, in which ζ_t , ε_t and θ are the vectors of endogenous variables, exogenous i.i.d. disturbances and unknown structural coefficients respectively. To obtain an estimate of our two main objects of interest, potential and optimal GDP, (3.1) must also include the equilibrium conditions of the corresponding counterfactual economies with stable markups and optimal monetary policy, and ζ_t must also contain all the variables necessary to characterize the dynamics of these counterfactual economies, including potential and optimal output.

We then log-linearize (3.1) around the non-stochastic steady state and solve the resulting linear system of rational expectation equations by standard methods (for example Sims (2001)). This procedure yields the following system of transition equations

$$(3.2) \quad \hat{\xi}_t = G(\theta) \hat{\xi}_{t-1} + M(\theta) \varepsilon_t,$$

where the “hat” denotes log deviations from the steady state, $\hat{\xi}_t$ is an extended version of $\hat{\zeta}_t$ that also includes the expectational variables that are necessary to characterize the solution of the model, and $G(\theta)$ and $M(\theta)$ are conformable matrices whose elements are functions of θ .

3.1. Data and Measurement. We estimate the model using eight series of U.S. quarterly data: the inflation rate, the nominal interest rate, the logarithm of per-capita hours, the log-difference of real per-capita GDP, consumption and investment, and *two* measures of nominal hourly wage inflation.

The inflation rate is the quarterly log difference of the GDP deflator, while the nominal interest rate is the effective Federal Funds rate. We measure per-capita hours as the number of hours worked in the total economy, divided by the civilian non-institutional population (16 years and older).³ The series of hours for the total economy exhibits a less pronounced low frequency behavior than that for the non-farm business sector, because it accounts better for sectoral shifts, as shown by Francis and Ramey (2009). Real per-capita GDP is nominal GDP divided by population and the GDP deflator. The real series for per-capita consumption and investment are obtained in the same manner. Consumption corresponds to the sum of non-durables and services, while investment is constructed by adding consumer durables to total private investment, all in nominal terms.

As pointed out in the introduction, we match the wage variable in the model, W_t , with two data series, following the methodology proposed by Boivin and Giannoni (2006a) and recently also adopted by Gali, Smets, and Wouters (2011). The first series is nominal compensation per hour in the total economy, from NIPA. The second measure is the “average hourly earnings of production and non-supervisory employees,” which is computed by the Bureau of Labor Statistics from the Establishment Survey, and is the one preferred by Galí (2010). We assume that both series are imperfect measures of the wage in the model, and allow each of them to be affected by an i.i.d. observation error. Finally, we normalize to one the coefficient relating the model’s wage to the first series, while we estimate the loading (Γ) on the other series, as standard in factor analysis. We will return to the reasons for this approach to the measurement of wages, and to its positive and normative implications in section 6.

³ We are grateful to Shawn Sprague, of the Bureau of Labor Statistics, for providing the data on hours and the labor share in the total economy. Breaks in the civilian population series due to census-based population adjustments are smoothed by splicing them uniformly over a 10-year window.

Denoting by d_t the vector of the observable series at time t , we link it to the corresponding variables in the model through the system of observation equations

$$(3.3) \quad d_t = H(\theta) \hat{\xi}_t + d(\theta) + e_t,$$

where e_t is a Gaussian *i.i.d.* measurement error that only enters the equations corresponding to the two wage series. The estimation sample starts in 1964:II, due to limited availability of the wage data, and ends in 2009:IV. We do not demean or detrend any series.

3.2. Bayesian Inference and Priors. We characterize the posterior distribution of the model's coefficients by combining the likelihood function with prior information (see An and Schorfheide (2007) for a survey of Bayesian estimation of DSGE models). The likelihood function can be evaluated by applying the Kalman Filter to the linear and Gaussian state space representation of the model, given by the state equation (3.2) and the measurement equation (3.3). Conditional on the sample information, the Kalman filter and smoother can also be used to estimate the historical path of the model's endogenous variables, $\{\hat{\xi}_t\}_{t=1}^T$, which include potential and optimal output. In the rest of this section we briefly discuss the specification of the priors, which is reported in Table 1.

Two parameters are fixed using level information not contained in our dataset: the quarterly depreciation rate of capital (δ) to 0.025 and the steady state ratio of government spending to GDP ($1 - 1/g$) to 0.2, which corresponds to the average value of G_t/X_t in our sample. Also due to lack of identification, we set the steady state net wage markup to 25 percent. The priors on the other coefficients are fairly diffuse and broadly in line with those adopted in previous studies, such as Smets and Wouters (2007) and Del Negro, Schorfheide, Smets, and Wouters (2007).

The prior distribution of all but two persistence parameters is a Beta, with mean 0.6 and standard deviation 0.2. The two exceptions are the autocorrelation of TFP shocks—whose prior is centered at 0.4, since this process already includes a unit root—and the autocorrelation of the inflation target shock, which we fix at 0.995. This value reflects the view that the exogenous movements of the inflation target should account for the very low frequency behavior of inflation. In reduced form, these secular movements might reflect the slow evolution of policymakers' beliefs and the consequent changes in the conduct of monetary policy, as suggested for instance by Cogley and Sargent (2004) and Primiceri (2006).

The intertemporal preference, price and wage markup shocks are normalized to enter with a unit coefficient in the consumption, price inflation and wage equations respectively (see appendix A for details). The priors on the standard deviations of the innovations to these normalized shocks are quite disperse and chosen to generate volatilities for the variables they impact directly broadly in line with those in the data. The covariance matrix of the vector of shocks is diagonal.

3.3. Posterior Estimates of the Parameters. Table 1 summarizes the posterior estimates of the parameters in our baseline specification. The data are quite informative about these parameters and the estimates we obtain are generally in line with those of previous studies. For this reason, and given the focus of our paper on the implications of these estimates for the inefficiency of the economy, and for optimal policy, we only briefly comment on the coefficients related to nominal rigidities and to monetary policy. The posterior distributions of the parameters ξ_p and ξ_w imply that prices and wages are re-optimized approximately every year and every three quarters respectively, while ι_p and ι_w indicate very low levels of backward indexation. As for monetary policy, it is fairly inertial, with ρ_R around 0.7, and it exhibits a substantial degree of activism, with interest rates responding with a long-run coefficient of more than 2 to inflation and of almost 1 to output growth.

4. POTENTIAL OUTPUT AND THE OUTPUT GAP

With an estimated structural model in hand, we are now ready to explore the relationship between the actual economy, as observed over the past fifty years, and its unobserved efficient frontier, as inferred from the model. In our environment, the observed macroeconomic outcomes deviate from those under perfect competition, and thus from efficiency, due to the presence of monopoly power in goods and labor markets. This monopoly power, which stems from the imperfect substitutability of intermediate goods and of specialized labor services, allows firms to price their output above marginal cost and households to price their labor above the marginal rate of substitution.

In the aggregate, these markups create a wedge in the intratemporal efficiency condition, the equality of the marginal rates of substitution (MRS) and transformation (MRT) between

TABLE 1. Prior distributions and posterior parameter estimates in the baseline model.

		Prior			Posterior					
		Dist	Mean	SE	Mode	Mean	SE	5%	Median	95%
α	Capital share	N	0.30	0.05	0.16	0.16	0.01	0.15	0.16	0.17
ι_p	Price indexation	B	0.50	0.15	0.15	0.16	0.07	0.06	0.15	0.29
ι_w	Wage indexation	B	0.50	0.15	0.04	0.04	0.02	0.02	0.04	0.08
γ	SS tech. growth	N	0.50	0.03	0.47	0.47	0.02	0.43	0.47	0.51
h	Habit formation	B	0.60	0.10	0.80	0.82	0.04	0.75	0.82	0.88
λ_p	SS price markup	N	0.15	0.05	0.25	0.25	0.04	0.18	0.25	0.32
$\log L_{ss}$	SS log hours	N	0.00	0.50	0.00	0.01	0.49	-0.79	0.00	0.80
$100(\pi - 1)$	SS inflation	N	0.62	0.10	0.24	0.24	0.07	0.13	0.24	0.35
$100(\beta^{-1} - 1)$	Discount factor	G	0.25	0.10	0.17	0.18	0.05	0.10	0.18	0.27
ν	Inverse Frisch	G	2.00	0.75	2.36	2.73	0.76	1.62	2.67	4.12
ξ_p	Price stickiness	B	0.66	0.10	0.84	0.84	0.03	0.79	0.84	0.89
ξ_w	Wage stickiness	B	0.66	0.10	0.72	0.73	0.04	0.66	0.73	0.80
χ	Elasticity util. cost	G	5.00	1.00	5.10	5.31	1.01	3.77	5.25	7.08
S''	Invest. adj. costs	G	4.00	1.00	3.64	3.93	0.78	2.73	3.87	5.30
ϕ_π	Reaction infl.	N	1.70	0.30	2.35	2.32	0.21	1.97	2.33	2.66
ϕ_X	Reaction GDP gr.	N	0.40	0.30	0.82	0.85	0.14	0.63	0.84	1.10
Γ	Loading coeff.	N	1.00	0.50	0.65	0.65	0.03	0.59	0.65	0.70
ρ_R	Auto. mp	B	0.60	0.20	0.69	0.70	0.05	0.60	0.70	0.77
ρ_z	Auto. tech.	B	0.40	0.20	0.13	0.13	0.06	0.04	0.13	0.23
ρ_g	Auto. gov. spending	B	0.60	0.20	1.00	1.00	0.00	0.99	1.00	1.00
ρ_μ	Auto. investment	B	0.60	0.20	0.69	0.69	0.05	0.60	0.69	0.77
ρ_p	Auto. price markup	B	0.60	0.20	0.25	0.28	0.15	0.08	0.25	0.56
ρ_φ	Auto. labor supply	B	0.60	0.20	0.98	0.97	0.01	0.95	0.98	0.99
ρ_b	Auto. intertemporal	B	0.60	0.20	0.68	0.63	0.08	0.49	0.64	0.76
$100\sigma_R$	Std mp	IG1	0.15	1.00	0.22	0.22	0.01	0.20	0.22	0.25
$100\sigma_z$	Std tech.	IG1	1.00	1.00	0.86	0.87	0.05	0.79	0.86	0.95
$100\sigma_g$	Std gov. spending	IG1	0.50	1.00	0.36	0.36	0.02	0.34	0.36	0.40
$100\sigma_\mu$	Std investment	IG1	0.50	1.00	6.98	7.56	1.46	5.39	7.45	10.17
$100\sigma_p$	Std price markup	IG1	0.15	1.00	0.17	0.16	0.02	0.12	0.16	0.19
$100\sigma_\varphi$	Std labor supply	IG1	1.00	1.00	3.86	4.73	1.58	2.88	4.49	7.23
$100\sigma_b$	Std intertemporal	IG1	0.10	1.00	0.03	0.03	0.01	0.02	0.03	0.05
$100\sigma_w$	Std wage markup	IG1	0.15	1.00	0.05	0.06	0.02	0.04	0.06	0.09
$100\sigma_\pi$	Std inflation target	IG1	0.05	0.03	0.05	0.05	0.01	0.04	0.05	0.07
$100\sigma_{1,mew}$	Std meas. error 1	IG1	0.15	1.00	0.49	0.49	0.03	0.45	0.49	0.55
$100\sigma_{2,mew}$	Std meas. error 2	IG1	0.15	1.00	0.28	0.28	0.02	0.24	0.28	0.32

consumption and leisure. More formally, define the aggregate price markup as

$$\mu_t^p = \frac{P_t}{MC_t},$$

and the wage markup as

$$\mu_t^w = \frac{W_t}{P_t MRS_t},$$

so that

$$MRS_t \mu_t^w \mu_t^p = MRT_t,$$

where we used the fact that the nominal marginal cost $MC_t = W_t/MPL_t$, and the marginal product of labor MPL_t is also the marginal rate of transformation between labor and final consumption (see also Gali, Gertler, and Lopez-Salido (2007)).

The equilibrium price and wage markups, μ_t^p and μ_t^w , vary over time for two reasons. First, exogenous shifts in the substitutability of goods and labor services, and thus in the elasticity of their demand, affect firms and workers' market power and desired markups. These shifts are captured by the stochastic processes $\Lambda_{p,t}$ and $\Lambda_{w,t}$. The second source of equilibrium markup variation depends on the presence of nominal rigidities, which prevent firms and workers from achieving their desired markups at any given point in time. The resulting movements in average markups are endogenous and can be triggered by any shock hitting the economy.

Given these distortions, we define potential output as the counterfactual level of output that would be observed if equilibrium price and wage markups were constant at their (positive) steady state levels Λ_p and Λ_w .⁴ This level of output would occur in equilibrium if prices and wages were flexible, and desired markups did not vary over time. Observe that potential output is the same as efficient output—equilibrium output in a competitive economy—up to a shift in the steady state that has no implications for the approximate log-linear dynamics.⁵ The reason for focusing on potential rather than on efficient output is that the current version of our model abstracts from distortions that are likely to be particularly relevant in steady

⁴ In the model, there is a small discrepancy between output and GDP, due to the presence of capital utilization costs. In the text, we sometimes refer to GDP as “output,” even if this usage is slightly imprecise.

⁵ The dynamic equivalence between the potential and efficient economies requires that the fixed cost F does not change with the degree of competition. If F , which is introduced in the model to “absorb” the steady state profits generated by monopolistic competition, is assumed to vanish as the economy becomes competitive, and those profits go to zero, the dynamic properties of the efficient equilibrium will be affected. However, we verified that this effect is quantitatively small, so we ignore it here.

FIGURE 1. GDP, potential GDP and the output gap in the baseline model. In panel (b), the blu line is the median of the posterior distribution of the output gap, while the four shades of green denote 60, 70, 80, and 90 percent posterior probability intervals.

state, such as taxes. Incorporating these other distortions into the analysis of the model's steady state, and also of its dynamics, is an important avenue for future research.

Figure 1a plots the logarithm of U.S. GDP and the posterior median of *potential output* from the DSGE model. The latter is on average somewhat more stable than actual output, especially starting in the mid-eighties, but it also displays sizable fluctuations. To focus more closely on cyclical variation, Figure 1b plots the posterior distribution of the model-implied *output gap*, which is simply the difference between output and the DSGE-based measure of potential. The shaded areas correspond to the NBER recessions. The DSGE-based output gap displays a pronounced cyclical behavior: it peaks at the end of most expansions and declines during recessions.

More specifically, figure 1b shows that the output gap declined steadily during the “twin” recessions of the early 1980s, which brought the economy back to its potential equilibrium after the “overheating” of the 1970s, particularly evident towards the end of that decade. This evolution of the output gap is consistent with the conventional wisdom that points to the importance of monetary and other nominal factors in driving real macroeconomic outcomes during that historical period. However, the level of the output gap remains positive throughout the two recessions, given the record levels it had reached just prior to the Volcker disinflation. Our estimates also indicate a slightly negative gap during most of the 1990s, implying that actual output was relatively slow in catching up with a surge in potential. This historical reconstruction squares well with the view that this period was characterized by a considerable acceleration in productivity. Finally, we observe that the last recession is associated with a substantial fall in output relative to the DSGE-based potential. This decline is driven by large negative investment shocks, whose origin can be traced to financial disturbances further propagated by nominal rigidities, as argued in Justiniano, Primiceri, and Tambalotti (2010).

The pattern of deviations of GDP from potential just described implies that output fluctuations in the postwar period would have been less pronounced if markets had been competitive. To quantify this statement, we focus on business cycle dynamics and decompose the deviation of GDP from a Hodrick-Prescott (HP) trend into two parts: the deviations of potential GDP and of the output gap from their respective HP trends. This is an exact decomposition, since the HP filter is linear and thus additive. Figure 2 shows the results of this exercise. First, note how the business cycle component of the DSGE gap peaks and troughs close to the NBER dates, with some delay in the last three “slow” recoveries. Second, the decomposition highlights that the DSGE model’s estimate of potential output is roughly consistent with the Real Business Cycle view of fluctuations, in the sense that its efficient frontier displays movements over the business cycle of similar magnitude to those of actual output (the correlation between detrended potential and output is equal to 0.75).⁶ This volatility in the efficient frontier, however, coexists with a substantial volatility of the HP-detrended output gap, which is about 70 percent of that of detrended GDP (the correlation between the detrended gap and output is equal to 0.23). This evidence is the basis for our first important

⁶ Our neutral technology shock has a smaller impact on fluctuations than in the RBC literature, due to the presence of investment adjustment costs and to a low estimate of the elasticity of labor supply. However, other shocks are also responsible for movements in the economy’s efficient frontier.

FIGURE 2. GDP, potential GDP and the output gap in the baseline model, HP filtered.

conclusion, that time variation in the economy's degree of inefficiency is a crucial factor in macroeconomic fluctuations.

This conclusion is consistent with the evidence in Galí, Gertler, and López-Salido (2007) and Sala, Söderström, and Trigari (2008), but stands in contrast to the results of many other studies in the empirical DSGE literature, as pointed out by Walsh (2005) and Mishkin (2007). This literature has typically estimated small DSGE output gaps, with little cyclical variation, but its findings can often be traced back to modeling assumptions that seem at odds with the data. For example, Andrés, López-Salido, and Nelson (2005) calibrate the coefficient of backward price indexation to one and do not include in their model markup shocks or any other disturbance that directly affects the Phillips curve. With these tight restrictions, marginal costs are forced to explain the high frequency fluctuations in the first difference of inflation. To make this possible, the slope of the Phillips curve must be quite large, erroneously suggesting that nominal rigidities are nearly irrelevant for cyclical fluctuations. Edge, Kiley, and LaForte (2008) instead estimate a large scale model, with data starting in

the early 1980s, but without an inflation target shock. Since the inflation rate exhibits a pronounced downward trend during this period, the inferred output gap inherits a similar trend. Finally, Levin, Onatski, Williams, and Williams (2005) detrend the real series prior to estimation, and assume that monetary policy responds directly to the output gap. Since their estimate of the output gap is small and countercyclical, this policy rule is at odds with the observed behavior of the Federal Reserve. If we impose the same restrictions assumed by these authors on our model, we can approximately replicate their results.

5. OPTIMAL OUTPUT AND THE (LACK OF A) POLICY TRADE-OFF

The previous section documented sizable movements in the degree of inefficiency of the U.S. economy in the last fifty years. Why did stabilization policy not counteract these inefficient fluctuations? One possible reason is that these movements in the output gap represent the price policymakers had to pay to minimize other distortions. And indeed, in our economy, policymakers face a trade-off between output gap stabilization, on the one hand, and price and wage stabilization, on the other.

This trade-off stems from the fact that, in the equilibrium in which output is stabilized around potential, desired prices and wages—those that agents would set in the absence of nominal rigidities—are in general not constant. For example, desired prices change mechanically in response to changes in desired markups due to markup shocks. A more subtle reason for movements in desired prices is the coexistence of price and wage stickiness. With sticky wages, an increase in the marginal product of labor, due for example to an increase in productivity, will result in a fall in firms' marginal costs, and hence in their desired price, since wages cannot fully adjust to absorb the productivity shock (Erceg, Henderson, and Levin (2000)).

Due to these movements in desired prices, workers and firms that reprice at different times will charge different prices. The resulting cross-sectional price and wage dispersion, whose movements over time go hand in hand with instability in price and wage inflation, is inefficient, since it forces workers and firms with identical tastes and technologies to supply different amounts of hours and of intermediate goods. Given that the aggregation of these inputs into aggregate utility and final output is concave, society would be better off if production and labor effort were shared equally across agents.

We can illustrate the inefficiency of an asymmetric distribution of the intermediate goods supply caused by price dispersion by aggregating the production functions (2.2) across firms i . Using demand (2.1), this aggregation produces a production function for the final good of the form

$$\Delta_t Y_t = A_t^{1-\alpha} K_t^\alpha L_t^{1-\alpha} - A_t F,$$

where

$$\Delta_t \equiv \int_0^1 \left(\frac{P_t(i)}{P_t} \right)^{-\frac{1+\Lambda_{p,t}}{\Lambda_{p,t}}} di$$

is a measure of price dispersion. Increases in price dispersion have the same effect as a fall in aggregate productivity, lowering the output of the final good, for any given level of the inputs. A similar dispersion term for wages directly reduces the utility of the average household, for any given level of the homogenous labor input L_t , due to the concavity of the labor aggregator (2.3).

To summarize, a stable output gap is in general incompatible with the absence of cross-sectional dispersion in prices and quantities in the intermediate goods and labor market, and therefore with stable price and wage inflation. As a result, stabilization policy faces a trade-off between these three objectives. In fact, in our economy, unlike in that of Erceg, Henderson, and Levin (2000) for instance, optimal policy needs to balance more than just these three goals, due to the presence of capital accumulation, and of all the other “frictions” needed to improve the model’s ability to fit the data.

For a more comprehensive and quantitative insight into these trade-offs, we turn to the numerical study of the model’s optimal equilibrium. This is computed by maximizing the utility of the average household, subject to the constraints represented by the optimal behavior of private agents, following the approach developed by Benigno and Woodford (2006) (see also Schmitt-Grohé and Uribe (2007) or Christiano, Ilut, Motto, and Rostagno (2010)). In the optimization, we assume that the short-term nominal interest rate is the only instrument available to the planner, which thus defines this problem as one of optimal monetary policy.

Figure 3a compares actual, potential and optimal output. To do so in a parsimonious way, and to maintain comparability with figure 1 above, actual and optimal output are both presented in deviation from potential output. Figure 3b and c plot instead observed price and wage inflation, as well as their counterfactual evolution under the optimal policy. The picture speaks for itself. Optimal and potential output move extremely closely together—the

FIGURE 3. Optimal output, price and wage inflation in the baseline model.

optimal output gap is virtually zero at all times—but this stabilization of the output gap is also consistent with a significant reduction in the volatility of price and, especially, wage inflation. The surprising implication of these findings, and the central result of the paper, is that stabilization policy appears to have faced a negligible trade-off among its three main objectives over our sample, and that much of the inefficient variation in output documented in section 4 could have been avoided, while at the same time reaping the benefit of more stable inflation.

6. WAGE MARKUP SHOCKS: FACT OR FICTION?

The key to the absence of a significant trade-off between output and inflation stabilization documented above is the treatment of wages in the estimation of the model. As described in section 3.1, and following the general methodology proposed by Boivin and Giannoni (2006a), our estimation procedure uses two wage series in the measurement equation, with the assumption that each series matches only imperfectly the model’s concept of “the wage”.

FIGURE 4. Two measures of nominal wage inflation.

As shown in figure 4, these two series behave similarly at medium and low frequencies, but differ markedly at high frequencies, reflecting some conceptual differences in what they attempt to measure, and more in general the well-known difficulties in measuring aggregate wages (e.g. Abraham, Spletzer, and Stewart (1999), Bosworth and Perry (1994)). Given these large high-frequency discrepancies, the use of two wage indicators helps to distinguish idiosyncratic measurement errors—high frequency wage movements uncorrelated across the two indicators—from macroeconomic factors, which should produce fluctuations in both series.

To highlight the importance of this measurement choice, this section presents estimates of our model based on only one wage series with no measurement error, as in most empirical DSGE exercises.⁷ Table 2 reports the posterior estimates for this alternative empirical specification, which are generally very similar to those for the baseline. Two relevant exceptions are steady state inflation (which in the baseline is affected by the lower mean of the inflation

⁷ We use the NIPA series of hourly compensation for the total economy, which is the more volatile of the two. This choice of data also allows us to estimate the model on a longer sample, starting in 1954. The typical estimation of DSGE models (for example Smets and Wouters (2007)) is conducted using compensation per hour in the non-farm business sector, which is even more volatile.

rate in the compensation series), and the persistence of the price markup shock, which is now substantially higher.

The crucial difference in the estimates, however, is that the standard deviation of the wage markup shock is now almost six times larger than in the baseline. In addition, the variance decomposition of the model reveals that the explanatory power of this shock is large for wage inflation, especially at high frequencies, but very small for the real variables.⁸ These observations, combined with the dubious microfoundations of these shocks (Chari, Kehoe, and McGrattan (2009), Shimer (2009)), suggest that they are likely to reflect mostly measurement error, rather than large variations in workers' monopoly power from one quarter to the next.

The extreme volatility of wage markup shocks in the model with only one wage measure, which we argued is empirically implausible, also has very misleading normative implications. This is because these large shocks directly affect workers' desired wages, thus generating wide dispersion in relative wages when nominal adjustments are staggered. Given the large welfare costs due to wage dispersion, the planner is willing to pay a high price to reduce it as much as possible. Figure 5a shows that, in this model estimated with a single wage series, this price is paid in the form of extremely volatile optimal output, which is consistent with Rotemberg and Woodford (1999). To reduce the variability of wage and price inflation, as shown in figures 5b and c, the planner must accept a path for optimal output that is significantly more volatile than that of actual output, at both business cycle and higher frequencies.

The striking implication of our calculations is that, in this case, optimal policy should have *de*-stabilized aggregate real activity, in order to reduce the volatility of price and, especially, wage inflation. In comparison, recall that this tension between output and inflation stabilization virtually disappears in the model with multiple wage indicators, in which the variation in wage markups is negligible.

Figure 6 confirms that the excess volatility of optimal output in this model is due to implausibly large wage markup shocks. The three panels plot GDP, price and wage inflation in the optimal allocation of an economy identical to the one depicted in Figure 5, but in which the wage markup shocks have been (arbitrarily) set to zero. In this economy, the deviations of optimal output from potential are much smaller than in figure 5, and more in line with

⁸ Wage markup shocks explain only 2 percent of GDP and hours volatility over the business cycle and 1 percent of consumption and investment.

TABLE 2. Prior distributions and posterior parameter estimates in the model with only one observable wage series (compensation).

		Prior			Posterior					
		Dist	Mean	SE	Mode	Mean	SE	5%	Median	95%
α	Capital share	N	0.30	0.05	0.17	0.17	0.01	0.16	0.17	0.18
ι_p	Price indexation	B	0.50	0.15	0.05	0.06	0.03	0.03	0.06	0.11
ι_w	Wage indexation	B	0.50	0.15	0.08	0.08	0.03	0.04	0.08	0.13
γ	SS tech. growth	N	0.50	0.03	0.49	0.49	0.02	0.45	0.49	0.53
h	Habit formation	B	0.60	0.10	0.84	0.84	0.03	0.78	0.84	0.89
λ_p	SS price markup	N	0.15	0.05	0.29	0.29	0.04	0.23	0.29	0.35
$\log L_{ss}$	SS log hours	N	0.00	0.50	0.03	0.04	0.49	-0.76	0.04	0.86
$100(\pi - 1)$	SS inflation	N	0.62	0.10	0.62	0.62	0.10	0.46	0.62	0.78
$100(\beta^{-1} - 1)$	Discount factor	G	0.25	0.10	0.12	0.13	0.04	0.07	0.13	0.20
ν	Inverse Frisch	G	2.00	0.75	1.96	2.32	0.69	1.32	2.24	3.56
ξ_p	Price stickiness	B	0.66	0.10	0.77	0.78	0.03	0.73	0.78	0.83
ξ_w	Wage stickiness	B	0.66	0.10	0.73	0.74	0.05	0.66	0.74	0.81
χ	Elasticity util. cost	G	5.00	1.00	4.95	5.17	1.03	3.61	5.10	6.97
S''	Invest. adj. costs	G	4.00	1.00	3.58	3.86	0.74	2.73	3.81	5.16
ϕ_π	Reaction infl.	N	1.70	0.30	2.10	2.17	0.20	1.86	2.17	2.51
ϕ_X	Reaction GDP gr.	N	0.40	0.30	0.94	0.94	0.13	0.73	0.94	1.17
ρ_R	Auto. mp	B	0.60	0.20	0.78	0.78	0.03	0.73	0.78	0.82
ρ_z	Auto. tech.	B	0.40	0.20	0.08	0.09	0.04	0.03	0.09	0.18
ρ_g	Auto. gov. spending	B	0.60	0.20	1.00	0.99	0.00	0.99	0.99	1.00
ρ_μ	Auto. investment	B	0.60	0.20	0.74	0.72	0.05	0.63	0.73	0.81
ρ_p	Auto. price markup	B	0.60	0.20	0.90	0.88	0.05	0.78	0.88	0.95
ρ_φ	Auto. labor supply	B	0.60	0.20	0.98	0.98	0.01	0.96	0.98	0.99
ρ_b	Auto. intertemporal	B	0.60	0.20	0.44	0.45	0.11	0.28	0.45	0.64
$100\sigma_R$	Std mp	IG1	0.15	1.00	0.21	0.22	0.01	0.20	0.22	0.24
$100\sigma_z$	Std tech.	IG1	1.00	1.00	0.92	0.93	0.05	0.85	0.93	1.01
$100\sigma_g$	Std gov. spending	IG1	0.50	1.00	0.37	0.37	0.02	0.34	0.37	0.40
$100\sigma_\mu$	Std investment	IG1	0.50	1.00	7.03	7.69	1.37	5.66	7.55	10.13
$100\sigma_p$	Std price markup	IG1	0.15	1.00	0.10	0.10	0.01	0.08	0.10	0.12
$100\sigma_\varphi$	Std labor supply	IG1	1.00	1.00	2.99	3.53	1.09	2.11	3.35	5.55
$100\sigma_b$	Std intertemporal	IG1	0.10	1.00	0.06	0.06	0.02	0.04	0.06	0.09
$100\sigma_w$	Std wage markup	IG1	0.15	1.00	0.30	0.30	0.02	0.27	0.30	0.33
$100\sigma_\pi$	Std inflation target	IG1	0.05	0.03	0.03	0.04	0.01	0.03	0.04	0.06

FIGURE 5. Optimal output, price and wage inflation, in the model with only one observable wage series (compensation).

those in the baseline model. The same is true for price and wage inflation, although the former is somewhat more volatile. This is because, in the one-wage model, the unconditional volatility of price markup shocks is higher than in the baseline estimation.

The conclusion we draw from these experiments is that ignoring measurement error in wages might lead to the erroneous normative implication that the output-inflation stabilization trade-off is substantial, and that, among their conflicting objectives, policymakers should prioritize the stabilization of wage inflation, even if this choice entails destabilizing output.

7. THE INTERPRETATION OF LABOR SUPPLY SHOCKS

As illustrated in section 2, the wage markup and labor disutility shocks are in general observationally equivalent in our empirical framework, and the only way to disentangle them is to posit a priori that they have different spectral profiles and hence contribute to fluctuations at different frequencies. This is the approach we have followed, by assuming that the wage markup shock is *i.i.d.*, while the labor disutility shock is a fairly persistent AR(1) process.

FIGURE 6. Optimal output, price and wage inflation in the model with only one observable wage series (compensation) and no wage markup shocks.

This identification strategy captures the idea that the large low frequency variation in labor supply identified by the model is mostly attributable to shifts in households' preferences for market work, such as the secular change in attitudes towards female labor participation that started in the 1960s, or other slow moving demographic developments. However, if we were to attribute all or part of this low frequency variation in hours to changes in the competitiveness of the labor market, our estimate of potential output would change significantly, since this shock directly affects the distance of the economy from its efficient frontier.

Figure 7 illustrates this point quantitatively, by comparing the estimates of potential output arising from two alternative interpretations of the source of low frequency fluctuations in hours, as in Sala, Söderström, and Trigari (2010). The first interpretation corresponds to our baseline assumption that these low frequency movements are mainly due to labor disutility shocks. Under the alternative interpretation, there are no taste shocks, and the low-frequency fluctuations in hours originate from changes in the monopoly power of workers, which do not

FIGURE 7. Effect of alternative assumptions on the origin of labor supply shocks on potential and optimal GDP. The blue line represents the difference between potential output in the model with all wage markup shocks and the baseline model with both wage markup and disutility of labor shocks. The red dashed line represents the difference in optimal outputs across the same two models.

affect the efficient equilibrium.⁹ The solid line in Figure 7 represents the difference between the estimates of potential output under the two scenarios. This difference is very large—ranging from -10 to $+15$ percent—but most of its variation is at low frequencies. This should not be too surprising, given that the shifts in the disutility of labor we estimate must be highly persistent and very large to explain the wide secular swings in total hours observed in the sample.

We just showed that the behavior of potential output changes considerably with respect to our baseline estimates under the extreme assumption that labor supply shocks only reflect changes in workers' market power. However, this is not the case for optimal output, as seen

⁹ Given the observational equivalence of these two alternative assumptions, the model does not need to be re-estimated.

from the dashed line in figure 7, which depicts the difference between the estimates of optimal output under the two alternative scenarios. This difference is an order of magnitude smaller than that between the estimates of potential. Optimal output is relatively invariant to the source of labor supply shifts because the planner has a strong distaste for wage dispersion, and thus chooses a very stable path for wage inflation, as in the baseline case of section 5. As a result, quantities—hours and output in particular—bear the brunt of the adjustment to labor supply shocks in the optimal equilibrium, regardless of these shocks’ origin.

According to the results in figure 7, the output gap is subject to large swings in the optimal equilibrium if labor supply shocks reflect changes in workers’ market power. Consequently, the policy trade-off between real and inflation stabilization is significant in this case. However, the fact that most of the variation in the output gap is concentrated at low frequencies, and is caused by changes in market power, suggests that monetary policy might not be the most appropriate tool to address this distortion, and that more targeted fiscal or microeconomic policies could be more effective in eliminating it.

Moreover, the fact that optimal output is roughly invariant in the two scenarios considered in figure 7 implies that researchers interested in characterizing the equilibrium implications of optimal policy do not need to take a strong stance on the ultimate sources of labor supply shocks, an issue that remains unsettled in the literature (e.g. Gali, Smets, and Wouters (2011) and Christiano, Trabandt, and Walentin (2010)). One consequence of this finding is that medium scale New-Keynesian models like the one presented here might be more useful for policy analysis than suggested for instance by Chari, Kehoe, and McGrattan (2009).

8. CONCLUSIONS

We estimated a New-Neoclassical Synthesis model of the U.S. economy. In this model, workers and firms have some monopoly power, and markups of price over marginal cost and of wages over the marginal rate of substitution fluctuate due to exogenous changes in monopoly power, as well as to the presence of sticky prices and wages. According to our estimates, the movements in these markups were associated with large fluctuations in U.S. GDP away from its efficient frontier. Moreover, we find that monetary policy could have virtually eliminated these inefficient output movements, reducing at the same time the volatility of price and wage inflation, and thus achieving an improvement in welfare for the model’s representative agent.

The key factor driving this surprising absence of a meaningful trade-off between output (gap) and inflation stabilization in our model is that we estimate little exogenous variation in desired wage markups, unlike in most previous empirical DSGE studies. This is because we match the model’s “wage” with two measures of earnings and compensation, which allows us to better isolate high frequency idiosyncrasies specific to each series, from a common component that is more likely to represent genuine macroeconomic factors. Indeed, the same model estimated with only one wage series finds implausibly large high frequency fluctuations in desired wage markups, which generate a strong tension between output and inflation stabilization. In this model, optimal policy would be forced to resolve this conflict by inducing higher output volatility than observed in the data, in order to obtain a smoother path of inflation.

Our results point to a significant discrepancy between the model’s equilibrium under historical monetary policy, as described by the estimated interest rate rule, and the optimal equilibrium. What are the reasons for this discrepancy? One possibility is that the model’s welfare function is not a good representation of the actual objectives of U.S. monetary policy. For example, wage distortions loom large in the utility of the representative agent in the model, but they are seldom mentioned as a direct preoccupation by policymakers. A second possibility is that we might have overlooked some relevant constraint in setting interest rates. The most obvious one in this respect is the zero lower bound on nominal interest rates. However, a preliminary exploration of this issue suggests that this bound is not violated often in the optimal equilibrium. Finally, actual monetary policy might have been misguided, at least at some points in the past. This is a fairly common conclusion among researchers, especially for the period roughly between 1965 and 1980 (e.g. Clarida, Gali, and Gertler (2000), Cogley and Sargent (2004) and Primiceri (2006)). The fact that we assume an invariant policy rule across the entire span of our long sample prevents us from addressing this possibility with finer historical detail, although doing so would be useful.

APPENDIX A. NORMALIZATION OF THE SHOCKS

As in Smets and Wouters (2007), some of the exogenous shocks are re-normalized by a constant term. In particular, we normalize the price and wage markups shocks and the intertemporal preference shock.

More specifically, the log-linearized Phillips curve is

$$\hat{\pi}_t = \frac{\beta}{1 + \beta\iota_p} E_t \hat{\pi}_{t+1} + \frac{1}{1 + \beta\iota_p} \hat{\pi}_{t-1} + \kappa \hat{s}_t + \kappa \hat{\lambda}_{p,t}.$$

The normalization consists of defining a new exogenous variable, $\hat{\lambda}_{p,t}^* \equiv \kappa \hat{\lambda}_{p,t}$, and estimating the standard deviation of the innovation to $\hat{\lambda}_{p,t}^*$ instead of $\hat{\lambda}_{p,t}$. We do the same for the wage markup and the intertemporal preference shock, for which the normalizations are

$$\begin{aligned} \hat{\lambda}_{w,t}^* &= \left(\frac{(1 - \beta\xi_w)(1 - \xi_w)}{\left(1 + \nu\left(1 + \frac{1}{\lambda_w}\right)\right)(1 + \beta)\xi_w} \right) \hat{\lambda}_{w,t} \\ \hat{b}_t^* &= \left(\frac{(1 - \rho_b)(e^\gamma - h\beta\rho_b)(e^\gamma - h)}{e^\gamma h + e^{2\gamma} + \beta h^2} \right) \hat{b}_t. \end{aligned}$$

These normalizations are chosen so that these shocks enter their equations with a coefficient of one. In this way, it is easier to choose a reasonable prior for their standard deviation. Moreover, the normalization is a practical way to impose correlated priors across coefficients, which is desirable in some cases. For instance, imposing a prior on the standard deviation of the innovation to $\hat{\lambda}_{p,t}^*$ corresponds to imposing priors that allow for correlation between κ and the standard deviation of the innovations to $\hat{\lambda}_{p,t}$. Often, these normalizations improve the convergence properties of the MCMC algorithm.

REFERENCES

- ABRAHAM, K. G., J. R. SPLETZER, AND J. C. STEWART (1999): “Why Do Different Wage Series Tell Different Stories?,” *American Economic Review*, 89(2), 34–39.
- AN, S., AND F. SCHORFHEIDE (2007): “Bayesian Analysis of DSGE Models,” *Econometric Reviews*, 26(2-4), 113–172.
- ANDRÉS, J., D. LÓPEZ-SALIDO, AND E. NELSON (2005): “Sticky-Price Models and the Natural Rate Hypothesis,” *Journal of Monetary Economics*, 52(5), 1025–53.
- BENIGNO, P., AND M. WOODFORD (2006): “Linear-Quadratic Approximation of Optimal Policy Problems,” NBER Working Papers 12672, National Bureau of Economic Research, Inc.
- BLANCHARD, O., AND J. GALÍ (2007): “Real Wage Rigidities and the New Keynesian Model,” *Journal of Money, Credit and Banking*, 39(s1), 35–65.
- BOIVIN, J., AND M. GIANNONI (2006a): “DSGE Models in a Data-Rich Environment,” NBER working paper No. 12772.
- (2006b): “Has Monetary Policy Become More Effective?,” *Review of Economics and Statistics*, 88(3), 445–462.
- BOSWORTH, B., AND G. L. PERRY (1994): “Productivity and Real Wages: Is There a Puzzle?,” *Brookings Papers on Economic Activity*, 25(1994-1), 317–343.

- CALVO, G. (1983): “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12(3), 383–98.
- CHARI, V., P. J. KEHOE, AND E. R. MCGRATTAN (2009): “New Keynesian Models Are Not Yet Useful for Policy Analysis,” *American Economic Journal: Macroeconomics*, 1(1), 242–266.
- CHRISTIANO, L., C. L. ILUT, R. MOTTO, AND M. ROSTAGNO (2010): “Monetary Policy and Stock Market Booms,” NBER Working Papers 16402.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal Rigidities and the Dynamic Effect of a Shock to Monetary Policy,” *The Journal of Political Economy*, 113(1), 1–45.
- CHRISTIANO, L. J., M. TRABANDT, AND K. VALENTIN (2010): “Involuntary Unemployment and the Business Cycle,” NBER Working Papers 15801.
- CLARIDA, R., J. GALI, AND M. GERTLER (2000): “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *The Quarterly Journal of Economics*, 115(1), 147–180.
- COGLEY, T., AND T. J. SARGENT (2004): “The Conquest of U.S. Inflation: Learning, Model Uncertainty, and Robustness,” mimeo, University of California, Davis.
- COGLEY, T., AND A. M. SBORDONE (2008): “Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve,” *American Economic Review*, 98(5), 2101–26.
- DEL NEGRO, M., F. SCHORFHEIDE, F. SMETS, AND R. WOUTERS (2007): “On the Fit and Forecasting Performance of New Keynesian Models,” *Journal of Business and Economic Statistics*, 25(2), 123–162.
- EDGE, R. M., M. T. KILEY, AND J.-P. LAFORTE (2008): “Natural Rate Measures in an Estimated DSGE Model of the U.S. Economy,” *Journal of Economic Dynamics and Control*, forthcoming.
- ERCEG, C. J., D. W. HENDERSON, AND A. T. LEVIN (2000): “Optimal Monetary Policy with Staggered Wage and Price Contracts,” *Journal of Monetary Economics*, 46(2), 281–313.
- FRANCIS, N. R., AND V. A. RAMEY (2009): “Measures of Hours Per Capita and their Implications for the Technology-Hours Debate,” *Journal of Money, Credit, and Banking*, 41(6), 1071–1097.
- GALI, J. (2008): *Monetary Policy, Inflation and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, Princeton, NJ.
- GALÍ, J. (2010): “The Return of the Wage Phillips Curve,” NBER Working Papers 15758, National Bureau of Economic Research, Inc.
- GALI, J., M. GERTLER, AND D. LOPEZ-SALIDO (2007): “Markups, Gaps and the Welfare Costs of Business Fluctuations,” *Review of Economics and Statistics*, 89(1), 44–59.
- GALI, J., F. SMETS, AND R. WOUTERS (2011): “Unemployment in an Estimated New Keynesian Model,” mimeo, Universitat Pompeu Fabra.
- GREENWOOD, J., Z. HERCOWITZ, AND P. KRUSELL (1997): “Long Run Implications of Investment-Specific Technological Change,” *American Economic Review*, 87(3), 342–362.
- HALL, R. E. (1997): “Macroeconomic Fluctuations and the Allocation of Time,” *Journal of Labor Economics*, 15(2), 223–250.
- IRELAND, P. N. (2007): “Changes in the Federal Reserve’s Inflation Target: Causes and Consequences,” *Journal of Money, Credit, and Banking*, 39(8), 1851–1882.

- JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2010): "Investment Shocks and the Relative Price of Investment," *Review of Economic Dynamics*, forthcoming.
- LEVIN, A. T., A. ONATSKI, J. C. WILLIAMS, AND N. WILLIAMS (2005): "Monetary Policy Under Uncertainty in Micro-Founded Macroeconometric Models," in *NBER Macroeconomics Annual*, pp. 229–312.
- LUBIK, T. A., AND F. SCHORFHEIDE (2004): "Testing for Indeterminacy: An Application to U.S. Monetary Policy," *American Economic Review*, 94(1), 190–217.
- MISHKIN, F. S. (2007): "Estimating Potential Output," Speech, May 24, 2007.
- NELSON, K. S. N. A. E. (2005): "Inflation Dynamics, Marginal Cost, and the Output Gap: Evidence from Three Countries," *Journal of Money, Credit, and Banking*, 37(6), 1019–45.
- PRIMICERI, G. E. (2006): "Why Inflation Rose and Fell: Policymakers' Beliefs and US Postwar Stabilization Policy," *The Quarterly Journal of Economics*, 121(3), 867–901.
- ROTEMBERG, J. J., AND M. WOODFORD (1995): "Dynamic General Equilibrium Models with Imperfectly Competitive Product Markets," in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, chap. 9, pp. 243–293. Princeton University Press, Princeton, NJ.
- (1997): "An Optimization-Based Econometric Model for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 12, 297–346.
- (1999): "Interest Rate Rules in an Estimated Sticky Price Model," in *Monetary Policy Rules*, NBER Chapters, pp. 57–126. National Bureau of Economic Research, Inc.
- SALA, L., U. SODERSTROM, AND A. TRIGARI (2008): "Monetary Policy under Uncertainty in an Estimated Model with Labor Market Frictions," *Journal of Monetary Economics*, forthcoming.
- SALA, L., U. SÖDERSTRÖM, AND A. TRIGARI (2010): "The Output Gap, the Labor Wedge, and the Dynamic Behavior of Hours," CEPR Discussion Papers 8005.
- SCHMITT-GROHE, S., AND M. URIBE (2004): "Optimal Operational Monetary Policy in the Christiano-Eichenbaum-Evans Model of the U.S. Business Cycle," NBER working paper 10724, National Bureau of Economic Research, Inc, mimeo, Duke University.
- SCHMITT-GROHÉ, S., AND M. URIBE (2007): "Optimal Inflation Stabilization in a Medium-Scale Macroeconomic Model," in *Monetary Policy under Inflation Targeting*, ed. by F. S. Mishkin, K. Schmidt-Hebbel, N. L. S. Editor), and K. Schmidt-Hebbel, vol. 11 of *Central Banking, Analysis, and Economic Policies Book Series*, chap. 5, pp. 125–186. Central Bank of Chile.
- SCHMITT-GROHE, S., AND M. URIBE (2010): "The Optimal Rate of Inflation," in *Handbook of Monetary Economics*, ed. by B. M. Friedman, and M. Woodford. Elsevier, forthcoming.
- SHIMER, R. (2009): "Convergence in Macroeconomics: The Labor Wedge," *American Economic Journal: Macroeconomics*, 1(1), 280–97.
- SIMS, C. A. (2001): "Solving Linear Rational Expectations Models," *Journal of Computational Economics*, 20(1-2), 1–20.
- SMETS, F., AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian Approach," *American Economic Review*, 97(3), 586–606.

- TAYLOR, J. B. (1993): “Discretion versus Policy Rules in Practice,” *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.
- WALSH, C. E. (2005): “Comments on “Monetary Policy Under Uncertainty in Micro-Founded Macroeconomic Models,” by A. Levin, A. Onatski, J. Williams and N. Williams,” in *NBER Macroeconomics Annual*, ed. by M. Gertler, and K. Rogoff. MIT Press, Cambridge, MA.
- WOODFORD, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ.

FEDERAL RESERVE BANK OF CHICAGO

NORTHWESTERN UNIVERSITY, CEPR AND NBER

FEDERAL RESERVE BANK OF NEW YORK