

Optimal Inflation Targeting Rules *

Marc P. Giannoni
Columbia University

Michael Woodford
Princeton University

February 24, 2003

Abstract

This paper characterizes optimal monetary policy for a range of alternative economic models, applying the general theory developed in Giannoni and Woodford (2002a). The rules computed here have the advantage of being optimal regardless of the assumed character of exogenous additive disturbances, though other aspects of model specification do affect the form of the optimal rule.

In each case, optimal policy can be implemented through a flexible inflation targeting rule, under which the central bank is committed to adjust its interest-rate instrument so as to ensure that projections of inflation and other variables satisfy a target criterion. The paper shows which additional variables should be taken into account, in addition to the inflation projection, and to what extent, for any given parameterization of the structural equations. It also explains what relative weights should be placed on projections for different horizons in the target criterion, and the manner and degree to which the target criterion should be history-dependent.

The likely quantitative significance of the various factors considered in the general discussion is then assessed by estimating a small, structural model of the U.S. monetary transmission with explicit optimizing foundations. An optimal policy rule is computed for the estimated model, and shown to correspond to a multi-stage inflation-forecast targeting procedure. The degree to which actual U.S. policy over the past two decades has conformed to the optimal target criteria is then considered.

*We would like to thank Jean Boivin, Rick Mishkin, Ed Nelson, and Lars Svensson for helpful discussions, Brad Strum for research assistance, and the National Science Foundation for research support through a grant to the NBER.

An increasingly popular approach to the conduct of monetary policy, since the early 1990s, has been inflation-forecast targeting. Under this general approach, a central bank is committed to adjust short-term nominal interest rates periodically so as to ensure that its projection for the economy’s evolution satisfies an explicit target criterion — for example, in the case of the Bank of England, the requirement that the RPIX inflation rate be projected to equal 2.5 percent at a horizon two years in the future (Vickers, 1998). Such a commitment can overcome the inflationary bias that is likely to follow from discretionary policy guided solely by a concern for social welfare, and can also help to stabilize medium-term inflation expectations around a level that reduces the output cost to the economy of maintaining low inflation.

Another benefit that is claimed for such an approach (*e.g.*, King, 1997; Bernanke *et al.*, 1999)— and an important advantage, at least in principle, of inflation targeting over other policy rules, such as a k -percent rule for monetary growth, that should also achieve a low average rate of inflation — is the possibility of combining reasonable stability of the inflation rate (especially over the medium to long term) with optimal short-run responses to real disturbances of various sorts. Hence Svensson (1999) argues for the desirability of “flexible” inflation targeting, by which it is meant¹ that the target criterion involves not only the projected path of the inflation rate, but one or more other variables, such as a measure of the output gap, as well.

We here consider the question of what sort of additional variables ought to matter — and with what weights, and what dynamic structure — in a target criterion that is intended to implement optimal policy. We wish to use economic theory to address questions such as which measure of inflation is most appropriately targeted (an index of goods prices only, or wage inflation as well?), which sort of output gap, if any, should justify short-run departures of projected inflation from the long-run target rate (a departure of real GDP from a smooth

¹Svensson discusses two alternative specifications of an inflation-targeting policy rule, one of which (a “general targeting rule”) involves specification of a loss function that the central bank should use to evaluate alternative paths for the economy, and the other of which (a “specific targeting rule”) involves specification of a target criterion. We are here concerned solely with policy prescriptions of the latter sort. On the implementation of optimal policy through a “general targeting rule,” see Svensson and Woodford (2003).

trend path, or from a “natural rate” that varies in response to a variety of disturbances?), and how large a modification of the acceptable inflation projection should result from a given size of projected output gap. We also consider how far in the future the inflation and output projections should extend upon which the current interest-rate decision is based, and the degree to which an optimal target criterion should be *history-dependent*, i.e., should depend on recent conditions, and not simply on the projected paths of inflation and other target variables from now on.

In a recent paper (Giannoni and Woodford, 2002a), we expound a general approach to the design of an optimal target criterion. We show, for a fairly general class of linear-quadratic policy problems, how it is possible to choose a target criterion that will satisfy several desiderata. First, the target criterion has the property that insofar as the central bank is expected to ensure that it holds at all times, this expectation will imply the existence of a determinate rational-expectations equilibrium. Second, that equilibrium will be optimal, from the point of view of a specified quadratic loss function, among all possible rational-expectations equilibria, given one’s model of the monetary transmission mechanism.² Thus the policy rule implements the optimal state-contingent evolution of the economy, in the sense of given it a reason to occur if the private sector is convinced of the central bank’s commitment to the rule and fully understands its implications.

Third, the rule is robustly optimal, in the sense that the same target criterion brings about an optimal state-contingent evolution of the economy regardless of the assumed statistical properties of the exogenous disturbances, despite the fact that the target criterion makes no explicit reference to the particular types of disturbances that may occur (except insofar as these may be involved in the definition of the target variables — the variables appearing

²Technically, the state-contingent evolution that is implemented by commitment to the policy rule is optimal from a “timeless perspective” of the kind proposed in Woodford (1999b), which means that it would have been chosen as part of an optimal commitment at a date sufficiently far in the past for the policymaker to fully internalize the implications of the anticipation of the specified policy actions, as well as their effects at the time that they are taken. This modification of the concept of optimality typically used in Ramsey-style analyses of optimal policy commitments allows a time-invariant policy rule to be judged optimal, and eliminates the time inconsistency of optimal policy. See Giannoni and Woodford (2002a) and Svensson and Woodford (2003) for further discussion.

in the loss function which defines the stabilization objectives). This robustness greatly increases the practical interest in the computation of a target criterion that is intended to implement optimal state-contingent responses to disturbances; for actual economies are affected by an innumerable variety of types of disturbances, and central banks always have a great deal of specific information about the ones that have most recently occurred. The demand that the target criterion be robustly optimal also allows us to obtain much sharper conclusions as to the form of an optimal target criterion. For while there would be a very large number of alternative relations among the paths of inflation and other variables that are equally consistent with the optimal state-contingent evolution in the case of a particular type of assumed disturbances, only relations of a very special sort continue to describe the optimal state-contingent evolution even if one changes the assumed character of the exogenous disturbances affecting the economy.

Our general characterization in Giannoni and Woodford (2002a) is in terms of a fairly abstract notation, involving eigenvectors and matrix lag polynomials. Here we offer examples of the specific character of the optimally flexible inflation targets that can be derived using that theory. Our results are of two sorts. First, we illustrate the implications of the theory in the context of a series of simple models that incorporate important features of realistic models of the monetary transmission mechanism. Such features include wage and price stickiness, inflation inertia, habit persistence, and predeterminedness of pricing and spending decisions. In the models considered, there is a tension between two or more of the central bank's stabilization objectives, that cannot simultaneously be achieved in full; in the simplest case, this is a tension between inflation and output-gap stabilization, but we also consider models in which it is reasonable to seek to stabilize interest rates or wage inflation as well. These results in the context of very simple models are intended to give insight into the way in which the character of the optimal target criterion should depend on one's model of the economy, and should be of interest even to readers who are not persuaded of the empirical realism of our estimated model.

Second, we apply the theory to a small quantitative model of the U.S. monetary transmis-

sion mechanism, the numerical parameters of which are fit to VAR estimates of the impulse responses of several aggregate variables to identified monetary policy shocks. While the model remains an extremely simple one, this exercise makes an attempt to judge the likely quantitative significance of the types of effects that have previously been discussed in more general terms. It also offers a tentative evaluation of the extent to which U.S. policy over the past two decades has differed from what an optimal inflation-targeting regime would have called for.

1 Model Specification and Optimal Targets

Here we offer a few simple examples of the way in which the optimal target criterion will depend on the details of one's model of the monetary transmission mechanism. (The optimal target criterion also depends, of course, on one's assumed stabilization objectives. But here we shall take the view that the appropriate stabilization objectives follow from ones assumptions about the way in which policy affects the economy, though the welfare-theoretic stabilization objectives implied by our various simple models are here simply asserted rather than derived.) The examples that we select illustrate the consequences of features that are often present in quantitative optimizing models of the monetary transmission mechanism. They are also features of the small quantitative model presented in section 2; hence our analytical results in this section are intended to provide intuition for the numerical results presented for the empirical model in section 3.

The analysis of Giannoni and Woodford (2002a) derives a robustly optimal target criterion from the first-order conditions that characterize the optimal state-contingent evolution of the economy. Here we illustrate this method by directly applying it to our simple examples, without any need to recapitulate the general theory.

1.1 An Inflation-Output Stabilization Tradeoff

We first consider the central issue addressed in previous literature on flexible inflation targeting, which is the extent to which a departure from complete (and immediate) stabilization of

inflation is justifiable in the case of real disturbances that prevent joint stabilization of both inflation and the (welfare-relevant) output gap.³ We illustrate how this question would be answered in the case of a simple optimizing model of the monetary transmission mechanism that allows for the existence of such “cost-push shocks” (to use the language of Clarida *et al.*, 1999).

As is well known, a discrete-time version of the optimizing model of staggered price-setting proposed by Calvo (1983) results in a log-linear aggregate supply relation of the form

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1} + u_t, \tag{1.1}$$

sometimes called the “New Keynesian Phillips curve” (after Roberts, 1995).⁴ Here π_t denotes the inflation rate (rate of change of a general index of goods prices), x_t the output gap (the deviation of log real GDP from a time-varying “natural rate”, defined so that stabilization of the output gap is part of the welfare-theoretic stabilization objective⁵), and the disturbance term u_t is a “cost-push shock”, collecting all of the exogenous shifts in the equilibrium relation between inflation and output that do not correspond to shifts in the welfare-relevant “natural rate” of output. In addition, $0 < \beta < 1$ is the discount factor of the representative household, and $\kappa > 0$ is a function of a number of features of the underlying structure, including both the average frequency of price adjustment and the degree to which Ball-Romer (1990) “real rigidities” are important.

We shall assume that the objective of monetary policy is to minimize the expected value

³Possible sources of disturbances of this sort are discussed in Giannoni (2000), Steinsson (2002), and Woodford (2003, chap. 6).

⁴See Woodford (2003, chap. 3) for a derivation in the context of an explicit intertemporal general equilibrium model of the transmission mechanism. Equation (1.1) represents merely a log-linear approximation to the exact equilibrium relation between inflation and output implied by this pricing model; however, under circumstances discussed in Woodford (2003, chap. 6), such an approximation suffices for a log-linear approximate characterization of the optimal responses of inflation and output to small enough disturbances. Similar remarks apply to the other log-linear models presented below.

⁵See Woodford (2003, chaps. 3 and 6) for discussion of how this variable responds to a variety of types of real disturbances. Under conditions discussed in chapter 6, the “natural rate” referred to here corresponds to the equilibrium level of output in the case that all wages and prices were completely flexible. However, our results in this section apply to a broader class of model specifications, under an appropriate definition of the “output gap”.

of a loss function of the form

$$W = E_0 \left\{ \sum_{t=0}^{\infty} \beta^t L_t \right\}, \quad (1.2)$$

where the discount factor β is the same as in (1.1), and the loss each period is given by

$$L_t = \pi_t^2 + \lambda(x_t - x^*)^2, \quad (1.3)$$

for a certain relative weight $\lambda > 0$ and optimal level of the output gap $x^* > 0$. Under the same microfoundations as justify the structural relation (1.1), one can show (Woodford, 2003, chap. 6) that a quadratic approximation to the expected utility of the representative household is a decreasing function of (1.2), with

$$\lambda = \kappa/\theta \quad (1.4)$$

(where $\theta > 1$ is the elasticity of substitution between alternative differentiated goods) and x^* a function of both the degree of market power and the size of tax distortions. However, we here offer an analysis of the optimal target criterion in the case of any loss function of the form (1.3), regardless of whether the weights and target values are the ones that can be justified on welfare-theoretic grounds or not. (In fact, a quadratic loss function of this form is frequently assumed in the literature on monetary policy evaluation, and is often supposed to represent the primary stabilization objectives of actual inflation-targeting central banks in positive characterizations of the consequences of inflation targeting.)

The presence of disturbances of the kind represented by u_t in (1.1) creates a tension between the two stabilization goals reflected in (1.3) of inflation stabilization on the one hand and output-gap stabilization (around the value x^*) on the other; under an optimal policy, the paths of both variables will be affected by cost-push shocks. The optimal responses can be found by computing the state-contingent paths $\{\pi_t, x_t\}$ that minimize (1.2) with loss function (1.3) subject to the sequence of constraints (1.1).⁶ The Lagrangian for this problem,

⁶Note that the aggregate-demand side of the model does not matter, as long as a nominal interest-rate path exists that is consistent with any inflation and output paths that may be selected. This is true if, for example, the relation between interest rates and private expenditure is of the form (1.15) assumed below, and the required path of nominal interest rates is always non-negative. We assume here that the non-negativity constraint never binds, which will be true, under the assumptions of the model, in the case of any small enough real disturbances $\{u_t, r_t^n\}$.

looking forward from any date t_0 , is of the form

$$\mathcal{L}_{t_0} = E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \left\{ \frac{1}{2} [\pi_t^2 + \lambda_x (x_t - x^*)^2] + \varphi_t [\pi_t - \kappa x_t - \beta \pi_{t+1}] \right\}, \quad (1.5)$$

where φ_t is a Lagrange multiplier associated with constraint (1.1) on the possible inflation-output pairs in period t . In writing the constraint term associated with the period t AS relation, it does not matter that we substitute π_{t+1} for $E_t \pi_{t+1}$; for it is only the conditional expectation of the term at date t_0 that matters in (1.5), and the law of iterated expectations implies that

$$E_{t_0} [\varphi_t E_t \pi_{t+1}] = E_{t_0} [E_t (\varphi_t \pi_{t+1})] = E_{t_0} [\varphi_t \pi_{t+1}]$$

for any $t \geq t_0$.

Differentiating (1.5) with respect to the levels of inflation and output each period, we obtain a pair of first-order conditions

$$\pi_t + \varphi_t - \varphi_{t-1} = 0, \quad (1.6)$$

$$\lambda(x_t - x^*) - \kappa \varphi_t = 0, \quad (1.7)$$

for each period $t \geq t_0$. These conditions, together with the structural relation (1.1), have a unique non-explosive solution⁷ for the inflation rate, the output gap, and the Lagrange multiplier (a unique solution in which the paths of these variables are bounded if the shocks u_t are bounded), and this solution (which therefore satisfies the transversality condition) indicates the optimal state-contingent evolution of inflation and output.

As an example, Figure 1, plots the impulse responses to a positive cost-push shock, in the simple case that the cost-push shock is purely transitory, and unforecastable before the period in which it occurs (so that $E_t u_{t+j} = 0$ for all $j \geq 1$). Here the assumed values of $\beta, \kappa,$

⁷Obtaining a unique solution requires the specification of an initial value for the Lagrange multiplier φ_{t_0-1} . See Woodford (2003, chap. 7) for the discussion of alternative possible choices of this initial condition and their significance. Here we note simply that regardless of the value chosen for φ_{t_0-1} , the optimal responses to cost-push shocks in period t_0 and later are the same.




Figure 1: Optimal responses to a positive cost-push shock under commitment, in the case of Calvo pricing.

and λ are those given in Table 1,⁸ and the shock in period zero is of size $u_0 = 1$; the periods represent quarters, and the inflation rate is plotted as an annualized rate, meaning that what is plotted is actually $4\pi_t$. As one might expect, in an optimal equilibrium inflation is allowed to increase somewhat in response to a cost-push shock, so that the output gap need not fall as much as would be required to prevent any increase in the inflation rate. Perhaps less intuitively, the figure also shows that under an optimal commitment, monetary policy remains tight even after the disturbance has dissipated, so that the output gap returns to

⁸These parameter values are based on the estimates of Rotemberg and Woodford (1997) for a slightly more complex variant of the model used here and in section 1.3. The coefficient λ here corresponds to λ_x in the table. Note also that the value of .003 for that coefficient refers to a loss function in which π_t represents the *quarterly* change in the log price level. If we write the loss function in terms of an annualized inflation rate, $4\pi_t$, as is conventional in numerical work, then the relative weight on the output-gap stabilization term would actually be $16\lambda_x$, or about .048. Of course, this is still quite low compared the relative weights often assumed in the *ad hoc* stabilization objectives used in the literature on the evaluation of monetary policy rules.

Table 1: Calibrated parameter values for the examples in section 1.

Structural parameters	
β	0.99
κ	.024
θ^{-1}	0.13
σ^{-1}	0.16
Shock processes	
ρ_u	0
ρ_r	0.35
Loss function	
λ_x	.003
λ_i	.236

zero only much more gradually. As a result of this, while inflation overshoots its long-run target value at the time of the shock, it is held *below* its long-run target value for a time following the shock, so that the unexpected increase in prices is subsequently undone. In fact, as the bottom panel of the figure shows, under an optimal commitment, the price level eventually returns to exactly the same path that it would have been expected to follow if the shock had not occurred.

This simple example illustrates a very general feature of optimal policy once one takes account of forward-looking private-sector behavior: optimal policy is almost always *history-dependent*. That is, it depends on the economy's recent history and not simply on the set of possible state-contingent paths for the target variables (here, inflation and the output gap) that are possible from now on. (In the example shown in the figure, the set of possible rational-expectations equilibrium paths for inflation and output from period t onward depends only on the value of u_t ; but under an optimal policy, the actually realized inflation rate and output gap depend on past disturbances as well.) This is because a commitment to respond later to past conditions can shift expectations at the earlier date in a way that helps to achieve the central bank's stabilization objectives. In the present example, if price-setters are forward-looking, the anticipation that a current increase in the general price level will

predictably be “undone” soon gives suppliers a reason not to increase their own prices currently as much as they otherwise would. This leads to smaller equilibrium deviations from the long-run inflation target at the time of the cost-push shock, without requiring such a large change in the output gap as would be required to stabilize inflation to the same degree without a change in expectations regarding future inflation. (The impulse responses under the best possible equilibrium that does not involve history-dependence are shown by the dashed lines in the figure.⁹ Note that a larger initial output contraction is required, even though both the initial price increase and the long-run price increase caused by the shock are greater.)

It follows that no *purely forward-looking* target criterion — one that involves only the projected paths of the target variables from the present time onward, like the criterion that is officially used by the Bank of England — can possibly determine an equilibrium with the optimal responses to disturbances. Instead, a history-dependent target criterion is necessary, as stressed by Svensson and Woodford (2003).

A target criterion that works is easily derived from the first-order conditions (1.6) – (1.7). Eliminating the Lagrange multiplier, one is left with a linear relation

$$\pi_t + \phi(x_t - x_{t-1}) = 0, \tag{1.8}$$

with a coefficient $\phi = \lambda/\kappa > 0$, that the state-contingent evolution of inflation and the output gap must satisfy. Note that this relation must hold in an optimal equilibrium regardless of the assumed statistical properties of the disturbances. One can also show that a commitment to ensure that (1.8) holds each period from some date t_0 onward implies the existence of a determinate rational-expectations equilibrium,¹⁰ given any initial output gap x_{t_0-1} . In this equilibrium, inflation and output evolve according to the optimal state-contingent evolution

⁹See Woodford (2003, chap. 7) for derivation of this “optimal non-inertial plan.” In the example shown in Figure 1, this optimal non-inertial policy corresponds to the Markov equilibrium resulting from discretionary optimization by the central bank. That equivalence would not obtain, however, in the case of serially correlated disturbances.

¹⁰The characteristic equation that determines whether the system of equations consisting of (1.1) and (1.8) has a unique non-explosive solution is the same as for the system of equations solved above for the optimal state-contingent evolution.

characterized above.

This is the optimal target criterion that we are looking for: it indicates that deviations of the projected inflation rate π_t from the long-run inflation target (here equal to zero) should be accepted that are proportional to the degree to which the output gap is projected to decline over the same period that prices are projected to rise. Note that this criterion is history-dependent, because the acceptability of a given projection (π_t, x_t) depends on the recent past level of the output gap; it is this feature of the criterion that will result in the output gap's returning only gradually to its normal level following a transitory cost-push shock, as shown in Figure 1.

How much of a projected change in the output gap is needed to justify a given degree of departure from the long-run inflation target? If λ is assigned the value that it takes in the welfare-theoretic loss function, then $\phi = \theta^{-1}$, where θ is the elasticity of demand faced by the typical firm. The calibrated value for this parameter given in Table 1 (based on the estimates of Rotemberg and Woodford, 1997) implies that $\phi = .13$. If we express the target criterion in terms of the annualized inflation rate ($4\pi_t$) rather than the quarterly rate of price change, the relative weight on the projected quarterly change in the output gap will instead be 4ϕ , or about 0.51. Hence a projection of a decline in real GDP of two percentage points relative to the natural rate of output over the coming quarter would justify an increase in the projected (annualized) rate of inflation of slightly more than one percentage point.

1.2 Inflation Inertia

A feature of the “New Keynesian” aggregate-supply relation (1.1) that has come in for substantial criticism in the empirical literature is the fact that past inflation rates play no role in the determination of current equilibrium inflation. Instead, empirical models of the kind used in central banks for policy evaluation often imply that the path of the output gap required in order to achieve a particular path for the inflation rate from now onward depends on what rate of inflation has already been recently experienced; and this aspect of one's model is of obvious importance for the question of how rapidly one should expect

that it is optimal to return inflation to its normal level, or even to “undo” past unexpected price-level increases, following a cost-push shock.

A simple way of incorporating inflation inertia of the kind that central-bank models often assume into an optimizing model of pricing behavior is to assume, as Christiano *et al.* (2001) propose, that individual prices are indexed to an aggregate price index during the intervals between re-optimizations of the individual prices, and that the aggregate price index becomes available for this purpose only with a one-period lag. When the Calvo model of staggered price-setting is modified in this way, the aggregate-supply relation (1.1) takes the more general form¹¹

$$\pi_t - \gamma\pi_{t-1} = \kappa x_t + \beta E_t[\pi_{t+1} - \gamma\pi_t] + u_t, \quad (1.9)$$

where the coefficient $0 \leq \gamma \leq 1$ indicates the degree of automatic indexation to the aggregate price index. In the limiting case of complete indexation ($\gamma = 1$), the case assumed by Christiano *et al.* and the case found to best fit US data in our own estimation results below, this relation is essentially identical to the aggregate-supply relation proposed by Fuhrer and Moore (1995), which has been widely used in empirical work.

The welfare-theoretic stabilization objective corresponding to this alternative structural model is of the form (1.2) with the period loss function (1.3) replaced by

$$L_t = (\pi_t - \gamma\pi_{t-1})^2 + \lambda(x_t - x^*)^2, \quad (1.10)$$

where $\lambda > 0$ is again given by (1.4), and $x^* > 0$ is similarly the same function of underlying microeconomic distortions as before.¹² (The reason for the change is that with the automatic indexation, the degree to which the prices of firms that re-optimize their prices and those that do not are different depends on the degree to which the current overall inflation rate π_t differs from the rate at which the automatically adjusted prices are increasing, *i.e.*, from $\gamma\pi_{t-1}$.) If we consider the problem of minimizing (1.2) with loss function (1.10) subject to

¹¹See Woodford (2003, chap. 3) for a derivation from explicit microeconomic foundations.

¹²See Woodford (2003, chap. 6) for derivation of this loss function as an approximation to expected utility.




Figure 2: Optimal responses to a positive cost-push shock under commitment, for alternative degrees of inflation inertia.

the sequence of constraints (1.9), the problem has the same form as in the previous section, except with π_t everywhere replaced by the quasi-differenced inflation rate

$$\pi_t^{qd} \equiv \pi_t - \gamma\pi_{t-1}. \quad (1.11)$$

The solution is therefore also the same, with this substitution.

Figure 2 shows the impulse responses of inflation, the output gap, and the price level to the same kind of disturbance as in Figure 1, under optimal policy for economies with alternative values of the indexation parameter γ . (The values assumed for β , κ , and λ are again as in Table 1.) Once again, under an optimal commitment, the initial unexpected increase in prices is eventually undone, as long as $\gamma < 1$; and this once again means that inflation eventually undershoots its long-run level for a time. However, for any large enough value of γ , inflation remains greater than its long-run level for a time even after the disturbance

has ceased, and only later undershoots its long-run level; and the larger is γ , the longer this period of above-average inflation persists. In the limiting case that $\gamma = 1$, the undershooting never occurs; inflation is simply gradually brought back to the long-run target level.¹³ In this last case, a temporary disturbance causes a permanent change in the price level, even under optimal policy. However, the *inflation rate* is eventually restored to its previously anticipated long-run level under an optimal commitment, even though the rate of inflation (as opposed to the rate of *acceleration* of inflation) is not welfare-relevant in this model. (Note that the optimal responses shown in Figure 2 for the case $\gamma = 1$ correspond fairly well to the conventional wisdom of inflation-targeting central banks; but our theoretical analysis allows us to compute an optimal rate at which inflation should be projected to return to its long-run target value following a disturbance.)

As in the previous section, we can derive a target criterion that implements the optimal responses to disturbances regardless of the assumed statistical properties of the disturbances. This optimal target criterion is obtained by replacing π_t in (1.8) by π_t^{qd} , yielding

$$\pi_t - \gamma\pi_{t-1} + \phi(x_t - x_{t-1}) = 0, \tag{1.12}$$

where $\phi > 0$ is the same function of model parameters as before. This indicates that the acceptable inflation projection for the current period should depend not only on the projected change in the output gap, but also (insofar as $\gamma > 0$) on the recent past rate of inflation: a higher existing inflation rate justifies a higher projected near-term inflation rate, in the case of any given output-gap projection.

In the special case that $\gamma = 1$, the optimal target criterion adjusts the current inflation target one-for-one with increases in the existing rate of inflation — the target criterion actually involves only the rate of acceleration of inflation. But this does not mean that disturbances are allowed to permanently shift the inflation rate to a new level, as shown in Figure 2. In fact, in the case of full indexation, an alternative target criterion that also leads

¹³Note that the impulse response of inflation (for $\gamma = 1$) in panel 1 of Figure 2 is the same as the impulse response of the price level (under optimal policy) in panel 3 of Figure 1. The scales are different because the inflation rate plotted is an annualized rate, $4\pi_t$ rather than π_t .

to the optimal equilibrium responses to cost-push shocks is the simpler criterion

$$\pi_t + \phi x_t = \bar{\pi}, \tag{1.13}$$

where again $\phi > 0$ is the same coefficient as in (1.12), and the value of the long-run inflation target $\bar{\pi}$ is arbitrary (but not changing over time). Note that (1.8) is just a first-differenced form of (1.13), and a commitment to ensure that (1.12) holds in each period $t \geq t_0$ is equivalent to a commitment to ensure that (1.13) holds, for a particular choice of $\bar{\pi}$, namely $\bar{\pi} = \pi_{t_0-1} + \phi x_{t_0-1}$. But the choice of $\bar{\pi}$ has no effect on either the determinacy of equilibrium or the equilibrium responses of inflation and output to real disturbances (only on the long-run average inflation rate), and so any target criterion of the form (1.13) implements the optimal responses to disturbances.¹⁴ Note that this optimal target criterion is similar in form to the kind that Svensson (1999) suggests as a description of the behavior of actual inflation-targeting central banks, except that the inflation and output-gap projections in (1.13) are not so far in the future (they refer only to the coming quarter) as in the procedures of actual inflation targeters.

The result that the long-run inflation target associated with an optimal target criterion is indeterminate depends, of course, on the fact that we have assumed a model in which no distortions depend on the inflation rate, as opposed to its rate of change. This is logically possible, but unlikely to be true in reality. (Distortions that depend on the level of nominal interest rates, considered in the next section, would be one example of a realistic complication that would break this result, even in the case of full indexation.) Because the model considered here with $\gamma = 1$ does not determine any particular optimal long-run inflation target (it need *not* vary with the initially existing inflation rate, for example), even a small perturbation of these assumptions is likely to determine an optimal long-run inflation

¹⁴Any such policy rule is also optimal from a timeless perspective, under the definition given in Giannoni and Woodford (2002a). Note that alternative rules, that result in equilibria that differ only in a transitory, deterministic component of the path of each of the target variables, can each be considered optimal in this sense. This ambiguity as to the initial behavior of the target variables cannot be resolved if our concept of optimal policy is to be time-consistent. In the present case, ambiguity about the required initial behavior of the target variable, inflation acceleration, implies ambiguity about the required long-run average level of the inflation rate, though there is no ambiguity about how inflation should respond to shocks.

target, and this will generally be independent of the initially existing rate of inflation. (The monetary frictions considered in the next subsection provide an example of this.)

It is worth noting that even though the optimal dynamic responses shown in Figure 2 for the case of large γ confirm the conventional wisdom of inflation-targeting central bankers with regard the desirability of a gradual return of the inflation rate to its long-run target level following a cost-push shock, the optimal target criterion for this model does *not* involve a “medium-term” inflation forecast rather than a shorter-run projection. Even in the case that we suppose that the central bank will often have advance information about disturbances that will shift the aggregate-supply relation only a year or more in the future, the robust description of optimal policy is one that indicates how short-run output-gap projections should modify the acceptable short-run inflation projection, rather than one that checks only that some more distant inflation forecast is still on track. Of course, a commitment to the achievement of the target criterion (1.12) each period does imply that the projection of inflation several quarters in the future should never depart much from the long-run inflation target; but the latter stipulation is not an equally useful guide to what should actually be done with interest rates at a given point in time.

1.3 An Interest-Rate Stabilization Objective

The policy problems considered above assume that central banks care only about the paths of inflation and the output gap, and not about the behavior of nominal interest rates that may be required to bring about a given evolution of inflation and output that is consistent with the aggregate-supply relation. However, actual central banks generally appear to care about reducing the volatility of nominal interest rates as well (Goodfriend, 1991). Such a concern can also be justified in terms of microeconomic foundations that are consistent with the kind of aggregate-supply relations assumed above, as discussed in Woodford (2003, chap. 6).

For example, the transactions frictions that account for money demand imply a distortion that should be an increasing function of the nominal interest rate, as stressed by Friedman

(1969); the deadweight loss resulting from a positive opportunity cost of holding money should also be a convex function of the interest rate, at least for interest rates close enough to the optimal one (the interest rate paid on base money). Alternatively, the existence of a zero lower bound on nominal interest rates can make it desirable to accept somewhat greater variability of inflation and the output gap for the sake of reducing the required variability of nominal interest rates, given that the smaller range of variation in the nominal interest rate allows the average nominal interest rate (and hence the average inflation rate) to be lower. A quadratic penalty for deviations of the nominal interest rate from a target level may then be justified as a proxy for a constraint that links the feasible average level of nominal interest rates to the variability of the nominal interest rate.

For any of these reasons, we may be interested in a policy that minimizes a loss function of the form

$$L_t = \pi_t^2 + \lambda_x(x_t - x^*)^2 + \lambda_i(i_t - i^*)^2, \quad (1.14)$$

where $\lambda_x > 0$ is the same function of underlying parameters as λ in (1.3), i_t is a short-term nominal interest rate, $\lambda_i > 0$ for one of the reasons discussed above, and i^* is the level around which the nominal interest rate would ideally be stabilized. In this case, the aggregate-supply relation is not the only relevant constraint in our optimal policy problem; it also matters what interest-rate path is required in order to induce a given evolution of aggregate demand.

In a simple optimizing model that has been used in many recent analyses of optimal monetary policy (e.g., McCallum and Nelson, 1999; Clarida *et al.*, 1999; and Woodford, 1999a), the aggregate-supply relation (1.1) is combined with an intertemporal Euler equation for the timing of private expenditure of the form

$$x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1} - r_t^n), \quad (1.15)$$

where $\sigma > 0$ represents the intertemporal elasticity of substitution and r_t^n exogenous variation in Wicksell's "natural rate of interest." Real disturbances that cause the natural rate of interest to vary are now another reason why (if $\lambda_i > 0$) it will be impossible for the central bank to completely stabilize all of its target variables simultaneously, and hence for transitory




Figure 3: Optimal responses to an increase in the natural rate of interest.

variations in the inflation rate to be optimal, even in the absence of cost-push shocks.

This leads us to consider the problem of finding the state-contingent evolution of inflation, output and interest rates to minimize the expected discounted value of (1.14) subject to the constraints (1.1) and (1.15). A similar Lagrangian method as in section 1.1 leads to first-order conditions of the form

$$\pi_t - \beta^{-1}\sigma\varphi_{1t-1} + \varphi_{2t} - \varphi_{2t-1} = 0, \quad (1.16)$$

$$\lambda_x(x_t - x^*) + \varphi_{1t} - \beta^{-1}\varphi_{1t-1} - \kappa\varphi_{2t} = 0, \quad (1.17)$$

$$\lambda_i(i_t - i^*) + \sigma\varphi_{1t} = 0, \quad (1.18)$$

where φ_{1t} is the multiplier associated with constraint (1.15) and φ_{2t} the one associated with constraint (1.1). We can once again solve this system of equations for unique bounded paths for the endogenous variables in the case of any bounded processes for the exogenous disturbances $\{r_t^n, u_t\}$. The implied optimal responses to an exogenous increase in the natural rate of interest are shown in Figure 3. Here the model parameters are calibrated as in Table 1, and the natural rate of interest is assumed to be a first-order autoregressive process with serial correlation coefficient $\rho_r = 0.35$.¹⁵

A notable feature of Figure 3 is that once again optimal policy must be history-dependent, for the optimal responses to the disturbance are more persistent than the disturbance itself. As discussed in Woodford (1999a), optimal interest-rate policy is *inertial*, in the sense that interest rates are both raised only gradually in response to an increase in the natural rate of interest, and then returned to their normal level more gradually than the natural rate itself as well. (The impulse response of the natural rate is shown by the dotted line in panel 1 of the figure.) Because spending responds to expected future interest rates and not only current short rates, it is possible to achieve a given degree of stabilization of demand (relative to the natural rate) in response to disturbances with less volatility of short-term interest rates if short rates are moved in a more inertial fashion. (The optimal responses among those achievable using a purely forward-looking target criterion are shown, for purposes of comparison, by the dashed lines in the figure.)

A history-dependent target criterion that can bring about the desired impulse responses, again regardless of the statistical properties of the disturbances r_t^n and u_t (including any assumptions about the degree of correlation between these disturbances), can be derived once more from the first-order conditions (1.16) – (1.18). Using the last two equations to substitute for the two Lagrange multipliers in the first equation, we are left with a linear

¹⁵The real disturbances that cause the natural rate of interest to vary are assumed to create no variation in the cost-push term u_t ; that is, they shift the equilibrium relation between inflation and output only through possible shifts in the natural rate of output. A variety of examples of real disturbances with this property are discussed in Woodford (2003, chap. 6).

relation of the form

$$A(L)(i_t - i^*) = \phi_\pi \pi_t + \phi_x (x_t - x_{t-1}) \quad (1.19)$$

that must be satisfied each period under an optimal policy. Here the coefficients of the lag polynomial are

$$A(L) \equiv 1 - \left(1 + \frac{\kappa\sigma}{\beta}\right) L - \beta^{-1} L(1 - L),$$

and the inflation and output response coefficients are

$$\phi_\pi = \frac{\kappa\sigma}{\lambda_i} > 0, \quad \phi_x = \frac{\sigma\lambda_x}{\lambda_i} > 0. \quad (1.20)$$

One can furthermore show that this is not only a necessary feature of an optimal equilibrium, but also suffices to characterize it, in the sense that the system consisting of equation (1.19) together with the structural equations (1.1) and (1.15) has a unique non-explosive solution, in which the equilibrium responses to shocks are optimal.¹⁶

Requirement (1.19) can be interpreted as an inertial Taylor rule, as discussed in Giannoni and Woodford (2002b). However, this requirement can also be equivalently expressed in a forward-integrated form, that more directly generalizes the optimal target criterion derived in section 1.1. It is easily seen that our sign assumptions on the model parameters imply that $A(L)$ can be factored as

$$A(L) \equiv (1 - \lambda_1 L)(1 - \lambda_2 L),$$

where $0 < \lambda_1 < 1 < \lambda_2$. It then follows that (1.19) is equivalent to

$$(1 - \lambda_1 L)(i_{t-1} - i^*) = -\lambda_2^{-1} E_t[(1 - \lambda_2^{-1} L^{-1})^{-1}(\phi_\pi \pi_t + \phi_x \Delta x_t)], \quad (1.21)$$

in the sense that bounded stochastic processes $\{i_t, \pi_t, x_t\}$ satisfy (1.19) for all $t \geq t_0$ if and only if they satisfy (1.21) for all $t \geq t_0$.¹⁷ Hence a commitment to ensure that (1.21) is satisfied at all times implies a determinate rational-expectations equilibrium in which the responses to shocks are optimal. This conclusion is once again independent of any assumption about the statistical properties of the disturbances, so that (1.21) is a robustly optimal target criterion.

¹⁶See Giannoni and Woodford (2002b), Proposition 6.

¹⁷See Giannoni and Woodford (2002b), Proposition 7.

This optimal target criterion can be expressed in the form

$$F_t(\pi) + \phi F_t(x) = \theta_x x_{t-1} - \theta_i (i_{t-1} - i^*) - \theta_\Delta \Delta i_{t-1}, \quad (1.22)$$

where for each of the variables $z = \pi, x$ we use the notation $F_t(z)$ for a conditional forecast

$$F_t(z) \equiv \sum_{j=0}^{\infty} \alpha_{z,j} E_t z_{t+j}$$

involving weights $\{\alpha_{z,j}\}$ that sum to one. Thus the criterion specifies a time-varying target value for a weighted average of an inflation forecast and an output-gap forecast, where each of these forecasts is in fact a weighted average of forecasts at various horizons, rather than a projection for a specific future date. The coefficients of this representation of optimal policy are given by

$$\phi = \theta_x = (1 - \lambda_2^{-1}) \frac{\lambda_x}{\kappa} > 0,$$

$$\theta_i = \lambda_2 (1 - \lambda_1) (1 - \lambda_2^{-1}) \frac{\lambda_i}{\kappa \sigma} > 0,$$

$$\theta_\Delta = \lambda_1 \lambda_2 (1 - \lambda_2^{-1}) \frac{\lambda_i}{\kappa \sigma} > 0,$$

while the optimal weights in the conditional forecasts are

$$\alpha_{\pi,j} = \alpha_{x,j} = (1 - \lambda_2^{-1}) \lambda_2^{-j}.$$

Thus the optimal conditional forecast is one that places positive weight on the projection for each future period, beginning with the current period, with weights that decline exponentially as the horizon increases. The mean distance in the future of the projections that are relevant to the target criterion is equal to

$$\sum_{j=0}^{\infty} \alpha_{z,j} j = (\lambda_2 - 1)^{-1}$$

for both the inflation and output-gap forecasts.

In the case of the calibrated parameter values in Table 1, the rate at which these weights decay per quarter is $\lambda_2^{-1} = .68$, so that the mean forecast horizon in the optimal target criterion is 2.1 quarters. Thus while the optimal target criterion in this case involves projections of inflation and output beyond the current quarter, the forecast horizon remains quite short compared to the actual practice of inflation forecast-targeting central banks. For these same parameter values, the optimal relative weight on the output-gap forecast is $\phi = .04$,¹⁸ indicating that the target criterion is largely an inflation target. The remaining optimal coefficients are $\theta_x = .04$, $\theta_i = .24$, and $\theta_\Delta = .51$, indicating a substantial degree of history-dependence of the optimal flexible inflation target. The fact that $\theta_x = \phi$ indicates that it is the forecasted increase in the output gap *relative* to the previous quarter's level, rather than the absolute level of the gap, that should modify the inflation target, just as in section 1.1. The signs of θ_i and θ_Δ imply that policy will be made tighter (in the sense of demanding a lower modified inflation forecast) when interest rates have been high and/or increasing in the recent past; this is a way of committing to interest-rate inertia of the kind shown in Figure 3.

Note that in the limiting case in which $\lambda_i = 0$, this target criterion reduces to (1.8). In that limit, θ_i , θ_Δ and the decay factor λ_2^{-1} become equal to zero, while ϕ and θ_x have a well-defined (common) positive limit. Thus in this limiting case, the optimal targeting rule is one in which the inflation target must be modified in proportion to the projected change in the output gap, but it is no longer also dependent on lagged interest rates, and the relevant inflation and output-gap projections do not involve periods beyond the current one. This will also be nearly true in the case of small enough positive values of λ_i .

We may similarly introduce an interest-rate stabilization objective in the case of the model with inflation inertia considered in section 1.2. In this case, the loss function (1.10) is generalized to

$$L_t = (\pi_t - \gamma\pi_{t-1})^2 + \lambda_x(x_t - x^*)^2 + \lambda_i(i_t - i^*)^2, \quad (1.23)$$

¹⁸If we write the target criterion in terms of a forecast for the annualized inflation rate ($4\pi_t$), the relative weight on the output-gap forecast will instead be 4ϕ , or about .15.

for some $\lambda_i > 0$ and some desired interest rate i^* . In this generalization of the problem just considered, the first-order condition (1.16) becomes instead

$$\pi_t^{qd} - \beta\gamma E_t \pi_{t+1}^{qd} - \beta^{-1}\sigma\varphi_{1t-1} - \beta\gamma E_t \varphi_{2,t+1} + (1 + \beta\gamma)\varphi_{2t} - \varphi_{2t-1} = 0, \quad (1.24)$$

where π_t^{qd} is again defined in (1.11). Conditions (1.17) – (1.18) remain as before.¹⁹

Again using the latter two equations to eliminate the Lagrange multipliers, we obtain a relation of the form

$$E_t[A(L)(i_{t+1} - i^*)] = -E_t[(1 - \beta\gamma L^{-1})q_t] \quad (1.25)$$

for the optimal evolution of the target variables. Here $A(L)$ is a cubic lag polynomial

$$A(L) \equiv \beta\gamma - (1 + \gamma + \beta\gamma)L + (1 + \gamma + \beta^{-1}(1 + \kappa\sigma))L^2 - \beta^{-1}L^3, \quad (1.26)$$

while q_t is a function of the projected paths of the target variables, defined by

$$q_t \equiv \frac{\kappa\sigma}{\lambda_i} \left[\pi_t^{qd} + \frac{\lambda_x}{\kappa} \Delta x_t \right].$$

The lag polynomial $A(L)$ can be factored as $A(L) = (1 - \lambda_1 L)L^2 B(L^{-1})$, where $B(L^{-1})$ is a quadratic polynomial, and under our sign assumptions one can further show²⁰ that $0 < \lambda_1 < 1$, while both roots of $B(L)$ are outside the unit circle. Relation (1.25) is then equivalent²¹ to a relation of the form

$$(1 - \lambda_1 L)(i_{t-1} - i^*) = -E_t[B(L^{-1})^{-1}(1 - \beta\gamma L^{-1})q_t], \quad (1.27)$$

which generalizes (1.21) to the case $\gamma \neq 0$.

This provides us with a robustly optimal target criterion that can be expressed in the form

$$F_t(\pi) + \phi F_t(x) = \theta_\pi \pi_{t-1} + \theta_x x_{t-1} - \theta_i (i_{t-1} - i^*) - \theta_\Delta \Delta i_{t-1}, \quad (1.28)$$

¹⁹One easily sees that in the case that $\gamma = 1$, the only long-run average inflation rate consistent with these conditions is $\bar{\pi} = i^* - \bar{r}$, where \bar{r} is the unconditional mean of the natural rate of interest. This is true for any $\lambda_i > 0$, no matter how small. Hence even a slight preference for lower interest-rate variability suffices to break the indeterminacy of the optimal long-run inflation target obtained for the case $\gamma = 1$ in section 1.2.

²⁰See Giannoni and Woodford (2002b), Proposition 8.

²¹See Giannoni and Woodford (2002b), Proposition 11.

generalizing (1.22). Under our sign assumptions, one can show²² that

$$\phi = \theta_x > 0,$$

$$0 < \theta_\pi \leq 1,$$

and

$$\theta_i, \theta_\Delta > 0.$$

Furthermore, for fixed values of the other parameters, as $\gamma \rightarrow 0$, θ_π approaches zero and the other parameters approach the non-zero values associated with the target criterion (1.22). Instead, as $\gamma \rightarrow 1$, θ_π approaches 1, so that the target criterion involves only the projected change in the rate of inflation relative to its already existing level, just as we found in section 1.2 when there was assumed to be no interest-rate stabilization objective.

The effects of increasing γ on the coefficients of the optimal target criterion (1.28) is illustrated in Figure 4, where the coefficients are plotted against γ , assuming the same calibrated values for the other parameters as before. It is interesting to note that each of the coefficients indicating history-dependence ($\theta_\pi, \theta_x, \theta_i$, and θ_Δ) increases with γ (except perhaps when γ is near one). Thus if there is substantial inflation inertia, it is even more important for the inflation-forecast target to vary with changes in recent economic conditions. It is also worth noting that the degree to which the inflation target should be modified in response to changes in the output-gap projection (indicated by the coefficient ϕ) increases with γ . While our conclusion for the case $\gamma = 0$ above ($\phi = .04$) might have suggested that this sort of modification of the inflation target is not too important, we find that a substantially larger response is justified if γ is large. The optimal coefficient is $\phi = 0.13$, as in sections 1.1 and 1.2, if $\gamma = 1$; and once again this corresponds to a weight of 0.51 if the inflation target is expressed as an annualized rate.

The panels of Figure 5 correspondingly show the relative weights $\alpha_{z,j}/\alpha_{z,0}$ on the forecasts at different horizons in the optimal target criterion (1.28), for each of several alternative values of γ . As above, the inclusion of an interest-rate stabilization objective makes the

²²See Giannoni and Woodford (2002b), Proposition 10.




Figure 4: Coefficients of the optimal targeting rule (1.28) as functions of γ .

optimal target criterion more forward-looking than was the case in section 1.2. Indeed, we now find, at least for high enough values of γ , that the optimal target criterion places non-negligible weight on forecasts more than a year in the future. But it is not necessarily true that a greater degree of inflation inertia justifies a target criterion with a longer forecast horizon. Increases in γ increase the optimal weights on the current-quarter projections of both inflation and the output gap (normalizing the weights to sum to one), and instead make the weights on the projections for quarters more than two quarters in the future *less* positive. At least for low values of γ (in which case the weights are all non-negative), this makes the optimal target criterion *less* forward-looking.

For higher values of γ , increases in γ do increase the absolute value of the weights on forecasts for dates one to two years in the future (these become more negative). But even




Figure 5: Relative weights on forecasts at different horizons in the optimal criterion (1.28).

in this case, the existence of inflation inertia does not justify the kind of response to longer-horizon forecasts that is typical of inflation-targeting central banks. An increase in the forecast level of inflation and/or the output gap during the second year of a bank's current projection should justify a *loosening* of current policy, in the sense of a policy intended to *raise* projected inflation and/or the output gap in the next few quarters. This is because in the model with large γ , welfare losses result from inflation variation rather than high inflation as such; a forecast of higher inflation a year from now is then a reason to accept somewhat higher inflation in the nearer term than one otherwise would.

1.4 Wages and Prices Both Sticky

A number of studies have found that the joint dynamics of real and nominal variables are best explained by a model in which wages as well as prices are sticky (e.g., Amato and Laubach, 2001b; Christiano *et al.*, 2001; Smets and Wouter, 2002; Altig *et al.*, 2002; and

Woodford, 2003, chap. 3). This is often modeled in the way suggested by Erceg *et al.* (2000), with monopolistic competition among the suppliers of different types of labor, and staggered wage setting analogous to the Calvo (1983) model of price setting. The structural equations of the supply side of this model can be written in the form

$$\pi_t = \kappa_p(x_t + u_t) + \xi_p(w_t - w_t^n) + \beta E_t \pi_{t+1}, \quad (1.29)$$

$$\pi_t^w = \kappa_w(x_t + u_t) + \xi_w(w_t^n - w_t) + \beta E_t \pi_{t+1}^w, \quad (1.30)$$

together with the identity

$$w_t = w_{t-1} + \pi_t^w - \pi_t, \quad (1.31)$$

generalizing the single equation (1.1) for the flexible-wage model. Here π_t^w represents nominal wage inflation, w_t is the log real wage, w_t^n represents exogenous variation in the “natural real wage”, and the coefficients $\xi_p, \xi_w, \kappa_p, \kappa_w$ are all positive. The coefficient ξ_p indicates the sensitivity of goods-price inflation to changes in the average gap between marginal cost and current prices; it is smaller the stickier are prices. Similarly, ξ_w indicates the sensitivity of wage inflation to changes in the average gap between households’ “supply wage” (the marginal rate of substitution between labor supply and consumption) and current wages, and measures the degree to which wages are sticky.²³

We note furthermore that $\kappa_p \equiv \xi_p \omega_p$ and $\kappa_w \equiv \xi_w (\omega_w + \sigma^{-1})$, where $\omega_p > 0$ measures the elasticity of marginal cost with respect to the quantity supplied, at a given wage; $\omega_w > 0$ measures the elasticity of the supply wage with respect to quantity produced, holding fixed households’ marginal utility of income; and $\sigma > 0$ is the same intertemporal elasticity of substitution as in (1.15). In the limit of perfectly flexible wages, ξ_w is unboundedly large, and (1.30) reduces to the contemporaneous relation $w_t - w_t^n = (\omega_w + \sigma^{-1})(x_t + u_t)$. Using this to substitute for w_t in (1.29), the latter relation then reduces to (1.1), where

$$\kappa \equiv \xi_p (\omega_p + \omega_w + \sigma^{-1}) \quad (1.32)$$

²³For further discussion of these coefficients, and explicit formulas for them in terms of the frequency of wage and price adjustment, see section 2 below.

and the cost-push shock u_t has been rescaled.

Given the proposed microeconomic foundations for these relations, Erceg *et al.* show that the appropriate welfare-theoretic stabilization objective is a discounted criterion of the form (1.2), with a period loss function of the form

$$L_t = \lambda_p \pi_t^2 + \lambda_w \pi_t^{w2} + \lambda_x (x_t - x^*)^2. \quad (1.33)$$

Here the relative weights on the various stabilization objectives are given by

$$\lambda_p = \frac{\theta_p \xi_p^{-1}}{\theta_p \xi_p^{-1} + \theta_w \phi^{-1} \xi_w^{-1}} > 0, \quad \lambda_w = \frac{\theta_w \phi^{-1} \xi_w^{-1}}{\theta_p \xi_p^{-1} + \theta_w \phi^{-1} \xi_w^{-1}} > 0, \quad (1.34)$$

$$\lambda_x = \lambda_p \frac{\kappa}{\theta_p} > 0, \quad (1.35)$$

as functions of the underlying model parameters. Note that we have normalized the weights so that $\lambda_p + \lambda_w = 1$, and that (1.35) generalizes the previous expression (1.4) for the flexible-wage case.

Here we again abstract from the motives for interest-rate stabilization discussed in the previous section. As a result, we need not specify the demand side of the model. We then wish to consider policies that minimize the criterion defined by (1.2) and (1.33), subject to the constraints (1.29) – (1.31).

The Lagrangian method illustrated above now yields a system of first-order conditions

$$\lambda_p \pi_t + \varphi_{pt} - \varphi_{p,t-1} + v_t = 0, \quad (1.36)$$

$$\lambda_w \pi_t^w + \varphi_{wt} - \varphi_{w,t-1} - v_t = 0, \quad (1.37)$$

$$\lambda_x (x_t - x^*) - \kappa_p \varphi_{pt} - \kappa_w \varphi_{wt} = 0, \quad (1.38)$$

$$v_t = \xi_p \varphi_{pt} - \xi_w \varphi_{wt} + \beta E_t v_{t+1}, \quad (1.39)$$

where φ_{pt} , φ_{wt} , v_t are the Lagrange multipliers associated with constraints (1.29), (1.30) and (1.31) respectively. We can again use three of the equations to eliminate the three Lagrange multipliers, obtaining a target criterion of the form

$$(\kappa_w - \kappa_p) \pi_t^{asym} + (\xi_p + \xi_w) q_t + (\kappa_w - \kappa_p) \{E_t[\beta q_{t+1} - q_t] - E_{t-1}[\beta q_t - q_{t-1}]\} = 0, \quad (1.40)$$

where

$$\pi_t^{asym} \equiv \lambda_p \xi_p \pi_t - \lambda_w \xi_w \pi_t^w$$

is a measure of the asymmetry between price and wage inflation,

$$\pi_t^{sym} \equiv \frac{\lambda_p \kappa_p \pi_t + \lambda_w \kappa_w \pi_t^w}{\lambda_p \kappa_p + \lambda_w \kappa_w}$$

is a (weighted) average of the rates of price and wage inflation, and

$$q_t \equiv (\lambda_p \kappa_p + \lambda_w \kappa_w) \left[\pi_t^{sym} + \frac{\lambda_x}{\lambda_p \kappa_p + \lambda_w \kappa_w} (x_t - x_{t-1}) \right]. \quad (1.41)$$

In the special case that $\kappa_w = \kappa_p = \kappa > 0$, which empirical studies such as that of Amato and Laubach (2001b) find to be not far from the truth,²⁴ the optimal target criterion (1.40) reduces simply to $q_t = 0$, or

$$\pi_t^{sym} + \phi(x_t - x_{t-1}) = 0, \quad (1.42)$$

with $\phi = \lambda_x / \kappa$ as in section 1.1.²⁵ More generally, the optimal target criterion is more complex, and slightly more forward-looking (as a result of the inertia in the real-wage dynamics when both wages and prices are sticky²⁶). But it still takes the form of an output-adjusted inflation target, involving the projected paths of both price and wage inflation; and since all terms except the first one in (1.40) are equal to zero under a commitment to ensure that $q_t = 0$ at all times, the target criterion (1.42) continues to provide a fairly good approximation to optimal policy even when κ_w is not exactly equal to κ_p .

This is of the same form as the optimal target criterion (1.8) for the case in which only prices are sticky, with the exception that the index of goods price inflation π_t is now replaced by an index π_t^{sym} that takes account of both price and wage inflation. Of course, the weight

²⁴See the discussion in Woodford (2003), chapter 3. In this case, the structural equations (1.29) – (1.30) imply that the real wage will be unaffected by monetary policy, instead evolving as a function of the real disturbances alone. Empirical studies often find that the estimated response of the real wage to an identified monetary policy shock is quite weak, and not significantly different from zero. Indeed, it is not significantly different from zero in our own analysis in section 2, though the point estimates for the impulse response function suggest that wages are not as sticky as prices.

²⁵Here we assume a normalization of the loss function weights in (1.33) in which $\lambda_p + \lambda_w = 1$, corresponding to the normalization in (1.3).

²⁶This only affects the optimal target criterion, of course, to the extent that the evolution of the real wage is endogenous, which requires that $\kappa_w \neq \kappa_p$.

that should be placed on wages in the inflation target depends on the relative weight on wage stabilization in the loss function (1.33). If one assumes a “traditional” stabilization objective of the form (1.3), so that $\lambda_w = 0$, then (1.42) is again identical to (1.8). However, one can show that expected utility maximization corresponds to minimization of a discounted loss criterion in which the relative weight on wage-inflation stabilization depends on the relative stickiness of wages and prices, as discussed by Erceg *et al.* (2000).²⁷

1.5 Habit Persistence

In the simple models thus far, the intertemporal IS relation (1.15) implies that aggregate demand is determined as a purely forward-looking function of the expected path of real interest rates and exogenous disturbances. Many empirical models of the monetary transmission mechanism instead imply that the current level of aggregate real expenditure should depend positively on the recent past level of expenditure, so that aggregate demand should change only gradually even in the case of an abrupt change in the path of interest rates. A simple way of introducing this is to assume that private expenditure exhibits “habit persistence” of the sort assumed in the case of consumption expenditure by authors such as Fuhrer (2000), Edge (2000), Christiano *et al.* (2001), Smets and Wouters (2002), and Altig *et al.* (2002).

Here, as in the models above, we model all interest-sensitive private expenditure as if it were non-durable consumption; that is, we abstract from the effects of variations in private expenditure on the evolution of productive capacity.²⁸ Hence we assume habit persistence in the level of aggregate private expenditure, and not solely in consumption, as in the models of Amato and Laubach (2001a) and Boivin and Giannoni (2003). This might seem odd, given that we do not really interpret the “ C_t ” in our model as referring mainly to consumption expenditure. But quantitative models that treat consumption and investment spending separately often find that the dynamics of investment spending are also best captured by

²⁷See also Woodford (2002xx), chap. 6, which modifies the derivation of Erceg *et al.* to take account of the discounting of utility.

²⁸See McCallum and Nelson (1999) and Woodford (2002, chap. 4) for further discussion of this simplification.

specifications of adjustment costs that imply inertia in the rate of investment spending (e.g., Edge, 2000; Christiano *et al.*, 2001; Altig *et al.*, 2002; Basu and Kimball, 2002). The “habit persistence” assumed here should be understood as a proxy for adjustment costs in investment expenditure of that sort, and not solely (or even primarily) as a description of household preferences with regard to personal consumption.²⁹

Following Boivin and Giannoni (2003), let us suppose that the utility flow of any household h in period t depends not only on its real expenditure C_t^h in that period, but also on that household’s level of expenditure in the previous period.³⁰ Specifically, we assume that the utility flow from expenditure is given by a function of the form

$$u\left(C_t^h - \eta C_{t-1}^h; \xi_t\right),$$

where ξ_t is a vector of exogenous taste shocks, $u(\cdot; \xi)$ is an increasing, concave function for each value of the exogenous disturbances, and $0 \leq \eta \leq 1$ measures the degree of habit persistence. (Our previous model corresponds to the limiting case $\eta = 0$ of this one.) The household’s budget constraint remains as before.

In this extension of our model, the marginal utility for the representative household of additional real income in period t is no longer equal to the marginal utility of consumption in that period, but rather to

$$\lambda_t = u_c(C_t - \eta C_{t-1}; \xi_t) - \beta \eta E_t[u_c(C_{t+1} - \eta C_t; \xi_{t+1})]. \quad (1.43)$$

The marginal utility of income in different periods continues to be linked to the expected return on financial assets in the usual way, so that equilibrium requires that

$$\lambda_t = \beta E_t[\lambda_{t+1}(1 + i_t)P_t/P_{t+1}]. \quad (1.44)$$

Using (1.43) to substitute for the λ ’s in (1.44), we obtain a generalization of the usual Euler equation for the intertemporal allocation of aggregate expenditure given expected rates of return.

²⁹For further discussion, see Woodford (2003, chapter 5, sec. 1.2).

³⁰Note that the consumption “habit” is assumed here to depend on the household’s own past level of expenditure, and not on that of other households.

Log-linearization of this Euler equation yields a generalization of our previous IS relation (1.15), of the form

$$\tilde{x}_t = E_t \tilde{x}_{t+1} - \varphi^{-1} [i_t - E_t \pi_{t+1} - r_t^n], \quad (1.45)$$

where

$$\begin{aligned} \tilde{x}_t &\equiv (x_t - \eta x_{t-1}) - \beta \eta E_t (x_{t+1} - \eta x_t), \\ \varphi^{-1} &\equiv (1 - \beta \eta) \sigma > 0, \end{aligned}$$

and $\sigma \equiv -u_c / (\bar{Y} u_{cc})$ as before. Here x_t is again the log gap between actual output and the flexible-price equilibrium level of output in the absence of markup fluctuations, and r_t^n is again the flexible-price equilibrium real interest rate in the absence of markup fluctuations, *i.e.*, the real interest rate associated with an equilibrium in which $x_t = 0$ at all times. Note that when $\eta = 0$, φ reduces to σ^{-1} , \tilde{x}_t reduces to x_t , and (1.45) reduces to (2.7). In the general case, the log marginal utility of real income is negatively related to \tilde{x}_t , rather than to x_t , which is why \tilde{x}_t appears in the generalized IS relation (1.45).

This modification of preferences changes the form of the aggregate-supply relation (1.1) as well. (For simplicity, we here consider only the case of a model with flexible wages and Calvo pricing.) In the derivation of (1.1), we have assumed that the log marginal utility of real income (which affects real supply costs owing to its effect on real wage demands) can be replaced by a linear function of x_t ; but just as in the case of the IS relation, this now must be written as a linear function of \tilde{x}_t instead. We then obtain an aggregate-supply relation of the form

$$\pi_t = \xi_p [\omega x_t + \varphi \tilde{x}_t] + \beta E_t \pi_{t+1} + u_t, \quad (1.46)$$

where $\xi_p > 0$ is the same coefficient as in (1.29), and $\omega \equiv \omega_p + \omega_w > 0$. The relation can equivalently be rewritten in the form

$$\pi_t = \kappa [(x_t - \delta x_{t-1}) - \beta \delta E_t (x_{t+1} - \delta x_t)] + \beta E_t \pi_{t+1} + u_t, \quad (1.47)$$

where $0 \leq \delta \leq \eta$ is the smaller root of the quadratic equation

$$\eta \varphi (1 + \beta \delta^2) = [\omega + \varphi (1 + \beta \eta^2)] \delta, \quad (1.48)$$

and³¹

$$\kappa \equiv \xi_p \eta \varphi / \delta > 0. \quad (1.49)$$

Again taking a second-order Taylor series expansion of the expected utility of the representative household,³² we again obtain a discounted criterion of the form (1.2), but now with a period loss function of the form

$$L_t = \pi_t^2 + \lambda(x_t - \delta x_{t-1} - \hat{x}^*)^2, \quad (1.50)$$

generalizing (1.3). Here λ is again defined as in (1.4), the parameters κ, δ are the same as in the aggregate-supply relation (1.47), and the size of $\hat{x}^* > 0$ depends once more on both the degree of market power and the size of tax distortions. As in the analysis of Amato and Laubach (2001a), habit persistence implies that the period loss function should depend on the lagged output gap as well as the present gap. However, we note that both the inflationary pressures indicated in (1.47) and the deadweight losses measured by (1.50) depend on the quasi-differenced output gap $x_t - \delta x_{t-1}$, where δ is the smaller root of (1.48). And while δ is an increasing function of η , it may be much smaller than it; if ω is large relative to φ , then δ may be quite small even in the presence of substantial habit persistence. This is the case that our estimates below suggest is empirically realistic: while the best empirical fit is obtained for the extreme value $\eta = 1$, the implied value of δ is only 0.14.

An optimal target criterion is easily derived, even in the presence of habit persistence, in the case that there are no transactions frictions, nor any other grounds for an interest-rate stabilization objective. In this case an optimal policy seeks to minimize the discounted sum of losses (1.50) subject to the sequence of constraints (1.47). The same Lagrangian method as above yields first-order conditions

$$\pi_t + \varphi_t - \varphi_{t-1} = 0, \quad (1.51)$$

$$\lambda(x_t - \delta x_{t-1} - \hat{x}^*) - \kappa \varphi_t + \delta \kappa \varphi_{t-1} = 0, \quad (1.52)$$

³¹In the limiting case in which $\eta = 0$, $\delta = 0$, while δ/η approaches the well-defined limit $\varphi/(\omega + \varphi)$, so that $\kappa = \xi_p(\omega + \varphi) = \xi_p(\omega + \sigma^{-1})$. Thus in this limit, (1.47) reduces to (1.1), where κ is defined as in (1.32).

³²For details of the calculation, see the derivation in the appendix for the full model, incorporating habit persistence, that is introduced in section 2.

generalizing (1.6) – (1.7). An optimal target criterion is again obtained by eliminating the Lagrange multiplier. In the case that $\delta < 1$, as is necessarily true (even in the extreme case where $\eta = 1$) given $\omega > 0$, (1.52) implies that a time-invariant way of identifying the Lagrange multiplier is

$$\varphi_t = (\lambda/\kappa)(x_t - x^*),$$

where $x^* \equiv \hat{x}^*/(1 - \delta)$. Substituting this into (1.51), we obtain

$$\pi_t + \frac{\lambda_x}{\kappa}(x_t - x_{t-1}) = 0. \quad (1.53)$$

Thus the optimal target criterion is *exactly the same* as in our baseline model, and is unaffected by the estimated value of η . The estimated degree of habit persistence does matter for the central bank’s judgment about which inflation/output paths are feasible, and also about the interest-rate path that will be necessary in order to achieve them. But it has no consequences for the target criterion that should be used to judge whether a given inflation/output projection is acceptable.

The degree of habit persistence does matter for the optimal target criterion in the case of an interest-rate stabilization objective. Suppose that the loss function (1.50) is generalized to the form

$$L_t = \pi_t^2 + \lambda_x(x_t - \delta x_{t-1} - \hat{x}^*)^2 + \lambda_i(i_t - i^*)^2, \quad (1.54)$$

where $\lambda_i > 0$ for any of the reasons discussed in section 1.3. In this case the relevant constraints on possible equilibrium paths of the target variables include both (1.45) and (1.47) each period. In the resulting system of first-order conditions, (1.16) and (1.18) are again exactly as in section 1.3, but (1.17) generalizes to

$$\lambda_x E_t[(1 - \beta\delta L^{-1})^{-1}(1 - \delta L)(x_t - x^*)] + E_t[B(L)\varphi_{1,t+1}] - \kappa E_t[(1 - \beta\delta L^{-1})^{-1}(1 - \delta L)\varphi_{2t}] = 0, \quad (1.55)$$

where

$$B(L) \equiv (1 - \beta^{-1}L)(1 - \eta L)(L - \beta\eta).$$

Using two of these relations to eliminate the Lagrange multipliers from the other, we obtain a target criterion of the form

$$(1 - \delta L)[\phi_\pi \pi_t + \phi_x(x_t - x_{t-1})] = (1 - L)E_t[(1 - \beta\delta L^{-1})^{-1}B(L)i_{t+1}] - \frac{\kappa}{\beta\varphi}(1 - \delta L)(i_{t-1} - i^*), \quad (1.56)$$

generalizing (1.19), where the definitions of ϕ_π and ϕ_x are as in (1.20), but with φ replacing σ^{-1} in the previous expressions. Here we see that the presence of habit persistence introduces additional dynamics into the form of the optimal target criterion. Nonetheless, it is interesting to note that once again, the optimal target criterion involves only the rate of change of the output gap, rather than its absolute level, even when the utility-based stabilization objective instead indicates a concern to stabilize the value of $x_t - \delta x_{t-1}$.

2 A Small Quantitative Model of the U.S. Economy

We now turn to the question of the likely quantitative importance of the various considerations discussed in section 1 in the actual conduct of monetary policy. In order to do this, we first estimate the numerical parameters of a model that, while still very stylized, is intended to capture important features of the monetary transmission mechanism in the U.S. economy. We present an updated version of the analysis in Rotemberg and Woodford (1997), incorporating a number of additional complications — habit persistence, wage stickiness, and inflation inertia — that have been argued in the subsequent empirical literature to afford important improvements in the realism of this sort of optimizing model of the transmission mechanism, as discussed in section 1. The model that we use is similar the one estimated by Boivin and Giannoni (2003), extended to allow for sticky wages.

Our approach to estimation of the model parameters follows the lines proposed in Rotemberg and Woodford (1997) and also used in Boivin and Giannoni (2003). First, we estimate an unconstrained vector autoregression model of a small number of U.S. aggregate time series. This VAR is used (along with weak identifying assumptions) both to identify the coefficients of the Fed's reaction function in the historical period, and to estimate the im-

pulse responses of our variables to an identified monetary policy shock under that historical policy. In a second step, we develop a simple optimizing model that can replicate the effects of identified monetary policy shocks, as implied by the VAR. We estimate the structural parameters of the model by minimizing the weighted distance between the estimated VAR impulse responses to a monetary policy shock and the model’s predicted responses to the same shock. We are then able to recover the historical sequence of structural disturbances and to estimate a law of motion for them, which we use for certain exercises in section 3. However, for purposes of the sort of characterization of optimal policy offered here (as opposed to those proposed by Rotemberg and Woodford, 1997, 1999), our conclusions about the character of the historical disturbance processes are much less important than our conclusions about the coefficients of the structural relations that relate the endogenous variables to one another.

In a third step, discussed in section 3, we derive a welfare-theoretic loss function for the evaluation of alternative monetary policy rules, by computing a second-order approximation to the expected utility of the representative household in our model. We then proceed along the lines of Giannoni and Woodford (2002a, 2002b) to derive a robustly optimal inflation-targeting rule for monetary policy.

2.1 The Effects of Monetary Disturbances

Here we briefly present the VAR that we use to estimate the actual monetary policy rule as well as the effects of monetary policy disturbances. We assume that the recent U.S. monetary policy can be described by the following feedback rule for the Federal funds rate

$$i_t = \bar{i} + \sum_{k=1}^{n_i} \phi_{ik} (i_{t-k} - \bar{i}) + \sum_{k=0}^{n_w} \phi_{wk} \hat{w}_{t-k} + \sum_{k=0}^{n_\pi} \phi_{\pi k} (\pi_{t-k} - \bar{\pi}) + \sum_{k=0}^{n_y} \phi_{yk} \hat{Y}_{t-k} + \varepsilon_t \quad (2.1)$$

where i_t is the Federal funds rate in period t , π_t denotes the rate of inflation between periods $t-1$ and t , \hat{w}_t is the deviation of the log real wage from trend at date t , \hat{Y}_t is the deviation of log real GDP from trend, and $\bar{i}, \bar{\pi}$ are long-run average values of the respective variables.³³ The disturbances ε_t represent monetary policy “shocks” and are assumed to be serially

uncorrelated. Estimated policy rules often omit real wages, but we include them in (2.1) for generality; the VAR that we use below to estimate impulse responses is then completely unrestricted (except as to number of lags).

To identify the monetary policy shocks and estimate the coefficients in (2.1), we assume as in the studies of Bernanke and Blinder (1992), Rotemberg and Woodford (1997), Bernanke and Mihov (1998), and Christiano *et al.* (2001), among others, that a monetary policy shock at date t has no effect on inflation, output or the real wage in that period. It follows that (2.1) can be estimated by OLS, and that the residuals of the estimated equation will represent a historical sequence of monetary policy shocks.

We model the dynamics of the vector $Z_t = [i_t, \hat{w}_{t+1}, \pi_{t+1}, \hat{Y}_{t+1}]'$ by a structural VAR of with three lags. This can then be written in companion form as

$$T\bar{Z}_t = a + A\bar{Z}_{t-1} + \bar{e}_t \quad (2.2)$$

where $\bar{Z}_t \equiv [Z'_t, Z'_{t-1}, Z'_{t-2}]'$ and T is a lower triangular matrix with ones on the diagonal and nonzero off-diagonal elements only in the first four rows, the first four rows of the vector a contain constants, and A contains estimated coefficients from the VAR in the first four rows, and an identity matrix in the lower rows. The first row of the estimated system (2.2) corresponds to the estimated monetary policy rule (2.1).

To estimate the VAR, we consider quarterly U.S. data on the sample period 1980:1 – 2002:2. As in Rotemberg and Woodford (1997), and Amato and Laubach (2001b), we begin the sample in the first quarter of 1980 because several empirical studies have identified a significant change in monetary policy around that period (see, *e.g.*, Clarida, Gali and Gertler, 2000; Boivin, 2001; Boivin and Giannoni, 2003; Cogley and Sargent, 2001, 2002).³⁴

³³Specifically, \hat{Y}_t is the log of real GDP minus a linear trend. Inflation is computed as the quarterly growth of the GDP deflator (chain-type), annualized. The interest rate i_t is the quarterly average of the Federal funds rate, annualized. The real wage is the log of wages and salaries in the compensation of employees published by the Bureau of Economic Analysis, divided by the GDP deflator; a linear trend is then subtracted from the log real wage to obtain \hat{w}_t .

³⁴Some studies suggest that monetary policy has changed again around the mid-1980's. However, Boivin and Giannoni (2003), following the approach proposed by Bernanke, Boivin and Elias (2002), show that impulse response functions to monetary policy disturbances in a factor-augmented VAR are similar to the ones reported here, when estimated both on the 1980-2002 and 1984-2002 sample periods.

ϕ_{i1}	ϕ_{i2}	ϕ_{i3}	ϕ_{w0}	ϕ_{w1}	ϕ_{w2}
0.572	-0.085	0.192	0.365	-0.008	-0.406
(0.104)	(0.127)	(0.090)	(0.202)	(0.302)	(0.191)
$\phi_{\pi0}$	$\phi_{\pi1}$	$\phi_{\pi2}$	ϕ_{y0}	ϕ_{y1}	ϕ_{y2}
0.071	0.146	0.472	0.333	-0.038	-0.118
(0.098)	(0.115)	(0.115)	(0.176)	(0.241)	(0.169)
R^2	DW				
0.956	2.033				

Standard errors are in parenthesis

Table 2: Estimated Monetary Policy Rule (1980:1 - 2002:2)

Table 2 reports the coefficients of the estimated policy rule. While these coefficients are difficult to interpret as such, we note that the estimated rule implies that the interest rate would eventually increase by 2.14 percentage points in the long run, in response to a one percentage point permanent increase in inflation, and that it would increase by 0.55 percentage point in response to a one percent permanent increase in output. These are similar long-run response coefficients to those obtained by authors such as Taylor (1993, 1999), Judd and Rudebusch (1998), and Clarida *et al.* (2000). The estimated real-wage response coefficients at different lags are close to cancelling; the estimated reaction function is quite similar to one in which the central bank responds only to the rate of real-wage growth, rather than to the level of real wages. The response to real wage growth is strongly positive, indicating that increases in wages lead to a stronger and more immediate increase in nominal interest rates than do increases in prices of the same magnitude. While wages are not often included as an explanatory variable in estimated Fed reaction functions, our results here suggest that wage growth is also an important explanatory variable.

Figure 6 shows the estimated impulse response functions of output, the real wage, inflation, and the interest rate. Here the dashed lines indicate 90% confidence intervals, obtained using Kilian's (1998) bootstrap procedure. Because of our identifying assumption, output, inflation, and the real wage remain unchanged in the period of the shock. In the quarter

following the shock, output still barely moves, while inflation and the real wage start declining. Output falls substantially in the second quarter after the shock and then returns progressively back to its initial level. In contrast, inflation and the real wage both reach their lowest levels only five quarters after the shock.

2.2 A Quantitative Model of the Transmission Mechanism

We now describe a simple optimizing model that we use to explain the effects of monetary policy on output, inflation, the real wage, and interest rates. While the model is still very stylized, it contains several ingredients that allow it to replicate important features of the impulse response functions estimated using our VAR. We assume that there exists a continuum of households indexed by h and distributed uniformly on the $[0, 1]$ interval. Each household h seeks, at date t , to maximize a lifetime expected utility of the form

$$E_t \left\{ \sum_{T=t}^{\infty} \beta^{T-t} \left[u \left(C_T^h - \eta C_{T-1}^h; \xi_T \right) - v \left(H_T^h; \xi_T \right) \right] \right\} \quad (2.3)$$

where $\beta \in (0, 1)$ is the household's discount factor (assumed to be equal for each household), C_t^h is a Dixit-Stiglitz (1977) index of the household's consumption of each of the differentiated goods supplied at time t , P_t is the corresponding price index, and H_t^h is the amount of labor (of type h) that household h supplies at date t . Here we assume that each household specializes in the supply of one type of labor, and that each type of labor is supplied by an equal number of households. The parameter $0 \leq \eta \leq 1$ represents the degree of habit formation, as in section 1.5. The stationary vector ξ_t represents exogenous disturbances to preferences. For each value of ξ , the function $u(\cdot; \xi)$ is assumed to be increasing and concave, while $v(\cdot; \xi)$ is increasing and convex.

2.2.1 Optimal Consumption Decisions

While the optimal allocation consumption at date t is chosen at date t , and is determined by the usual Dixit-Stiglitz demand relations, we assume as in Rotemberg and Woodford (1997) that households must choose their index of consumption C_t^h at date $t - 2$. Equivalently, we

assume that C_t^h is determined at the beginning of period $t - 1$, i.e., before the monetary policy shock in $t - 1$ is known. We assume that financial markets are complete so that risks are efficiently shared. As a result, each household faces a single intertemporal budget constraint.

The first-order conditions for optimal timing of consumption by the representative household require that

$$E_{t-2} \{u_c(C_t - \eta C_{t-1}; \xi_t) - \beta \eta u_c(C_{t+1} - \eta C_t; \xi_{t+1})\} = E_{t-2} \{\lambda_t\} \quad (2.4)$$

for each date $t \geq 2$ and each possible state at date $t - 2$, generalizing (1.43), where again λ_t represents the representative household's marginal utility of real income at date t .³⁵ The marginal utilities of income at different dates and in different states must furthermore satisfy

$$\lambda_t Q_{t,T} / P_t = \beta^{T-t} \lambda_T / P_T \quad (2.5)$$

for any possible state at any date $T \geq t$, where $Q_{t,T}$ is the stochastic discount factor that defines the market valuations of alternative random income streams. Noting that the riskless one-period nominal interest rate i_t must satisfy $(1 + i_t)^{-1} = E_t Q_{t,t+1}$, we obtain once again (1.44) as an equilibrium relation linking interest rates to the evolution of the marginal utility of income. We assume furthermore that the government purchases a Dixit-Stiglitz aggregate G_t of the form of all goods in the economy, so that aggregate demand Y_t satisfies $Y_t = C_t + G_t$.

We make use of log-linear approximations of these relationships about a steady state equilibrium in which there is no inflation. Log-linearization of (1.44) yields

$$\hat{\lambda}_t = E_t[\hat{\lambda}_{t+1} + \hat{i}_t - \pi_{t+1}]. \quad (2.6)$$

where $\hat{\lambda}_t \equiv \log\left(\frac{\lambda_t}{\bar{\lambda}}\right)$, $\hat{i}_t \equiv \log\left(\frac{1+i_t}{1+\bar{i}}\right)$, and $\pi_t \equiv \log(P_t/P_{t-1})$. Using this, and log-linearizing (2.4) we obtain an equation of the form

$$\tilde{Y}_t = \check{g}_t + E_{t-2}(\tilde{Y}_{t+1} - \check{g}_{t+1}) - \varphi^{-1} E_{t-2}(\hat{i}_t - \pi_{t+1}) - \beta \eta (E_t \hat{Y}_{t+1} - E_{t-2} \hat{Y}_{t+1}) \quad (2.7)$$

³⁵Because the problem is the same for each household h (the initial level of wealth is assumed to differ for any two households in a way that compensates for any difference in their expected labor incomes, and complete financial markets allow complete pooling of idiosyncratic labor income risk thereafter), all households choose identical state-contingent plans for consumption.

where φ is defined as in (1.45), \check{g}_t represents exogenous demand shocks including preference shocks and fluctuations in government expenditure, and $\tilde{Y}_t \equiv (\hat{Y}_t - \eta\hat{Y}_{t-1}) - \beta\eta(E_t\hat{Y}_{t+1} - \eta\hat{Y}_t)$, $\hat{Y}_t \equiv \log(Y_t/\bar{Y})$. Equation (2.7) generalizes the intertemporal IS relation (1.45).

For our welfare analysis, it is convenient to rewrite this relation in terms of the output gap

$$x_t \equiv \hat{Y}_t - \hat{Y}_t^n$$

where \hat{Y}_t^n indicates log deviations in the natural rate of output, by which we mean the equilibrium level of output under flexible prices, flexible wages, constant levels of distorting taxes and of desired markups in the labor and product markets, and with wages, prices and spending decisions predetermined by only one period.³⁶

Expressing (2.7) in terms of the output gap, we obtain

$$E_{t-2}\tilde{x}_t = E_{t-2}\tilde{x}_{t+1} - \varphi^{-1}E_{t-2}(\hat{i}_t - \pi_{t+1} - \hat{r}_t^n) \quad (2.8)$$

where $\tilde{x}_t \equiv (x_t - \eta x_{t-1}) - \beta\eta(E_t x_{t+1} - \eta x_t)$ and \hat{r}_t^n is an exogenous variable that represents the deviation from steady state of the natural rate of interest, *i.e.*, the equilibrium real rate of interest in the ideal situation defined above. The actual output gap relates furthermore to the expected output gap through

$$\tilde{x}_t = E_{t-2}\tilde{x}_t + (\check{g}_t - \tilde{Y}_t^n) - E_{t-2}(\check{g}_t - \tilde{Y}_t^n) - \beta\eta[E_t(x_{t+1} + \hat{Y}_{t+1}^n) - E_{t-2}(x_{t+1} + \hat{Y}_{t+1}^n)]. \quad (2.9)$$

³⁶Up to the log-linear approximation used in our estimation of the model, \hat{Y}_t^n defined in this way is just the conditional expectation at date $t-1$ of the log deviation of the equilibrium level of output when none of these variables are predetermined at all. Because wages and prices are both predetermined a period in advance, it is only the component of the output gap that is forecastable a period in advance that matters in any event for these equations. It is similarly only the variation in the forecastable component of the output gap that need be considered when evaluating welfare under alternative policies, since the unforecastable component of the output gap (defined relative to a concept of the “natural rate” that is not predetermined) would in any event be both exogenous and uncorrelated with the forecastable component. It then simplifies notation to *define* the output gap as the gap between actual output and the *forecastable component* of the natural rate. In this way, x_t becomes a predetermined state variable.

2.2.2 Optimal Wage and Price Setting

As in Erceg et al. (2000), Amato and Laubach (2001b), and Woodford (2003, chap. 3), we assume that there is a single economy-wide labor market. The producers of all goods hire the same kinds of labor and face the same wages. Firm z is a monopolistic supplier of good z , which it produces according to the production function

$$y_t(z) = A_t F(\bar{K}, H_t(z)) \equiv A_t f(H_t(z))$$

where $f' > 0$, $f'' < 0$, the variable $A_t > 0$ is an exogenous technology factor, and capital is assumed to be fixed so that labor is the only variable input. The labor used to produce each good z is a CES aggregate

$$H_t(z) \equiv \left[\int_0^1 H_t^h(z)^{\frac{\theta_w-1}{\theta_w}} dh \right]^{\frac{\theta_w}{\theta_w-1}} \quad (2.10)$$

for some elasticity of substitution $\theta_w > 1$, where $H_t^h(z)$ is the labor of type h that is hired to produce a given good z . The demand for labor of type h by firm z is again of the Dixit-Stiglitz form $H_t^h(z) = H_t(z) \left(\frac{w_t(h)}{W_t} \right)^{-\theta_w}$, where $w_t(h)$ is the nominal wage of labor of type h , and W_t is a wage index.

We assume that the wage for each type of labor is set by the supplier of that type, who is in a situation of monopolistic competition and who is ready to supply as many hours of work as may be demanded at that wage. We assume that each wage is reoptimized with a fixed probability $1 - \alpha_w$ each period. However, as in Woodford (2002, ch. 3), if a wage is not reoptimized, it is adjusted according to the indexation rule

$$\log w_t(h) = \log w_{t-1}(h) + \gamma_w \pi_{t-1}$$

for some $0 \leq \gamma_w \leq 1$. A worker of type h who chooses a new wage $w_t(h)$ at date t , expects to have a wage $w_t(h) (P_{T-1}/P_{t-1})^{\gamma_w}$ with probability α_w^{T-t} at any date $T \geq t$. We assume furthermore that the newly chosen wage that comes into effect in period t , w_t^* , is chosen at the end of period $t - 1$, i.e., on the basis of information available at date $t - 1$.

As shown in Woodford (2002, ch. 3), this setup yields as a first-order approximation, a wage inflation equation of the form

$$(\pi_t^w - \gamma_w \pi_{t-1}) = \xi_w E_{t-1} (\omega_w x_t + \varphi \tilde{x}_t) - \xi_w E_{t-1} \mu_t + \xi_w E_{t-1} (w_t^n - w_t) + \beta E_{t-1} (\pi_{t+1}^w - \gamma_w \pi_t), \quad (2.11)$$

generalizing (1.30) to allow for indexation to the lagged price index, habit persistence, and predetermined wage-setting and spending decisions. Here π_t^w denotes nominal wage inflation, w_t is the log real wage, and w_t^n is an exogenous variable representing the log of the “natural real wage”, *i.e.*, the equilibrium real wage when both wages and prices are fully flexible and consumption is not predetermined. The parameter

$$\xi_w \equiv \frac{(1 - \alpha_w)(1 - \alpha_w \beta)}{\alpha_w(1 + \nu \theta_w)} > 0 \quad (2.12)$$

is a function of the degree of wage stickiness, the elasticity of marginal disutility of labor supply at the steady-state, $\nu \equiv \frac{v_{hh} \bar{H}}{v_h}$, and the elasticity of substitution for different types of labor. The parameter $\omega_w \equiv \nu \phi > 0$ indicates the degree to which higher economic activity increases workers’ desired wages for given prices. (Once again, $\phi \equiv f / (\bar{H} f') > 0$ is the elasticity of the required labor input with respect to output variations.)

Integrating (2.11) forward, we note that nominal wages at date t tend to increase (above lag inflation) when expected future positive output gaps are positive and when real wages are expected to be below their natural rate. The variable $\mu_t \equiv \hat{\lambda}_t - \varphi E_t (\tilde{g}_t - \tilde{Y}_t)$, which corresponds to the discrepancy between the (log) marginal utility of real income and the (log) marginal utility of consumption satisfies

$$E_{t-1} \mu_t = E_{t-1} (\hat{i}_t - \pi_{t+1}) + \varphi E_{t-1} [(\tilde{g}_{t+1} - \tilde{g}_t) - (\tilde{Y}_{t+1} - \tilde{Y}_t)]. \quad (2.13)$$

The presence of $E_{t-1} \mu_t$ in (2.11) indicates a moderating effect on nominal wage inflation of an expectation at date $t - 1$ of real rates of return between t and $t + 1$ that are higher than those that were anticipated at $t - 2$, *i.e.*, at the time that consumption decisions were made for period t . In fact, unexpectedly high real rates of return increase the value of income in period t and thus lower average wage demands.

Similarly, we assume that the suppliers of goods are in monopolistic competition and that each price is reoptimized with a fixed probability $1 - \alpha_p$ each period. However, as in Woodford (2002, ch. 3), if a price is not reoptimized, it is again adjusted according to the indexation rule

$$\log p_t(z) = \log p_{t-1}(z) + \gamma_p \pi_{t-1}$$

for some $0 \leq \gamma_p \leq 1$. Again following the development in Woodford (2002, ch. 3), we can show that optimal pricing decisions result in an aggregate supply relation of the form

$$\pi_t - \gamma_p \pi_{t-1} = \xi_p \omega_p E_{t-1} x_t + \xi_p E_{t-1} (w_t - w_t^n) + \beta E_{t-1} (\pi_{t+1} - \gamma_p \pi_t), \quad (2.14)$$

generalizing (1.29) to allow for indexation to the lagged price index and predetermination of pricing decisions. Here

$$\xi_p \equiv \frac{(1 - \alpha_p)(1 - \alpha_p \beta)}{\alpha_p(1 + \omega_p \theta_p)} > 0 \quad (2.15)$$

is a function of the degree of price stickiness, the elasticity of substitution for different goods $\theta_p > 1$, and $\omega_p > 0$ which measures the degree to which higher economic activity increases producers' prices for given wages. Integrating (2.14) forward, we observe that inflation tends to increase (relative to past inflation) when agents expect positive future output gaps and/or expect that real wages will be above their natural rate.

Finally, the evolution of the real wage is linked to wage inflation and price inflation through the identity (1.31). Our structural model can then be summarized by a demand block (2.8) – (2.9) and a supply block consisting of (2.11) – (2.14) together with (1.31). We finally close the model with an equation such as (2.1) that characterizes the behavior of the central bank. These equations then allow us to determine the equilibrium evolution of the variables of interest: π_t , π_t^w , x_t , \hat{i}_t , and w_t .

2.3 Estimated Parameter Values

We turn now to the estimation of the parameters of the structural model just set out. As mentioned above, we are looking for structural parameters that allow the model to describe as well as possible the transmission mechanism of monetary policy. Following Rotemberg and

Woodford (1997), we choose the structural parameters that minimize the distance between the estimated VAR impulse response functions to a monetary policy shock and the model's predicted response to the same shock. As discussed in Amato and Laubach (2001b), Boivin and Giannoni (2003) and in Christiano *et al.* (2001), this is quite generally an estimation procedure that allows for statistical inference on the model's estimated structural parameters. Note also that the model that we consider is constructed so as to be consistent with the identifying assumptions made for the estimation of the VAR impulse response functions. In particular, both the model and the VAR have the feature that output, inflation and the real wage respond to unexpected changes in the interest rate with a lag of at least one quarter. In addition, to the extent that we estimate the structural parameters on the basis of impulse responses to monetary shocks, our estimation method has the advantage of providing parameter estimates that are robust to potential misspecifications of the remaining shock processes in the model. This is because in order to compute the impulse responses, we don't need to specify the stochastic process of the shocks such as $\tilde{g}_t, \hat{Y}_t^n, \hat{\omega}_t^n, \hat{r}_t^n$.

As in the studies mentioned above, we set $\beta = 0.99$ so that β^{-1} corresponds approximately to the steady-state real gross rate of interest which is about 1.01. In addition, we calibrate the elasticity $\omega_p \equiv -f''\bar{Y}/(f')^2$ to 0.33 as in Rotemberg and Woodford (1997). This would be implied by a Cobb-Douglas production function in which the elasticity of output with respect to hours is 0.75. Such a production function would yield a share of wages in the value of output of $0.75/\mu_p$ where $\mu_p \equiv \theta_p/(\theta_p - 1)$ is the average gross markup of prices over marginal cost due to market power in the goods markets. (This means a labor share of 0.74, given the markup estimate reported below.)

We estimate the vector of the remaining seven structural parameters $v \equiv [\varphi, \eta, \xi_p, \xi_w, \omega_w, \gamma_p, \gamma_w]'$ by minimizing the distance

$$D(v) = [\hat{f}_V - f_M(\hat{\phi}, v)]' V [\hat{f}_V - f_M(\hat{\phi}, v)]$$

where \hat{f}_V is a vector that contains the VAR-based impulse response functions of output, inflation, the real wage, and the interest rate to an unexpected monetary policy shock, and




Figure 6: Estimated and predicted impulse responses to a monetary policy shock.

$f_M(\hat{\phi}, v)$ is vector containing the corresponding impulse response functions generated by the model, for a given vector of structural parameters v and the vector of policy rule coefficients $\hat{\phi}$ estimated in section 2.1. In fact to the extent that we estimated consistently the policy rule of the form (2.1) when estimating the VAR, we do not need to estimate again its coefficients at this stage. The positive definite weighting matrix V that we use in our estimation is a diagonal matrix, with the inverse of the variance of the estimate of each impulse response as the corresponding diagonal element. This allows us to weight the various impulse responses according to the degrees of precision with which each is estimated.³⁷ We estimate the structural parameters by matching model-based and VAR-based impulse responses of output,

³⁷The use of the inverse of the complete variance-covariance matrix of impulse responses as a weighting matrix would be more attractive, as this would yield efficient estimates. But such a weighting matrix appears to hinder the stability of the minimization algorithm. The matrix that we propose has the advantage of reducing the weight on responses about which we are less sure, in addition to making our results independent of the units in which we happen to measure the various series.

	Baseline	No habit $\eta = 0$	No indexation $\gamma_p = \gamma_w = 0$	Flexible wages $\xi_w^{-1} = 0$
Estimated parameters				
$\psi \equiv \frac{\varphi^{-1}}{1+\beta\eta^2}$	0.6715 (0.3330)	4.3144 (1.0253)	1.5026 (0.4221)	0.7564 (0.2823)
$\tilde{\eta} \equiv \frac{\eta}{1+\beta\eta^2}$	0.5025 (0.0692)*	0 (—)	0.5025 (0.1121)*	0.5025 (0.0515)*
ξ_p	0.0020 (0.0009)	0.0015 (0.0005)	0.0072 (0.0039)	0.0015 (0.0012)
ξ_w	0.0042 (0.1343)	0.0042 (0.0612)	0.0046 (0.0310)	$+\infty$ (—)
ω_w	19.551 (595.1)	19.991 (269.5)	19.072 (122.6)	0.5642 (0.1253)
γ_p	1 (0.3800)*	1 (0.3484)*	0 (—)	1 (0.5374)*
γ_w	1 (10.908)*	1 (12.4613)*	0 (—)	0 (—)
Implied parameters				
φ	0.7483	0.2318	0.3344	0.6643
η	1	0	1	1
$\kappa_p \equiv \xi_p \omega_p$	0.0007	0.0005	0.0024	0.0004
$\omega \equiv \omega_p + \omega_w$	19.884	20.325	19.405	0.8975
$\nu \equiv \omega_w / \phi$	14.663	14.994	14.304	0.4231
$\mu_p \equiv \frac{\theta_p}{\theta_p - 1}$	1.0039	1.0027	1.0143	1.0029
$\mu_w \equiv \frac{\theta_w}{\theta_w - 1}$	1.5361	1.5731	1.6113	—
Objective function value	13.110	15.886	16.580	18.837
Wald test (p -value)	—	0.000	0.000	0.000

Table 3: Estimated structural parameters for the baseline case and restricted models.

inflation, the real wage, and the interest rate on quarters 0 to 12 following a unexpected monetary policy shock. For consistency with the model, we constrain all parameters to be positive and impose an upper bound at 1 on η , γ_p and γ_w .

The estimated parameter values are shown in Table 3. Standard errors are in parentheses; an asterisk next to the reported standard error indicates that the standard error may not be reliable as the estimated parameter lies on the boundary of the allowed parameter space. Here we report estimates (with standard errors) for parameters $\psi \equiv \frac{\varphi^{-1}}{1+\beta\eta^2}$ and $\tilde{\eta} \equiv \frac{\eta}{1+\beta\eta^2}$ rather than for φ and η , as the former nonlinear transformations of these parameters can

be estimated with greater precision.³⁸ The values of φ and η implied by these estimates are shown in the second part of the table, along with the implied values for other model parameters, making use of the calibrated parameter values reported in Table 4.

While some of the model parameters cannot be estimated at all precisely, as indicated by the large standard errors, our estimation results are consistent with our theory insofar as we estimate positive values for the response coefficients φ , ξ_p , ξ_w , and ω_w in our structural equations. The values of ψ , measuring the interest-sensitivity of aggregate expenditure,³⁹ and ξ_p , measuring the response of inflation to the real-wage gap, are both significantly positive, though the estimates of ξ_w and ω_w are instead quite imprecise. We also find small enough standard errors on the standard errors of $\tilde{\eta}$, measuring the degree of habit persistence, and γ_p , measuring the degree of indexation of prices, to allow some inference about the magnitudes of those parameters (for example, both are significantly positive), while the value of γ_w is very imprecisely estimated. In general, the parameters of our wage equation are poorly estimated, while both our IS relation and our inflation equation are much better estimated.⁴⁰

The second through fourth columns of Table 3 report the corresponding estimates, using the same method, of various restricted versions of our model. In column 2, we assume zero habit persistence, as in the models of Rotemberg and Woodford (1997) and Amato and Laubach (2001b); in column 3, no inflation inertia (*i.e.*, no indexation of either wages or prices to the lagged price index), also like the two models just mentioned; and in column 4, flexible wages, as in the models of Rotemberg and Woodford (1997) and Boivin and Giannoni (2003).⁴¹ In each case, the objective function value is reported for the restricted model, *i.e.*,

³⁸Here ψ is estimated to be significantly positive, implying a significant effect of interest rates on aggregate demand, while the corresponding standard error for an estimate of φ would not allow us to judge that the latter coefficient was significantly positive. Similarly, $\tilde{\eta}$ is estimated to be significantly positive, implying habit persistence, even though the corresponding standard error for the estimated value of η is much greater than one.

³⁹The parameter ψ is called by Boivin and Giannoni (2003) the “pseudo-elasticity of substitution”; it measures the elasticity of expected output growth with respect to changes in the expected real rate of return, holding constant output growth in other periods.

⁴⁰A MATLAB program, to be made available on our webpages, will allow readers to check the extent to which our numerical characterization of optimal policy would be different in the case of alternative parameter values.

⁴¹The restricted model considered in column 4 corresponds to the model of Boivin and Giannoni, though

the weighted distance $D(\nu)$ defined above. The p -values reported on the last line refer to Wald tests of the null hypothesis that the restricted model is correct. In the last column, the parameter γ_w is set to zero as it is not identified in the case of flexible wages. We see that each of these restrictions assumed in earlier studies can be individually rejected, though the assumption of flexible wages is the one that would reduce the model's ability to fit the estimated impulse response functions to the greatest extent.⁴² Hence each of the complications introduced here are found to be justified; in this respect, our findings agree with those of Christiano *et al.* (2001), Altig *et al.* (2002), and Smets and Wouters (2002), though these authors all also introduce additional complications in order to explain a larger set of time series.

It is striking to note that the model fits the impulse responses best when the degree of inflation indexing (γ_p) and wage indexing to inflation (γ_w) reach their upper bound at 1. This corresponds to the assumption of full wage and price indexing made by Christiano *et al.* (2001). A value of $\gamma_p = 1$ is also roughly consistent with the weight on lagged inflation in the "hybrid" aggregate-supply relation estimated by Gali and Gertler (1999), and results in an aggregate supply relation quite similar to the one proposed by Fuhrer and Moore (1995).

The relatively small values of ξ_p , and ξ_w suggest that changes in the output gap and the real wage gap have a relatively small impact on price and wage inflation. However the estimated value of ω_w suggests that a one percent increase in economic activity increases workers' desired wages by nearly 20 percent, for given prices. The estimate of φ corresponds to an elasticity of intertemporal substitution (adjusted by the degree of habit formation) of $\varphi^{-1} = 1.3$. While authors such as Fuhrer (2000) and Christiano *et al.* (2001) among others have estimated substantial degrees of habit formation, our estimate lies at the upper bound

their method of estimation is different, in that they do not fit estimated impulse responses of the real wage along with those of the other three variables, and their model assumes a different form of monetary policy rule. They also calibrate the value of $\omega = \omega_w + \omega_p$, rather than only specifying a calibrated value for ω_p , and they assume a value of ω much smaller than our estimate. Nonetheless, the estimates for the other parameters reported in column 4 are similar to those obtained by Boivin and Giannoni, providing further evidence regarding the robustness of our conclusions here.

⁴²The implied impulse response functions are compared to the estimated ones in the case of each of the restricted models in the technical appendix to this paper.

β	ω_p	α_p	α_w	ϕ
0.99	1/3	2/3	2/3	4/3

Table 4: Additional calibrated parameter values.

of 1.

While the estimated parameter values for η , γ_p , and γ_w are significantly smaller when we estimate our model using impulse response functions over the six first quarters or less following the monetary shock, all parameter estimates are very similar to those reported in Table 2, when we use impulse response functions that extend longer than six quarters.⁴³ This suggests that in order to adequately capture the degree of persistence in the endogenous variables, we need to perform our estimation using long enough responses.

Assuming, as in Rotemberg and Woodford (1997) that $\alpha_p = 2/3$,⁴⁴ and similarly that $\alpha_w = 2/3$, together with the other parameter values already mentioned above, it is possible to infer the elasticities of substitution θ_p and θ_w from the estimated values of ξ_p and ξ_w respectively, using the definitions (2.12) and (2.15). The values of these elasticities implied by our estimates imply a gross markup of prices over marginal costs of only $\mu_p = \theta_p / (\theta_p - 1) = 1.004$ in the goods market, but a considerably higher gross markup of $\mu_w = \theta_w / (\theta_w - 1) = 1.54$ in the labor market. The fact that these implied markups are greater than one (*i.e.*, that the implied elasticities of substitution are greater than one) again indicates consistency of our estimates with our theoretical model.

Finally, our estimated value for ω_w can be used to derive an implied value of ν , the inverse of the Frisch elasticity of labor supply, using the definition $\omega_w = \nu\phi$ and a calibrated value for ϕ , the inverse of the elasticity of output with respect to the labor input. (The calibrated value of ϕ reported in Table 4 is implied by the same Cobb-Douglas production function as was used to calibrate the value of ω_p , discussed above.) The Frisch elasticity of labor supply implied by our estimates is thus only on the order of 0.07, less than one one-hundredth of the value implied by the estimates of Rotemberg and Woodford (1997),

⁴³Again, see the technical appendix for details.

⁴⁴Rotemberg and Woodford (1997) base this calibration on Blinder's (1994) survey evidence that prices are maintained constant for an average of 9 months, so that $1/(1 - \alpha_p)$ equals 3 quarters.

and much more consistent with many estimates in the empirical literature on labor supply. Because of the assumption of sticky wages, our model is able to account for non-negligible effects of a monetary disturbance on real activity without assuming that voluntary labor supply (under flexible wages) would be highly elastic. (Note that under the restriction of flexible wages, we would obtain estimates implying an elasticity of labor supply greater than 2.) While the values of these implied parameters do not matter for the ability of our model to fit the estimated impulse responses, they do matter for our welfare analysis below.

The solid lines in Figure 6 indicate the impulse response functions generated by our estimated model. Overall, it appears that the model is able to replicate quite well the impulse responses estimated by the VAR (circled lines), and the impulse responses remain consistently within the 90% confidence intervals. The model replicates in particular the estimated hump shaped output and real wage responses. While it does not capture the oscillations in the inflation response implied by the VAR, we note that this response is estimated quite imprecisely.

3 Optimal Policy for the Estimated Model

Now that we have an estimated structural model which allows us to account for at least certain aspects of the responses of output and of price and wage inflation to monetary disturbances, we turn to the characterization of optimal policy in the context of this model.

3.1 A Welfare-Theoretic Stabilization Objective

An advantage of having developed a structural model based on optimizing behavior is that it provides a natural objective for the monetary policy, namely maximization of the expected utility of the representative household. Following the method of Woodford (2003, chap. 6), we can express a second-order Taylor series approximation to this objective as a quadratic function of (wage and price) inflation, the output gap, and the nominal interest rate. The way in which various aspects of our model specification affect the appropriate welfare-theoretic stabilization objective in simple cases has already been discussed in section 1.

λ_p	λ_w	$16\lambda_x$	δ
0.9960	0.0040	0.0026	0.035

Table 5: Loss-function coefficients implied by our parameter estimates.

In the technical appendix to this paper, we show that for the model developed in section 2, the corresponding welfare-theoretic loss function, abstracting from any grounds for concern with interest-rate stabilization, is given by

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\lambda_p (\pi_t - \gamma_p \pi_{t-1})^2 + \lambda_w (\pi_t^w - \gamma_w \pi_{t-1}^w)^2 + \lambda_x (x_t - \delta x_{t-1} - \hat{x}^*)^2 \right]. \quad (3.1)$$

In this expression, the weights $\lambda_p, \lambda_w > 0$ are again defined as in (1.34); the weight $\lambda_x > 0$ is again defined as in (1.35), but using now the definition (1.49) for κ in the latter expression; the coefficient $0 \leq \delta \leq \eta$ is again the smaller root of (1.48); and $\hat{x}^* > 0$ is the same function of the microeconomic distortions affecting the efficiency of the steady-state output level as in (1.50).

This result combines features of several simpler cases discussed in section 1. Deadweight loss depends on squared deviations of both price and wage inflation (separately) from the rates that would minimize relative-price and relative-wage distortions, given that both wages and prices are sticky, as in (1.33). Due to the indexation of both prices and wages to a lagged price index, the loss-minimizing rates of wage and price inflation each period are determined by the lagged inflation rate and the indexation coefficients in each case, as in (1.10). And finally, the presence of habit persistence implies that deadweight loss depends not on squared deviations of the output gap from a constant value, but rather on squared deviations of $x_t - \delta x_{t-1}$ from a constant value, as in (1.50).

The numerical coefficients of the welfare-theoretic loss function implied by the estimated parameter values reported in Table 3 (for the baseline model) are reported in Table 5. Interestingly, our estimated model implies that it is optimal for the central bank to put a much larger weight on the stabilization of goods-price inflation than on the stabilization of wage inflation or of the output gap. Moreover, despite the fact that we estimate a very high degree of habit formation, which implies that household utility depends on the rate of change

of real expenditure rather than its level, the central bank's loss function does not involve the variability of the change in the output gap. Instead, it involves the variability of the level of the output gap relative to a small fraction of the lagged output gap.

These conclusions depend, of course, on our parameter estimates. It may seem surprising that the weight on wage inflation stabilization is so small, given that our estimates do not imply that wages are substantially more flexible than prices (for example, ξ_w is larger than ξ_p , but not by a large factor). The conclusion that λ_w is nonetheless very much smaller than λ_p reflects mainly the fact that our estimates imply a value for θ_p that is much larger than $\phi^{-1}\theta_w$. This in turn results from the fact that the estimated value of ω_w is much larger than the calibrated value of ω_p .⁴⁵ Because it is not plausible to assume a technology for which ω_p could be nearly as large as the estimated value of ω_w , we are led to assume a value of θ_p substantially larger than $\phi^{-1}\theta_w$. The result that λ_p greatly exceeds λ_w then follows, using (1.34).

The conclusion that λ_x is small follows, using (1.35), from the small value of κ_p and large value of θ_p implied by our parameter estimates. Since $\kappa_p \equiv \xi_p\omega_p$ and the value of θ_p is inferred from the value of ξ_p using (2.15), both of these conclusions depend crucially on the small estimated value for ξ_p . Essentially, the observed insensitivity of inflation to variations in output allows us to infer underlying microeconomic parameters that imply that variations in the output gap cause relatively modest distortions — this is the only way, in the context of our other assumptions, to explain the fact that inflation is not more strongly affected (*i.e.*, that the Phillips curve is not steeper).

Finally, the conclusion that δ is small (despite the fact that $\eta = 1$) follows, using (1.48), from the fact that the value of ω implied by our estimates is large relative to the estimated value of φ . Essentially, the observed sensitivity of wages to variations in real activity on the one hand (implying a large value for ω_w) and the sensitivity of aggregate expenditure to

⁴⁵If ξ_p and ξ_w were assigned equal values, then under our assumption of equal values for α_p and α_w , (2.12) and (2.15) would imply equal values for $\omega_p\theta_p$ and $\omega_w\phi^{-1}\theta_w$. (Here we recall that $\omega_w \equiv \nu\phi$.) The implied value of θ_p is then larger than $\phi^{-1}\omega_w$ by exactly the same factor as ω_w is larger than ω_p . In fact, our estimated value for ξ_p is smaller than our estimate for ξ_w , and this further increases the relative size of the implied value of θ_p .

interest-rate changes on the other (implying that φ cannot be too large) indicate preferences under which variations in the level of real activity will create greater distortions than variations in the rate of growth of real activity. Even when $\eta = 1$, the level of output matters to the representative household because of its consequences for the amount that the household must work; if the marginal disutility of output supply increases sharply with the level of real activity (as implied by a large value of ω), it will still be relatively more important to stabilize the level of real activity than its rate of change.⁴⁶

3.2 An Optimal Target Criterion

The method illustrated in section 1 for the derivation of optimal target criteria under alternative assumptions can be applied as well in the case of the empirical model described in section 2. Details of the relevant calculations are included in the technical appendix to this paper; here we simply present the quantitative implications of our estimated parameter values.

A first observation about optimal policy in our estimated model follows from the fact that wages, prices, and output are all predetermined for one quarter or longer in the model. It follows that in our structural equations, any variations in the short-term nominal interest rate i_t that are not forecastable a quarter earlier are irrelevant to the determination of wages, prices, or output. Hence this component of interest-rate policy cannot be relevant for welfare except through its consequences for the expected discounted value of the $\lambda_i(i_t - i^*)^2$ term that must be added to (3.1) if we take account of monetary frictions. But this last term is obviously minimized (in the case of any $\lambda_i > 0$) by a policy under which the nominal interest rate is completely forecastable a quarter in advance. Even in the case that $\lambda_i = 0$, there is no harm to any other stabilization objectives in eliminating unforecastable interest-rate variations; and so it seems plausible to assume at least some tiny concern with interest-rate

⁴⁶As discussed in section 1.3 above, it may also be desirable to reduce the variability of nominal interest rates; in this case, the loss function (3.1) should include an additional term, proportional to the squared deviation of the nominal interest rate from an optimal value. We do not take up this possible extension of the analysis here.

stabilization, so that it is optimal to suppress such variation in the interest rate.⁴⁷

Hence

$$i_t = E_{t-1} \dot{i}_t \tag{3.2}$$

is a requirement for optimal policy. This can be understood to say that all interest-rate changes should be signaled by the central well in advance of the date at which they take effect. The instrument that the central bank must adjust in period t in order to ensure that its period t target criterion will be projected to be satisfied is then not the period t interest rate i_t , but rather the bank's precommitted value $E_t i_{t+1}$ for the level of short-term nominal interest rates in the following period.⁴⁸ We turn now to the property that the bank's projections regarding period t endogenous variables should be made to satisfy through an appropriate commitment of this kind.

To simplify, we shall restrict attention to the case of a model in which $\gamma_p = \gamma_w = 1$, as assumed by Christiano *et al.* (2001), and as indicated by our estimates in section 2. In the appendix, we show that the first-order conditions for an optimal state-contingent evolution of the endogenous variables can be manipulated, after the fashion illustrated in section 1, to yield a characterization of optimal policy in terms of the projected paths of the target variables alone. However, in the present case, unlike the simpler ones discussed in section 1, the most convenient representation of these conditions is not in terms of a single target criterion, but two distinct ones. First of all, optimality requires that projections in any

⁴⁷For example, even if we assume that monetary frictions are of negligible quantitative significance, we may reasonably assume that the economy is a "cashless limiting economy" of the kind discussed in Woodford (1998), rather than a genuinely cashless economy. In this case, there should in fact exist tiny monetary frictions, that suffice to entail a preference for a completely forecastable nominal interest rate, in the absence of any offsetting benefit from variations in response to current shocks.

⁴⁸See further discussion in Svensson and Woodford (2003).

period t satisfy a condition of the form⁴⁹

$$F_t(\pi) + \phi_w[F_t(w) - w_t] = \bar{\pi}_t. \quad (3.3)$$

Here for each of the variables $z = \pi, w$, the expression $F_t(z)$ refers to a weighted average of forecasts of the variable z at various future horizons, conditional on information at date t ,

$$F_t(z) \equiv \sum_{k=1}^{\infty} \alpha_k^z E_t z_{t+k}, \quad (3.4)$$

where the weights α_k^z sum to one. Thus the coefficient ϕ_w is actually the sum of the weights on real-wage forecasts at different horizons k . We observe that the target criterion can be thought of as a wage-adjusted inflation target. In addition to the correction for the projected growth of real wages in the future, the acceptable rate of projected future inflation also varies due to time variation in the target $\bar{\pi}_t$. Optimality further requires that $\bar{\pi}_t$ be a function only of information available at date $t - 1$, and hence that

$$\bar{\pi}_t = E_{t-1}[F_t(\pi) + \phi_w(F_t(w) - w_t)]. \quad (3.5)$$

In general, this optimal target will not be constant over time.

In addition to the above requirement (which amounts to the condition that the left-hand-side of (3.3) be forecastable a quarter in advance), optimality also requires that projections at date t satisfy another condition as well, of the form

$$F_t^*(\pi) + \phi_w^*[F_t^*(w) - w_t] + \phi_x^*[F_t^*(x) - x_t] = \pi_t^*, \quad (3.6)$$

where the expressions $F_t^*(z)$ are again weighted averages of forecasts at different horizons (but with relative weights α_k^{z*} that may be different in this case), and π_t^* is another time-varying target value, once again a predetermined variable. In this case the criterion specifies a target for a wage- and output-adjusted inflation projection.⁵⁰

⁴⁹The target criterion could equivalently be expressed in the form $\phi_p F_t(\pi) + \phi_w F_t(\pi^w) = \bar{\pi}_t$, in which case the target criterion would refer solely to projected inflation of different sorts (both price and wage inflation). This would be a representation analogous to the one given in section 1.4 above, and would make clear that only the projected future paths of target variables (variables that enter the loss function) matter. We feel, however, that the representation proposed here allows a more convenient numerical summary of the content of the target criterion, by collecting the central bank's projections regarding the future level of nominal quantities in a single variable, the projected future price level.

⁵⁰As with (3.3), we could equivalently express this criterion in terms of a linear function of projections for price inflation, wage inflation, and the output gap.

Optimality requires that the target value be given by an expression of the form

$$\begin{aligned} \pi_t^* &= F_{t-1}^1(\pi) + \phi_w^*[F_{t-1}^1(w) - w_t] + \phi_x^*[F_{t-1}^1(x) - x_t] \\ &+ \theta_\pi^*[F_{t-1}^1(\pi) - F_{t-2}^2(\pi)] + \theta_w^*[F_{t-1}^1(w) - F_{t-2}^2(w)] + \theta_x^*[F_{t-1}^1(x) - F_{t-2}^2(x)], \end{aligned} \quad (3.7)$$

where the expressions $F_t^j(z)$, for $j = 1, 2$, are still other weighted averages of forecasts at different horizons, with relative weights α_k^{zj} that again sum to one. Here, as with (3.5), the optimal target value depends on previous quarters' forecasts of the economy's subsequent evolution; this is a further example of the history-dependence of optimal target criteria, already observed in simpler cases in section 1.

It is worth noting that in the case of each of the three variables — the inflation rate, the real wage, and the output gap — it is only the *differences* between the projections for the variables and their current values, or between current projections and projections for the same variables at an earlier date, that matter for the optimal target criterion.⁵¹ What this means is that we could alternatively express the target criterion entirely in terms of projections of the *first differences* of these three variables, with no dependence on the *absolute levels* of any of the variables.⁵² The reasons are somewhat different, however, in each case. The inflation projections involve only the projected first difference of the inflation rate because in the loss function (3.1), it is only the values of $\pi_t - \pi_{t-1}$ and $\pi_t^w - \pi_{t-1}$ that matter,⁵³ rather than the absolute rate of either price or wage inflation. Projections for these two target variables can alternatively be expressed in terms of projections for $\pi_t - \pi_{t-1}$ and $w_t - w_{t-1}$. Because real wage projections enter the target criterion only for this reason (*i.e.*, as a substitute for wage-inflation projections), it is only projections for the first difference of

⁵¹Note that the same is true of the short-run target criterion (3.3) – (3.5).

⁵²We have chosen not to do so, because this representation seems to allow a more convenient description of the numerical character of the optimal target criterion. On the other hand, we could simplify the algebraic statement of the target criterion by eliminating the $-\phi_w^*w_t - \phi_x^*x_t$ terms that appear both on the left-hand side of (3.6) and the right-hand side of (3.7). But we have not done this, in order to make it clear that this target criterion does not actually require an estimate of either the absolute level of the output or of the deviation of the real wage from its trend path, but only of output-gap and real-wage *growth*. This is an important practical consideration, given that real-time estimates of the absolute levels of both variables are much more controversial.

⁵³The result thus depends on the fact that our estimates in section 2 imply that $\gamma_p = \gamma_w = 1$.

the real wage that matter as well. The observation that the optimal target criterion involves only the projected rate of change of the output gap, rather than its level, is instead not a consequence of the form of the loss function — for δ may be much less than one, as it is for our estimated parameter values. Nor is this a consequence of our estimated value $\eta = 1$. The optimality conditions just stated apply regardless of the value of η , and we note once again that $\eta = 1$ does not imply $\delta = 1$ in any event. The result is instead an illustration of the generality of the conclusion, first obtained in section 1, that the optimal target criterion involves the rate of change of the output gap rather than its absolute level, regardless of the degree of habit persistence.

It may be wondered how we can specify optimal policy in terms of two distinct target criteria involving different linear combinations of projections, when the central bank has only one instrument at its disposal. The key to this is to observe that the target criterion specified by (3.3) – (3.5) restricts only the *surprise* components of the quarter t projections, *i.e.*, the way in which they may differ from the projections that were made in quarter $t - 1$ for the same variables. Hence it is only the *surprise* component of the central bank’s interest-rate decision — the difference between the $E_t i_{t+1}$ announced in quarter t and $E_{t-1} i_{t+1}$ — that can be determined by this criterion for optimal policy. The evolution of the (two-period-ahead) predetermined component of policy, $E_{t-2} i_t$, can instead be chosen so as to ensure that the second target criterion, specified by (3.6) – (3.7), is satisfied each period.

We may thus imagine the implementation of the optimal targeting rule to occur in the following way.⁵⁴ First, in each quarter t , the central bank intervenes in the money markets (through open-market operations, repurchases, standing facilities in the interbank market for central-bank balances, etc.) so as to implement the interest-rate target i_t announced in quarter $t - 1$. Second, as part of the quarter t decision cycle, the bank must choose an operating target i_{t+1} to announce for the following quarter. This is chosen in order to imply

⁵⁴Because our empirical model is quarterly, it is simplest to discuss the policy process as if a policy decision is also made once per quarter, even though in reality most central banks reconsider their operating targets for overnight interest rates somewhat more frequently than this. Our discussion should not be taken to imply that it is optimal for the policy committee to meet only once per quarter; this would follow from our analysis only if (as in our model) all other markets were also open only once per quarter.

a projected evolution of (wage and price) inflation from quarter $t + 1$ onward that satisfies the target criterion (3.3), where $\bar{\pi}_t$ is a target value that had been determined in quarter $t - 1$. Third, it is also necessary, as part of the quarter t decision cycle, for the central bank to choose the target $\bar{\pi}_{t+1}$ for the following quarter. This is chosen so as to ensure that future policy will be conducted in a way that allows the bank to project (conditional on its current information) that the target criterion (3.6) – (3.7) should be satisfied. In practice, this means that the central bank should use its model of the transmission mechanism to determine the future evolution of the economy under the assumption that (3.6) – (3.7) will hold in all future periods; this forecast then determines the target value $\bar{\pi}_{t+1}$ using (3.5).⁵⁵

Algebraic expressions for each of the coefficients in the optimal target criteria, as functions of the underlying model parameters, are given in the appendix. Here we discuss only the numerical coefficients implied by our estimated parameter values. In the case of the short-term criterion (3.3), the coefficient ϕ_w is equal to 0.565.⁵⁶ Thus if unexpected developments in quarter t are projected to imply a higher future level of real wages than had previously been anticipated, policy must ensure that projected future price inflation is correspondingly reduced. This is because of a desire to stabilize (nominal) wage inflation as well as price inflation, and under circumstances of expected real wage growth, inflation must be curbed in order for nominal wage growth to not be even higher.

The relative weights that this criterion places on projections at different future horizons are shown in Figure 7. The two panels plot the coefficients α_k^π, α_k^w respectively, as functions of the horizon k . Note that the quarter for which the projections receive greatest weight is one quarter in the future, in each case. However, while the real-wage projection that matters is primarily the projected growth in real wages between the present quarter and the next one, substantial weight is also placed on projected inflation farther in the future;

⁵⁵See Svensson and Woodford (2003) for further discussion of the sort of calculations involved in a forecast-targeting decision procedure.

⁵⁶Here and below, we present the coefficients for a target criterion where the inflation rate is measured in annualized percentage points, rather than as a quarterly rate of change as in the model of section 2. When the variables are defined as in the model, the coefficients multiplying the real-wage and output-gap terms are only 1/4 as large as those given here and below.




Figure 7: Relative weights on projections at different horizons in the short-run target criterion (3.3). The horizontal axis indicates the horizon k in quarters.

in fact, the mean lead $\sum_k \alpha_k^\pi k$ is between 10 and 11 quarters in the future in the case of the inflation projection $F_t(\pi)$. Thus the short-run target criterion is a (time-varying) target for the average rate of inflation that is projected over the next several years, adjusted to take account of expected wage growth, mainly over the coming quarter. Roughly speaking, optimal policy requires the central bank to choose $E_t i_{t+1}$ in quarter t so as to head off any change in the projected average inflation rate over the next several years that due to any developments not anticipated in quarter $t-1$ (and hence reflected in the current target $\bar{\pi}_{t-1}$). This is a criterion in the spirit of inflation-forecast targeting as currently practiced at central banks such as the Bank of England, except that projected wage growth matters as well as price inflation, and that the target shifts over time.

In the case of the long-term criterion (3.6), instead, the numerical coefficients of the




Figure 8: Relative weights on projections at different horizons in the long-run target criterion. Panels in the first row indicate the projections in (3.6), while the other two rows indicate the projections from previous quarters that define the target value π_t^* .

target criterion are given by

$$\phi_w^* = 0.258, \quad \phi_x^* = 0.135.$$

In this case, output-gap projections matter as well; a higher projected future output gap will require a reduction in the projected future rate of inflation, just as will a higher projected future real wage. The numerical size of the weight placed on the output-gap projection may appear modest; but as we shall see in the next section, the degree of variability of output-gap projections in practice are likely to make this a quite significant correction to the path of the target criterion.

The relative weights on forecasts at different horizons in this criterion are plotted in the panels in the first row of Figure 8. We observe that in the case of this criterion, the

projections that mainly matter are those for two quarters in the future; the criterion is nearly independent of projections regarding the quarter after the current one. Hence it makes sense to think of this criterion as the one that should determine the policy that the central bank plans on in periods two or more quarters in the future (and hence its choice in quarter t of the target $\bar{\pi}_{t+1}$ to constrain its choice in the following period of $E_{t+1}i_{t+2}$), but not as a primary determinant of whether the bank's intended policy in period $t + 1$ is on track.

Finally, the coefficients of the rule (3.7) determining the target value for the long-term criterion are given by

$$\theta_{\pi}^* = 0.580, \quad \theta_w^* = 0.252, \quad \theta_x^* = 0.125.$$

The weights in the projections (conditional on information in the previous two quarters) at various horizons are plotted in rows 2 and 3 of Figure 8. Here too, it is primarily projections for two quarters in the future that matter in each case. Roughly speaking, then, the target value for the wage- and output-adjusted inflation projection two quarters in the future is high when a similar adjusted inflation projection (again for a time two quarters in the future) was high in the previous quarter, or when the adjusted inflation projection in the previous quarter was high, not in absolute terms, but *relative* to a similar adjusted inflation projection two quarters ago.

Thus we find that forecasting exercises, in which the central bank projects the evolution of both inflation and real variables many years into the future under alternative hypothetical policies on its own part, play a central role in a natural approach to the implementation of optimal policy. A forecast of inflation several years into the future is required in each (quarterly) decision cycle in order to check whether the intended interest-rate operating target for the following quarter is consistent with the criterion (3.3). In addition, the time-varying medium-term inflation target $\bar{\pi}_t$ must be chosen each period on the basis of yet another forecasting exercise. While the long-run target criterion (3.6) primarily involves projections for a time only two quarters in the future, the choice of $\bar{\pi}_{t+1}$ requires that the central bank solve for a projected path of the economy in which (3.6) is satisfied not only in

the current period, but in all future periods as well. Hence this exercise as well requires the construction of projected paths for inflation and real variables extending many years into the future. The relevant paths, however, will not be constant-interest-rate projections (of the kind currently published by the Bank of England), but rather projections of the economy's future evolution given how policy is expected to evolve. Indeed, the projections are used to select constraints upon the bank's own actions in future decision cycles (by choosing both the interest-rate operating target $E_t i_{t+1}$ and the adjusted inflation target $\bar{\pi}_{t+1}$ in period t).

3.3 A Comparison with Actual U.S. Policy

An interesting question about this policy rule is the extent to which it would prescribe policy different from that which the Fed has actually pursued during our sample period. A simple way of considering this is to ask to what extent, under actual policy, projections of the evolution of inflation and output have satisfied the optimal target criteria stated above. Answering this question requires, of course, that we estimate what the projected future paths of the target variables should have been at various past dates. However, our VAR characterization of the data over our sample period provides one way of generating such projections. Here we propose to appraise how close actual policy has been to being optimal by asking to what extent projections based on the VAR would have satisfied the target criterion.

In our characterization of optimal policy above, there are actually three criteria that must be satisfied each period — one relating to the component of interest-rate policy that cannot be forecasted even a quarter in advance, one relating to the component of policy that is forecastable a quarter in advance but not earlier, and one relating to the component of policy that can be anticipated two quarters in advance. The first criterion, that the evolution of interest rates satisfy (3.2) each period, is simplest to check, as long as we are willing to assume that our VAR forecasts fully capture public information in a given quarter. Figure 9 shows a plot of the actual (quarterly average) path of the federal funds rate over our sample period, together with the VAR forecast using the previous quarter's information




Figure 9: Actual and forecastable variation in the U.S. federal funds rate.

set.⁵⁷ This allows a test of the degree to which condition (3.2) has been satisfied in practice. We find that under actual U.S. policy, variation in the U.S. federal funds rate has been largely predictable; the gap between the two series in Figure 9 has a standard deviation of only 65 basis points.⁵⁸ This means that the identified monetary policy shocks, according to the VAR analysis discussed in section 2, have been relatively small. This is what one should expect, in a period in which the conduct of monetary policy has been fairly sensible.

⁵⁷Note that here and below, the “quarter t information set” is taken to include π_{t+1} , w_{t+1} , and \hat{Y}_{t+1} , as well as all variables dated t or earlier, on the ground that prices, wages, and output are all predetermined variables according to our model. See Rotemberg and Woodford (1997) for further discussion.

⁵⁸Of course, we are judging the forecastability of the funds rate using a VAR that has been fit to this data set, rather than considering the out-of-sample forecasting ability of a regression model estimated using only data prior to the quarter for which the funds rate is being forecasted. We are also including variables in the quarter t information set the values of which are not announced in quarter t (indeed, not even during quarter $t+1$, though the measurements are made during that quarter), which also exaggerates the information actually available in quarter t . But it must also be recognized that decisionmakers have access to a great deal of information in quarter t that is not included in our data set, that might well allow better forecasting of the funds rate than is possible on the basis of only the variables included in our VAR.




Figure 10: Testing whether actual U.S. policy has satisfied the short-run target criterion: the adjusted inflation projection (3.3) compared with the optimal target given by (3.5).

The next condition for optimality that we consider is the short-term target criterion (3.3) – (3.5). Figure 10 shows a plot of the historical path of the wage-adjusted inflation projection that is targeted under this criterion, using the VAR forecasts to form this projection each quarter, together with the path for the target value $\bar{\pi}_t$ given by (3.5), also using the VAR forecasts for the projections in the previous quarter. Figure 11 decomposes the variation in both the adjusted inflation projection (3.3) and the time-varying target $\bar{\pi}_t$ into the parts that are due to variation in the inflation projections (at various horizons) on the one hand and the parts that are due to variation in the real wage projections. We observe that a substantial part of the quarter-to-quarter variation in the adjusted inflation projection is in fact due, over this historical period, to variation in the real wage projection,⁵⁹ though variation in the real wage projection a quarter earlier appears to be less important as a source of variation

⁵⁹We have not attempted to quantify the share since the two components are not orthogonal.




Figure 11: Decomposition of the variation in the short-run target criterion into parts due to variation in inflation projections and real wage projections respectively.

in the optimal target value.

Once again, the data are fairly consistent with this criterion for optimal policy. While the wage-adjusted inflation projection has varied (according to the VAR) over a range of a few percentage points, these variations have been fairly forecastable based on the previous quarter’s information set, as required by the target criterion. The gap between the projection and the target value has a standard deviation of only 48 basis points over this sample.

Of course, passing this test requires only that wage and price inflation, like the federal funds rate, be highly forecastable a quarter in advance. It may accordingly be felt that it is the inertial character of wage and price inflation that is confirmed by Figure 10, rather than something that depends much on monetary policy. It should also be noted that the “target” series plotted in the figure only indicates how the right hand side of (3.5) has

varied over the sample period, under actual U.S. monetary policy, rather than the way in which the target $\bar{\pi}_t$ *would have* evolved under optimal monetary policy, given that the inflation projections that determine this target would have been different under a different sort of monetary policy. This latter sort of exercise would require that we solve for the counterfactual equilibrium paths the endogenous variables under optimal policy, given the historical sequence of exogenous shocks, as undertaken by Rotemberg and Woodford (1997). We do not attempt such an exercise here.

Testing the extent to which the historical data have satisfied the long-run target criterion (3.6) – (3.7) is more complicated, because it requires the construction of projections for the path of the output gap. The output gap is not directly observed, and our approach to the estimation of the model in section 2 does not require us to commit ourselves to an empirical proxy for the gap, despite the appearance of this variable in the model structural equations. For in order to estimate the model parameters needed for our calculations thus far, we had only to be able to compute the predicted impulse responses of prices, wages, output and interest rates to a monetary disturbance. For this purpose, we could rely on the fact that, according to our model, the output gap should equal \hat{Y}_t (detrended log output) minus a term that is unaffected by monetary disturbances; there was no need to identify the time variation in that latter term. Yet in order to evaluate the long-run target criterion at each date, we need to be able to do so.

One possible approach is to use our estimated structural equations to infer the historical sequence of disturbances from the residuals of the structural equations, using VAR forecasts of the endogenous variables as proxies for the expectation terms in these equations, as do Rotemberg and Woodford (1997). This approach can be used, however, only under strong assumptions of debatable validity. The “natural rate of output” process that we are able to infer from the residuals of our structural equations corresponds to the equilibrium level of output under complete wage and price flexibility.⁶⁰ But this may or may not be the concept

⁶⁰To be precise, it corresponds to the component of this variable that is forecastable a quarter in advance. This is all that can be reconstructed from the paths of the endogenous variables, given that wages, prices and output are all predetermined according to our model, but this is also what is relevant for the construction of

of exogenously given potential output that should be used to define the welfare-relevant “output gap” that appears in the loss function (3.1).

Under certain assumptions that are made precise in the appendix (and that have been tacitly maintained thus far in our exposition), the “output gap” that appears in the structural equations (2.11) and (2.14) as a source of inflationary pressure — without any additional “cost-push shock” term of the kind routinely included in the models of section 1 — is exactly the same variable as the distortion measure appearing in (3.1). Yet this need not be true in general; time variation in distorting taxes or in the degree of market power in either labor markets or goods markets, for example, will result in a time-varying wedge between the flexible-wage-and-price equilibrium level of output and the efficient level of output, with the result that the relevant output gap for the two purposes ceases to be the same.⁶¹ We can allow for this extension of our framework by letting the gap between actual output and the flexible-wage-and-price equilibrium output be denoted $x_t + u_t$, as in equations (1.29) – (1.30) above, where x_t is the welfare-relevant output-gap concept (the variable that appears in the welfare-theoretic loss function), while u_t is a “cost-push disturbance” term.

In the case of the extended model, the method of Rotemberg and Woodford allows us to construct an empirical proxy for the evolution of the series $x_t + u_t$, as this is what appears in the wage- and price-setting equations. However, the projections that are required for checking whether the target criterion is satisfied are projections for x_t , the variable that appears in the loss function (3.1). Further assumptions must be made in order to infer what the projected variations in the welfare-relevant output gap should have been. These assumptions are not testable within the context of the model and the small set of time series used here.

One simple, though extreme, assumption, would be that the welfare-relevant concept of potential output is a smooth trend, so that cyclical variation in \hat{Y}_t^n should be almost entirely attributed to transitory variation in the cost-push term u_t .⁶² In this case, it should

the variable x_t that appears in our loss function (3.1), and hence the target criterion stated in the previous section.

⁶¹See Giannoni (2000) or Woodford (2003, chap. 6) for further discussion in the context of simpler models.




Figure 12: Testing whether actual U.S. policy has satisfied the long-run target criterion: the adjusted inflation projection (3.6) compared with the optimal target given by (3.7).

be more accurate to identify the welfare-relevant output gap with \hat{Y}_t , detrended output, than with the series $x_t + u_t$ inferred from the residuals of the structural equations. Under this assumption, we can construct our output-gap projections using the VAR alone, without any need reconstruct disturbances using the equation residuals.

We first consider the conformity of historical policy with the optimal target criteria when detrended output is considered an adequate proxy for the output gap. In Figure 12, we plot the historical series for the wage- and output-adjusted inflation projection that is targeted under the long-term criterion (3.6) over our sample period, using the VAR forecasts for inflation, the real wage, and detrended output. Figure 13 similarly decomposes both the projection and its optimal target value into their components due to variation over time

⁶²This view is implicit in the output-gap measures commonly used in the literature on empirical central-bank reaction functions.




Figure 13: Decomposition of the variation in the long-run target criterion into parts due to variation in inflation, real-wage, and output projections respectively.

in inflation projections, real-wage projections, and output projections. Note that when the output gap is measured in this way, the projected change in the output gap over a two-quarter horizon is modest enough that terms of this kind are not responsible for too much of the variation from quarter to quarter in either the adjusted inflation projection or in its optimal target value. Instead, the target criterion is largely a function of the inflation and real-wage projections (or alternatively, projected price and wage inflation).

This alternative (longer-run) adjusted inflation projection has also been relatively stable over our historical sample, and once again, the gap between the target and the current projection has never been large; the standard deviation of target misses in the case of this criterion is only 69 basis points. Furthermore, target misses under this criterion have been quite transitory; their autocorrelation is actually negative (-0.57), indicating that a target




Figure 14: Alternative version of figure 12, using equation residuals to infer the variation in the natural rate of output.

overshoot one quarter tends to have its sign reversed in the next quarter. One therefore cannot identify any periods of time of even a year over which policy was consistently too loose or too tight under this criterion. Furthermore, in the case of this criterion, the target is not simply the forecast a quarter earlier of what the projection would be, so that the lack of serial correlation in the target misses is not an automatic consequence of the way in which we have constructed our projections, as in the other two cases. The historical record is therefore fairly consistent with this target criterion as well.

If, instead, we use the residuals from our structural equations to infer the evolution of the output gap, the plots corresponding to Figures 12 and 13 instead look like those shown in Figures 14 and 15. In this case, historical paths of both the adjusted inflation projection and its optimal target value are considerably more volatile. The change is due to the much greater (and much more transitory) volatility of the output gap process that is inferred




Figure 15: Alternative version of figure 13, using equation residuals to infer the variation in the natural rate of output.

in this manner. As shown in Figure 15, in this case the quarter-to-quarter variation in projected growth of the output gap is the most important factor resulting in variation in the adjusted inflation projection and in the target value. Of course, the extreme volatility of (and high-frequency variation in) this series may well suggest that it reflects mainly specification error in the structural equations of our wage-price block, rather than actual variation in the welfare-relevant output gap.⁶³

In this case, the gap between the adjusted inflation projection and its optimal target value is also found to often be fairly large. The standard deviation of the gap between the two series shown in Figure 14 is nearly 3.12 percentage points. However, once again the target misses are extremely transitory (the gap series has an autocorrelation of -0.69), so that there

⁶³The fact that our model does relatively poorly at matching the dynamics of the estimated response of inflation, as shown in Figure 6, does not give us much confidence in this regard.

are no periods of time over which policy can be identified as having been consistently too tight or too loose.

Overall, a comparison between U.S. time series over the past twenty years and the criteria for optimal policy discussed in the previous section do not indicate any gross discrepancy. However, this may simply mean that the diagnostics proposed here are not very useful as a way of diagnosing deviations from optimal policy in the historical record. We have plotted only the time-variation in the optimal target criteria that would be implied by the variation in lagged projections that *has* occurred, given the actual evolution of the U.S. time series, rather than attempting to determine the variation in the target values that *would have* occurred under optimal policy, given the historical disturbance processes. These two ways of judging the historical time series might yield quite different pictures. For our optimal target criteria demand that certain adjusted inflation projections not be too different than similar projections have been in the quarter before; this will result in plots of projections and target values that look fairly similar, regardless of the paths of the U.S. time series, as long as each of our four variables has been relatively smooth (as is the case). Nonetheless, inflation and other variables might have wandered for years at some distance from the levels that they would have had under fully optimal responses to the historical disturbances.

4 Conclusions

We have shown that it is possible to derive robustly optimal monetary policy rule for optimizing models of the monetary transmission mechanism that incorporate a number of common features of recent empirical models: staggered wage- and price-setting, inflation inertia resulting from automatic indexation of wages and prices to a lagged price index, predetermined wage-setting, pricing and spending decisions, and habit persistence in the level of real private expenditure. In this way, we have sought to show that the approach to the design of optimal policy rules proposed by Giannoni and Woodford (2002a) can be applied to models of practical interest.

In each of the cases that we have discussed, the optimal policy rule is a modified inflation-

forecast targeting rule. The optimal rule differs from a simple (or “strict”) inflation target in that projections of the future paths of variables other than goods-price inflation also receive some weight in the target criterion — in particular, wage inflation, a measure of the output gap, and nominal interest rates. Nonetheless, according to our numerical analysis in the case of an estimated model of the U.S. monetary transmission mechanism, the weight on the inflation projection (in each of the two target criteria involved in our characterization of optimal policy for that model) is strong enough that it makes sense to speak of optimal policy as a (flexible) inflation-forecast targeting procedure.

In our examples, the optimal rule also differs from a simple inflation target (and even from many simple examples of “flexible inflation targeting” rules discussed in the literature) in that the optimal target value for the modified inflation forecast should vary over time, depending on current and recent past macroeconomic conditions. We have illustrated the possible degree of history-dependence of an optimal inflation target by showing how our two optimal target criteria would have varied in the U.S. over the past two decades, given our VAR characterization of the U.S. time series and the parameters of our estimated structural model. Even when we use detrended output as our proxy for the output gap (which results in a less volatile output-gap series than the one implied by the residuals of our structural equations), and even over the relatively uneventful period 1984-2000, our analysis implies that the optimal target criterion has varied from quarter to quarter over a range of several percentage points.

Finally, we have shown, in the context of our empirical model, that an optimal policy may be too complex in structure to be conveniently described by a single target criterion. Our estimated model of the U.S. monetary transmission mechanism implies that optimal policy must satisfy three distinct criteria — one that governs the way that interest rates in a given quarter should respond to unexpected developments during that quarter; one that governs the way in which the central bank’s commitment regarding interest rates in that quarter, announced the quarter before, should respond to unexpected developments in the quarter when the commitment is made; and still a third criterion that determines the component

of interest-rate policy that can be anticipated two quarters in advance. Nonetheless, the decision procedure takes the form of an inflation-forecast targeting procedure, in which (i) the instrument used to ensure satisfaction of the target criterion is the central bank's commitment regarding its interest-rate operating target for the following quarter, and (ii) the inflation target each quarter is itself the product of a policy decision in the previous quarter, also aimed at ensuring that a certain adjusted inflation projection satisfies a target criterion.

Our optimal target criteria are a good bit more complex than the sort used by actual inflation-targeting central banks, which typically specify a time-invariant inflation target and a particular horizon at which it is to be reached (for example, RPIX inflation of 2.5 percent at a horizon of 8 quarters in the future, in the case of the Bank of England). Our advocacy of a more complex form of targeting rule is not meant to deny the desirability of having a medium-term inflation target that remains the same even if the actual inflation rate may depart from it temporarily. In the examples that we have considered, optimal policy almost always involves a well-defined long-run inflation target, to which the inflation rate should be expected to return after each disturbance; and it is surely desirable for a central bank to be explicit about this aspect of its policy commitment, in order to anchor the public's medium-term inflation expectations.

Rather, we wish to suggest that it is *insufficient* to specify no more of a policy commitment than this. The mere fact that a central bank wishes to see inflation return to a rate of 2.5 percent at a horizon two years in the future is not sufficient to say which of the various possible transition paths that reach that endpoint should be preferred. There will always be a range of possible scenarios consistent with the terminal condition: for example, looser policy this year to be compensated by tighter policy next year, or alternatively the reverse.

In practice, the Bank of England, like many other forecast-targeting banks, deals with this problem by demanding that a *constant-interest-rate* forecast satisfy the terminal condition. That is, the current level of overnight interest rates is held to be justified if a projection under the assumption that that level of interest rates will be maintained implies that RPIX

inflation should equal 2.5 percent eight quarters in the future. However, this implies no commitment to actually maintain interest rates at the current level over that period, or even that interest rates are currently expected to remain at that level on average. (It is frequently the case that the published constant-interest-rate projection would itself imply that interest rates will need to be changed over the coming year, in order for the target criterion to be satisfied by a constant-interest-rate projection under the conditions that are forecasted to obtain by then.) It is thus hard to see how basing policy decisions on a forecast-targeting exercise of this particular kind can be expected to serve the goals of making monetary policy more transparent, or improving the degree to which policy is correctly anticipated by the private sector.

The conceptually superior approach, surely, is to base policy on a projection that is computed under the assumption that policy will be made in accordance with the targeting rule in the future as well,⁶⁴ so that the projection that is used to justify current policy will correspond to the bank's own best forecast of how it should act in the future, as in the case of the projections used to justify policy decisions by the Reserve Bank of New Zealand. It will, of course, be necessary to stress that the bank's only commitment is to the *rule* embodied in this projection, not to the particular time path of interest rates indicated as most likely. But given the use of "fan charts" to show that a variety of possible future scenarios can be envisioned, depending on how various types of uncertainty happen to be resolved, it is not clear why it should not be possible to talk about probability distributions for future interest rates along with those for inflation and real activity, without giving rise to the appearance of a more specific commitment than is intended.

Once this is done, however, it becomes necessary to specify a target criterion that can determine the appropriate short-run dynamics for the economy, and not simply a terminal condition for a date some years in the future. Such a criterion will accordingly place substantial weight on projections of the target variables over the coming year, as in the case of the optimal target criteria derived in this chapter. It will also have to take a stand as to

⁶⁴See Svensson and Woodford (2003) for further discussion of what this would mean in practice.

the kinds of projected departures of real variables from their long-run average values that justify short-run departures of the inflation projection from its long-run target value; it will no longer suffice simply to specify what the (unchanging) long-run inflation target is. None of the inflation-targeting central banks actually believe that it is desirable to keep inflation as close as possible to the long-run target value at all times; this is why forecast-targeting procedures only seek to ensure that inflation is projected to return to the target value after many quarters.⁶⁵ But by formulating no explicit doctrine as to the way in which one should choose among alternative transition paths to that medium-term goal, they avoid having to clarify the nature of acceptable trade-offs among competing stabilization goals.⁶⁶

A coherent approach — and in particular, one that could be justified as seeking to implement the conditions for optimal policy discussed in this chapter — would instead have to make explicit the kind of projections for output and other real variables that should justify a modification of the short-run inflation target, and the degree to which they should affect it. In all likelihood, the inflation-targeting banks have shied away from such explicitness out of a suspicion that the types of circumstances that might reasonably justify short-term departures from the inflation target are too various to be catalogued. But the theory developed here has sought to show that it is possible to state *short-run* target criteria (criteria that apply to the shortest horizon at which current policy decisions can still have an effect) that will be *robustly* optimal, meaning that the same criterion continues to determine the correct degree of short-run departure from the long-run inflation target regardless of the nature of the disturbance that may have occurred.

Much work remains to be done, of course, before a quantitative characterization of optimal policy of the kind that we offer in section 3 could be used in practical policy deliberations. One of the most obvious issues requiring further study concerns the way in which a

⁶⁵On this point, see, e.g., Bernanke *et al.*, 1999, or Svensson, 1999.

⁶⁶The fact that a real GDP projection is always included along with the projection for RPIX inflation in the introduction to the Bank of England's *Inflation Report* — and in fact, is always discussed *first* — suggests that the some attention is paid to the projected path of output in deciding upon the appropriateness of the current level of interest rates. But the Bank's official target criterion, involving only the constant-interest-rate projection of RPIX inflation at the eight-quarter horizon, does not make explicit the way in which the output projection should be taken into account.

central bank should take account of uncertainty about the correct model of the transmission mechanism, as well as uncertainty in its evaluation of current macroeconomic conditions. Uncertainty about the current state of the economy is relatively straightforward to deal with, at least in principle. One can allow for partial information on the part of the central bank in characterizing the optimal equilibrium responses to shocks, using methods similar to those employed here, and derive an optimal target criterion that is valid in the presence of partial information (Svensson and Woodford, 2002a, 2002b; Giannoni and Woodford, 2002b). Because of the principle of certainty-equivalence in linear-quadratic policy problems of this kind (discussed in detail by Svensson and Woodford), the optimal target criterion (once correctly expressed) involves coefficients that are independent of the degree of uncertainty in central-bank estimates of the current state of the economy; however, the target may involve variables that are not directly observed by the central bank, and must instead be estimated using a Kalman filter.

Dealing with uncertainty about the numerical values of structural parameters (to say nothing of more fundamental doubts about model specification) is a much harder problem, for which few general guidelines exist at present. Giannoni (2001, 2002) illustrates one approach to the problem, for the case of uncertainty about the numerical values of the elasticities κ and σ in a model similar to our baseline model (but in which an interest-rate stabilization objective is assumed). For the particular kind of parameter uncertainty considered, Giannoni finds that a concern for robustness (in the sense of guarding against bad outcomes in the least-favorable case) should lead a central bank to choose a Taylor-style interest-rate rule with stronger response coefficients than it would choose on the basis of its preferred estimates of the model parameters; this means allowing *less* variability of inflation in equilibrium, at the cost of greater variability in nominal interest rates. This suggests that a concern for robustness might justify targeting rules that are even closer to strict inflation targeting than the optimal rules obtained in this paper; the question is surely one that deserves further analysis.

References

- Altig, David, Lawrence J. Christiano, Martin S. Eichenbaum, and Jesper Linde, "Technology Shocks and Aggregate Fluctuations," unpublished, Federal Reserve Bank of Cleveland, 2002.
- Amato, Jeffery D., and Thomas Laubach, "Implications of Habit Formation for Optimal Monetary Policy," FEDS paper no. 2001-58, Federal Reserve Board, August 2001a.
- —, and — —, "Estimation and Control of an Optimization-Based Model with Sticky Wages and Prices," unpublished, October 2001b. (Forthcoming, *Journal of Economic Dynamics and Control*.)
- Ball, Laurence, and David Romer, "Real Rigidities and the Non-Neutrality of Money," *Review of Economic Studies* 57: 183-203 (1990).
- Basu, Susanto, and Miles S. Kimball, "Investment Planning Costs and the Effects of Fiscal and Monetary Policy," unpublished, Univ. of Michigan, November 2002.
- Bernanke, Ben S., and Alan S. Blinder, "The Federal Funds Rate and the Transmission of Monetary Policy," *American Economic Review* 82: 901-921 (1992).
- —, Thomas Laubach, Frederic S. Mishkin, and Adam S. Posen, *Inflation Targeting*, Princeton NJ: Princeton University Press (1999).
- —, Jean Boivin and Piotr Elias, "Factor Augmented VAR and the Analysis of Monetary Policy," unpublished manuscript, Princeton University, 2002.
- — and Ilian Mihov, "Measuring Monetary Policy," *Quarterly Journal of Economics* 113(3): 869-902 (1998).
- Blinder, Alan S., "On Sticky Prices: Academic Theories Meet the Real World," in N. G. Mankiw, ed., *Monetary Policy*, Chicago: University of Chicago Press (1994).
- Boivin, Jean, "The Conduct of Monetary Policy: Has it Changed and Does it Matter?" unpublished manuscript, Columbia University, 2001.
- — and Marc P. Giannoni, "Has Monetary Policy Become More Effective?" NBER Working Paper no. 9459, January 2003.
- Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics* 12(3): 383-98 (1983).

- Christiano, Lawrence J., Martin S. Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," NBER Working Paper no. 8403, July 2001.
- Clarida, Richard, Jordi Galí, and Mark Gertler, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature* 37: 1661-1707 (1999).
- —, — —, and — —, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics* 115: 147-180 (2000).
- Cogley, Timothy, and Thomas Sargent (2001), "Evolving Post-World War II U.S. Inflation Dynamics," *NBER Macroeconomics Annual* 2001.
- —, and — —, "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.," unpublished manuscript, August 2002.
- Dixit, Avinash K., and Joseph E. Stiglitz, "Monopolistic Competition and Optimum Product Diversity," *American Economic Review* 67: 297-308 (1977).
- Edge, Rochelle, "Time-to-Build, Time-to-Plan, Habit-Persistence, and the Liquidity Effect," IF discussion paper 2000-673 (2000).
- Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin, "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics* 46: 281-313 (2000).
- Friedman, Milton, "The Optimum Quantity of Money," in *The Optimum Quantity of Money and Other Essays*, Chicago: Aldine, 1969.
- Fuhrer, Jefferey C., "Habit Formation in Consumption and Its Implications for Monetary-Policy Models," *American Economic Review* 90(3): 367-390 (2000).
- —, and Geoffrey R. Moore, "Inflation Persistence," *The Quarterly Journal of Economics* 110(1): 127-159 (1995b).
- Galí, Jordi, and Mark Gertler, "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics* 44: 195-222 (1999).
- Giannoni, Marc P., "Optimal Interest-Rate Rules in a Forward-Looking Model, and Inflation Stabilization versus Price-Level Stabilization," unpublished, Princeton University, 2000.
- —, "Robust Optimal Monetary Policy in a Forward-Looking Model with Parameter and Shock Uncertainty," unpublished, Federal Reserve Bank of New York, 2001.

- —, “Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model,” *Macroeconomic Dynamics* 6(1): 111-144 (2002).
- —, and Michael Woodford, “Optimal Interest-Rate Rules: I. General Theory,” NBER Working Paper no. 9419, December 2002a.
- —, and — —, “Optimal Interest-Rate Rules: II. Applications,” NBER Working Paper no. 9420, December 2002b.
- Goodfriend, Marvin, “Interest Rate Smoothing in the Conduct of Monetary Policy,” *Carnegie-Rochester Conference Series on Public Policy* 34: 7-30 (1991).
- Judd, John F., and Glenn D. Rudebusch, “Taylor’s Rule and the Fed: 1970-1997,” *Federal Reserve Bank of San Francisco Economic Review*, 1998(3): 3-16.
- Kilian, Lutz, “Small-Sample Confidence Intervals for Impulse Response Functions,” *Review of Economics and Statistics*, 80(2): 218-230 (May 1998).
- King, Mervyn, “Changes in UK Monetary Policy: Rules and Discretion in Practice,” *Journal of Monetary Economics* 39: 81-97 (1997).
- McCallum, Bennett T. and Edward Nelson, “Performance of Operational Policy Rules in an Estimated Semi-Classical Structural Model,” in J.B. Taylor, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press, 1999.
- Roberts, John M., “New Keynesian Economics and the Phillips Curve,” *Journal of Money, Credit and Banking* 27(4): 975-984 (1995).
- Rotemberg, Julio J., and Michael Woodford, “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy,” *NBER Macroeconomics Annual* 12: 297-346 (1997).
- — and — —, “Interest-Rate Rules in an Estimated Sticky-Price Model,” in J.B. Taylor, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press, 1999.
- Smets, Frank, and Raf Wouters, “Sources of Business Cycle Fluctuations in the U.S.: A Bayesian DSGE Approach,” seminar presentation, Princeton University, November 1, 2002.
- Steinsson, Jon, “Optimal Monetary Policy in an Economy with Inflation Persistence,” unpublished, Harvard University, May 2002.
- Svensson, Lars E.O., “Inflation Targeting as a Monetary Policy Rule,” *Journal of Monetary Economics* 43: 607-654 (1999).

- —, and Michael Woodford, “Indicator Variables for Optimal Policy, ” Princeton University, January 2002a. Forthcoming, *Journal of Monetary Economics*.)
- —, and — —, “Indicator Variables for Optimal Policy under Asymmetric Information,” Princeton University, August 2002b. (Forthcoming, *Journal of Economic Dynamics and Control*.)
- — and — —, “Implementing Optimal Policy through Inflation-Forecast Targeting,” January 2003. (Forthcoming in B.S. Bernanke and M. Woodford, eds., *Inflation Targeting*, University of Chicago Press.)
- Taylor, John B., “Discretion versus Policy Rules in Practice,” *Carnegie-Rochester Conference Series on Public Policy* 39: 195-214 (1993).
- —, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press, 1999.
- Vickers, John, “Inflation Targeting in Practice: The UK Experience,” *Bank of England Quarterly Bulletin* 38, 368-375 (1998).
- Woodford, Michael, “Doing Without Money: Controlling Inflation in a Post-Monetary World,” *Review of Economic Dynamics* 1: 173-219 (1998).
- —, “Optimal Monetary Policy Inertia,” NBER Working Paper no. 7261, August 1999a.
- —, “Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?” in *New Challenges for Monetary Policy*, Kansas City: Federal Reserve Bank of Kansas City, 1999b.
- —, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, forthcoming 2003.