MATTHEW B. CANZONERI
ROBERT E. CUMBY
BEHZAD T. DIBA

The Cost of Nominal Rigidity in NNS Models

We present a model with Calvo wage and price setting, capital formation, and estimated rules for government spending and monetary policy. Our model captures many aspects of U.S. data, including the volatility that has been observed in various efficiency gaps. We estimate the cost of nominal rigidity—welfare under flexible wages and prices minus welfare with nominal rigidities—to be as much as 3% of consumption each period. Since there are interest rate rules that virtually eliminate this cost, our model suggests that—contrary to Lucas's (2003) assertion—there is considerable room for improvement in demand management policy.

THE NEW NEOCLASSICAL SYNTHESIS (NNS) is characterized by monopolistic competition, wage, and/or price stickiness, and demand determination of output and employment. The NNS has been used to revisit the central issues of stabilization policy, and a number of theoretical insights have emerged. Rotemberg and Woodford (1997) showed that smoothing output—which was strongly emphasized in traditional Keynesian analyses—can lower household welfare in a model driven by productivity shocks. A number of papers have shown that the trade-offs for monetary policy can depend on the type of nominal rigidity that is postulated. For example,

MATTHEW B. CANZONERI, ROBERT E. CUMBY, and BEHZAD T. DIBA *are professors of economics, Georgetown University* (*E-mail*: canzonem@georgetown.edu, cumbyr@georgetown.edu, dibab@georgetown.edu).

© 2007 The Ohio State University

King and Wolman (1999) showed that there was no inflation–output trade-off in a model with staggered price setting; the optimal policy in their model was to stabilize the price level. By contrast, Erceg, Henderson, and Levin (2000) (EHL) showed that an inflation–output trade-off can emerge in a model with both staggered wage setting and staggered price setting.[1]

But, are these theoretical insights of any practical import? A challenge hanging over this new literature on monetary policy is Lucas's claim that the macroeconomic stabilization problem has been solved: "Taking U.S. performance over the past 50 years as a benchmark, the potential for welfare gains from better long-run, supply-side policies exceeds *by far* the potential from further improvements in short-run demand management" (Lucas 2003, p. 1).

In this paper, we calibrate a standard NNS model—with wage and price stickiness, capital formation, and empirically estimated rules for public spending and the cental bank's interest rate policy—to U.S. data, and we use the model to calculate the welfare cost of wage and price rigidities. We express the cost as the consumption an average household would be willing to give up each period to obtain the flexible wage/price solution. Our estimate of the cost of nominal rigidities is large: between 1% and 3% of consumption. We think that a welfare cost of this magnitude is worth caring about. We also show that there exist simple interest rate rules that would nearly eliminate this welfare cost. This suggests to us that there may be considerable room for improvement in demand management policies.

NNS models typically envision a production economy with complete consumption risk sharing. In such an environment, it is natural to associate the welfare cost of nominal rigidity with variations in the gap between the marginal product of labor (MPL) and the marginal rate of substitution (MRS) between consumption and leisure. EHL motivated their analysis in this way, but the welfare losses they found were quite small, suggesting that the variations in this gap generated by their model would also be small.[2] By contrast, Gali, Gertler, and Lopez-Salido (Forthcoming) (GGLS) developed empirical proxies for the MRS–MPL gap, and found that the gap was quite volatile, much more volatile than output. GGLS did not present a model, so it was not clear that an NNS model could generate the gap volatility they observed in the data.

There are several possible interpretations of the contrasting results of EHL and GGLS. Productivity shocks are the only source of uncertainty in EHL's model, and the monetary policies that EHL consider are all reasonably good in NNS models that are driven by productivity shocks. So, one interpretation is that the U.S. monetary policy was quite a bit worse than the policies EHL studied, or that the EHL model

---

1. An inflation–output trade-off also emerges in Giannoni's (2000) model with "inefficient" supply shocks (e.g., markup shocks). Our model does not include these shocks.

2. In Canzoneri, Cumby, and Diba (2005), we found small welfare losses for a number of variations on the EHL model; our variations included dropping fiscal subsidies, modeling capital accumulation, assuming fixed firm-specific capital, and adding habit persistence in consumption. We generate big losses here by adding demand shocks, assuming a more inelastic labor supply, and estimating a more realistic monetary policy rule.

is missing some of the shocks that have caused the MRS–MPL gap to fluctuate in the data. We have already alluded to the other interpretation: NNS models may simply be incapable of generating the volatility in the gap that GGLS found in the data.[3]

We calibrate our model, and we show that our model captures many aspects of the quarterly U.S. data. Moreover, we show that our model is capable of generating most of the volatility in the MRS–MPL gap that has been observed in the data. Finally, we calculate welfare—with and without nominal rigidities—using a second-order approximation to both the model and the welfare function. The difference between the two is what we call the cost of nominal rigidity. The difference between the cost of nominal rigidity across two specifications of policy represents the welfare gain (under nominal rigidity) of moving from one policy to the other.

Our estimate of the cost of nominal rigidity depends crucially on what we assume about the Frisch elasticity of labor supply; an inelastic labor supply reflects a rapidly increasing marginal disutility of work, and this in turn implies large utility costs for fluctuations in work effort. Real business cycle (RBC) models have to assume a highly elastic labor supply curve—much more elastic than estimates coming from the labor economics literature—to generate the volatility in hours worked that is observed in the data. In our NNS model, wages are sticky, and work effort is demand determined. The elasticity of notional labor supply is essentially a free parameter in our model; its value has little to do with the model's ability to fit moments in the data.[4] Consequently, we can choose the parameter to conform with the labor economists' estimates, and this will be seen to result in rather large welfare costs for business cycles.[5]

Our estimate of the cost of nominal rigidity also depends importantly on what we assume about the central bank's interest rate policy. The interest rate rule we use to characterize existing monetary policy has lagged interest rates, inflation, an output gap, and a residual (which we call the interest rate shock). In estimating this rule, we define the output gap to be the difference between actual output and the CBO's estimate of potential output. While this interest rate rule and its estimation are quite conventional, it is not so clear how the rule should be interpreted in our NNS model. There is no CBO in the model to provide estimates of potential output.

So, in our model we consider several alternative specifications of monetary policy: in our benchmark case (which seems to make the model fit the data best) the output gap is defined as the difference between actual output and steady-state output; in a

---

3. Of course, any business cycle model can "explain" the MRS–MPL gap by allowing for sufficiently large preference shocks. The question we address here is whether an NNS model—with wage rigidity but no preference shocks—can explain the observed volatility in the gap.

4. We have learned (in private conversation) from Frank Smets, that the Bayesian estimation procedure used in Smets and Wouters (2003a, 2003b) is also essentially silent on the value of the Frisch elasticity.

5. A subtle question, beyond the scope of this paper, is whether we can trust the labor economists' estimates of the Frisch elasticity and postulate, at the same time, that workers are not on their labor supply curves. All we do to address the issue in the present paper is to consider a range of values for the elasticity to calculate how large (or small) the welfare costs may be.

second specification, the output gap is defined as the difference between actual output and output in the flexible wage/price solution, as defined by Neiss and Nelson (2003). We refer to these two specifications as the "benchmark" interest rate rule and the "N&N" interest rate rule. We also consider two other specifications of the policy rule. In the first, we re-estimate the interest rate rule replacing the CBO's potential output with a log-linear trend. This estimate of the rule may be more consistent with our use of steady-state output in the model's gap term. We call this specification the "trend-based gap" policy. Finally, we consider a rule in which the interest rate simply reacts to nominal wage inflation. We refer to this specification as the "wage targeting" policy.

The bottom line of our welfare calculations is that the average household in our model would be willing to give up 1% to 3% of consumption each period to be free of the effects of wage and price stickiness. This is a conservative estimate: increasing certain model parameters to levels that have been used elsewhere in the literature can easily raise the cost of nominal rigidities to 5% or 6% of consumption. These estimates assume that our benchmark policy is in operation; replacing it with the trend-based gap policy, our estimates are a little higher. On the other hand, replacing the benchmark policy with the N&N policy results in a big welfare gain: the cost of nominal rigidities falls significantly; replacing the benchmark policy with the wage-targeting policy virtually eliminates the cost of nominal rigidities.[6]

If the N&N policy is an accurate description of monetary policy in practice, then our results would be consistent with Lucas's (2003) claim that the short-run demand management problem has already been solved. But if the benchmark policy is a better description of existing monetary policy, then our results suggest that there is considerable room for improvement. We are inclined toward the latter view since estimating potential output is very difficult in practice and our model seems to fit the data better with the benchmark policy. The welfare improvements in going to the N&N rule or the wage-targeting rule illustrate the potential gains from instituting better policies. Since the wage-targeting rule virtually eliminates the cost of nominal rigidity in our model, this cost—1% to 3% of consumption each period—is a good measure of the room for improvement.

The rest of our paper is organized as follows. In Section 1, we outline our NNS model and describe our calculation of the welfare of the average household in it. In Section 2, we discuss the calibration of our benchmark model and its fit with the data. In Section 3, we report our benchmark estimation of the welfare cost of nominal rigidities. In Section 4, we consider alternative policy rules and give our estimate of the room for improvement in demand management policy. In Section 5, we perform two robustness exercises, one of which suggests that our benchmark estimates of the cost of nominal rigidities and the room for improvement in demand management policy are really quite conservative. In Section 6, we discuss the relevance of our results and

6. EHL suggested targeting both wage and price inflation; here, we find that wage inflation targeting is by far the more important of the two. Canzoneri, Cumby, and Diba (2005) and Levin et al. (2005) report similar results.

some potential weaknesses of our model; we also highlight some directions for future work. The appendices referred to below are available on Matthew Canzoneri's web page.

1. AN NNS MODEL WITH PRICE AND WAGE RIGIDITIES AND CAPITAL FORMATION

Like other NNS models, our model is characterized by optimizing agents, capital accumulation, monopolistic competition, and nominal rigidities. Staggered price setting leads to a dispersion in the firms' prices that creates an inefficiency in household consumption decisions, staggered wage setting leads to a dispersion in the households' wages that creates an inefficiency in firm hiring decisions. Our purpose here is to get an idea of the magnitude of these inefficiencies.

1.1 Firms

There is a continuum of firms indexed by f on the unit interval. At time t, each firm rents capital $K_{t-1}(f)$ at the rate R_t, hires a labor bundle $N_t(f)$ (to be defined below) at the rate W_t (also defined below), and produces a differentiated product using the Cobb–Douglass technology

$$Y_t(f) = Z_t K_{t-1}(f)^\nu N_t(f)^{\nu-1}, \tag{1}$$

where $0 < \nu < 1$, and Z_t is an economy wide productivity shock. Z_t follows a simple autoregressive process—$\log(Z_t) = \rho \log(Z_{t-1}) + \epsilon_{z,t}$; we describe our estimation of this process in Appendix B (available online). The firm's cost minimization implies

$$R_t/W_t = [\nu/(1-\nu)](N_t(f)/K_{t-1}(f)), \tag{2}$$

and the firm's marginal cost can be expressed as (see Appendix A, available online).

$$MC_t(f) = \left[\nu^\nu(1-\nu)^{(1-\nu)}\right]^{-1} R_t^\nu W_t^{1-\nu}/Z_t. \tag{3}$$

The modeling of monopolistic competition is now standard in the NNS literature. Household (and government) preferences for the firms' products are represented by a constant elasticity aggregator,

$$Y_t = \left[\int_0^1 Y_t(f)^{(\phi p - 1)/\phi p} df\right]^{\phi p/(\phi p - 1)}, \tag{4}$$

where $\phi_p > 1$. The demand for an individual firm's product is

$$Y_t^d(f) = (P_t/P_t(f))^{\phi p} Y_t, \tag{5}$$

where

$$P_t = \left[\int_0^1 P_t(f)^{1-\phi p} df \right]^{1/(1-\phi p)} \tag{6}$$

is the price of the composite good, and can also be interpreted as the aggregate price level.

Following Calvo (1983), firms set prices in staggered "contracts" of random duration. In any period t, each firm gets to announce a new price with probability $(1 - \alpha)$; otherwise, the old contract, and its price, remains in effect.[7] With this scheme, the average duration of a price contract is $(1 - \alpha)^{-1}$ periods (quarters, in what follows).

If a firm gets to announce a new contract in period t, it chooses a new price $P_t^*(f)$ to maximize the value of its profit stream over states of nature in which the new price holds

$$E_t \sum_{j=t}^{\infty} (\alpha\beta)^{j-t} \lambda_j \left[P_t^*(f) Y_j(f) - TC_j(f) \right], \tag{7}$$

where $TC(f)$ is the firm's total cost, β is the households' discount factor, and λ_j is the households' marginal utility of nominal wealth (to be defined below). The firm's first-order condition is

$$P_t^* = \mu_p (PB_t / PA_t), \tag{8}$$

where $\mu_p = \phi_p / (\phi_p - 1)$ is a monopoly markup factor, and

$$PB_t = E_t \sum_{j=t}^{\infty} (\alpha\beta)^{j-t} \lambda_j MC_j(f) P_j^{\phi p} Y_j$$
$$= \alpha\beta E_t PB_{t+1} + \lambda_t MC_t(f) P_t^{\phi p} Y_t \tag{9}$$

$$PA_t = E_t \sum_{j=t}^{\infty} (\alpha\beta)^{j-t} \lambda_j P_j^{\phi p} Y_j = \alpha\beta E_t PA_{t+1} + \lambda_t P_t^{\phi p} Y_t. \tag{10}$$

As $\alpha \to 0$, all firms reset their prices each period (the flexible price case), and $P_t^*(f) \to \mu_p MC_t(f)$.

1.2 Households

There is a continuum of households indexed by h on the unit interval. Each household supplies a differentiated labor service to all of the firms in the economy. Firm

---

7. We set steady-state inflation equal to zero. But, our results would be the same if we let the contract price rise with a non-zero steady-state rate of inflation.

preferences (or technology) for these labor services are represented by a constant elasticity of aggregator,

$$N_t = \left[\int_0^1 L_t(h)^{(\phi w - 1)/\phi w} dh \right]^{\phi w/(\phi w - 1)}, \tag{11}$$

where $\phi_w > 1$. The demand for an individual household's services is

$$L_t^d(h) = (W_t/W_t(h))^{\phi w} N_t, \tag{12}$$

where

$$W_t = \left[\int_0^1 W_t(h)^{1-\phi w} dh \right]^{1/(1-\phi w)} \tag{13}$$

is the price of the composite bundle of labor services and can be interpreted as the aggregate wage.

The utility of an individual household is

$$U_t(h) = E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} \big[(1-\theta)^{-1} C_\tau(h)^{1-\theta} - (1+\chi)^{-1} L_\tau(h)^{1+\chi} \big], \tag{14}$$

where $C_t(h)$ is the household's consumption of the composite good Y_t. In what follows, we set θ equal to 1; our welfare results are not very sensitive to changes in the value of θ between 1 and 4. And in keeping with much of the NNS literature, we consider a cashless economy.

The household's budget constraint is

$$E_t[\Delta_{t,t+1} B_{t+1}(h)] + P_t[C_t(h) + I_t(h) + T_t] = B_t(h) + W_t(h)L_t^d(h)$$
$$+ R_t K_{t-1}(h) + D_t(h), \tag{15}$$

where the first term on the left-hand side is a portfolio of contingent claims,[8] I_t is the household's investment in capital, T_t is a lump sum tax (used by the government to balance its budget constraint each period), and the last three terms on the right-hand side are the household's wage, rental and dividend income. The household's accumulation of capital is governed by

$$K_t(h) = (1-\delta)K_{t-1}(h) + I_t(h) - \tfrac{1}{2}\psi[(I_t(h)/K_{t-1}(h)) - \delta]^2 K_{t-1}(h), \tag{16}$$

where δ is the depreciation rate, and the last term is the cost of adjusting the capital stock.

8. $B_{t+1}(h)$ is the number of (period $t+1$) dollar claims in the portfolio, contingent on a given state, $\Delta_{t,t+1}$ is the stochastic discount factor (the price of a dollar claim divided by the probability of the state), and E_t is an expectation over all the states of nature. See Cochrane (2001).

Households set wages in staggered "contracts" of random duration. In any period t, each household gets to announce a new wage with probability $(1 - \omega)$; otherwise, the old contract, and its wage, remains in effect. The average duration of a wage contract is $(1 - \omega)^{-1}$ periods.

If the household gets to announce a new contract in period t, it chooses the new wage

$$W_t^{*1+\phi_w\chi} = \mu_w(WB_t/WA_t), \tag{17}$$

where $\mu_w = \phi_w/(\phi_w - 1)$ is a monopoly markup factor, and

$$\begin{aligned} WB_t &= E_t \sum_{j=t}^{\infty} (\omega\beta)^{j-t} N_j^{1+\chi} W_j^{\phi_w(1+\chi)} \\ &= \omega\beta E_t WB_{t+1} + N_t^{1+\chi} W_t^{\phi_w(1+\chi)}, \end{aligned} \tag{18}$$

$$\begin{aligned} WA_t &= E_t \sum_{j=t}^{\infty} (\omega\beta)^{j-t} \lambda_j N_j W_j^{\phi_w} \\ &= \omega\beta E_t WA_{t+1} + \lambda_t N_t W_t^{\phi_w}, \end{aligned} \tag{19}$$

where λ_j is the household's marginal utility of nominal wealth (to be defined below). As $\omega \to 0$, all households get to reset their wages each period (the flexible wage case), and $W_t^*(h) \to \mu_w N_t^{\chi}/\lambda_t$; that is, the wage is a markup over the marginal disutility of work. Note that $1/\chi$ is the Frisch (or constant λ_t) elasticity of labor supply; this parameter will play a prominent role in the next section.

When wages are sticky ($\omega > 0$), the households' work efforts will differ. But our assumption of complete contingent claims markets makes households identical in terms of their consumption and investment decisions. In equilibrium, aggregate consumption will be equal to household consumption, and the same is true of the aggregate capital stock. So, we can write the first-order conditions for consumption and investment in terms of aggregate values:

$$1/P_t C_t = \lambda_t, \tag{20}$$

$$\beta E_t[\lambda_{t+1}/\lambda_t] = E_t[\Delta_{t,t+1}] = (1 + i_t)^{-1} \tag{21}$$

$$\lambda_t P_t = \xi_t - \xi_t \psi[(I_t/K_{t-1}) - \delta], \tag{22}$$

$$\begin{aligned} \xi_t = \beta E_t\{\lambda_{t+1} R_{t+1} &+ \xi_{t+1}[(1 - \delta) - \tfrac{1}{2}\psi[(I_{t+1}/K_t) - \delta]^2 \\ &+ \psi[(I_{t+1}/K_t) - \delta](I_{t+1}/K_t)\}, \end{aligned} \tag{23}$$

where λ_t and ξ_t are the Lagrange multipliers for the households' budget and capital accumulation constraints, and i_t is the return on a "risk-free" bond; $1/(1 + i_t) = E_t[\Delta_{t,t+1}]$.

1.3 The Aggregate Price and Wage Levels, Aggregate Employment, and Aggregate Output

The aggregate price level can be written as

$$P_t = \left[\int_0^1 P_t(f)^{1-\phi p}df\right]^{1/(1-\phi p)}$$

$$= \left[\sum_{j=0}^{\infty}(1-\alpha)\alpha^j\left(P_{t-j}^*(f)\right)^{1-\phi p}\right]^{1/(1-\phi p)}, \tag{24}$$

since the law of large numbers implies that $(1 - \alpha)\alpha^j$ is the fraction of firms that set their prices $t - j$ periods ago, and have not gotten to reset them since. It is straightforward to show that

$$P_t^{1-\phi p} = (1 - \alpha)P_t^{*1-\phi p} + \alpha(P_{t-1})^{1-\phi p}. \tag{25}$$

Similarly, the aggregate wage (defined by equation (12)) can be written as

$$W_t^{1-\phi w} = (1 - \omega)W_t^{*1-\phi w} + \omega(W_{t-1})^{1-\phi w}. \tag{26}$$

These calculations illustrate the beauty of the Calvo scheme for wage and price setting.

In Appendix A (available online), we show that aggregate output can be written as

$$Y_t = Z_t K_{t-1}^v N_t^{1-v}/DP_t, \tag{27}$$

where $N_t = \int_0^1 N_t(f)\,df$ is aggregate employment, $K_{t-1} = \int_0^1 K_{t-1}(h)\,dh = \int_0^1 K_{t-1}(f)\,df$ is the aggregate capital stock, and $DP_t = \int_0^1 (P_t/P_t(f))^{\phi p}df$ is a measure of price dispersion across firms; DP_t can be written as

$$DP_t = (1 - \alpha)\left(P_t/P_t^*(f)\right)^{\phi p} + \alpha(P_t/P_{t-1})^{\phi p}DP_{t-1}. \tag{28}$$

Price dispersion leads to inefficiency in consumption choices. Since each firm has the same marginal cost, consumers should choose equal amounts of the firms' products to maximize the consumption good aggregator for a given resource cost. If prices are flexible ($\alpha = 0$), then $P_t(f) = P_t$ for all f, $DP_t = 1$, and this efficiency condition will be met; if prices are sticky ($\alpha > 0$), then product prices will differ, $DP_t > 1$, and consumption decisions will be distorted.

1.4 Monetary and Fiscal Policy

Monetary policy and government spending are given empirical specifications. We use a standard interest rate rule to describe monetary policy:

$$i_t = 0.222 + 0.824 i_{t-1} + 0.356 \pi_t + 0.032 (\text{output gap})_t + \epsilon_{i,t}, \tag{29}$$

where $\pi_t = \log(P_t/P_{t-1})$ and the standard error of the interest rate shock, $\epsilon_{i,t}$, is 0.00245. We estimated this rule over the Volcker and Greenspan years (1979.3–2003.2); a full description is given in Appendix B (available online). Here, it is important to note that for estimation purposes we defined the output gap to be actual GDP minus the Congressional Budget Office's "potential" GDP.[9]

We estimated an autoregressive process for government spending:

$$\log(G_t) = \zeta + 0.973 \log(G_{t-1}) + \epsilon_{g,t}, \tag{30}$$

where the standard error of the fiscal shock, $\epsilon_{g,t}$, is 0.01; once again, Appendix B provides the details. The intercept term, ζ, will be chosen to make $G/Y = 0.20$ in our model's steady state.

1.5 Welfare

Our measure of national welfare is

$$U_t = E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} \big[\log(C_\tau) - (1+\chi)^{-1} AL_\tau \big], \tag{31}$$

where $C_t (= \int_0^1 C_t(h)\,dh = C_t(h))$ is per capita consumption, and $AL_t = \int_0^1 L_t(h)^{1+\chi}\,dh$ is the average disutility of work. If wages are flexible ($\omega = 0$), then $W_t(h) = W_t$, and firms hire the same amount of work from each household; $AL_t = L_t(h)^{1+\chi} \int_0^1 dh = L_t(h)^{1+\chi}$. In this special case, households are identical, and national welfare, U_t, reduces to individual household utility $U_t(h)$ (equation (14)).

If wages are sticky ($\omega > 0$), then there is a dispersion of wages that makes firms hire different amounts of work from each household. This creates an inefficiency similar to the one created by price dispersion: the composite labor service used by firms—$N_t = \int_0^1 [L_t(h)^{(\phi_w-1)/\phi_w} dh]^{\phi_w/(\phi_w-1)}$—will not be maximized for a given aggregate labor input, $\int_0^1 L_t(h)\,dh$. In Appendix A (available online), we show that

$$AL_t = N_t^{1+\chi} DW_t \tag{32}$$

$$DW_t = (1-\omega)\big(W_t^*(h)/W_t\big)^{-\phi_w(1+\chi)} + \omega(W_{t-1}/W_t)^{-\phi_w(1+\chi)} DW_{t-1}. \tag{33}$$

---

9. Both variables were measured in real per capita terms and expressed in logarithms.

TABLE 1
PARAMETERS FOR THE BENCHMARK CALIBRATION

$1/\chi$	ρ	σ	α	ω	ϕ_p	ϕ_w	δ	ψ	v	β
0.33	0.923	0.008	0.67	0.75	7	7	0.025	8	0.25	0.99

NOTES: The parameters are defined as follows: $1/\chi$ is the Frisch elasticity of labor supply; ρ and σ are the autoregressive coefficient and the standard deviation of the innovation in the process for productivity; α is the probability that a firm does not reset its price in a given quarter; ω is the probability at a worker's wage is not reset in a given quarter; ϕ_p and ϕ_w are the elasticities of substitution across goods and across worker, respectively; δ is the quarterly rate of depreciation of the capital stock; ψ is the capital stock adjustment cost parameter; $1 - v$ is labor share of income; and β is the discount factor. The choice of these values is described in an appendix available online.

where $DW_t = \int_0^1 (W_t(h)/W_t)^{-\phi w(1+\chi)} dh$ is a measure of wage dispersion analogous to DP_t.

In summary, nominal rigidity creates two distortions: staggered price setting creates a price dispersion that distorts households' consumption decisions, and staggered wage setting creates a wage dispersion that distorts firms' hiring decisions. These distortions interact with the distortions created by monopolistic competition to create what we call the welfare cost of nominal rigidity.

2. THE BENCHMARK CALIBRATION

Here, we show that our rather rudimentary NNS model, with only a few easily measured shocks, is capable of capturing many features of the U.S. data. We also show that it captures much of the volatility of the efficiency gaps emphasized by EHL and GGLS. The benchmark calibration of our model uses the estimated interest rate rule equation (29) with the gap defined as output minus steady-state output. In the sections that follow, we will discuss alternatives interest rate rules.

Table 1 specifies the parameters we use in our benchmark calibration. In Appendix B (available online), we discuss our choice of parameter values, our estimation of the interest rate rule and the stochastic processes for productivity and government spending. Here, we focus attention on just three of the parameters: $1/\chi$, the Frisch (or constant λ_t) elasticity of labor supply; ρ, the autoregressive parameter in the stochastic process for productivity; and σ, the standard deviation of the innovation in the productivity process.

The Frisch elasticity will be important in our welfare analysis: an inelastic supply implies a high cost of nominal rigidity. Empirical estimates of the Frisch elasticity range from 0.05 to 0.35.[10] So, our benchmark specification—$1/\chi = 0.33$—is quite conservative for the purposes of our welfare analysis. In the next section, we will consider a range of values from 0.14 to 1.00; they are all are conservative in the sense that they fall in the upper half of the estimated range and beyond.

10. See Bayoumi, Laxton, and Pesenti (2004) and GGLS for a discussion of these studies.

TABLE 2

THE BENCHMARK CALIBRATION OF THE MODEL

Std / Cor	$1/\chi = 1$	$1/\chi = 0.33$	$1/\chi = 0.20$	$1/\chi = 0.14$	Actual data
Output	0.014	0.014	0.014	0.014	0.016
	1.000	1.000	1.0000	1.000	1.000
Consumption	0.832	0.839	0.840	0.840	0.799
	0.963	0.962	0.961	0.961	0.871
Investment	3.112	3.133	3.111	3.118	3.122
	0.990	0.992	0.992	0.992	0.893
Hours	0.965	0.972	0.972	0.972	0.894
	0.633	0.635	0.636	0.636	0.857
Real wage	0.497	0.497	0.493	0.500	0.470
	0.577	0.553	0.547	0.543	0.243
Output gap	1	1	1	1	0.962
	1	1	1	1	0.965

NOTES: The first four columns of the table contain the volatilities of the six variables (top figure) and their correlation with output (bottom figure) from the calibration of the model for several values of the Frisch elasticity of labor supply. All other parameters are set at their benchmark values. The final column contains the volatilities of the six variable and their correlations with output from the U.S. data from 1960:1 to 2003:2. Both model data and actual data are in logarithms and have been Hodrick–Prescott-filtered. Model data were generated by Dynare, using a first-order approximation. For output, the volatility is the standard deviation. For the other variables, the volatilities are the standard deviations relative to standard deviation of output. As hours and real wages are for the non-farm business sector, we normalize their standard deviations by the standard deviation of real GDP of the non-farm business sector.

The persistence and volatility of productivity shocks will also be important in our welfare analysis: high values of ρ and σ imply high costs of nominal rigidity. As explained in Appendix B, we have three different estimates of the productivity process, each of which has some merit. Our benchmark—$(\rho, \sigma)_2 = (0.930, 0.008)$—comes from an estimate of the 1960–2002 data (with a log-linear trend). Our alternative specifications are $(\rho, \sigma)_3 = (0.979, 0.007)$, which comes from King and Rebelo's (1999) review of RBC models, and $(\rho, \sigma)_1 = (0.843, 0.007)$, which comes from an estimate of the 1980–2002 data (again, with a log-linear trend).[11]

Table 2 compares results from our calibrated model with quarterly data from the U.S. economy. The model's variables are expressed as log deviations from a deterministic steady state; moments are calculated from a linear approximation. The U.S. data are also in logs, and both the model data and the actual data are Hodrick–Prescott-filtered. Beginning with the row for output and the column headed $1/\chi = 0.33$ (the benchmark specification), 0.014 is the model's standard deviation of output; it is slightly smaller than the standard deviation of output in the data, 0.016, which is given in the last column. Proceeding to the row for consumption, 0.839 is the ratio of the standard deviation of consumption to the standard deviation of output in the model, and 0.962 is the correlation between consumption and output. These are close to the corresponding statistics in the data. The following three rows provide the same statistics for investment, hours, and real wages.[12] The model comes fairly close to

11. This data period corresponds to the one for our interest rate rule. The two data periods produced very similar estimates for the government spending process (see Appendix B).

12. We set the adjustment cost coefficient ψ to make the standard deviation of investment (relative to the standard deviation of output) in the model mirror the ratio in the data.

matching the data for most of these variables, but real wages and output are more positively correlated in the model than they are in the data.

More to the point for us, Table 2 suggests that the value of $1/\chi$ has almost no effect on the model's moments. RBC models need a very elastic labor supply curve to generate the employment volatility that is observed in the data. Employment and output are demand determined in NNS models, and workers may be off their notional labor supply curves for long periods of time. We do not need an elastic labor supply to match the volatility of employment.

The output gap in the benchmark calibration of the model is defined as actual output minus steady-state output, which explains the ones in the last row of Table 2. The interesting thing to note here is that the output gap in the data—output minus the CBO's "potential" output—comes very close to matching these statistics.

The work of EHL and (especially) GGLS suggests that we may also be able to test our model against the data in a way that is more directly related to welfare. Economic efficiency requires that the MRS between consumption and work be equal to the MPL. Neglecting constant terms, the efficiency gap is

$$\text{gap}_t = \log(\text{MRS}_t) - \log(\text{MPN}_t) = \log(C_t) + \chi \log(N_t)$$
$$- [\log(Y_t) - \log(N_t)], \tag{34}$$

This gap can be partitioned into a wage gap,

$$\text{wgap}_t = \log(\text{MRS}_t) - \log(W_t/P_t) = \log(C_t) + \chi \log(N_t) - \log(W_t/P_t), \tag{35}$$

and a price gap,

$$\text{pgap}_t = \log(W_t/P_t) - \log(\text{MPN}_t) = \log(W_t/P_t) - [\log(Y_t) - \log(N_t)]. \tag{36}$$

GGLS calculate the volatility of these gaps, using the U.S. data, and make inferences about the cost to welfare that this volatility implies. The overall gap is very volatile, and GGLS claim that the welfare cost of U.S. business cycles is high.

Table 3 shows the volatility in these efficiency gaps, both in our model and in the U.S. data. The first number in each cell is the standard deviation of the gap; the second number is the correlation between the gap and output. Our NNS model is capable of generating much of the volatility that is observed in the data. The wage gap is much more volatile than the price gap, suggesting that it is the major source of the welfare costs, as GGLS claim.

In summary, our benchmark calibration matches a number of the variances and covariances in the U.S. data, and the volatility of its efficiency gaps matches as the volatility found by GGLS.

TABLE 3

THE BENCHMARK CALIBRATION OF THE EFFICIENCY GAPS

Std Correlation	Model MRS–MPL gap	Model wage gap	Model price gap	Data MRS–MPL gap	Data wage gap	Data price gap
$1/\chi = 1$	0.025	0.022	0.008	0.029	0.028	0.011
	0.580	0.741	−0.188	0.730	0.844	−0.230
$1/\chi = 0.33$	0.053	0.050	0.008	0.064	0.063	0.011
	0.611	0.689	−0.196	0.807	0.862	−0.230
$1/\chi = 0.20$	0.082	0.078	0.008	0.099	0.098	0.011
	0.659	0.705	−0.151	0.826	0.862	−0.230
$1/\chi = 0.14$	0.110	0.106	0.008	0.134	0.133	0.011
	0.623	0.661	−0.198	0.834	0.861	−0.230

NOTES: The table presents the standard deviations (top figure) and correlations with output (bottom figure) of three different "efficiency" gaps for alternative values of the Frisch elasticity of labor supply. The MRS–MPL gap is defined in equation (34) and difference between the marginal rate of substitution between consumption and work and the marginal product of labor. The wage gap is defined in equation (35) and is the difference between the marginal rate of substitution and the real wage. The price gap is defined in equation (36) and is the difference between the real wage and the marginal product of labor. The first three columns report the values generated from a first-order approximation of the model by Dynare. The final three columns are estimated using the U.S. data from 1960:1 to 2003:2. Both model data and actual data are in logarithms and have been Hodrick–Prescott-filtered.

3. MEASURING THE COST OF NOMINAL RIGIDITY IN THE BENCHMARK MODEL

Let V_t be the value function for aggregate welfare in period t. In light of equation (31), V_t is given by

$$V_t = \log(C_t) - (1 + \chi)^{-1} AL_t + \beta E_t[V_{t+1}]. \tag{37}$$

In this section, we use Dynare to calculate a second-order approximation of V_t under various assumptions about nominal rigidity and other key parameters in the model. Assuming state variables are at their deterministic steady-state values at time 0, let $V_0(\alpha, \omega)$ represent aggregate utility for an economy with a given type of nominal rigidity (characterized by α and ω). The welfare cost of nominal rigidity in this economy is

$$CC(\alpha, \omega) = V_0(0, 0) - V_0(\alpha, \omega). \tag{38}$$

CC is a cardinal number, and its units are hard to understand. However, following Lucas (2003), we can interpret CC in terms of consumption equivalents. We have assumed a log specification for the utility of consumption; so, our $CC(\alpha, \omega)$ can be interpreted as the percentage of consumption households would be willing to give up each period to be free of a particular type of nominal rigidity, *assuming that work effort is held constant*.[13]

---

13. Lucas' utility function did not include leisure. Using U.S. data, he calculated that the welfare cost of fluctuations in consumption (about trend) was only about one-twentieth of 1% of consumption each period. Although our thought experiment differs from his, the spirit of his analysis suggests that our $CC(\alpha, \omega)$ should be quite small.

To see this, let $\{C_j^*\}$ and $\{AL_j^*\}$ be consumption and the average disutility of work in the flexible wage/price solution, let $\{C_j\}$ and $\{AL_j\}$ be consumption and the average disutility of work in the solution with nominal rigidity, and let ξ solve:

$$V_0(0, 0) = E_0 \sum_{j=0}^{\infty} \beta^j \left[\log\left(C_j^*\right)(1 + \chi)^{-1} AL_j^* \right]$$

$$= E_0 \sum_{j=0}^{\infty} \beta^j \left[\log((1 + \xi)C_j) - (1 + \chi)^{-1} AL_j \right]$$

$$= \xi/(1 - \beta) + E_0 \sum_{j=0}^{\infty} \beta^j \left[\log(C_j) - (1 + \chi)^{-1} AL_j \right]$$

$$= \xi/(1 - \beta) + V_0(\alpha, \omega) \tag{39}$$

or

$$\xi = (1 - \beta)[V_0(0, 0) - V_0(\alpha, \beta)] = .01 * [V_0(0, 0) - V_0(\alpha, \omega)], \tag{40}$$

for our assumed value of β. Our $CC(\alpha, \omega) = 100 * \xi$, which expresses the cost as a percentage of consumption (instead of a fraction) each quarter.

Table 4 presents our estimates of this cost in the benchmark model. In our benchmark parameterization—$1/\chi = 0.33$ and $(\rho, \sigma)_2$—the cost is 1.03% of consumption. Recall that $1/\chi = 0.33$ is at the upper end of the range of empirical estimates for the Frisch elasticity of labor supply. If we use a value closer to the middle of the estimated range—$1/\chi = 0.14$—the cost of nominal rigidity is 2.13% of consumption; if we use a value at the bottom of the range, the cost goes up to 5 or 6%. If we keep the benchmark elasticity of labor supply, but use a productivity process typical of the RBC literature—$(\rho, \sigma)_3$—the cost is 3.5%.

In summary, our benchmark model conservatively estimates the cost of nominal rigidity to be between 1% and 3% of consumption each quarter. If we raise the Frisch

TABLE 4
CONSUMPTION COST OF THE BENCHMARK WAGE AND PRICE RIGIDITIES

	$(\rho, \sigma)_1 = (0.843, 0.007)$	$(\rho, \sigma)_2 = (0.923, 0.0086)$	$(\rho, \sigma)_3 = (0.979, 0.007)$
$1/\chi = 1$	0.15	0.47	1.39
$1/\chi = 0.33$	0.28	1.02	3.54
$1/\chi = 0.20$	0.42	1.57	5.71
$1/\chi = 0.14$	0.55	2.11	7.88
$1/\chi = 0.10$	0.58	2.94	11.15
$1/\chi = 0.05$	1.46	5.72	22.19

NOTES: The table reports the welfare costs—expressed as a percent of consumption—of wage and price stickiness for several alternative Frisch elasticities of labor supply ($1/\chi$) and for alternative stochastic processes for productivity (ρ, σ). The probabilities of no change in prices and wages are set at our benchmark values of 0.67 and 0.75. The middle column is our benchmark specification for productivity and is estimated using linearly de-trended U.S. data from 1960:1 to 2003:2. The first column uses a shorter sample (1979:3–2003:2). The third column uses a process (taken from King and Rebelo 1999) that is typical in the RBC literature. The welfare costs are computed from a second-order approximation of the model and the welfare function by Dynare.

elasticity of labor supply or the persistence of productivity shocks to levels that have been used elsewhere in the literature, the cost can easily rise to 5% or 6%. This cost would certainly seem to be significant.

4. ROOM FOR IMPROVEMENT IN DEMAND MANAGEMENT POLICY

The cost of nominal rigidity may be high, but this does not necessarily mean that there is much room for improvement in demand management policy. It may be that monetary policy is incapable of doing much about this welfare loss. In this section, we show that this is not the case. Simple monetary policy rules can virtually eliminate the cost of nominal rigidities. Of course, the model's fit generally deteriorates under these alternative rules, but that is just another indication of the fact that our benchmark rule seems to capture existing monetary policy fairly well.

Rotemberg and Woodford (1997) have shown that a policy of smoothing output can lower household welfare in a model that is driven by productivity shocks. Productivity shocks play an important role in our NNS model; for example, (infinite horizon) variance decompositions reveal that productivity shocks account for 52% of the variation in output and 95% of the variation in inflation. Consequently, our benchmark interest rate rule, which defines the output gap as the difference between actual output and steady-state output, might be expected to do rather poorly in terms of welfare. Moreover, we alluded in the introduction to studies which suggest that the major welfare losses in NNS models come from the nominal wage rigidity rather than price rigidity. The benchmark interest rate rule responds to movements in price inflation, and not wage inflation. In this section, we consider two alternative interest rate rules that might be expected to do better.

The first alternative is the N&N interest rate rule, which replaces the CBO's measurement of potential output in our estimated rule, equation (29), with the model's own flexible wage/price output. Using the N&N gap might be expected to address the concern raised by Rotemberg and Woodford (1997). The second alternative is the wage-targeting rule:

$$i_t = -\log(\beta) + 2.0\pi_{w,t}, \tag{41}$$

where $\pi_{w,t}$ is aggregate wage inflation.

It is straightforward to use the consumption cost, $CC(\alpha, \omega)$, that we have already defined to make normative comparisons between alternative policy rules. Since $V_0(0, 0)$ does not depend on monetary policy, the welfare gain from moving to a new policy is just the difference between the value of $CC(\alpha, \omega)$ associated with the benchmark policy and the value of $CC(\alpha, \omega)$ associated with the alternate policy. Once again, the welfare gain is expressed as the percent of consumption an average household would be willing to give up each quarter to have the alternate policy instead of the benchmark policy.

TABLE 5
CONSUMPTION COSTS OF ALTERNATIVE INTEREST RATE RULES

	$1/\chi = 1$	$1/\chi = 0.33$	$1/\chi = 0.14$
Benchmark rule	0.47	1.02	2.11
Benchmark (no shock)	0.43	0.92	1.92
N&N rule	0.12	0.23	0.46
N&N rule (no shock)	0.07	0.14	0.26
Wage inflation rule	0.06	0.08	0.12
Wage inflation (no shock)	0.06	0.07	0.11

NOTES: The table reports the welfare costs—expressed as a percent of consumption—of wage and price stickiness for alternative Frisch elasticities of labor supply ($1/\chi$) and for alternative monetary policy rules. The probabilities of no change in prices and wages are set at our benchmark values of 0.67 and 0.75. All other parameters are also set at their values. The first row reports the welfare costs obtained with our benchmark specification of monetary policy, equation (29), which uses the non-stochastic steady state value of output to compute the output gap. The third row reports the costs using the same parameter values as the benchmark rule but computes potential output using uses the using the level of employment and the capital stock of the corresponding economy with flexible wages and prices to compute the output gap. We refer to this rule as the Neiss and Nelson rule. The fifth row reports the welfare costs using the rule equation (41) in which interest rates respond only to wage inflation. The rows labeled "no shock" report the welfare costs using the corresponding policy rule but with the standard deviation of the shock set to zero. The welfare costs are computed from a second-order approximation of the model and the welfare function by Dynare.

Table 5 reports the consumption cost of nominal rigidity associated with the three policy rules;[14] the first row of Table 5 picks up the corresponding elements of the third column in Table 4. We also report the cost of nominal rigidity when the interest rate shocks are eliminated, since some have questioned our interpretation of the residuals in the estimated policy rule; in particular, these residuals may be the Fed's use of other indicator variables, or differences between CBO and Fed estimates of potential output. Table 6 shows the model's fit under the three policy rules; here again, the first row of Table 6 picks up the corresponding elements of the third column in Table 2. The model's fit deteriorates markedly without the interest rate shocks; we do not report these results.

The benchmark policy's definition of the output gap does matter. Under the N&N policy, the consumption cost of nominal rigidity is about a quarter of what it was in the benchmark case. Put another way, if the Frisch labor supply elasticity is one-third, the welfare gain from replacing the benchmark policy with the N&N policy is a little over .75% of consumption per period; and if the interest rate shocks can be eliminated in the N&N policy, the welfare gain rises to almost 1%. If (as we think likely) the Frisch elasticity is lower, say 0.14, then the welfare gain is 1.5% of consumption per period.

Table 6 shows however that N&N policy does a poor job of fitting the observed output gap (i.e., actual output minus the CBO's estimate of potential output). The N&N gap is not volatile enough, and its correlation with output is too low. The benchmark gap is a much better fit. This fact, and the rather extreme informational requirements of the N&N policy, lead us to conclude that the benchmark policy is a better description of existing policy. The welfare gain associated with the N&N

14. For the wage inflation-targeting rule, we add a shock of the same size as the shock we estimated for the benchmark interest rate rule.

TABLE 6

THE MODEL CALIBRATION UNDER ALTERNATIVE INTEREST RATE RULES

Std / Cor	Consumption	Investment	Output	Employment	Real wages	Output gap
Benchmark rule	0.839	3.133	0.014	0.972	0.497	1
	0.962	0.992	1.000	0.635	0.553	1
N&N rule	0.846	3.044	0.018	0.808	0.368	0.593
	0.976	0.993	1.000	0.817	0.608	0.808
Wage inflation rule	0.538	1.622	0.008	0.0687	0.490	0.0573
	0.903	0.971	1.000	−0.000	0.917	1
Actual data	0.799	3.122	0.016	0.894	0.470	0.962
	0.871	0.893	1.000	0.857	0.243	0.965

NOTES: The table contains the volatilities of the six variables (top figure) and the correlation of those variables with output (bottom figure) from the benchmark calibration of the model for the three policy rules from Table 5. The bottom row contains the volatilities of the six variables and their correlations with output from the U.S. data from 1960:1 to 2003:2. Both model data and actual data are in logarithms and have been Hodrick–Prescott-filtered. Model data were generated by Dynare, using first-order approximations. For output, the volatility is that standard deviation. For the other variables, the volatilities are the standard deviations relative to standard deviation of output. As hours and real wages are for the non-farm business sector, we normalize their standard deviations by the standard deviation of real GDP of the non-farm business sector. All parameters, including the Frisch elasticity of labor supply, are set at their benchmark values.

policy is one measure of the potential room for improvement in demand management policy.

The simple wage-targeting rule does even better than the N&N policy. The welfare gain from replacing the benchmark policy with the wage-targeting policy is almost as large as our original cost of nominal rigidity (associated with the benchmark policy).[15] The wage-targeting rule (with or without shocks) virtually eliminates the welfare cost associated with wage and price stickiness. Of course, the wage-targeting policy does a very bad job of fitting the data, as can be seen in Table 6.

In summary, our model suggests that our benchmark estimate of the cost of nominal rigidity is itself a good measure of the potential room for improvement in demand management policy.[16]

5. TWO ROBUSTNESS EXERCISES

We have already done a number of robustness exercises. We checked the model's fit, and the implied welfare costs, of alternative monetary policy rules. And we discussed the implications of different values of the Frisch labor supply elasticity $(1/\chi)$ and the persistence of productivity processes. Here, we perform two additional exercises.

15. This is also the case for lower Frisch labor supply elasticities and more inertia-ridden productivity processes.

16. The wage-targeting rule is not the best policy. Benigno and Woodford (2005) show that in an economy like ours, optimal policy responds to wage and price inflation as well as a measure of the output gap. A referee suggested that with wage and price rigidity, a Ramsey planner would not do as well as the flexible wage/price solution and therefore would not do much better than our wage-targeting rule.

TABLE 7
CALIBRATION WITH LESS NOMINAL RIGIDITY ($\alpha = \omega = 0.5$)

Std / Correlation	$1/\chi = 0.33$	$1/\chi = 0.14$	Actual data
Output	0.014	0.014	0.016
	1.000	1.000	1.000
Consumption	0.840	0.860	0.799
	0.964	0.962	0.871
Investment	3.118	3.042	3.122
	0.990	0.991	0.893
Hours	0.743	0.741	0.894
	0.634	0.630	0.857
Real wage	0.597	0.601	0.470
	0.757	0.754	0.243
Output gap	1	1	0.962
	1	1	0.965

NOTES: The first two columns of the table contain the volatilities of the six variables (top figure) and their correlations with output (bottom figure) from the calibration of the model for two values of the Frisch elasticity of labor supply. The probabilities of no change in prices and wages are both reduced to 0.5 so that both wages and prices are more flexible than in the benchmark calibration. All other parameters are set at their benchmark values. The final column contains the volatilities of the six variables and their correlations with output from the U.S. data from 1960:1 to 2003:2. Both model data and actual data are in logarithms, and have been Hodrick–Prescott-filtered. Model data were generated by Dynare, using first-order approximations. For output, the volatility is that standard deviation. For the other variables, the volatilities are the standard deviations relative to standard deviation of output. As hours and real wages are for the non-farm business sector, we normalize their standard deviations by the standard deviation of real GDP of the non-farm business sector.

Our first exercise will be to see what happens to the model's fit, and the implied welfare costs, when we vary the degree of wage and price stickiness. The benchmark values of the Calvo parameters are standard in the literature, but since our benchmark model does not generate enough variability in wage and price inflation, it seems natural to ask how our results would change if we lowered the degree of nominal rigidity. The second exercise will be to re-estimate the benchmark interest rate policy, replacing the CBO's estimate of potential output with a log-linear trend (as Rotemberg and Woodford 1997 and others have done).[17] We outline our estimation procedure and give our data sources in Appendix B (available online). This estimate might be more consistent with our model's benchmark policy, which defines the output gap as the deviation of actual output from its steady-state value. We call this specification the "trend-based gap" rule.

5.1 Lowering the Degree of Wage and Price Stickiness

In the benchmark calibration, we set the Calvo parameters at $\alpha = 0.67$ and $\omega = 0.75$; the average durations of price and wage contracts are, respectively, three and four quarters. Table 7 reports the same calibration exercise, but with the Calvo parameters (α and ω) set equal to 0.5; that is, we reduce the average duration of both price and wage contracts to two quarters.

Lowering the degree of nominal rigidity helps the model fit the data in some ways: relative price and wage inflation volatility in the model (0.35 and 0.15) are closer

17. We are grateful to a referee for suggesting this.

TABLE 8

CONSUMPTION COSTS WITH ALTERNATIVE DEGREES OF WAGE AND PRICE RIGIDITY

	$1/\chi = 0.33$	$1/\chi = 0.14$
$CC(0.67, 0.75)$	1.02	2.11
$CC(0.50, 0.50)$	0.45	0.95
$CC(0.67, 0)$	0.13	0.22
$CC(0, 0.75)$	1.13	2.40

NOTES: The table reports the welfare costs—expressed as a percent of consumption—of wage and price stickiness for two alternative Frisch elasticities of labor supply $(1/\chi)$ and for four alternative values for price and wage stickiness. The first row sets the probabilities of no change in prices and wages are set at our benchmark values of 0.67 and 0.75. The second row sets both probabilities to 0.5. The third row sets the probability of no price changes to our benchmark value but allows wages to be fully flexible. The final row sets the probability of no change in wages to our benchmark value but allows prices to be fully flexible. All other parameters are set at their benchmark values.

to what is observed in the data (0.36 and 0.30). But other aspects of the model's fit deteriorate; most notably, hours worked now have less volatility than in the data.

Table 8 presents the consumption costs of different degrees of nominal rigidity, under our conservative benchmark assumption that the Frisch elasticity is $1/\chi = 0.33$, and under the more typical estimated value of $1/\chi = 0.14$. The first row of Table 8 is taken from Table 4; it is the cost of nominal rigidity in our benchmark case $-(\alpha, \omega) = (0.67, 0.75)$. The second row gives the cost of the wage and price rigidity specified in Table 5 $-(\alpha, \omega) = (0.50, 0.50)$. The costs are now about half of what they were in the benchmark case. But for $1/\chi = 0.14$, which is in the middle of the range of estimates, the cost of nominal rigidity is still about 1% of consumption.

The third and fourth rows of Table 8 examine the relative importance of wage and price rigidities for welfare. The third row gives the cost of price rigidity with flexible wages $-(\alpha, \omega) = (0.67, 0)$, and the last row gives the cost of wage rigidity with flexible prices $-(\alpha, \omega) = (0, 0.75)$. Wage rigidity appears to be much more costly than price rigidity. This result is consistent with the fact that the volatility of the wage gap in Table 3 is much larger than the volatility of the price gap.

In summary, cutting the degree of wage and price stickiness helps the model's fit in some ways and hurts it in other ways. Cutting the degree of wage stickiness decreases the cost of nominal rigidities considerably.[18] But cutting the degree of price stickiness—while leaving the degree of wage stickiness the same—may actually increase the cost of nominal rigidity.[19]

5.2 The "Trend-Based" Version of the Benchmark Case

Table 9 reports the second calibration exercise, with the "trend-based" estimate of the interest rate rule replacing the original estimate. Once again, the model's gap is

18. This is consistent with the results of Paustian (2005) and Levin et al. (2005): wage contracting schemes that imply less wage dispersion than the Calvo framework lead to smaller estimates of the cost of wage rigidity.

19. There are a number of distortions in our NNS model. Eliminating just one—price rigidity—need not increase welfare.

TABLE 9
THE "TREND OUTPUT" POLICY RULE

Std\Cor	$1/\chi = 0.33$		
	Benchmark rule	Trend output rule	Actual data
Output	0.014	0.013	0.016
	1.000	1.0000	1.000
Consumption	0.839	0.836	0.799
	0.962	0.956	0.871
Investment	3.133	3.142	3.122
	0.992	0.991	0.893
Hours	0.972	1.067	0.894
	0.635	0.592	0.857
Real wage	0.497	0.537	0.470
	0.553	0.528	0.243
Output gap	1	1	0.962
	1	1	0.965
Welfare cost of nominal rigidity baseline productivity	1.02	1.21	
Welfare cost $(\rho, \sigma)_1 = (0.843, 0.007)$	0.28	0.29	
Welfare cost $(\rho, \sigma)_3 = (0.979, 0.007)$	3.54	5.02	

NOTES: The table summarizes the model calibration and the welfare costs of price and wage stickiness using our estimated benchmark monetary policy rule (which uses the CBO's measure of potential output) and a rule for a similar rule for monetary policy estimated using trend output rather than the CBO measure of potential output in computing the output gap. The top six rows report the volatilities of six variables and their correlations with output. The final three rows report the welfare costs of price and wage stickiness for the three alternative stochastic processes for productivity described in Table 4. All other parameters are set at their benchmark values.

defined as actual output minus steady state output. The model's fit with the newly estimated policy rule is worse than the fit of the original benchmark model, but only slightly. Table 9 reports the implied welfare costs for the three productivity processes we considered in Section 3 (and with the Frisch labor supply elasticity set at $1/\chi = 1/3$). The consumption cost of nominal rigidity rises significantly in each case, and especially in the case of the productivity process that is typical of the RBC literature.

6. CONCLUSION

In this paper, we have conservatively estimated the welfare cost of nominal rigidity to be between 1% and 3% of the "average" household's consumption each period. Wage stickiness accounted for most of this welfare cost. The estimates are conservative in the sense that if we decrease the Frisch labor supply elasticity or increase the persistence of productivity shocks to levels that have been used elsewhere in the literature, the estimates could easily double. The estimates are also conservative in the sense that we have not modeled unemployment and assumed consumption risk sharing.

To gauge the room for improvement in demand management policy, we showed that there exist simple interest rate rules that would—if implemented—raise welfare in our model by an amount that is roughly equal to the original cost of nominal rigidity. Replacing the output gap in the benchmark rule (actual output minus steady-state output) with the N&N output gap (actual output minus flexible wage/price output)

would increase welfare by about three-fourths of the benchmark cost of nominal rigidity. And simply targeting wage inflation would virtually offset the cost of nominal rigidity. The model may not fit the data as well with these alternative policies, but that is just a reflection of the fact that our benchmark policy is a better description of existing monetary policy.

There is of course a good reason why central banks may not have implemented these alternative interest rate policies: each has a demanding informational requirement. It is very difficult for central banks to estimate the level of potential output. And similarly, it would be difficult to get an accurate measure of wage inflation since it is hard to get data for salaried and self-employed workers. However, central banks are currently investing significant resources in the measurement of productivity and potential output, and our results suggest that these resources are being well spent.

These statements take our model at face value, and even we do not actually do that. The cost of nominal rigidity and the welfare gain from alternative monetary policies are model dependent, and our results hinge on theoretical and empirical issues that are the subject of ongoing research. In addition, central banks may be responding to a myriad of shocks or concerns that we have not modeled. For example, our finding that the N&N interest rate rule performs so much better than our benchmark rule may reflect the fact that we do not model what Giannoni (2000) calls inefficient supply shocks—shocks that make the flexible wage/price level of output inefficient and may thereby change the welfare relevant measure of the output gap. In the same spirit, we suspect that some demand shocks may either be absent or incorrectly specified in our model: inflation and output are negatively correlated in our model, while they appear to be positively correlated in the U.S. data; similarly, interest rates and output are negatively correlated in the model, while they appear to be positively correlated in the data.[20] This may be driving some of our results: it is the predominance of "efficient" productivity shocks in our model that makes getting potential output right so important in the central bank's interest rate rule.

More fundamentally, Goodfriend and King (2001) have questioned the relevance of observed wage rigidity on theoretical grounds: "There is a fundamental asymmetry between product and labor markets. The labor market is characterized by long-term relationships where there is opportunity for firms and workers to neutralize the allocative effects of temporarily sticky nominal wages. . . . (However), spot transactions predominate in product markets where there is much less opportunity for the effects of sticky nominal prices to be privately neutralized." On the other hand, a growing empirical literature suggests that wage rigidity helps NNS models explain the U.S. data: Christiano, Eichenbaum, and Evans (2005) found that wage stickiness helps explain persistence in the effects of monetary shocks, and Smets and Wouters (2003b) showed that wage stickiness also helps explain other features of the data. The resolution of this debate may well matter for our claim that targeting wage inflation would be a major improvement on current policy.

20. In conventional Keynesian models, an increase in aggregate demand would be expected to raise output and inflation, and via the central bank's interest rate rule, short-term interest rates.

LITERATURE CITED

Bayoumi, Tamim, Douglas Laxton, and Paolo Pesenti. (2004) "Benefits and Spillovers of Greater Competition in Europe: A Macroeconomic Assessment." NBER Working Paper No. 10416.

Benigno, Pierpaolo, and Michael Woodford. (2005) "Optimal Stabilization Policy When Wages and Prices are Sticky: The Case of a Distorted Steady State." In *Models and Monetary Policy: Research in the Tradition of Dale Henderson, Richard Porter, and Peter Tinsley*, edited by Jon Faust, Athanasios Orphanides, and David Reifschneider. Board of Governors of the Federal Reserve System.

Calvo, Guillermo. (1983) "Staggered Prices in a Utility Maximizing Framework." *Journal of Monetary Economics,* 12, 383–98.

Canzoneri, Matthew, Robert Cumby, and Behzad Diba. (2005) "Price and Wage Inflation Targeting: Variations on a Theme by Erceg, Henderson and Levin." In *Models and Monetary Policy: Research in the Tradition of Dale Henderson, Richard Porter, and Peter Tinsley*, edited by Jon Faust, Athanasios Orphanides, and David Reifschneider. Board of Governors of the Federal Reserve System.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. (2005) "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy,* 113, 1–45.

Cochrane, John. (2001) *Asset Pricing*. Princeton and Oxford: Princeton University Press.

Erceg, Christopher, Dale Henderson, and Andrew Levin. (2000) "Optimal Monetary Policy with Staggered Wage and Price Contracts." *Journal of Monetary Economics,* 46, 281–313.

Gali, Jordi, Mark Gertler, and David Lopez-Salido. (Forthcoming) "Markups, Gaps, and the Welfare Costs of Business Fluctuations." *Review of Economics and Statistics.*

Giannoni, Marc. (2000) "Optimal Interest-Rate Rules in a Forward-Looking Model, and Inflation Stabilization versus Price-Level Stabilization." Mimeo, Columbia University, October 2000.

Goodfriend, Marvin, and Robert King. (2001) "The Case for Price Stability." NBER Working Paper # 8423.

King, Robert, and Sergio Rebelo. (1999) "Resuscitating Real Business Cycles." In *Handbook of Macroeconomic*, Volume 1B, edited by John Taylor and Michael Woodford. Amsterdam: Elsevier.

King, Robert, and Alexander Wolman. (1999) "What Should the Monetary Authority Do When Prices are Sticky?" In *Monetary Policy Rules*, edited by John Taylor. Chicago: University of Chicago Press.

Levin, Andrew, Alexei Onatski, John Williams, and Noah Williams. (2005) "Monetary Policy under Uncertainty in Micro-Founded Macroeconometric Models." In *NBER Macroeconomics Annual 2005*, edited by Mark Gertler and Kenneth Rogoff.

Lucas, Robert. (2003) "Macroeconomic Priorities." American Economic Association Presidential Address. *American Economic Review,* 93:1, 1–14.

Neiss, Katharine S., and Edward Nelson. (2003) "The Real-Interest-Rate Gap as an Inflation Indicator." *Macroeconomic Dynamics,* 7:2, 239−62.

Paustian, Mathias. (2005) "The Role of Contracting Schemes for the Welfare Costs of Nominal Rigidities over the Business Cycle." Mimeo, April 2005.

Rotemberg, Julio J., and Michael Woodford. (1997) "An Optimization Based Econometric Framework for the Evaluation of Monetary Policy." In *NBER Macroeconomics Annual 1997,* edited by Ben S. Bernanke, and Julio J. Rotemberg, pp. 297–346.

Smets, Frank, and Raf Wouters. (2003a) "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association,* 1, 1123–75.

Smets, Frank, and Raf Wouters. (2003b) "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." Mimeo, May 2003.