

NBER WORKING PAPER SERIES

INFLATION STABILIZATION AND WELFARE

Michael Woodford

Working Paper 8071
<http://www.nber.org/papers/w8071>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2001

I would like to thank Marvin Goodfriend, Jinill Kim, Julio Rotemberg, and Marco Vega for helpful comments on an earlier draft, Gauti Eggertson for research assistance, and the National Science Foundation and the John Simon Guggenheim Memorial Foundation for research support. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

© 2001 by Michael Woodford. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Inflation Stabilization and Welfare
Michael Woodford
NBER Working Paper No. 8071
January 2001
JEL No. D60, E52

ABSTRACT

This paper derives loss functions for monetary policy that are grounded in the welfare of private agents, for optimizing models with nominal price rigidities. Inflation stabilization enhances welfare, insofar as variable inflation results in real distortions when prices are not adjusted throughout the economy in a perfectly synchronized fashion. The exact sense in which inflation variability matters for welfare, however, depends upon the details of price-setting behavior.

Conditions are described under which optimal policy involves complete stabilization of the price level. This may be optimal even in the presence of “supply shocks” of several possible sorts, and even in the presence of distortions that imply that the optimal output gap is positive (despite existence of a long-run Phillips curve). At the same time, a variety of reasons are discussed why complete price-level stabilization is not optimal in more complicated (and probably more realistic) settings.

Michael Woodford
Department of Economics
Princeton University
Princeton, NJ 08544
and NBER
woodford@princeton.edu

What should be the goal of monetary stabilization policy? There is thus a fair amount of consensus in the academic literature that a desirable monetary policy is one that achieves a low expected value of a discounted loss function, where the losses each period are a weighted average of terms quadratic in the deviation of inflation from a target rate and in some measure of output relative to potential. But even agreement upon this general form of the objective still allows considerable scope for disagreement about details, that may well matter a great deal for the design of an optimal policy. First of all, obviously, there is the question of the relative weight to be placed upon inflation stabilization and output stabilization. But this is hardly the only ambiguity in the conventional prescription. For instance, which kind of output measure should be stabilized? In particular, should one seek to stabilize output relative to a concept of “potential output” that varies in response to real disturbances that shift the short-run aggregate supply curve, or should one seek to stabilize output relative to a smooth trend?¹

Similarly, in which sense should price stability be pursued? Should one seek to stabilize deviations of the *price level* from a deterministic target path (as proposed, for example, by Hall and Mankiw, 1994), so that unexpected inflation in excess of one’s target rate should subsequently be deliberately counteracted, in order to bring the price level back to its target path? Or should one seek to stabilize deviations of the *inflation rate* from its target level (as assumed, for example, by Svensson, 1997), so that – assuming that the variance of the unforecastable component of inflation cannot be reduced by policy – one should *not* seek to counteract past inflation fluctuations, in order to minimize variation in the forecastable component of inflation? Should greater priority perhaps be given to reducing the variability of *unforecastable* inflation, on the ground that this is what causes unexpected modifications of the real consequences of pre-existing nominal contracts, while forecastable variations in inflation can simply be incorporated into contracts? Or should greater priority be given to stabilization of *forecastable* inflation, on the ground that expected inflation

¹Different answers to this question lead Bean (1983) and West (1986) to reach diametrically opposite conclusions about the case in which nominal GDP targeting would be preferable to money-supply targeting.

distorts incentives (like an anticipated tax), while unforecastable inflation has no incentive effects (like an unanticipated wealth levy)?

The aim of this paper is to show how economic analysis can be brought to bear upon these questions. An important advantage of using a model founded upon private-sector optimization to analyze the consequences of alternative policy rules is that there is a natural welfare criterion in the context of such a model, provided by the preferences of private agents that are displayed in the structural relations that determine the effects of alternative policies. Such a utility-based approach to welfare analysis has long been standard in the theory of public finance. It is not yet too common in analyses of monetary policy, perhaps because it is believed that the main concerns of monetary stabilization policy are assumed away in models with explicit micro-foundations. But we have seen that models founded on individual optimization can be constructed that, thanks to the presence of nominal rigidities, allow for realistic effects of monetary policy upon real variables. Here we shall see those same nominal rigidities provide welfare-economic justification for central bankers' traditional concern for price stability.

Individuals are not assumed, of course, to care directly about prices; their economic welfare depends directly only upon the goods they consume and the amount of effort they expend upon production. But just as taxes can cause deadweight losses because of their effects upon the equilibrium allocation of resources, so can inflation. In a model with nominal rigidities – more specifically, in one in which it is recognized that prices are not adjusted in perfect synchronization with one another (which requires, but is stronger than, the observation that they are not all adjusted continually) – instability of the general price level leads to unnecessary and undesired variation in the relative prices of goods whose prices are adjusted at different times. These relative price distortions result in deadweight losses, just as in the case of distorting taxes. We shall see that this effect can justify not only a loss function that penalizes inflation variations, but indeed – if one assumes parameter values implied by the apparent degree of nominal rigidity in actual economies – a much larger relative weight on inflation variation than upon output variation than is assumed in the loss functions used in

many monetary policy evaluation exercises.

Derivation of a utility-based welfare criterion in this way can not only allow us to justify a general concern with price stability, but can furthermore provide exact answers to the questions raised above about the precise formulation of the appropriate loss function. These answers, of course, depend upon the assumptions we make about the structure of the economy; for example, they depend crucially upon the nature of the nominal rigidities that are present. Insofar as the correct structural relations of our model of the economy remain controversial, the proper welfare criterion to use in evaluating policy will remain controversial as well; and our goal here is more to illustrate a method than to reach final conclusions. But insofar as particular parameter values are found to be empirically justified in that they are required in order for our structural equations to fit historical data, they will contain important information about the proper welfare criterion as well.

1 Approximation of Loss Functions and of Optimal Policies

The method that we shall employ in the analysis below derives a quadratic loss function, that represents a quadratic (second-order Taylor series) approximation to the level of expected utility of the representative household in the rational expectations equilibrium associated with a given policy. There are several reasons for our resort to this approximation. One is simply mathematical convenience; with a quadratic approximation to our objective function and linear approximations to our structural equations, we can address the nature of optimal policy within a linear-quadratic optimal control framework that has been extensively studied, and numerical computation of optimal policy is relatively simple. This convenience is especially great when we turn to questions such as the optimal use of indicator variables under circumstances of partial information.

But there are other advantages as well. One is comparability of our results with those of the traditional literature on monetary policy evaluation, which almost always assumes a quadratic loss function of one sort or another. Casting our own results in this familiar form

allows us to discuss similarities and differences between our utility-based welfare criterion and those assumed in other studies without letting matters be obscured by superficial differences in functional form that may have relatively little consequence for the results obtained.

And finally, it does not make sense to be concerned with a higher-order approximation to our welfare criterion if we do not plan to characterize the effects of alternative policies with a degree of precision sufficient to allow computation of those higher-order terms. In the first part of this study, we have shown how to derive a log-linear approximation to the equilibrium fluctuations in inflation and output under alternative policies, using a log-linear approximation to the exact structural equations of our model. Using this method, we compute the equilibrium fluctuations in these variables only up to a residual of order $\mathcal{O}(\|\xi\|^2)$, where $\|\xi\|$ is a bound on the amplitude of the exogenous disturbances. Given that we do not compute the terms of second order in $\|\xi\|$ in characterizing equilibrium fluctuations, we cannot expect to compute the terms of third or higher order in $\|\xi\|$ in evaluating the expected utility of the representative household.² Of course, one might also wish to undertake a more accurate approximation of the predicted evolution of the endogenous variables under alternative regimes. However, such a study would introduce a large number of additional free parameters, to which numerical values would have to be assigned for purposes of computation; and there is likely to be little empirical basis for the assignment of such values in most cases, given the degree to which the empirical study of macroeconomic time series makes use of linear models.³

However, even a second-order approximation to utility can be computed on the basis

²There is thus no obvious advantage to the approach sometimes adopted in utility-based welfare analyses, such as Ireland (1997) or Collard *et al.* (1998), of evaluating an exact utility function but using a log-linear approximation to the model's structural equations in order to compute the equilibrium.

³A common approach in the quantitative equilibrium business cycle literature, of course, is to assume special functional forms for preferences and technology that allow the higher derivatives of these functions to be inferred from the same small number of parameters as determine the lower-order derivatives, which may then be inferred from first and second moments of the time series alone. This approach often obscures the relation between the properties of the time series and the model parameters that are identified by them, and allows "identifications" that are in fact quite sensitive to the arbitrary functional form assumption. We prefer instead to assume special functional forms as little as possible, but to be clear about the order of approximation that is involved in our calculations.

of a merely linear approximation to our model structural equation only under special circumstances. We shall assume that these hold in our calculations here, but it is important to be clear about the scope of validity of our results. Let x represent a vector of endogenous variables, and suppose that we wish to evaluate $E[U(x; \xi)]$ under alternative policies, where ξ is a vector of random exogenous disturbances, and $U(\cdot; \xi)$ is a concave function for each possible realization of ξ , and at least twice differentiable. Now suppose that we are able to compute a linear (or log-linear) approximation to the equilibrium responses of the endogenous variables, under a given policy regime, of the form

$$x = x^0 + a'\xi + \mathcal{O}(\|\xi\|^2), \quad (1.1)$$

where the vectors of coefficients x^{ss} and a may both depend upon policy. (This represents a first-order Taylor series approximation to the exact equilibrium responses $x(\xi)$, assumed to be nonlinear but differentiable, taken around the mean values $\xi = 0$. The conditions under which the solution to the linearized structural equations yield a valid approximation of this kind to the solution to the exact structural equations are discussed in Appendix A.)

Under the assumption that the constant term x^0 in (1.1) is itself of at most order $\mathcal{O}(\|\xi\|)$, we can take a similar Taylor series expansion of the utility function $U(x)$, and be confident that terms that are of at most order $\mathcal{O}(\|x\|^3)$ are of at most order $\mathcal{O}(\|\xi\|^3)$. We then can write

$$U(x; \xi) = \bar{U} + U_x \tilde{x} + U_\xi \xi + \frac{1}{2} \tilde{x}' U_{xx} \tilde{x} + \tilde{x}' U_{x\xi} \xi + \frac{1}{2} \xi' U_{\xi\xi} \xi + \mathcal{O}(\|\xi\|^3), \quad (1.2)$$

where $\bar{U} \equiv U(\bar{x}; 0)$, $\tilde{x} \equiv x - \bar{x}$, and all partial derivatives of U are evaluated at $(\bar{x}; 0)$. We wish to use approximation (1.1) to the equilibrium fluctuations in x to compute the terms of second or lower order in approximation (1.2) to utility. But the term $U_x \tilde{x}$ in (1.2) will generally contain terms of second order that depend upon the neglected second order terms in (1.1). In order to be able to neglect these terms, we must also assume that $U_x(\bar{x}; 0)$ is at most of order $\mathcal{O}(\|\xi\|)$. In that case, the neglected terms of order $\mathcal{O}(\|\xi\|^2)$ contribute only to terms of order $\mathcal{O}(\|\xi\|^3)$ in $U_x \tilde{x}$.

Assuming this, substitution of (1.1) minus the residual into (1.2) minus the residual yields

a correct quadratic approximation to $U(x; \xi)$. Taking the expected value of this expression, and using the fact that we normalize ξ so that $E(\xi) = 0$, we obtain the approximate welfare criterion

$$E[U] = \bar{U} + U_x E[\tilde{x}] + \frac{1}{2} \text{tr}\{U_{xx} \text{var}[x]\} + \text{tr}\{U_{x\xi} \text{cov}(\xi, \tilde{x})\} + \frac{1}{2} \text{tr}\{U_{\xi\xi} \text{var}[\xi]\} + \mathcal{O}(\|\xi\|^3). \quad (1.3)$$

Here we use the notation $E[z]$ for the expectation of a random vector z , $\text{var}[z]$ denotes the variance-covariance matrix, and $\text{cov}(z_1, z_2)$ the matrix of covariances between two random vectors z_1, z_2 . In expression (1.3) it is understood that the various first and second moments are those that one computes using the linear approximation (1.1).

The validity of this last expression, when the first and second moments are computed using (1.1), thus depends upon two special assumptions. These are that x^0 is only of order $\mathcal{O}(\|\xi\|)$, and that $U_x(\bar{x}; 0)$ is similarly only of order $\mathcal{O}(\|\xi\|)$. Technically, we shall suppose that $\|\xi\|$ is a bound *both* upon the amplitude of the exogenous disturbances, and upon the size of x^0 and $U_x(\bar{x}; 0)$. Our approximation result then refers to a sequence of economies in which $\|\xi\|$ eventually becomes arbitrarily small; as we progress along this sequence, both the distribution of the disturbances, and certain other parameters of the model that determine x^0 and U_x , are varied so as to respect the changing bound, while keeping the specification of the policy rule the same. What the Taylor theorem guarantees is then that if (1.3) minus the residual yields a higher value for one policy than for another, it will be true for all economies far enough out in this sequence that the first policy yields higher expected utility than the other in the equilibrium of the exact model.

The stipulation regarding x^0 is an assumption about the kind of policy regime which we seek to evaluate, while the stipulation regarding $U_x(\bar{x}; 0)$ is an assumption about the point around which we choose to compute the Taylor expansion in (1.3). The latter assumption implies that we expand around a state of affairs \bar{x} that is close to being *optimal*, not simply in the sense of being the best we can do using the set of policies under consideration, but in the sense of being near the maximum of $U(x; 0)$ over all possible values of x .⁴ Of course, it

⁴One way to guarantee this would be to stipulate that \bar{x} is in fact the value that maximizes $U(x; 0)$.

is really only necessary that U_x be small in directions in which it is possible for the average value of x to differ under alternative policies. Thus it is not necessary for households to be nearly satiated both in consumption and in leisure in order for \bar{x} to be optimal in the necessary sense; it is enough that it not be possible to greatly increase utility by varying both consumption and work effort in a way that is feasible given the economy's production function. But it is somewhat delicate to draw conclusions about the directions in which it is possible for policy to vary the second-order terms in x without actually computing the second-order terms in (1.1), and so we prefer to substitute constraints of this kind into our definition of the objective function $U(x)$, and then require all elements of the vector U_x to be small.

Kim and Kim (1999) provide an example of a problem, relating to the welfare gains from risk-sharing, where this requirement for validity of a welfare calculation based upon the linear approximation (1.1) is not satisfied. They consider the expected utility $E[U(C_i)]$ obtained by a household i in each of two cases. In the first case (autarchy), each household consumes its own random income Y_i , while in the second case (perfect risk-sharing), two households pool their incomes, so that $C_i = (Y_1 + Y_2)/2$ for each. Kim and Kim consider the validity of a log-linear approximation to the relation between consumption and income. In the case of autarchy, the log-linear relation (which is exact) is given by

$$\hat{C}_i = \hat{Y}_i,$$

where as usual hatted variables denote deviations of the logs from the value $\log \bar{Y}$ around which one log-linearizes. In the case of perfect risk-sharing, the log-linear approximation is instead

$$\hat{C}_i = (\hat{Y}_1 + \hat{Y}_2)/2.$$

Substitution of these two log-linear expressions into a quadratic approximation to the utility

We do not wish, however, to insist upon this. In some cases, those “first-best” values do not correspond to a possible equilibrium, even in the absence of disturbances. We could linearize the model's structural equations around such values nonetheless, but we prefer to follow convention in linearizing around values that represent a particular equilibrium in the absence of shocks. One advantage of this convention is that our linearized structural equations always have zero constant terms.

function,

$$U(C_i) = U(\bar{Y}) + \bar{Y}U'(\bar{Y}) \left[\hat{C}_i + \frac{1-\gamma}{2}\hat{C}_i^2 \right],$$

where $\gamma > 0$ is the coefficient of relative risk aversion (evaluated at consumption level \bar{Y}), does not yield a correct quadratic approximation to utility as a function of \hat{Y}_1 and \hat{Y}_2 . Indeed, if $0 < \gamma < 1$, this approximation implies that expected utility is higher under autarchy.

The problem is that the partial derivative of U with respect to $\log C_i$ (which is equal to $\bar{Y}U'(\bar{Y})$) is non-zero, so that the correct quadratic approximation to expected utility in the case of risk-sharing involves quadratic terms in the Taylor series expansion for \hat{C}_i . (The omitted terms raise expected utility in that case, since less variable consumption means a higher expected value for log consumption, by Jensen's inequality.) It does not make sense to assume that this derivative can be made arbitrarily close to zero as we make the bound $\|\xi\|$ on the amplitude of income variations smaller, either. This could be done only by varying preferences and/or average income as we make $\|\xi\|$ smaller in such a way that proportional variations in consumption cease to matter much; but this would mean that in the limit, no comparisons between alternative consumption processes would be possible. This difficulty does not arise in the case of our analysis of the welfare gains from macroeconomic stabilization, below, as long as we log-linearize around a level of economic activity (in each sector) that is sufficiently near to being efficient; since an interior optimum does exist in our case, this is possible. But it is important that we check that the derivatives in question are indeed small, under our assumptions, and that the qualifications that this requires to our results be noted.

The assumption that x^0 is small means that the policies considered are all ones with the property that in the absence of shocks, the equilibrium value of x would be near the linearization point \bar{x} — or alternatively, that in equilibrium the *mean* value of x is near \bar{x} . Given our assumption about \bar{x} , this means that the policies considered are ones under which the equilibrium value x is nearly *optimal*, in the sense discussed above. As we are primarily interested in whether our approximate welfare criterion correctly identifies the optimal policy, the essential requirement is that our model (and the family of available policies) be such

that the best available policy can achieve an outcome that is sufficiently close to being fully optimal. Thus the *unavoidable* frictions – the ones that cannot be ameliorated through an appropriate choice of policy – must be small, even if there exist frictions that imply that outcomes under a *bad* policy could be significantly worse.

In the case of the class of monetary models treated below, the only frictions that prevent equilibrium from being efficient are (i) the market power possessed by suppliers of goods, as a result of monopolistic competition, and (ii) the failure to adjust all goods prices each period. We log-linearize our structural equations around the steady state with zero inflation each period, that represents a possible equilibrium in the absence of real disturbances. In this equilibrium, the failure to adjust prices constantly results in no distortion of the allocation of resources; this allocation is thus nearly optimal as long as the distortion due to market power is sufficiently small. In our calculations below, we assume that it is.

The other assumption required for the validity of the quadratic approximation obtained from our log-linear structural equations is that the policy rules considered be ones under which the equilibrium rate of inflation in the absence of shocks would in fact be near zero – or alternatively, that these policies be ones under which the average rate of inflation is low. This also is assumed; since we conclude that it is optimal for a country with characteristics like those of the U.S. to choose a policy under which the average rate of inflation is slightly, but only slightly, positive, this last assumption is relatively innocuous. However, it is important to realize that under other circumstances – say, an analysis of optimal monetary policy in the presence of a need for significant seignorage revenues – this assumption as well might be inappropriate.

Finally, it is important to note that the conditions required for validity of a quadratic approximation to welfare obtained from log-linear approximations to the structural equations do not relate solely to the structure of the economy; it also matters in which form we choose to express our approximate loss function. Alternative quadratic approximations to U , each equally valid second-order Taylor series expansions (but in terms of different variables), may not yield equally valid approximations to welfare when evaluated using a log-linear solution

(1.1) for the model's endogenous variables.

For example, consider a model like that of the next section. Each period's contribution to the utility of the representative household can be approximated by an expression of the form

$$U = a\hat{C} - b\hat{H} + Q_1 + R_1, \quad (1.4)$$

where \hat{C} , \hat{H} denote the percentage deviations in consumption and hours worked respectively, Q_1 is a set of quadratic terms in the log deviations, and the residual R_1 contains terms of third or higher order, or terms that are independent of the policy chosen, that can be neglected. Alternatively, one may eliminate hours using the necessary relation between aggregate consumption and aggregate hours implied by the production function, and obtain a Taylor series expansion of the form

$$U = c\hat{C} + Q_2 + R_2, \quad (1.5)$$

where Q_2 , R_2 are other quadratic terms and residual. Under the assumptions just described, the coefficient c is of order $\mathcal{O}(\|\xi\|)$. It follows that substitution into (1.5) of the solution to our log-linear structural equations yields a valid second-order approximation to utility.

However, the same is not true of substitution of the same solution into (1.4). The structural equations include a production-function relation between \hat{C} and \hat{H} , of the form

$$\hat{C} = f\hat{H} + Q_3 + R_3, \quad (1.6)$$

and one of the log-linear structural relations is given by the linear terms in this. In approximation (1.5), $c = a - bf^{-1}$, so that the first-order terms in the two approximations would have the same value. But the terms Q_2 are not equal to Q_1 , because of the presence of non-zero quadratic terms Q_3 in (1.6).⁵ Hence (1.4) will not yield a correct second-order approximation to welfare, if one substitutes the solutions for \hat{C} and \hat{H} implied by the log-linear structural relations, including the linear part of (1.6). In terms of the criterion set

⁵Even if the production function is of a constant-elasticity form, the log-linear approximation (1.6) will contain non-zero quadratic terms in the event of variations in government purchases.

out above, substitution of the log-linear approximate solutions into (1.4) yields an incorrect result because the coefficients a and b are not individually of order $\mathcal{O}(\|\xi\|)$, even though the linear combination c is. Thus it is not enough that one expand around a near-optimal equilibrium; the expansion must be written in a form that contains no first-order terms that do not involve coefficients of order $\mathcal{O}(\|\xi\|)$.

2 A Utility-Based Welfare Criterion

We turn now to the computation of a utility-based approximate welfare criterion, of the kind discussed in the previous section, for a class of optimizing monetary models with sticky prices. Each of the models is characterized by a representative household and monopolistically competitive supply of a continuum of differentiated goods. We abstract from monetary frictions,⁶ endogenous variations in productive capacity, factor market inefficiencies (such as sticky wages), and disturbances with asymmetric effects upon the demand for or cost of production of different goods, though we offer remarks in the final section on the extension of our methods to more general cases. The microeconomic foundations of these models, and the derivation of the aggregate supply relations resulting from alternative assumptions regarding the timing of price changes, are presented in greater detail in Woodford (2000).⁷

The natural welfare criterion in our models is the level of expected utility of the representative household. This can be written

$$E \left\{ \sum_{t=0}^{\infty} \beta^t U_t \right\}, \quad (2.1)$$

where U_t is the period t contribution to utility from the consumption and supply of the various differentiated goods. Equating the consumption of the representative household with aggregate production,⁸ and expressing labor effort as a function of quantity produced

⁶The welfare criteria derived here thus apply to the “cashless limit” of a monetary economy, in the sense introduced in Woodford (1998, 1999a). The modification required when monetary frictions are non-negligible is discussed in Woodford (1999c).

⁷Rotemberg and Woodford (1997, 1999) apply a similar, though not identical, method to a variant model with in the context of a variant model with additional decision lags. The method expounded here is applied to other more complex environments in Amato and Laubach (2000), Aoki (1999), Benigno (1999), Erceg *et al.* (1999), and Steinsson (2000).

using the production function, we can write U_t as a function of the quantity produced in period t of the various goods. We assume the specific form

$$U_t = u(Y_t; \xi_t) - \int_0^1 \tilde{v}(y_t(i); \tilde{\xi}_t) di, \quad (2.2)$$

where

$$Y_t \equiv \left[\int_0^1 y_t(i)^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}} \quad (2.3)$$

is a Dixit-Stiglitz index of aggregate demand, in which $y_t(i)$ is production (and consumption) in period t of differentiated good i ,⁹ and ξ_t is a vector of preference shocks.

The function $\tilde{v}(y; \tilde{\xi})$ indicates the disutility of supplying quantity y . If we assume a “yeoman farmer” model, as in Rotemberg and Woodford (1997, 1999), this can be interpreted directly as the household’s disutility of supplying output. Alternatively, if we assume firms and a labor market, as in Woodford (2000a), we can define

$$\tilde{v}(y; \tilde{\xi}) \equiv v(f^{-1}(y/A); \xi), \quad (2.4)$$

where $v(h; \xi)$ is the disutility of working h hours in any given production activity, and $Af(h)$ is the output produced using that labor input. Here $\tilde{\xi} \equiv (\xi, a)$ denotes the complete vector of exogenous disturbances, including both the preference shocks ξ ¹⁰ and the technology shock $a \equiv \log A$.¹¹ We assume that u is an increasing, concave function of Y for each possible value of ξ , while \tilde{v} is an increasing, convex function of y for each possible value of $\tilde{\xi}$. Thus (2.2) implies that U_t is a concave function of the entire vector of levels of production of the various goods.

⁸Recall that we can interpret the model as allowing for government purchases, treating these as a shift term in the utility function $u(Y; \xi)$.

⁹If there is trend growth in productivity and output, the variables Y_t , $y_t(i)$, and A_t (introduced in the next paragraph) should all be interpreted as having been deflated by a common exponential trend factor, to render them stationary.

¹⁰Note that the vector ξ contains many elements, so that the disturbances to the utility of consumption may or may not be correlated with the disturbances to the disutility of working.

¹¹We assume a normalization of the productivity measure A such that the unconditional expectation of a is zero. We include a rather than A in our definition of $\tilde{\xi}$ because it is a rather than A that is assumed to be always sufficiently close to zero (in order for our Taylor series approximation to be accurate), and because we wish to approximate the production-function relationship by one that is linear in a rather than linear in A .

Expression (2.2) assumes that the representative household produces and consumes all goods, though distinct types of effort (each with its own increasing marginal disutility) are required to produce each good. Nonetheless, we assume that there is no coordination of the pricing decisions of the suppliers of the different goods. One interpretation here is that separate firms supply the various goods, though the representative household supplies the appropriate type of labor to each of the firms. Alternatively, we may assume (as in standard “yeoman farmer” models) that each household specializes in the production of a single good, so that different goods are supplied by different households, but that households insure one another against the income risk associated with the differential effects of shocks upon the producers of different goods (that have set their prices at different times). Perfect risk sharing (the equilibrium with complete financial markets) then implies aggregate supply and demand decisions that are the same as if a representative household chooses for all of them, maximizing the average utility of the producers of the various goods subject to a pooled budget constraint. The objective of the stand-in representative household in the latter case will again be of the form (2.2). (Our interest in this latter interpretation explains the assumption of additive separability of the disutility of supplying the various goods.)

Our Taylor series expansions group terms of different powers in the elements of $\tilde{\xi}$, though we shall continue to use the notation $\|\xi\|$ for the bound on the magnitude of the entire vector of disturbances.

2.1 Output-Gap Stability and Welfare

It will be useful to express welfare as a function of the gap between output in each sector and the *natural rate* of output, by which we mean the equilibrium level of output under full price flexibility.¹² Note that the real marginal cost function of supplying any good i is given by

$$s(y(i), Y; \tilde{\xi}) = \frac{\tilde{v}_y(y(i); \tilde{\xi})}{u_c(Y; \xi)}. \quad (2.5)$$

¹²Under our symmetry assumptions, under flexible prices the equilibrium output of all goods would be the same, and so the *aggregate* natural rate also applies to each good individually. More generally, of course, there would be a good-specific natural rate.

The natural rate of output $Y_t^n = Y^n(\tilde{\xi}_t)$ is then defined implicitly by

$$s(Y_t^n, Y_t^n; \tilde{\xi}_t) = \frac{1 - \tau}{\mu} \equiv 1 - \Phi. \quad (2.6)$$

Here $\tau < 1$ is the constant proportional tax rate on sales proceeds (negative in the case of a subsidy), and $\mu \equiv \theta/(\theta - 1) > 1$ is the desired markup as a result of suppliers' market power under monopolistic competition. The parameter Φ then summarizes the overall distortion in the level of output that would exist under flexible prices, as a result of both taxes and market power.

It is plainly realistic to assume that $\Phi > 0$. However, we shall assume that Φ is small, specifically of order $\mathcal{O}(\|\xi\|)$. This is the assumption of near-efficiency of the steady state level of output with zero inflation that is made in order to allow us to use our log-linear approximations to the model structural equations in welfare comparisons, as explained in the previous section. Note that the introduction of the distorting tax rate τ allows us to contemplate a series of economies in which Φ is made progressively smaller, without this having to involve any change in the size of θ , a parameter that also affects the coefficients of the log-linearized equilibrium conditions.¹³

It might seem more natural to express welfare in terms of the gap between actual output and the *efficient* level of output, $Y^*(\tilde{\xi}_t)$, given by the solution to (2.6) when $\Phi = 0$. However, under our assumption that $\Phi = \mathcal{O}(\|\xi\|)$, we observe that

$$\log(Y_t^n/Y_t^*) = -(\omega + \sigma^{-1})^{-1}\Phi + \mathcal{O}(\|\xi\|^2), \quad (2.7)$$

where

$$\omega \equiv \frac{bar Y \tilde{v}_{yy}}{\tilde{v}_y} > 0$$

¹³Rotemberg and Woodford (1997, 1999) instead assume that τ is of exactly the (negative) size required to offset the distortion due to market power, so that $\Phi = 0$. The intention is to consider optimal monetary stabilization policy as part of a broader analysis of optimal policy, in which another instrument (tax policy) is assigned responsibility for achieving the optimal average level of economic activity, while monetary policy is used to ameliorate the economy's response to shocks. However, it is clear that monetary policy must in practice be chosen in an environment in which such an output subsidy does not, and probably cannot, exist. Furthermore, the fact that the "natural rate" of output is inefficiently low is of importance for certain issues, notably the inflationary bias associated with discretionary policymaking. Hence we here allow for $\Phi > 0$, while still assuming that Φ is of order $\mathcal{O}(\|\xi\|)$.

is the elasticity of real marginal cost with respect to a firm's own output, and

$$\sigma^{-1} \equiv - \frac{\bar{Y} u_{cc}}{u_c} > 0$$

is the elasticity with respect to aggregate output.¹⁴ Here $\bar{Y} \equiv Y^n(0)$ is the steady-state level of output under flexible prices (and also the steady-state level of output associated with zero inflation, in each of the sticky-price models considered here), and both elasticities are evaluated at $(\bar{Y}; 0)$, which is also the point around which we expand our Taylor series below. That is, the percentage difference between Y_t^n and Y_t^* is independent of the disturbances, to a first-order approximation. Thus the output gap relative to the natural rate differs from the gap relative to the efficient level only by a constant, if we neglect terms of order $\mathcal{O}(\|\xi\|^2)$. This accounts for the role of the output gap relative to the natural rate (a variable that plays a crucial role in the aggregate supply relations presented below, for reasons developed in Woodford, 2000) in our welfare expansions.

We now proceed to compute a quadratic Taylor series approximation to (2.2). The first term can be approximated as

$$\begin{aligned} u(Y_t; \xi_t) &= \bar{u} + u_c \tilde{Y}_t + u_\xi \xi_t + \frac{1}{2} u_{cc} \tilde{Y}_t^2 + u_{c\xi} \xi_t \tilde{Y}_t + \frac{1}{2} \xi_t' u_{\xi\xi} \xi_t + \mathcal{O}(\|\xi\|^3) \\ &= \bar{u} + \bar{Y} u_c \cdot (\hat{Y}_t + \frac{1}{2} \hat{Y}_t^2) + u_\xi \xi_t + \frac{1}{2} \bar{Y}^2 u_{cc} \hat{Y}_t^2 \\ &\quad + \bar{Y} u_{c\xi} \xi_t \hat{Y}_t + \frac{1}{2} \xi_t' u_{\xi\xi} \xi_t + \mathcal{O}(\|\xi\|^3) \\ &= \bar{Y} u_c \hat{Y}_t + \frac{1}{2} [\bar{Y} u_c + \bar{Y}^2 u_{cc}] \hat{Y}_t^2 - \bar{Y}^2 u_{cc} g_t \hat{Y}_t + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3) \\ &= \bar{Y} u_c \left\{ \hat{Y}_t + \frac{1}{2} (1 - \sigma^{-1}) \hat{Y}_t^2 + \sigma^{-1} g_t \hat{Y}_t \right\} + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \end{aligned} \quad (2.8)$$

Here the first line represents the usual Taylor expansion, in which $\bar{u} \equiv u(\bar{Y}; 0)$ and $\tilde{Y}_t \equiv Y_t - \bar{Y}$, and we assume that the fluctuations in \tilde{Y}_t are only of order $\mathcal{O}(\|\xi\|)$. The second line substitutes for \tilde{Y}_t in terms of $\hat{Y}_t \equiv \log(Y_t/\bar{Y})$, using the Taylor series expansion

$$Y_t/\bar{Y} = 1 + \hat{Y}_t + \frac{1}{2} \hat{Y}_t^2 + \mathcal{O}(\|\xi\|^3).$$

¹⁴This latter quantity is also the reciprocal of σ , the intertemporal elasticity of substitution of private expenditure.

The third line collects together in the term “t.i.p.” all of the terms that are independent of policy, as they involve only constants and exogenous variables, and introduces the notation

$$g_t \equiv -\frac{u_{c\xi}\xi_t}{\bar{Y}u_{cc}}$$

for the percentage variation in output required in order to keep the marginal utility of expenditure u_c at its steady-state level, given the preference shock. The final line collects terms in a useful way; note that the only part of this expression that differs across policies is the expression inside the curly braces.

We may similarly approximate $\tilde{v}(y_t(i); \xi_t)$ by

$$\begin{aligned}\tilde{v}(y_t(i); \xi_t) &= \bar{Y}\tilde{v}_y \left\{ \hat{y}_t(i) + \frac{1}{2}(1 + \omega)\hat{y}_t(i)^2 - \omega q_t \hat{y}_t(i) \right\} + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3) \\ &= \bar{Y}u_c \left\{ (1 - \Phi)\hat{y}_t(i) + \frac{1}{2}(1 + \omega)\hat{y}_t(i)^2 - \omega q_t \hat{y}_t(i) \right\} + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3),\end{aligned}$$

where $\hat{y}_t(i) \equiv \log(y_t(i)/\bar{Y})$, and

$$q_t \equiv -\frac{\tilde{v}_{y\xi}\xi_t}{\bar{Y}\tilde{v}_{yy}}$$

is the percentage variation in output required to keep the marginal disutility of supply \tilde{v}_y at its steady-state level, given the preference shock. The second line uses (2.5) and (2.6) to replace \tilde{v}_y by $(1 - \Phi)u_c$, and the assumption that Φ is of order $\mathcal{O}(\|\xi\|)$ to simplify. Note that the term premultiplying the expression in curly braces is now the same as in (2.8).

Integrating this expression over the differentiated goods i , we obtain

$$\begin{aligned}\int_0^1 \tilde{v}(y_t(i); \xi_t) &= \bar{Y}u_c \left\{ (1 - \Phi)E_i\hat{y}_t(i) + \frac{1}{2}(1 + \omega)[(E_i\hat{y}_t(i))^2 + \text{var}_i\hat{y}_t(i)] - \omega q_t E_i\hat{y}_t(i) \right\} \\ &\quad + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3) \\ &= \bar{Y}u_c \left\{ (1 - \Phi)\hat{Y}_t + \frac{1}{2}(1 + \omega)\hat{Y}_t^2 - \omega q_t \hat{Y}_t + \frac{1}{2}(\theta^{-1} + \omega)\text{var}_i\hat{y}_t(i) \right\} \\ &\quad + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3),\end{aligned}\tag{2.9}$$

using the notation $E_i\hat{y}_t(i)$ for the mean value of $\hat{y}_t(i)$ across all differentiated goods at date t , and $\text{var}_i\hat{y}_t(i)$ for the corresponding variance. In the second line, we use the Taylor series approximation to (2.3),

$$\hat{Y}_t = E_i\hat{y}_t(i) + \frac{1}{2}(1 - \theta^{-1})\text{var}_i\hat{y}_t(i) + \mathcal{O}(\|\xi\|^3),$$

to eliminate $E_i \hat{y}_t(i)$.

Combining (2.8) and (2.9), we finally obtain

$$\begin{aligned} U_t &= \bar{Y} u_c \left\{ \Phi \hat{Y}_t - \frac{1}{2}(\sigma^{-1} + \omega) \hat{Y}_t^2 + (\sigma^{-1} g_t + \omega q_t) \hat{Y}_t - \frac{1}{2}(\theta^{-1} + \omega) \text{var}_i \hat{y}_t(i) \right\} \\ &\quad + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3) \\ &= -\frac{\bar{Y} u_c}{2} \left\{ (\sigma^{-1} + \omega)(x_t - x^*)^2 + (\theta^{-1} + \omega) \text{var}_i \hat{y}_t(i) \right\} + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \end{aligned} \quad (2.10)$$

Here the second line rewrites the expression in terms of the *output gap* $x_t \equiv \hat{Y}_t - \hat{Y}_t^n$, where \hat{Y}_t^n denotes the deviation of the (log of the) *natural rate* of output from its steady-state level, or

$$\hat{Y}_t^n \equiv \log(Y_t^n / \bar{Y}) = \frac{\sigma^{-1} g_t + \omega q_t}{\sigma^{-1} + \omega}, \quad (2.11)$$

and in terms of the efficient level of the output gap, $x^* \equiv \log(Y^* / \bar{Y})$. (This expression for \hat{Y}_t^n follows from log-linearization of (2.6), while the expression for x^* follows from (2.7). Note that if Φ is positive and of order $\mathcal{O}(\|\xi\|)$, the same is true of x^* .)

Expression (2.10) represents a quadratic approximation to (2.2), under the assumption that Φ (and hence the inefficiency of the steady-state level of output) is of order $\mathcal{O}(\|\xi\|)$. It is interesting to observe that the preference and technology shocks $\tilde{\xi}_t$ matter, in this approximation, only through their effects upon a single exogenous state variable, the natural rate of output \hat{Y}_t^n . Furthermore, output variability as such does not matter for our utility-based welfare criterion; rather, it is the variability of the *output gap* that matters, and the measure of potential output with respect to which the gap should be measured for purposes of the welfare criterion is the *same* “natural rate” of output that (as shown in Woodford, 2000) determines the short-run relation between output and inflation. Thus we can already offer an answer to one question posed in the introduction to this paper: it is the output gap, rather than output relative to trend, that one should seek to stabilize, and (if distortions are small enough) the relevant output gap is the same one that appears in the short-run aggregate supply curve.

However, (2.10) implies that stabilization of the output gap should not be the sole concern of policy, since the dispersion of output levels across sectors matters as well.¹⁵ In fact, in

our baseline framework, there is no reason for equilibrium output to be different for different goods except as a result of relative price distortions that result from sticky prices in an environment where the overall price level is unstable. It is through this channel that price stability turns out to be relevant for welfare, in a way that goes beyond the mere association between inflation and the level of the aggregate output gap.

Specifically, our assumed CES (Dixit-Stiglitz) preferences over differentiated goods imply that each supplier faces a constant-elasticity demand curve of the form

$$\log y_t(i) = \log Y_t - \theta(\log p_t(i) - \log P_t).$$

It follows from this that

$$\text{var}_i \log y_t(i) = \theta^2 \text{var}_i \log p_t(i),$$

so that (2.10) may equivalently be written

$$U_t = -\frac{\bar{Y}u_c}{2} \left\{ (\sigma^{-1} + \omega)(x_t - x^*)^2 + \theta(1 + \omega\theta)\text{var}_i \log p_t(i) \right\} + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \quad (2.12)$$

Thus we find that, in addition to stabilization of the output gap, it is also appropriate for policy to aim to reduce price dispersion. In our framework, this is achieved by stabilizing the general price level; but the exact way in which fluctuations in the general price level affect price dispersion, and hence welfare, depend upon the details of price-setting.

2.2 Inflation and Relative-Price Distortions

The approximation (2.12) to the utility of the representative household applies to any model with no frictions other than those due to monopolistic competition and sticky prices, regardless of the nature of the delays involved in price-setting. The relation between the price dispersion term and the stability of the general price level depends, instead upon the details

¹⁵More generally, it is the dispersion of output gaps across sectors that matters, along with the aggregate output gap. We here assume that the only disturbances $\tilde{\xi}_t$ that affect the natural rate of output have identical effects upon all sectors, so that the dispersion of output gaps across sectors is identical to the dispersion of output levels. In section 3.4 below, we briefly discuss the consequences of allowing for shocks with asymmetric effects on different sectors.

of price-setting. Here we do not attempt a general treatment, but illustrate the form of the relation in three simple examples.

As a first example, we consider the case of an economy in which a fraction $0 < \gamma < 1$ of goods prices are fully flexible, while the remaining $1 - \gamma$ must be fixed a period in advance. In such an economy, as shown in Woodford (2000a), the aggregate supply relation takes the familiar “New Classical” form

$$\pi_t = \kappa x_t + E_{t-1} \pi_t, \quad (2.13)$$

where the slope coefficient is given by

$$\kappa \equiv \frac{\gamma}{1 - \gamma} \frac{\sigma^{-1} + \omega}{1 + \omega\theta} > 0.$$

In this model, in any period all flexible-price goods have the same price, p_t^1 , and all sticky-price goods have the same price, p_t^2 , which satisfies

$$\log p_t^2 = E_{t-1} \log p_t^1 + \mathcal{O}(\|\xi\|^2). \quad (2.14)$$

The Dixit-Stiglitz price index furthermore satisfies

$$\log P_t = \gamma \log p_t^1 + (1 - \gamma) \log p_t^2 + \mathcal{O}(\|\xi\|^2),$$

so that

$$\begin{aligned} \pi_t - E_{t-1} \pi_t &= \gamma [\log p_t^1 - E_{t-1} \log p_t^1] + \mathcal{O}(\|\xi\|^2) \\ &= \gamma [\log p_t^1 - \log p_t^2] + \mathcal{O}(\|\xi\|^2), \end{aligned}$$

using (2.14). It follows that under this assumption about pricing,

$$\begin{aligned} \text{var}_i \log p_t(i) &= \gamma(1 - \gamma)(\log p_t^1 - \log p_t^2)^2 \\ &= \frac{1 - \gamma}{\gamma} (\pi_t - E_{t-1} \pi_t)^2. \end{aligned}$$

As asserted above, equilibrium price dispersion is closely connected with the stability of the general price level; but in this special case, it is only the volatility of the *unexpected component* of inflation that matters.

Substituting this expression into (2.12), we obtain

$$U_t = -\Omega L_t + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3),$$

where Ω is a positive constant and L_t is a quadratic loss function of the form

$$L_t = (\pi_t - E_{t-1}\pi_t)^2 + \lambda(x_t - x^*)^2, \quad (2.15)$$

with a relative weight on output gap variability of $\lambda = \kappa/\theta$. We thus obtain precise conclusions regarding both the sense in which aggregate output and inflation variations matter for welfare (it is the output *gap* that matters, and the *unexpected* component of inflation), and the relative weight that should be placed upon the two concerns (the relative weight on output gap variations is proportional to the slope κ of the short-run Phillips curve).

In fact, in the context of this model, there is no tension between the goals represented by the two terms of (2.15). For (2.13) implies that the output gap is itself proportional to the surprise component of inflation. Thus we can simplify (2.15) further, and say that the sole goal of policy should be to minimize the variability of unexpected inflation, or alternatively, that the sole goal should be to stabilize the output gap (when properly measured).¹⁶

While we obtain a simple result in this case, the model is not a very realistic one, since, as is well known, it is unable to account for the persistence of the observed output effects of monetary disturbances. Let us consider instead, then, the consequences of the kind of staggered pricing assumed in another popular model, a discrete-time version of the Calvo (1983) pricing model. In this model, a fraction $0 < \alpha < 1$ of all prices remain unchanged each period, with the probability of a price change assumed to be independent of both the length of time since the price was last changed and of the degree to which that good's price is out of line with others. This implies that each period, the distribution of prices $\{p_t(i)\}$ consists of α times the distribution of prices in the previous period, plus an atom of size

¹⁶However, if we allow for disturbances to the short-run aggregate supply relation (2.13) that – unlike the preference, technology, or government-purchase shocks considered in section 3.2 below – do not shift the efficient level of output to the same extent, then the loss function (2.15) would still be correct, while the output gap that appears in this formula would no longer coincide perfectly with unexpected inflation. In that extension of the model, it would be quite important to know the correct relative weight λ to place on output-gap variations.

$(1 - \alpha)$ at the price p_t^* that is chosen at date t by all suppliers who choose a new price at that date. As shown in Woodford (2000a), the aggregate supply relation takes in this case the “New Keynesian” form

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1}, \quad (2.16)$$

where now the slope coefficient is given by

$$\kappa \equiv \frac{(1 - \alpha)(1 - \alpha\beta)(\sigma^{-1} + \omega)}{\alpha} \frac{1}{1 + \omega\theta} > 0. \quad (2.17)$$

Letting

$$\bar{P}_t \equiv E_i \log p_t(i), \quad \Delta_t \equiv \text{var}_i \log p_t(i),$$

we observe from the above recursive characterization of the distribution of prices at date t that

$$\begin{aligned} \bar{P}_t - \bar{P}_{t-1} &= E_i[\log p_t(i) - \bar{P}_{t-1}] \\ &= \alpha E_i[\log p_{t-1}(i) - \bar{P}_{t-1}] + (1 - \alpha)(\log p_t^* - \bar{P}_{t-1}) \\ &= (1 - \alpha)(\log p_t^* - \bar{P}_{t-1}). \end{aligned}$$

Similar reasoning about the dispersion measure Δ_t yields

$$\begin{aligned} \Delta_t &= \text{var}_i[\log p_t(i) - \bar{P}_{t-1}] \\ &= E_i\{[\log p_t(i) - \bar{P}_{t-1}]^2\} - (E_i \log p_t(i) - \bar{P}_{t-1})^2 \\ &= \alpha E_i\{[\log p_{t-1}(i) - \bar{P}_{t-1}]^2\} + (1 - \alpha)(\log p_t^* - \bar{P}_{t-1})^2 - (\bar{P}_t - \bar{P}_{t-1})^2 \\ &= \alpha \Delta_{t-1} + \frac{\alpha}{1 - \alpha} (\bar{P}_t - \bar{P}_{t-1})^2. \end{aligned}$$

Finally, substituting the log-linear approximation

$$\bar{P}_t = \log P_t + \mathcal{O}(\|\xi\|^2),$$

we obtain

$$\Delta_t = \alpha \Delta_{t-1} + \frac{\alpha}{1 - \alpha} \pi_t^2 + \mathcal{O}(\|\xi\|^3) \quad (2.18)$$

as a law of motion for the dispersion of prices. Note that price dispersion is again a function of the degree of instability of the general price level, though now the relation is a dynamic one. Note also that under this assumptions about pricing, both expected and unexpected inflation contribute *equally* to increases in price dispersion.

Integrating forward (2.18) starting from any initial degree of price dispersion Δ_{-1} in the period before the first period for which a new policy is contemplated, the degree of price dispersion in any period $t \geq 0$ under the new policy will be given by

$$\Delta_t = \alpha^{t+1} \Delta_{-1} + \sum_{s=0}^t \alpha^{t-s} \left(\frac{\alpha}{1-\alpha} \right) \pi_s^2 + \mathcal{O}(\|\xi\|^3).$$

Note that the first term will be independent of the policy that is chosen to apply in periods $t \geq 0$. Thus if we take the discounted value of these terms over all periods $t \geq 0$, we obtain

$$\sum_{t=0}^{\infty} \beta^t \Delta_t = \frac{\alpha}{(1-\alpha)(1-\alpha\beta)} \sum_{t=0}^{\infty} \beta^t \pi_t^2 + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3).$$

Substitution of this in turn into (2.12), we find that

$$\sum_{t=0}^{\infty} \beta^t U_t = -\Omega \sum_{t=0}^{\infty} \beta^t L_t + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3), \quad (2.19)$$

where in this case the normalized quadratic loss function is given by

$$L_t = \pi_t^2 + \lambda(x_t - x^*)^2. \quad (2.20)$$

Here the relative weight on output gap variability is again given by $\lambda = \kappa/\theta$, but now the value of κ referred to is that given in (2.17).¹⁷

The loss function (2.20) is in fact of a form widely assumed in the literature on monetary policy evaluation (and also in positive models of central bank behavior).¹⁸ Here, however, we

¹⁷Note that the values of Ω and λ obtained here are slightly different from those that follow from the derivation presented in Rotemberg and Woodford (1999). The reason is that we are here interested in approximating the expected value of the discounted sum of utilities, conditioning upon the pre-existing degree of price dispersion at date -1 , whereas they compute an unconditional expectation. Note that the loss measure that we compute here, for a given policy, will not depend upon the initial price dispersion Δ_{-1} . Nonetheless, it matters whether one conditions upon the value of Δ_{-1} in computing the expected utility. Computing the unconditional expectation, rather than conditioning upon the value of Δ_{-1} , penalizes policies that lead to higher average price dispersion also for the higher average value assumed for Δ_{-1} if one integrates over the unconditional distribution of values for Δ_t associated with a given stationary equilibrium.

¹⁸See, e.g., Walsh (1998, chap. 8) and Clarida *et al.* (1999).

are able to present a theoretical justification for the attention to variations in inflation (rather than, say, variations in the price level), as well as for the common assumption that inflation variations are equally costly whether forecastable or not, in terms of the relative-price distortions resulting from price-level instability in the Calvo model of staggered price-setting. We are also able to derive an optimal rate of inflation with respect to which deviations should be measured (namely, zero, as it is in this case that no relative-price distortions result from imperfect synchronization of price changes). And finally, we are again able to derive an optimal relative weight upon output-gap variation as opposed to inflation variation; this depends upon model parameters, but in a way that makes an estimate of the slope of the short-run aggregate supply curve directly informative about the proper size of this weight.¹⁹

As is discussed further in Woodford (1999c), the estimate of the slope of the short-run aggregate supply curve for the U.S. of Rotemberg and Woodford (1997) implies a value for λ_x on the order of .05, if the output gap is measured in percentage points and inflation is measured as an annualized percentage rate. This value is much lower than the value $\lambda_x = 1$ often assumed in the literature on evaluation of monetary policy rules, on a ground such as “giving equal weight to inflation and output” as stabilization objectives.²⁰ Our utility-based analysis implies instead that if one assumes the degree of price stickiness that is needed to account for the persistence of the real effects of monetary policy shocks, the distortions associated with inflation are more important than those associated with variation in the aggregate output gap.

Alternative assumptions about the timing of price changes would justify still other loss functions. As a single further example, suppose that a fraction $1 - \alpha$ of all goods change their prices each period, and that these are randomly chosen, but that among these, a fraction γ

¹⁹The size of this weight is of greater interest in the case of this model, since aggregate supply relation (2.16) does not imply that inflation and the output gap should perfectly co-vary under most circumstances. It is true that complete stabilization of one implies complete stabilization of the other, as we discuss further in the next section, and in this sense there is no tension between the two goals if (2.16) holds. But it may not be possible to achieve complete stabilization, e.g., because of the zero lower bound on nominal interest rates, or informational restrictions upon feasible policies; and in such cases optimal policy will generally depend upon the relative weight placed upon the two goals.

²⁰See, e.g., Rudebusch and Svensson (1999) and Williams (1999).

choose the new price that takes effect in period t at that time, while the remaining $1 - \gamma$ adopt a new price that was chosen in period $t - 1$ (or at any rate, using only public information as of date $t - 1$). As shown in Woodford (2000a), this assumption leads to an aggregate supply relation of the form

$$\pi_t = (1 - \psi)E_{t-1}\pi_t + \psi[\kappa x_t + \beta E_t \pi_{t+1}], \quad (2.21)$$

where $\psi \equiv \gamma\alpha/(1 - \gamma(1 - \alpha))$ is a positive fraction and κ is the same positive coefficient as in (2.16). Note that this aggregate supply relation reduces to the “New Classical” specification (2.13) in the limit as $\alpha \rightarrow 0$ (in which case $\psi \rightarrow 0$ but $\psi\kappa$ approaches a positive limit), and to the “New Keynesian” specification (2.16) in the limit as $\gamma \rightarrow 1$ (in which case $\psi \rightarrow 1$).

Under these assumptions, the distribution of prices in any period t is equal to α times the distribution in period $t - 1$, plus an atom of size $\gamma(1 - \alpha)$ at the new price p_t^1 charged by all suppliers who have a new price in period t and no decision lag, and another atom of size $(1 - \gamma)(1 - \alpha)$ at the new price p_t^2 charged by all suppliers with a new price in period t who are subject to the one-period decision lag. The two new prices are again related by (2.14), even though p_t^1 no longer corresponds to the price that would be chosen by a flexible-price supplier. A straightforward extension of the above calculations²¹ shows that in this case (2.18) generalizes to

$$\Delta_t = \alpha\Delta_{t-1} + \frac{\alpha}{1 - \alpha} \left[\pi_t^2 + \left(\frac{1 - \psi}{\psi} \right) (\pi_t - E_{t-1}\pi_t)^2 \right] + \mathcal{O}(\|\xi\|^3). \quad (2.22)$$

We thus obtain a quadratic approximation to our welfare criterion in which the normalized period loss function is now given by

$$L_t = \pi_t^2 + \lambda_u(\pi_t - E_{t-1}\pi_t)^2 + \lambda_x(x_t - x^*)^2, \quad (2.23)$$

where $\lambda_u = \psi^{-1} - 1$ and once again $\lambda_x = \kappa/\theta$. In this case, variations in inflation create distortions whether they are forecastable or not, but there are additional distortions associated with the unforecastable component of inflation, so that there is an additional positive weight on the squared inflation surprise, as in the criterion of Rotemberg and Woodford (1997, 1999).²²

²¹For details, see the appendix to Rotemberg and Woodford (1999).

3 The Case for Price Stability

While the loss measures derived above under the various assumptions about the timing of pricing decisions are each different in certain respects, they all share an important common property. This is that the deadweight losses due to relative price distortions can in each case be completely eliminated, in principle, by stabilizing the aggregate price level. The intuition for this result is simple. The aggregate price level is stabilized by creating an environment in which suppliers who choose a new price have no desire at any time to set a price different from the average of existing prices. But if this is so, the average of existing prices never changes, and so the new prices that are chosen at all times are always the same, and eventually all goods prices are equal to that same, constant value. Thus aggregate price stability is a sufficient condition for the absence of price dispersion in our simple framework. At the same time, in most cases, it is also a necessary condition. This is not true in the pure “New Classical” case, as in that case it is only necessary that there be no *unexpected* changes in the aggregate price level in order for there to be no price dispersion. But this is clearly a highly special case; in the hybrid case considered at the end of the previous section, $\pi_t = 0$ at all times is required for zero price dispersion, even in the case of a very small positive value for α .

Moreover, price stability is not only the case in which the distortions associated with inefficient output *composition* are eliminated. As we shall see, it is also the route to minimization of the distortions associated with an inefficient *level* of output; and so, in the context of the kind of simple model considered thus far, it is an unambiguously desirable goal for monetary policy. The argument for this is simplest in the case that the equilibrium level of output under flexible prices is optimal, so we take up this case first. But as we shall

²²The welfare criterion of Rotemberg and Woodford differs from (2.23) in two respects. First, they assume pricing decision lags of one and two quarters for the two groups of goods, rather than of zero and one quarters, as here. Thus in their model, π_t is entirely forecastable at date $t - 1$, and they obtain a loss function of the form (2.23) in which the inflation surprise component is instead defined as $\pi_t - E_{t-2}\pi_t$. Second, they evaluate the unconditional expectation of utility rather than the conditional expectation of discounted utility, as discussed above. Thus they compute a period loss function L_t with the property that $E[U_t] = -\Omega E[L_t]$ (plus terms that may be neglected). This results in a period loss function of the form (2.23), but with a slightly different expression for the weight λ_x .

see, our conclusions require only minor modification even when we allow for the possibility that the natural rate of output is inefficiently low.

3.1 The Case of an Efficient Natural Rate of Output

Here we assume not merely that the inefficiency wedge Φ defined in (2.6) is of order $\mathcal{O}(\|\xi\|)$, but that it is equal to zero (or at any rate, that it is of order $\mathcal{O}(\|\xi\|^2)$, so that we may neglect it in our quadratic approximation to expected utility). This implies that $\bar{Y} = Y^*$ (or at least that their log difference x^* is of order $\mathcal{O}(\|\xi\|^2)$), so that the steady-state level of output under flexible prices is efficient (at least to second order). Since we have already verified, above, that percentage *fluctuations* in the natural rate are equal (to second order) to the percentage fluctuations in the efficient level of output, this actually implies that (to second order) the natural rate of output coincides with the efficient level of output at all times.

In this case, we easily obtain a very simple conclusion about the nature of optimal monetary policy. For *each* of the individual terms in our quadratic loss function can be shown to achieve its minimum possible value, zero, if inflation is zero at all times. We have just discussed the fact that this is true of the terms that measure the deadweight loss due to an inefficient composition of output. But in the present case, $x^* = 0$, so that the term in the loss function that involves the aggregate output gap is also minimized (and equal to zero) if and only if $x_t = 0$ at all times. Each of the aggregate supply relations (2.13), (2.16) and (2.21) implies that this will be true in the case of zero inflation at all times. Furthermore, except in the pure “New Classical” case (where all prices are changed each period), zero inflation at all times is also *necessary* in order to minimize this term. (In the “New Classical” case, it is once again only necessary that there be no *unforecastable* inflation.) Thus there is no conflict between the goals of minimizing the losses represented by the two separate terms inside the curly braces in (2.10); and both are minimized by complete stabilization of the aggregate price level.

In the case that some new prices must be chosen in advance (*i.e.*, in the “New Classical” or

hybrid cases), these results are only precisely correct under the assumption that the initial rate of expected inflation, $E_{-1}\pi_0$, that is given as an initial condition at the time that a new policy is chosen, is equal to zero.²³ For if this does not equal zero, then a policy that achieves zero inflation from period zero onward in all states will not result in a zero output gap in period zero, though the output gap will be zero thereafter. (Nor will the relative-price distortions be completely eliminated, since the surprise component of inflation will be non-zero in period zero, though it is zero in all later periods.) Of course, if it is known in advance that such a policy will be followed from period zero onward, $E_{-1}\pi_0$ should indeed equal zero. But if we wish to consider optimal policy choice starting from arbitrary initial conditions, we need also to allow for the possibility that $E_{-1}\pi_0 \neq 0$.

In the pure “New Classical” case, the elimination of deadweight losses only requires that the inflation *surprise* be zero at all times, and this is possible regardless of the initial condition $E_{-1}\pi_0$. Thus we need only modify our description of optimal policy to state that inflation in the initial period should equal whatever rate has already been expected (or more precisely, is reflected in the price changes that have already been chosen for period zero), while it is possible (though not essential) to commit to zero inflation in all later periods. The problem is more complicated in the hybrid case; here it is really not possible to eliminate all distortions in the case of an initial condition $E_{-1}\pi_0 \neq 0$. Nonetheless, the optimal monetary policy commitment from date zero onward makes the inflation rate each period a linear function of the initial condition and the shocks realized to that period,²⁴

$$\pi_t = \pi^{ss} + \sum_{j=0}^t a'_{tj} \epsilon_{t-j} + b_t E_{-1}\pi_0,$$

where ϵ_t represents the vector of innovations in period t in a state-space model of the evolution of the exogenous disturbances ξ_t . Furthermore, in this representation of optimal policy, the long-run average (or steady-state) inflation rate $\pi^{ss} = 0$, and all of the response coefficients $a_{tj} = 0$ as well. Only the coefficients b_t are non-zero (though $b_t \rightarrow 0$ as t grows).

²³Technically, the predetermined state variable is the new price that has already been chosen for period zero for the goods whose prices must be set in advance, as a ratio to the aggregate price level in period -1.

²⁴A Lagrangian method that can be used to solve for the optimal commitment in problems of this kind is presented in Woodford (1999c).

Briefly, one sees that the terms a_{tj} must be zero, as any dependence of the path of inflation upon the innovations ϵ_t adds additional terms of positive expected value to the discounted loss criterion, in addition to the terms resulting from the deterministic part of the inflation path.²⁵ Hence the optimal commitment involves a deterministic path for inflation $\{\pi_t\}$. Using the aggregate supply relation (2.21) to eliminate the output gap x_t in the loss function (2.23), one finds that the optimal inflation sequence must minimize

$$\begin{aligned} \sum_{t=0}^{\infty} \beta^t L_t &= \pi_0^2 + (\kappa\theta)^{-1}[(\pi_0 - \beta\pi_1) + (\psi^{-1} - 1)(\pi_0 - E_{-1}\pi_0)]^2 \\ &\quad + (\psi^{-1} - 1)(\pi_0 - E_{-1}\pi_0)^2 + \sum_{t=1}^{\infty} \beta^t [\pi_t^2 + (\kappa\theta)^{-1}(\pi_t - \beta\pi_{t+1})^2]. \end{aligned}$$

This implies a sequence of first-order conditions

$$\beta\pi_{t+1} - (1 + \beta + \kappa\theta)\pi_t + \pi_{t-1} = 0 \quad (3.1)$$

for each $t \geq 2$. There are corresponding conditions for $t = 0$ and 1, each of which is of the form

$$c_i E_{-1}\pi_0 + d_i \pi_0 + e_i \pi_1 + f_i \pi_2 = 0, \quad (3.2)$$

for $i = 1, 2$. (The first-order condition for $t = 1$ is actually of the same form as (3.1), except that π_{t-1} is replaced by $\psi^{-1}\pi_0 - (\psi^{-1} - 1)E_{-1}\pi_0$.)

The characteristic equation associated with the homogeneous difference equation (3.1),

$$\beta\lambda^2 - (1 + \beta + \kappa\theta)\lambda + 1 = 0, \quad (3.3)$$

has two real roots, satisfying

$$0 < \lambda_1 < 1 < \beta^{-1} < \lambda_2.$$

Hence for any given value of π_1 , (3.1) has a unique bounded solution (which is accordingly the solution satisfying the transversality condition) given by $\pi_t = \pi_1 \lambda_1^{t-1}$. The remaining two

²⁵Note that neither the structural equation (2.21) nor the objective function (2.23) involve any of the stochastic disturbance terms explicitly, once written in terms of inflation and the output gap. Hence in the case of any feasible stochastic solution for inflation and the output gap, the associated *certainty-equivalent* solution, obtained by replacing each random variable by its expectation as of period zero, is also consistent with (2.21), and achieves a strictly lower expected value for (2.23).

first-order conditions (3.2) then provide a system of two inhomogenous linear equations that can be solved for π_0 and π_1 as multiples of $E_{-1}\pi_0$. (Note that $\pi_1 = \lambda_1(\psi^{-1}\pi_0 - (\psi^{-1} - 1)\pi_0$.) One thereby obtains a solution of the kind described above, where the coefficients b_t satisfy $b_t = [\psi^{-1}(b_0 - 1) + 1]\lambda_1^{t-1}$ for all $t \geq 1$.

Thus optimal policy involves a deterministic path for inflation that depends upon the initial condition; but regardless of the initial condition, the optimal inflation rate eventually approaches zero, and it never responds to any of the exogenous disturbances. Once the optimal policy has been in place long enough, inflation will be zero at all times, regardless of the recent history of shocks. If we abstract from issues relating to the transition to an optimal policy regime starting from initial conditions that may not ever be generated again under such a regime, and simply consider how inflation should depend upon the recent history of shocks once an optimal policy is in place,²⁶

then a simple conclusion is obtained: in each of the cases considered thus far, optimal policy involves zero inflation at all times, regardless of the shocks hitting the economy. Though we do not provide an explicit analysis here, the same conclusion holds for a much larger family of possible time-dependent pricing rules (such as overlapping price commitments of fixed length of the kind considered by authors such as Blanchard and Fischer, 1989, pp. 395-398, or King and Wolman, 1999); only the transition dynamics associated with optimal policy are different in each case.

3.2 Optimal Responses to Shocks

This strong conclusion regarding the optimality of complete price stability depends upon various details of our model, as we discuss further in section 3.4. Nonetheless, it is interesting to remark that it holds despite our having allowed for several different kinds of stochastic disturbances. In particular, our framework allows for exogenous disturbances to technology, to government purchases, to households' impatience to consume, to their willingness to

²⁶The appropriateness of this "timeless perspective" in evaluating alternative policy rules is argued in Woodford (1999d). See also Svensson and Woodford (1999) and McCallum and Nelson (2000) for further discussion.

supply labor, or to the transactions technology that determines their demand for money balances. In the face of each of these types of disturbance, it remains optimal, under the circumstances assumed here, for the general level of prices to be held fixed.

The generality of the conclusion results from a simple intuition, stressed by Goodfriend and King (1997). Under the circumstances assumed here, the failure of prices to be continually adjusted is the *only* distortion that prevents rational expectations equilibrium from achieving an optimal allocation of resources. Thus an optimal monetary policy is one that achieves *the same allocation of resources as would occur with flexible prices*, if this is possible. Flexible-price equilibrium models of aggregate fluctuations (*i.e.*, real business cycle models²⁷) are then of practical interest, not as descriptions of what aggregate fluctuations should be like *regardless* of the monetary policy regime, but as descriptions of what they would be like under an *optimal* policy regime. Finally, our models of optimal price-setting imply that price stickiness will have no effects upon equilibrium outcomes in the case that monetary policy keeps the general price level completely unchanged over time, since in this case suppliers of goods would not wish to change their prices more frequently even if it were costless for them to do so. Thus complete price stability achieves the optimal allocation of resources.

Verifying that it is in fact possible, in principle, to achieve this first-best allocation through suitable monetary policy requires that we verify that we can solve the equations of our model for the evolution of all variables (including the interest-rate instrument of the

²⁷Standard real business cycle models (King and Rebelo, 1999) differ from the flexible-price limit of the model assumed here in that product markets are competitive, rather than monopolistically competitive; in that all output is produced using inputs purchased from the same factor markets, so that there is a common level of marginal cost for all firms at any time; and in that the endogenous dynamics of the capital stock in response to shocks is modeled, and indeed emphasized (as the only endogenous propagation mechanism in simple RBC models). However, in the flexible-price limit of our baseline model, all goods prices move together, and similarly the levels of production of each good, so that marginal cost is in fact the same for all firms. If we assume, as in this section, that an output or employment subsidy offsets the distortion due to firms' market power, the flexible-price equilibrium is equivalent to that of a competitive model with a single good. Finally, if we extend the basic model expounded here to take account of capital-accumulation dynamics, as in Woodford (2000b), then the flexible-price dynamics of our model are fully equivalent to those of a standard RBC model. Note that these models, like the "cashless limit" of our model, abstract from real-balance effects upon consumption demand, labor supply, and so on.

central bank) under the assumption that $\pi_t = 0$ at all times. We observe that each of our alternative aggregate supply relations implies that this requires that the output gap be zero at all times,²⁸ so that aggregate output satisfies $\hat{Y}_t = \hat{Y}_t^n$ at all times.

Recalling our definitions of the composite exogenous disturbance terms, these can be expressed in terms of more fundamental disturbances as

$$g_t = \hat{G}_t + (1 - s_G)\tilde{c}_t,$$

$$q_t = (1 + \omega^{-1})a_t + \omega^{-1}\nu\tilde{h}_t.$$

Here \hat{G}_t denotes the deviation of government purchases from their steady-state level, measured as a percentage of steady-state output \bar{Y} , which shifts the level of private expenditure implied by any given level of aggregate demand \hat{Y}_t ; \tilde{c}_t denotes the percentage shift in the Frisch (constant marginal utility of income) consumption demand, due to a shift in the utility-of-consumption function; a_t is the multiplicative technology disturbance; and \tilde{h}_t is the percentage shift in the Frisch labor supply, due to a shift in the disutility-of-labor function v . (The exogenous shifts in the Frisch demand schedules are measured at the steady-state values of their arguments.) In addition, $s_G < 1$ is the steady-state fraction of total spending consisting of government purchases, and $\nu > 0$ is the inverse of the Frisch (or intertemporal) elasticity of labor supply. It then follows from (2.11) that

$$\hat{Y}_t^n = \frac{\sigma^{-1}}{\sigma^{-1} + \omega}(\hat{G}_t + (1 - s_G)\tilde{c}_t) + \frac{1}{\sigma^{-1} + \omega}((1 + \omega)a_t + \nu\tilde{h}_t).$$

We observe that *each* of the exogenous disturbances \hat{G}_t , \tilde{c}_t , a_t , and \tilde{h}_t increases the natural rate of output, and thus, under the optimal policy each of them is allowed to perturb the equilibrium level of economic activity \hat{Y}_t . Nonetheless, it is optimal to completely stabilize prices in each case. Thus the case for price stability does not depend upon an assumption that the only important shocks are ones that affect aggregate demand without shifting the

²⁸Here we abstract from transition issues. In the case of our “hybrid” aggregate supply specification (2.21), if one starts from an initial condition $E_{-1}\pi_0 \neq 0$, the optimal commitment, derived above, implies non-zero output gaps. However, the implied sequence for x_t is deterministic (*i.e.*, unaffected by shocks in period zero or later), and converges to zero asymptotically at the rate λ_1^t .

aggregate supply curve. Of course, it is also optimal to completely stabilize prices in response to pure “demand” disturbances, such as the exogenous shifts in money demand considered below.

Substituting the optimal price and output movements into the intertemporal Euler equation,

$$x_t = E_t x_{t+1} - \sigma[\hat{\lambda}_t - E_t \pi_{t+1} - \hat{r}_t^n], \quad (3.4)$$

we can solve for the required movements of the central bank’s interest-rate instrument. These are given by $\hat{\lambda}_t = \hat{r}_t^n$, *i.e.*, the interest rate must perfectly track the exogenous fluctuations in the Wicksellian natural rate of interest, as discussed in Woodford (2000b).²⁹ If this represents a feasible path for the interest rate, as will necessarily be true at least in the case of small enough shocks, then we have verified that it is possible in principle for monetary policy to achieve the optimal allocation of resources.

Noting that

$$\hat{r}_t^n \equiv \sigma^{-1}[(g_t - \hat{Y}_t^n) - E_t(g_{t+1} - \hat{Y}_{t+1}^n)],$$

and assuming (for simplicity) that each of the exogenous disturbances follows an independent first-order autoregressive process, we see that the required interest-rate variations are given by

$$\hat{r}_t^n = (\sigma + \omega^{-1})^{-1}[(1 - \rho_G)\hat{G}_t + (1 - s_G)(1 - \rho_c)\tilde{c}_t - (1 + \omega^{-1})(1 - \rho_a)a_t - \omega^{-1}\nu(1 - \rho_h)\tilde{h}_t],$$

where ρ_G , ρ_c , ρ_a , and ρ_h are the coefficients of serial correlation of the four exogenous disturbance processes. Since stationarity requires that $\rho_i < 1$ in each case, we observe that under this assumption, interest rates must increase in response to temporary increases in government purchases or in the impatience of households to consume, and decrease in response to temporary increases in productivity or in the willingness of households to supply labor. In each case, the required interest-rate changes under the optimal policy are larger the more temporary the disturbance (*i.e.*, the less positive the serial correlation).

²⁹As is discussed in that paper, this is not a prescription for a *policy rule* to achieve price stability, since a commitment to keep the nominal interest rate on this path does not lead to a determinate rational expectations equilibrium; it simply describes how interest rates must vary in the desired equilibrium.

It is worth noting that the required interest-rate variations in response to the various types of shocks cannot be achieved, in general, through a simple “Taylor rule” under which the nominal interest rate is a function solely of inflation and the deviation of output from trend. In the optimal equilibrium, inflation does not vary in response to the shocks at all, and so conveys no information about them. Output does vary in response to all of the shocks, but the desired interest-rate response is not proportion to the desired output response across the various types of shocks; indeed, one wants interest rates to vary procyclically in the case of government-purchase or consumption-demand shocks, but countercyclically in response to technology or labor-supply shocks. Thus the central bank will need additional information in order to implement the optimal policy.

Finally, by substituting the necessary equilibrium variations in prices, output and interest rates into a money demand equation,³⁰

$$\log M_t^s = \log P_t + \eta_y \hat{Y}_t - \eta_i \hat{r}_t + \epsilon_t^m, \quad (3.5)$$

where $\eta_y, \eta_i > 0$ and ϵ_t^m is an exogenous disturbance, we can learn how the money supply must be allowed to vary as part of the optimal policy. We obtain

$$\log M_t^s = \log P^* + \eta_y \hat{Y}_t^n - \eta_i \hat{r}_t^n + \epsilon_t^m,$$

where \hat{Y}_t^n and \hat{r}_t^n are the functions of the primitive disturbances just derived. We find that, in general, the money supply should be allowed to vary in response to all five types of exogenous disturbances, so that a constant money growth rate is certainly not the optimal policy.

Nor does the optimal evolution of the money supply necessarily involve “leaning against the wind”. For example, *procyclical* variations in the money supply are optimal in response to temporary fluctuations in productivity, as argued by Ireland (1996);³¹ for an increase

³⁰This can be derived as an additional equilibrium condition in a standard Sidrauski-Brock monetary model, under assumptions consistent with the other equations used above; see, e.g., Woodford (1999a). Here we normalize the steady-state level of real money balances \bar{m} to equal one.

³¹Aiyagari and Braun (1998) reach a similar conclusion, in the case of their model with sticky prices, though they assume convex costs of price changes, following Rotemberg (1995), rather than predetermined prices as in the models considered here. These authors also reach a similar conclusion with regard to government-purchase shocks, in the case of their numerical calibration of their model.

in a_t raises \hat{Y}_t^n while lowering \hat{r}_t^n , thus warranting an increase in $\log M_t^s$ at the same time as an increase in \hat{Y}_t . The same is true of temporary labor-supply shocks, and while the result depends upon parameter values, it is also true of government-purchase shocks and consumption-demand shocks, at least if these are sufficiently persistent. Furthermore, in the case of technology or labor-supply shocks, it is actually optimal for the money supply to be *more* procyclical than would be the case if interest rates were held unchanged; for one actually wants nominal interest rates to *decline* in response to a positive shock. In the case of the other two shocks, this is not true, but it is still possible that holding the nominal interest rate fixed is closer to the optimal response than holding the money supply fixed; in particular, this is necessarily true if the shocks are sufficiently persistent, as in that case the natural rate of interest is affected very little.

These results contrast with the classic analysis of Poole (1970), according to which it is optimal to accommodate money-demand shocks (fixing the nominal interest rate and letting the money supply vary), but not “IS shocks”. Our analysis similarly concludes that full accommodation of money-demand shocks is optimal. But our government-purchase or consumption-demand shocks presumably correspond to what Poole intends by “IS shocks”, yet even in the case of these shocks, some degree of accommodation is often optimal, and if the shocks are sufficiently persistent, the optimal degree of accommodation of the “IS shift” may be nearly 100 percent. The difference, of course, is that Poole assumes that output stabilization should be the goal of policy, whereas here we find that optimal policy stabilizes the output *gap* instead, and that the natural rate of output is affected by “IS shocks” among others. If we consider the possibility of technology or labor-supply shocks, neglected by Poole altogether, our results are even more different, and even more strongly support a presumption in favor of full accommodation.³²

³²Of course, even in the case that procyclical variation in the money supply is *not* optimal, it does not follow that there is any reason to target the money supply, or even to respond to it at all. Assuming that it is possible for the central bank to make its interest-rate instrument depend upon aggregate output and the aggregate price level, then any rule that puts a non-zero weight on the money supply is dominated by another rule that puts a zero weight on the money supply, and instead involves direct feedback from output and/or prices to the interest rate, without introducing the noise associated with the money-demand shocks ϵ_t^m .

Our results also differ, at least superficially, from those of Ireland (1996), who argues that one should use monetary policy to “insulate aggregate output” against “shocks to demand”, while accommodating “shocks to supply”. Many readers might assume that “shocks to demand” would include disturbances such as our government-purchase or consumption-demand shocks, but we have seen that it is not optimal to stabilize output in response to these shocks. In fact, in Ireland’s model, “shocks to demand” refer solely to money-demand shocks, which is the only exogenous disturbance other than a technology shock that he considers.

Our results also differ superficially from those of Clarida *et al.* (1999), who state (in their “Result 4”) that optimal policy involves “adjusting the interest rate to perfectly offset demand shocks,” while “perfectly accommodat[ing] shocks to potential output by keeping the nominal interest rate constant”. In fact, the variable (their “ g_t ”) here referred to as a “demand shock” corresponds to our natural rate of interest r_t^n .³³ What these authors mean by “perfectly offsetting” movements in this variable is that the central bank’s interest-rate instrument should move one-for-one with variations in the natural rate of interest. (Thus “perfectly offsetting” the shocks does not mean that output is insulated from them, but that the *output gap* is.) And what they mean by “perfectly accommodating shocks to potential output” is that, *given* the value of the natural rate of interest, the interest rate should be independent of the natural rate of output. That is, disturbances to the natural rate of output *that do not shift the natural rate of interest* should not affect nominal interest rates. Stated this way, there is no difference between these results and our own.³⁴ However, it is *not* true, in general, that optimal policy involves no interest-rate response to shocks that affect the natural rate of output, because (as shown above) such shocks almost always do affect the natural rate of interest to some extent.

³³The variable is evidently thought of as a “demand shock” because it is the disturbance term in the Euler equation (3.4). But because this condition has been written in terms of the *output gap* x_t rather than the level of output \hat{Y}_t , the composite disturbance \hat{r}_t^n , unlike our variable g_t , cannot properly be regarded as a pure demand shock, if one supposes that transitory disturbances to the natural rate of output occur.

³⁴Actually, the results referred to in Clarida *et al.* are characterizations of optimizing central bank policy under discretion, which is *not* in general optimal policy, in the sense of the policy that best achieves the central bank’s assumed objectives, as is emphasized in Woodford (1999b). However, under the special circumstances that we assume here, optimizing policy under discretion coincides with truly optimal policy, as optimal policy is time-consistent in this case.

3.3 Consequences of an Inefficient Natural Rate of Output

We now consider the extent to which the above conclusions must be modified in the case that (quite realistically) we assume that $\Phi > 0$, so that the equilibrium rate of output under flexible prices would be inefficiently low. The distortions represented by the coefficient Φ , *i.e.*, the market power resulting from monopolistic competition and the constant rate of distorting taxation τ , introduce a wedge between this natural rate of output and the efficient output level. However, this wedge is assumed to be constant over time, so that *percentage changes* in the natural rate still correspond precisely (in our log-linear approximation) to percentage changes in the efficient level of output. Thus, as shown above, the distortions associated with a suboptimal aggregate level of economic activity are still measured a quadratic function of the output gap, $\lambda(x_t - x^*)^2$, even if now the constant x^* is assumed to be positive.

While this difference matters for the optimal *average* levels of inflation and output – that is, for the deterministic part of our above description of the optimal policy commitment – it has *no* effect (in our log-linear approximation to optimal policy) upon the optimal responses to shocks. We first demonstrate this in the simple context of our “New Classical” model of price-setting. In this case, the normalized quadratic loss function (2.15) can be written

$$L_t = (\pi_t - E_{t-1}\pi_t)^2 - 2\lambda x^* x_t + \lambda x_t^2, \quad (3.6)$$

dropping the term λx^{*2} that is independent of policy. The second term on the right-hand side now indicates a welfare gain from an increase in the expected output gap in any period. However, under the assumption that Φ , and hence x^* , is of order $\mathcal{O}(|\xi|)$, a *first-order* approximation to the solution for x_t suffices to give us a *second-order* approximation to this term. Hence we may substitute using the aggregate supply relation (2.13), to obtain

$$L_t = (\pi_t - E_{t-1}\pi_t)^2 - 2\theta^{-1}x^*[\pi_t - E_{t-1}\pi_t] + \lambda x_t^2,$$

recalling that $\lambda = \kappa/\theta$.

Taking the expected discounted value of such terms (and dropping the term $E_{-1}\pi_0$ that

is independent of policy), we obtain the utility-based welfare criterion

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t L_t \right\} = -2\theta^{-1} x^* \pi_0 + E_0 \left\{ \sum_{t=0}^{\infty} \beta^t [(\pi_t - E_{t-1} \pi_t)^2 + \lambda x_t^2] \right\} \quad (3.7)$$

Note that each of the terms proportional to x^* has canceled, except the one indicating a welfare gain from surprise inflation at date zero, the time at which a new policy commitment is adopted. Because it is not possible to commit in advance to an inflation *surprise* at any later date, the corresponding terms for dates $t \geq 1$ do not matter. But this means that allowing for $x^* > 0$ has no effect upon the nature of the optimal policy commitment, except in the initial (transitional) period, when it is possible to take advantage of the fact that private sector expectations of period zero inflation are already given, before the policy is adopted.³⁵ It is arguable (see Woodford, 1999d) that it does not make sense to behave differently in this initial period than one commits to behave later, if one wants the commitment to be credible. But regardless of how one manages the transition to the optimal regime, it is optimal to commit to an eventual zero rate of inflation, and to a path for inflation that is unaffected by any stochastic disturbances.³⁶

It might be thought that this result depends upon the fact that in the special case in which all prices are changed every period (though some are committed a period in advance), only *unexpected* inflation has an effect upon output. Yet a similar conclusion is obtained in our baseline model, with Calvo price-setting. In this case, we can similarly substitute into (2.20) using the aggregate supply relation (2.16), to obtain

$$L_t = \pi_t^2 - 2\theta^{-1} x^* [\pi_t - \beta E_t \pi_{t+1}] + \lambda x_t^2,$$

noting again that $\lambda = \kappa/\theta$. Taking the expected discounted value, we obtain the utility-based welfare criterion

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t L_t \right\} = -2\theta^{-1} x^* \pi_0 + E_0 \left\{ \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda x_t^2] \right\}. \quad (3.8)$$

³⁵Of course, this different prescription in the case of the initial period shows that optimal policy is not *time-consistent* in this case. See Woodford (1999b) for further discussion.

³⁶Of course, in this model, there is no advantage of complete price stability over any other policy that makes inflation completely forecastable a period in advance. But in order to stress the similarity of the results obtained under the alternative aggregate supply specifications, it is worth noting that also in this case there is no *advantage* to any variation in inflation in response to shocks.

Once again all of the terms proportional to x^* cancel, except the one indicating welfare gains from a surprise inflation in period zero. Committing in advance to non-zero inflation in any later period does *not* produce any such effect. For the value of the increase in output in any period $t \geq 1$ resulting from higher inflation in period t must be offset by the cost of the *reduction* in output in period $t - 1$ as a result of *expectation* of that higher inflation in period t . From the standpoint of the discounted loss criterion (3.7), the costs resulting from the anticipation of the inflation are weighted more strongly (by a factor of $\beta^{-1} > 1$), as they occur earlier in time. On the other hand, the output effect of anticipated inflation, by shifting the short-run aggregate supply curve, is also smaller than the effect of current inflation, by exactly the factor $\beta < 1$, with the result that the two effects exactly cancel, to first order (which is to say, to second order when multiplied by x^*). Thus once again there is no welfare gain, up to our order of approximation, from a commitment to inflation that can be anticipated in advance. In particular, we find once again that except for transition effects, resulting from the different term in (3.8) for the initial period, it is again optimal to commit to zero inflation, independent of the shocks to the economy.

Nonetheless, the term in (3.8) that is linear in π_0 now affects the optimal commitment for periods later than π_0 as well. That is because of the intertemporal linkage implied by aggregate supply relation (2.16). The welfare gain from inflation at date zero can be obtained with less increase in the period zero output gap (and hence less increase in the λx_0^2 term) if it is accompanied by an increase in expected inflation at date one; and since the welfare loss from such inflation is merely quadratic, it is optimal to commit to some amount of such inflation. Thus the inflation associated with the transition to the optimal regime lasts for more than a single period in this case.

Let us briefly sketch a derivation of the optimal commitment in this case. As in section 3.1, we easily see that any dependence of the inflation rate upon the stochastic disturbances increases expected losses (3.8), so that we may restrict attention to deterministic inflation commitments. Again using (2.16) to eliminate x_t from (3.8), we see that the optimal com-

mitment involves an inflation sequence $\{\pi_t\}$ that minimizes

$$\sum_{t=0}^{\infty} \beta^t L_t = -2\theta^{-1}x^*\pi_0 + \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + (\kappa\theta)^{-1}(\pi_t - \beta\pi_{t+1})^2].$$

The first-order condition for the optimal choice of π_0 is given by

$$\beta\pi_1 - (1 + \kappa\theta)\pi_0 + \kappa x^* = 0, \tag{3.9}$$

while that for π_t is again given by (3.1) for each $t \geq 1$.

As shown above, the sequence of conditions (3.1) have a unique solution satisfying the transversality condition, given by

$$\pi_t = \pi_0 \lambda_1^t,$$

where λ_1 is again the smaller root of the characteristic equation (3.3). Substituting this into (3.9), we find that the optimal initial inflation rate equals

$$\pi_0 = (\lambda_1^{-1} - \beta)^{-1} \kappa x^* > 0.$$

Thus it is optimal to arrange an initial inflation, taking account of the fact that the decision to do so can have no effect upon expectations prior to date zero (if one is not bothered by the non-time-consistency of such a principle of action). The optimal policy involves positive inflation in subsequent periods as well, but there should be a commitment to reduce inflation to its optimal long-run value, of zero, asymptotically (as $\lambda_1 < 1$). And the rate at which inflation is committed to decline to zero should be completely unaffected by random disturbances to the economy in the meantime.³⁷ Thus the assumption of $\Phi > 0$ makes no difference for the conclusions of the previous section with regard to the optimal response to shocks. And if one takes the view (as we shall argue later) that one should actually conduct policy as one *would have optimally committed to do far in the past*, thus foregoing the temptation to exploit the private sector's failure to anticipate the new policy, then it is optimal simply to choose $\pi_t = 0$ at all times – *i.e.*, to completely stabilize the price level – just as in the previous section.

³⁷These results agree with those of King and Wolman (1999) in the context of a model with two-period overlapping price commitments in the style of Taylor (1980).

It is interesting to note that this result – that the optimal commitment involves a long-run inflation rate of zero, even when the natural rate of output is inefficiently low – does *not* depend upon the existence of a vertical “long-run Phillips curve” tradeoff. For the aggregate supply relation (2.16) in our baseline model implies an upward-sloping relation

$$x^{ss} = (1 - \beta)\kappa^{-1}\pi^{ss}$$

between steady-state inflation π^{ss} and the steady-state output gap x^{ss} . (This is because the expected-inflation term has a coefficient $\beta < 1$, unlike that of the “New Classical” relation (2.13).) It is sometimes supposed that the existence of a long-run Phillips-curve tradeoff, together with an inefficient natural rate, should imply that the Phillips curve should be exploited to some extent, resulting in positive inflation forever, even under commitment. But here that is not true, because the smaller coefficient on the expected-inflation term relative to that on current inflation – which results in the long-run tradeoff – is exactly the size of shift term in the short-run aggregate supply relation that is needed *to precisely eliminate any long-run incentive for non-zero inflation* under an optimal commitment. If one were instead to “simplify” the New Keynesian aggregate supply relation, putting a coefficient of one on expected inflation (as is done in some presentations,³⁸ presumably in order to conform to the conventional wisdom regarding the long-run Phillips curve), we would then fail to obtain such a simple result. The optimal long-run inflation rate would actually be found to be *negative*, as the stimulative effects of lower expected inflation would be judged to be worth more than the output cost of lower current inflation – even though there would actually be no long-run output increase as a result of the policy!

Similar conclusions can be obtained in the case of the “hybrid” aggregate supply specification (2.21), though the transition dynamics are slightly more complex. In this case we obtain instead

$$L_t = \pi_t^2 + \lambda_u(\pi_t - E_{t-1}\pi_t)^2 - 2\theta^{-1}x^*[\psi^{-1}\pi_t - (\psi^{-1} - 1)E_{-1}\pi_t - \beta E_t\pi_{t+1}] + \lambda_x x_t^2,$$

³⁸See, e.g., Roberts (1995) or Clarida *et al.* (1999).

so that

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t L_t \right\} = -2\theta^{-1}x^*[\psi^{-1}\pi_0 - (\psi^{-1} - 1)E_{-1}\pi_0] \\ + E_0 \left\{ \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda_u(\pi_t - E_{t-1}\pi_t)^2 + \lambda_x x_t^2] \right\},$$

generalizing (3.8). Again, all of the terms proportional to x^* cancel, except the term involving π_0 and $E_{-1}\pi_0$. The optimal commitment again involves a deterministic path for inflation, and the first-order conditions satisfied by this path are the same as in the problem considered in section 3.1. It follows once again that the optimal commitment is of the form

$$\pi_t = [\psi^{-1}\pi_0 - (\psi^{-1} - 1)E_{-1}\pi_0]\lambda_1^t$$

for all $t \geq 1$, so that the value of x^* affects only the relation between the optimal π_0 and the initial condition $E_{-1}\pi_0$. Thus once again the optimal commitment involves no dependence of inflation upon the exogenous disturbances, and a commitment to reduce inflation to zero eventually.

3.4 Caveats

We have seen that, within the class of sticky-price models discussed above, the optimality of a monetary policy that aims at complete price stability is surprisingly robust. Not only does this conclusion not depend upon the fine details of how many prices are set a particular time in advance or left unchanged for a particular length of time, but it remains valid in the case of a considerable range of types of stochastic disturbances, and in the case of an inefficient natural rate of output. Nonetheless, it is likely that some degree of deviation from full price stability is warranted in practice. Some of the more obvious reasons for this are sketched here.

First of all, complete price stability may not be *feasible*. We have just argued, in section 3.2, that in our baseline model, it is feasible, because we are able to solve for the required path of the central bank's nominal interest-rate instrument. This is correct, as long as the random disturbances are small enough in amplitude. But if they are larger, such a policy

might not be possible, because it might require the nominal interest rate to be *negative* at some times, which is not possible under any policy. Specifically, this will occur if it is ever the case that the natural rate of interest is negative. On average, it does not seem that it should be, and thus zero inflation *on average* would seem to be feasible; but it may be temporarily negative as a result of certain kinds of disturbances, and this is enough to make complete price stability infeasible. As a result, a policy will have to be pursued which involves less volatility of the short nominal interest rate in response to shocks, and some amount of price stability will have to be sacrificed for the sake of this.³⁹ The way in which optimal monetary policy is different in the presence of such a concern is an important concern of Woodford (1999b).⁴⁰

Varying nominal interest rates as much as the natural rate of interest varies may also be desirable as a result of the “shoe-leather costs” involved in economizing on money balances. As argued by Friedman (1969), the size of these distortions is measured by the level of nominal interest rates, and they are eliminated only if nominal interest rates are zero at all times.⁴¹ Taking account of these distortions – from which we have abstracted thus far in our welfare analysis – provides another reason for the equilibrium with complete price stability, even if feasible, not to be fully efficient; for as Friedman argues, a zero nominal interest rate will typically require expected *deflation* at a rate of at least a few percent per year.

One might think that this should make no more difference to our analysis of optimal policy than does the existence of an inefficient natural rate of output due to market power – that it may similarly affect the *deterministic* part of the optimal path for inflation without creating any reason for inflation to vary in response to random shocks. But monetary frictions do

³⁹In general, it will be optimal to back off from complete price stability both by allowing inflation to vary somewhat in response to disturbances, *and* by choosing an average rate of inflation that is somewhat greater than zero, as suggested by Summers (1991), in order to allow more room for interest-rate fluctuations consistent with the zero lower bound. However, the quantitative analysis undertaken in Rotemberg and Woodford (1997, 1999) and Woodford (1999b, 1999c) finds that the effect of the interest-rate lower bound on the optimal response of inflation to shocks is more significant than the effect upon the optimal average rate of inflation.

⁴⁰See also Woodford (1999c, section 4.2).

⁴¹See Woodford (1990) for justification of this relation in a variety of alternative models of the demand for money.

not have implications only for the optimal *average* level of nominal interest rates. As with distorting taxes, it is plausible that the deadweight loss is a convex function of the relative-price distortion, so that temporary increases in nominal interest rates are more costly than temporary decreases of the same size are beneficial. In short, monetary frictions provide a further reason for it to be desirable to reduce the *variability* of nominal interest rates, even if one cannot reduce their average level. (At the same time, reducing their average level will require less variable rates, because of the zero floor.) Insofar as these costs are important, they too will justify a departure from complete price stability, in the case of any kinds of real disturbances that cause fluctuations in the natural rate of interest, in order to allow greater stability of nominal interest rates. This tradeoff is treated more explicitly in Woodford (1999c, section 4.1).

Even apart from these grounds for concern with interest-rate volatility, it should be recognized that the class of sticky-price models analyzed above are still quite special in certain respects. One of the most obvious is that there are assumed to be no shocks as a result of which the relative prices of any of the goods with sticky prices would vary over time in an *efficient* equilibrium (*i.e.*, the shadow prices that would decentralize the optimal allocation of resources involve no variation in the relative prices of such goods). This is because we have assumed that only goods prices are sticky, that all goods enter the model in a perfectly symmetrical way, and that all random disturbances have perfectly symmetrical effects upon all sectors of the economy. These assumptions are convenient, but plainly an idealization. Yet it should be clear that they are relied upon in our conclusion that stability of the general price level suffices to eliminate the distortions due to price stickiness.

If an efficient allocation of resources requires relative price changes, due to asymmetries in the way that different sticky-price commodities are affected by shocks, this will not be true. We show, however, in Woodford (1999c), that even in the presence of asymmetric shocks, it is possible to define a symmetric case in which it is still optimal to completely stabilize the general price level, even though this does not eliminate all of the distortions resulting from price stickiness. But this holds exactly only in a special case, in which different goods

are similar, among other respects, in the *degree* of stickiness of their prices. If sectors of the economy differ in their degree of price stickiness (as is surely realistic), then complete stabilization of an aggregate price index will not be optimal. Stabilization of an appropriately defined *asymmetric* price index (that puts more weight on the stickier prices) is a better policy, as argued by Aoki (1999) and Benigno (1999), though even the best policy of this kind need not be fully optimal.

An especially important reason for disturbances to require relative price changes between sticky commodities with sticky prices is that wages are probably as sticky as are prices. Real disturbances almost inevitably require real wage adjustments in order for an efficient allocation of resources to be decentralized, and if *both* wages and prices are sticky, it will then not be possible to achieve all of the relative prices associated with efficiency simply by stabilizing the price level – specifically, the real wage will frequently be misaligned, as will be the relative wages of different types of labor if these are not set in perfect synchronization. In such circumstances, complete price stability may not be a good approximation at all to the optimal policy, as Erceg *et al.* (1999) show. Nonetheless, stabilization of an appropriately weighted average of prices and wages is still found to be a good approximation to optimal policy, and fully optimal in some cases (Woodford, 1999c). Thus concerns of this kind are not so much reasons not to pursue price stability as they are reasons why care in the choice of the index of prices (including wages) that one seeks to stabilize may be important.

Yet another qualification to our results in this section is that we have assumed a framework in which the natural rate of output – the equilibrium level in the case of flexible prices – is efficient, or at any rate differs from the efficient level only by a (small) constant factor. As we have seen, this assumption is compatible with the existence of a variety of types of economic disturbances, including technology shocks, preference shocks, and variations in government purchases. But it would not hold in the case of other sorts of disturbances, that cause *time variation in the degree of inefficiency* of the natural rate. These could include variation in the level of distorting taxes, variation in the degree of market power of firms or workers, or variation in the size of the wage premium that must be paid on efficiency-wage

grounds.⁴²

These are presumably the kinds of disturbances intended by the “cost-push shock” term in the aggregate supply equation of Clarida *et al.* (1999).⁴³ In the presence of such shocks, but with assumptions otherwise as in our baseline model, the aggregate supply relation (2.16) takes the form

$$\pi_t = \kappa[x_t + c_t] + \beta E_t \pi_{t+1}, \quad (3.10)$$

where c_t indicates the percentage deviation (from its steady-state level) of the amount by which the natural rate of output (the equilibrium level under flexible prices) exceeds the efficient level of output, and x_t is now the output gap relative to the efficient level of output, rather than the natural rate.⁴⁴ The normalized utility-based loss function is still given by (2.20), in terms of this notation, as it is the gap between actual output and the *efficient* level that matters for the computation of deadweight loss. (Here we continue to assume that both the efficient level of output and the natural rate are the same for all goods, so that the dispersion of efficiency gaps continues to be directly related to price dispersion.)

With this modification of the aggregate supply relation, complete stabilization of inflation is no longer sufficient for complete stabilization of the welfare-relevant output gap, and given that output-gap variability also affects the loss function, complete stabilization of inflation will not generally be optimal.⁴⁵ It is not obvious that stabilization of an alternative price index makes sense as a solution to the problem in this case, either, whereas some degree of

⁴²See Giannoni (2000) for further discussion of the origin and consequences for optimal policy of such disturbances.

⁴³The terminology is not entirely helpful, however, as there is no necessary connection between shocks that affect inflation by increasing costs of production and this kind of time variation in the degree of inefficiency of the natural rate of output. Technology shocks, energy price shocks, or variations in real wage demands may all shift the aggregate supply curve without implying any variation in Y_t^n/Y_t^* .

⁴⁴See Giannoni (2000). For consistency with the literature, *e.g.*, Clarida *et al.*, we technically define x_t as $\log(Y_t/Y_t^*)$ plus the steady-state value of $\log(Y_t^n/Y_t^*)$. This is somewhat illogical notation, but allows us to follow the literature, such as Clarida *et al.*, in representing a non-zero average gap between the natural rate and the efficient level by a non-zero target value x^* in the loss function, rather than by a constant in the aggregate-supply relation, while at the same time representing time variation in this gap by a random term in the aggregate-supply relation, rather than a time-varying target for the output gap that appears in the loss function.

⁴⁵See Clarida *et al.* (1999), Woodford (1999d), and Svensson and Woodford (1999) for discussions of optimal policy in the presence of this kind of disturbances.

concern for stabilization of the (appropriately measured) output gap is clearly appropriate, alongside a concern for inflation stabilization.

In fact, one can easily show (Clarida *et al.*, 1999; Woodford, 1999d) that minimization of the welfare-based loss function (2.20) subject to the sequence of constraints implied by the aggregate supply relation (3.10) requires stabilization of a hybrid index

$$\log P_t + \theta^{-1}x_t, \tag{3.11}$$

rather than stabilization of either the price index or the output gap alone. This is an example of the form of policy rule that Hall (1984) calls an “elastic price standard”. However, under a plausible calibration of the degree of market power in an economy like the U.S., the relative weight on the output gap in this index should only be about 0.1 — much smaller than the weights assumed, for example, by Hall.

Hence while it is true that the presence of inefficient supply disturbances justifies a departure from complete price stability for the sake of greater stabilization of the output gap, the optimal degree of response to variations in the output gap remains small. And the above calculation assumes perfect observability of both inflation and the output gap. In practice, real-time estimates of the output gap are much more uncertain than are measures of inflation, and taking this into account is likely to justify even less of a deviation from complete stabilization of the price level.⁴⁶

Thus the case for price-level stabilization is more robust than is often assumed. Complete stabilization of the price level is fully optimal only under relatively special circumstances — but as we have seen, this case still allows for a wide range of types of disturbances, and a modest degree of inefficiency of the economy for reasons other than the stickiness of prices. Furthermore, while there are many reasons why complete price stability is not likely to be

⁴⁶See, *e.g.*, Svensson and Woodford (2000) for an analysis of optimal policy in the case of inefficient supply disturbances, when potential output is observed only with noise. Because of a certainty-equivalence principle, it is still optimal to stabilize one’s best estimate of the current value of the index (3.11). However, the greater the degree of noise in the observation of potential output (*i.e.*, the estimate based on quantity data alone), the greater the extent to which the optimal estimate of the output gap (based on a Kalman filter) is a function primarily of fluctuations in prices. Hence the rule amounts to something quite close to price-level stabilization.

fully optimal, none of these seem likely to justify departures from price stability that are too great, quantitatively (at least if the price index to be stabilized is appropriately defined). This remains, of course, an important topic for investigation in quantitative models. Furthermore, the question of which price index it is most desirable to stabilize remains an important topic for further study, despite the promising initial results of Aoki (1999), Benigno (1999) and Erceg *et al.* (1999).

References

- Aiyagari, S. Rao, and R. Anton Braun, "Some Models to Guide Monetary Policymakers," *Carnegie-Rochester Conference Series on Public Policy* 48: 1-42 (1998).
- Amato, Jeffery D., and Thomas Laubach, "Rule-of-Thumb Behavior and Monetary Policy," unpublished, Bank for International Settlements, April 2000.
- Aoki, Kosuke, "Optimal Monetary Policy Responses to Relative Price Changes," unpublished, Princeton University, July 1999.
- Bean, Charles R., "Targeting Nominal Income: An Appraisal," *Economic Journal* 93: 806-819 (1983).
- Benigno, Pierpaolo, "Optimal Monetary Policy in a Currency Area," unpublished, Princeton University, May 1999.
- Blanchard, Olivier J., and Stanley Fischer, *Lectures on Macroeconomics*, Cambridge: M.I.T. Press, 1989.
- Blinder, Alan S., Elie R.D. Canetti, David E. Lebow, and Jeremy B. Rudd, *Asking About Prices: A New Approach to Understanding Price Stickiness*, New York: Russell Sage Foundation, 1998.
- Calvo, Guillermo, "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12: 383-98 (1983).
- Clarida, Richard, Jordi Gali and Mark Gertler, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature* 37: 1661-1707 (1999).
- Collard, Fabrice, Harris Dellas, and Guy Ertz, "Poole Revisited," unpublished, CEPREMAP (Paris), October 1998.
- Dupor, Bill, "Optimal Monetary Policy with Nominal Rigidities," unpublished, Wharton School, February 1999.
- Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin, "Optimal Monetary Policy with Staggered Wage and Price Contracts", unpublished, Federal Reserve Board, March 1999.
- Friedman, Milton, "The Optimum Quantity of Money," in *The Optimum Quantity of Money and Other Essays*, Chicago: Aldine, 1969.

- Giannoni, Marc, "Price-Level Rules versus Inflation Rules," unpublished, Federal Reserve Bank of New York, September 2000.
- Goodfriend, Marvin, and Robert G. King, "The New Neoclassical Synthesis and the Role of Monetary Policy," *NBER Macroeconomics Annual* 12: 231-283 (1997).
- Hall, Robert E., "Monetary Policy with an Elastic Price Standard," in *Price Stability and Public Policy*, Kansas City: Federal Reserve Bank of Kansas City, 1984.
- —, and N. Gregory Mankiw, "Nominal Income Targeting," in N.G. Mankiw, ed., *Monetary Policy*, Chicago: University of Chicago Press, 1994.
- Ireland, Peter N., "The Role of Countercyclical Monetary Policy," *Journal of Monetary Policy* 104: 704-723 (1996).
- —, "A Small, Structural, Quarterly Model for Monetary Policy Evaluation," *Carnegie-Rochester Conference Series on Public Policy* 47: 83-108 (1997).
- Kim, Jinill, and Sunghyun Henry Kim, "Spurious Welfare Reversals in International Business Cycle Models," unpublished, University of Virginia, October 1999.
- King, Robert G., and Sergio Rebelo, "Resuscitating Real Business Cycles," in J.B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: North Holland, forthcoming 1999.
- King, Robert G., and Alexander L. Wolman, "What Should the Monetary Authority Do when Prices are Sticky?" in J.B. Taylor, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press, forthcoming 1999.
- McCallum, Bennett T., and Edward Nelson, "Timeless Perspective vs. Discretionary Monetary Policy in Forward-Looking Models," NBER Working Paper no. 7915, September 2000.
- Poole, William, "Optimal Choice of Monetary Policy Instrument in a Simple Stochastic Macro Model," *Quarterly Journal of Economics* 84: 197-216 (1970).
- Roberts, John M., "New Keynesian Economics and the Phillips Curve," *Journal of Money, Credit and Banking* 27: 975-984 (1995).
- Rotemberg, Julio J., "Prices, Output and Hours: An Empirical Analysis Based on a Sticky Price Model," *Journal of Monetary Economics* 37: 505-533 (1995).
- —, and Michael Woodford, "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual* 12: 297-346

- (1997). [Expanded version available as NBER Technical Working Paper no. 233, May 1998.]
- —, and — —, “Interest-Rate Rules in an Estimated Sticky-Price Model,” in J.B. Taylor, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press, 1999.
- Rudebusch, Glenn D., and Lars E.O. Svensson, “Policy Rules for Inflation Targeting,” in J.B. Taylor, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press, forthcoming 1999.
- Steinsson, Jon, “Optimal Monetary Policy in an Economy with Inflation Persistence,” unpublished, Princeton University, May 2000.
- Summers, Lawrence, “How Should Long Term Monetary Policy Be Determined?” *Journal of Money, Credit and Banking*, 23, August 1991, 625-31.
- Svensson, Lars E.O., “Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets,” *European Economic Review* 41: 1111-1147 (1997).
- —, “Inflation Targeting as a Monetary Policy Rule,” *Journal of Monetary Economics*, forthcoming 1999a.
- — and Michael Woodford, “Implementing Optimal Policy through Inflation-Forecast Targeting,” unpublished, Princeton University, November 1999.
- — and — —, “Indicator Variables for Optimal Policy,” unpublished, Princeton University, September 2000.
- Taylor, John B., “Aggregate Dynamics and Staggered Contracts,” *Journal of Political Economy* 88: 1-24 (1980).
- Walsh, Carl E., *Monetary Theory and Policy*, Cambridge: M.I.T. Press, 1998.
- West, Kenneth D., “Targeting Nominal Income: A Note,” *Economic Journal* 96: 1077-1083 (1986).
- Williams, John C., “Simple Rules for Monetary Policy,” Finance and Economics Discussion Series paper no. 1999-12, Federal Reserve Board, February 1999.
- Woodford, Michael, “The Optimum Quantity of Money,” in B.M. Friedman and F.H. Hahn, eds. *Handbook of Monetary Economics*, vol. II, Amsterdam: North-Holland, 1990.

- —, “Doing without Money: Controlling Inflation in a Post-Monetary World,” *Review of Economic Dynamics* 1: 173-219 (1998).
- —, “Price-Level Determination under Interest-Rate Rules,” unpublished, Princeton University, April 1999a.
- —, “Optimal Monetary Policy Inertia,” NBER Working Paper no. 7261, July 1999b.
- —, “Inflation Stabilization and Welfare,” unpublished, Princeton University, August 1999c.
- —, “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, 1999d.
- —, “Optimizing Models with Nominal Rigidities,” unpublished, Princeton University, February 2000a.
- —, “A Neo-Wicksellian Framework for the Analysis of Monetary Policy,” unpublished, Princeton University, September 2000b.