

Optimal Monetary Stabilization Policy*

Michael Woodford
Columbia University

Revised October 2010

Contents

1	Optimal Policy in a Canonical New Keynesian Model	3
1.1	The Problem Posed	4
1.2	Optimal Equilibrium Dynamics	7
1.3	The Value of Commitment	13
1.4	Implementing Optimal Policy through Forecast Targeting	17
1.5	Optimality from a “Timeless Perspective”	24
1.6	Consequences of the Interest-Rate Lower Bound	31
1.7	Optimal Policy Under Imperfect Information	40
2	Stabilization and Welfare	44
2.1	Microfoundations of the Basic New Keynesian Model	44
2.2	Welfare and the Optimal Policy Problem	50
2.3	Local Characterization of Optimal Dynamics	54
2.4	A Welfare-Based Quadratic Objective	63
2.4.1	The Case of an Efficient Steady State	64
2.4.2	The Case of Small Steady-State Distortions	69
2.4.3	The Case of Large Steady-State Distortions	71
2.5	Second-Order Conditions for Optimality	74
2.6	When is Price Stability Optimal?	76
3	Generalizations of the Basic Model	79
3.1	Alternative Models of Price Adjustment	79
3.1.1	Structural Inflation Inertia	81
3.1.2	Sticky Information	88
3.2	Which Price Index to Stabilize?	92

*Prepared for the new (2010) volumes of the *Handbook of Monetary Economics*, edited by Benjamin M. Friedman and Michael Woodford. I would like to thank Ozge Akinci, Ryan Chahrour, V.V. Chari, Marc Giannoni and Ivan Werning for comments, Luminita Stevens for research assistance, and the National Science Foundation for research support under grant SES-0820438.

3.2.1	Sectoral Heterogeneity and Asymmetric Disturbances	93
3.2.2	Sticky Wages as Well as Prices	107
4	Research Agenda	111

This chapter reviews the theory of optimal monetary stabilization policy in New Keynesian models, with particular emphasis on developments since the treatment of this topic in Woodford (2003). The primary emphasis of the chapter is on methods of analysis that are useful in this area, rather than on final conclusions about the ideal conduct of policy (that are obviously model-dependent, and hence dependent on the stand that one might take on many issues that remain controversial), and on general themes that have been found to be important under a range of possible model specifications.¹ With regard to methodology, some of the central themes of this review will be the application of the method of Ramsey policy analysis to the problem of the optimal conduct of monetary policy, and the connection that can be established between utility maximization and linear-quadratic policy problems of the sort often considered in the central banking literature. With regard to the structure of a desirable decision framework for monetary policy deliberations, some of the central themes will be the importance of commitment for a superior stabilization outcome, and more generally, the importance of advance signals about the future conduct of policy; the advantages of history-dependent policies over purely forward-looking approaches; and the usefulness of a target criterion as a way of characterizing a central bank's policy commitment.

In this chapter, the question of monetary *stabilization* policy — *i.e.*, the proper monetary policy response to the various types of disturbances to which an economy may be subject — is somewhat artificially distinguished from the question of the optimal long-run inflation target, which is the topic of another chapter (Schmitt-Grohé and Uribe, 2010). This does not mean (except in section 1) that I simply take as given the desirability of stabilizing inflation around a long-run target that has been determined elsewhere; the kind of utility-based analysis of optimal policy expounded in section 2 has implications for the optimal long-run inflation rate as much as for the optimal response to disturbances, though it is the latter issue that is the focus of the discussion here. (The question of the optimal long-run inflation target is not entirely independent of the way in which one expects that policy should respond to shocks, either.) It is nonetheless reasonable to consider the two aspects of optimal policy in separate chapters, insofar as the aspects of the structure of the economy that are of greatest significance for the answer to one question are not entirely the same as those that matter most for the other. For example, the consequences of inflation for

¹Practical lessons of the modern literature on monetary stabilization policy are developed in more detail in the chapters by Taylor and Williams (2010) and by Svensson (2010) in this *Handbook*.

people’s incentive to economize on cash balances by conducting transactions in less convenient ways has been a central issue in the scholarly literature on the optimal long-run inflation target, and so must be discussed in detail by Schmitt-Grohé and Uribe (2010), whereas this particular type of friction has not played a central role in discussions of optimal monetary stabilization policy, and is abstracted from entirely in this chapter.²

Monetary stabilization policy is also analyzed here under the assumption (made explicit in the welfare-based analysis introduced in section 2) that a non-distorting source of government revenue exists, so that stabilization policy can be considered in abstraction from the state of the government’s budget and from the choice of fiscal policy. This is again a respect in which the scope of the present chapter has been deliberately restricted, because the question of the interaction between optimal monetary stabilization policy and optimal state-contingent tax policy is treated in another chapter of the *Handbook*, by Canzoneri *et al.* (2010). While the “special” case in which lump-sum taxation is possible might seem of little practical interest, I believe that an understanding of the principles of optimal monetary stabilization policy in the simpler setting considered in this chapter provides an important starting point for understanding the more complex problems considered in the literature reviewed by Canzoneri *et al.* (2010).³

In section 1, I introduce a number of central methodological issues and key themes of the theory of optimal stabilization policy, in the context of a familiar textbook example, in which the central bank’s objective is assumed to be the minimization of a conventional quadratic objective (sometimes identified with “flexible inflation targeting”), subject to the constraints implied by certain log-linear structural equations (sometimes called “the basic New Keynesian model”). In section 2, I then consider the connection between this kind of analysis and expected-utility-maximizing policy

²This does not mean that transactions frictions that result in a demand for money have no consequences for optimal stabilization policy; see *e.g.*, Woodford (2003, chap. 6, sec. 4.1) or Khan *et al.* (2003) for treatment of this issue. This is one of many possible extensions of the basic analysis presented here that are not taken up in this chapter, for reasons of space.

³From a practical standpoint, it is important not only to understand optimal monetary policy in an economy where only distorting sources of government revenue exist, but taxes are adjusted *optimally*, as in the literature reviewed by Canzoneri *et al.* (2010), but also when fiscal policy is sub-optimal owing to practical and/or political constraints. Benigno and Woodford (2007) offer a preliminary analysis of this less-explored topic.

in a New Keynesian model with explicit microfoundations. Methods that are useful in analyzing Ramsey policy and in characterizing the optimal policy commitment in microfounded models are illustrated in section 2 in the context of a relatively simple model that yields policy recommendations that are closely related to the conclusions obtained in section 1, so that the results of section 2 can be viewed as providing welfare-theoretic foundations for the more conventional analysis in section 1. However, once the association of these results with very specific assumptions about the model of the economy has been made, an obvious question is the extent to which similar conclusions would be obtained under alternative assumptions. Section 3 shows how similar methods can be used to provide a welfare-based analysis of optimal policy in several alternative classes of models, that introduce a variety of complications that are often present in empirical DSGE models of the monetary transmission mechanism. Section 4 concludes with a much briefer discussion of other important directions in which the analysis of optimal monetary stabilization policy can or should be extended.

1 Optimal Policy in a Canonical New Keynesian Model

In this section, I illustrate a number of fundamental insights from the literature on the optimal conduct of monetary policy, in the context of a simple but extremely influential example. In particular, this section shows how taking account of the way in which the effects of monetary policy depend on expectations regarding the future conduct of policy affects the problem of policy design. The general issues that arise as a result of forward-looking private-sector behavior can be illustrated in the context of a simple model in which the structural relations that determine inflation and output under given policy on the part of the central bank involve expectations regarding future inflation and output, for reasons that are not discussed until section 2. Here I shall simply take as given both the form of the model structural relations and the assumed objectives of stabilization policy, to illustrate the complications that arise as a result from forward-looking behavior, especially (in this section) the dependence of the aggregate-supply tradeoff at a point in time on the expected rate of inflation. I shall offer comments along the way about the extent to which the issues that arise in the analysis of this example are ones that occur in broader classes of stabilization

policy problems as well. The extent to which specific conclusions from this simple example can be obtained in a model with explicit microfoundations is then taken up in section 2.

1.1 The Problem Posed

I shall begin by recapitulating the analysis of optimal policy in the linear-quadratic problem considered by Clarida *et al.* (1999), among others.⁴ In a log-linear version of what is sometimes called the “basic New Keynesian model,” inflation π_t and (log) output y_t are determined by an aggregate-supply relation (often called the “New Keynesian Phillips curve”)

$$\pi_t = \kappa(y_t - y_t^n) + \beta E_t \pi_{t+1} + u_t \quad (1.1)$$

and an aggregate-demand relation (sometimes called the “intertemporal IS relation”)

$$y_t = E_t y_{t+1} - \sigma(i_t - E_t \pi_{t+1} - \rho_t). \quad (1.2)$$

Here i_t is a short-term nominal interest rate; y_t^n , u_t , and ρ_t are each exogenous disturbances; and the coefficients of the structural relations satisfy $\kappa, \sigma > 0$ and $0 < \beta < 1$. It may be wondered why there are two distinct exogenous disturbance terms in the aggregate-supply relation (the “cost-push shock” u_t in addition to allowance for shifts in the “natural rate of output” y_t^n); the answer is that the distinction between these two possible sources of shifts in the inflation-output tradeoff matters for the assumed stabilization objective of the monetary authority (as specified in (1.6) below).

The analysis of optimal policy is simplest if we treat the nominal interest rate as being directly under the control of the central bank, in which case equations (1.1)–(1.2) suffice to indicate the paths for inflation and output that can be achieved through alternative interest-rate policies. However, if one wishes to treat the central bank’s instrument as some measure of the money supply (perhaps the quantity of base money), with the interest rate being determined by the market given the central bank’s control of the money supply, one can do so by adjoining an additional equilibrium relation,

$$m_t - p_t = \eta_y y_t - \eta_i i_t + \epsilon_t^m, \quad (1.3)$$

⁴The notation used here follows the treatment of this model in Woodford (2003).

where m_t is the log money supply (or monetary base), p_t is the log price level, ϵ_t^m is an exogenous money-demand disturbance, $\eta_y > 0$ is the income elasticity of money demand, and $\eta_i > 0$ is the interest-rate semi-elasticity of money demand. Combining this with the identity

$$\pi_t \equiv p_t - p_{t-1},$$

one then has a system of four equations per period to determine the evolution of the four endogenous variables $\{y_t, p_t, \pi_t, i_t\}$ given the central bank's control of the path of the money supply.

In fact, the equilibrium relation (1.3) between the monetary base and the other variables should more correctly be written as a pair of inequalities,

$$m_t - p_t \geq \eta_y y_t - \eta_i i_t + \epsilon_t^m, \tag{1.4}$$

$$i_t \geq 0, \tag{1.5}$$

together with the complementary slackness requirement that at least one of the two inequalities must hold with equality at any point in time. Thus it is possible to have an equilibrium in which $i_t = 0$ (so that money is no longer dominated in rate of return⁵), but in which (log) real money balances exceed the quantity $\eta_y y_t + \epsilon_t^m$ required for the satiation of private parties in money balances — households or firms should be willing to freely hold the additional cash balances as long as they have a zero opportunity cost.

One observes that (1.5) represents an additional constraint on the possible paths for the variables $\{\pi_t, y_t, i_t\}$ beyond those reflected by the equations (1.1)–(1.2). However, if one assumes that the constraint (1.5) happens never to bind in the optimal policy problem, as in the treatment by Clarida *et al.* (1999),⁶ then one can not only replace the pair of relations (1.4)–(1.5) by the simple equality (1.3), one can furthermore neglect this subsystem altogether in characterizing optimal policy, and simply

⁵For simplicity, I assume here that money earns a zero nominal return. See, for example, Woodford, 2003, chaps. 2,4, for extension of the theory to the case in which the monetary base can earn interest. This elaboration of the theory has no consequences for the issues taken up in this section: it simply complicates the description of the possible actions that a central bank may take in order to implement a particular interest-rate policy.

⁶This is also true in the micro-founded policy problem treated in section 2, in the case that all stochastic disturbances are small enough in amplitude. See, however, section 1.6 below for an extension of the present analysis to the case in which the zero lower bound may temporarily be a binding constraint.

analyze the set of paths for $\{\pi_t, y_t, i_t\}$ consistent with conditions (1.1)–(1.2). Indeed, one can even dispense with condition (1.2), and simply analyze the set of paths for the variables $\{\pi_t, y_t\}$ consistent with the condition (1.1). Assuming an objective for policy that involves only the paths of these variables (as assumed in (1.6) below), such an analysis would suffice to determine the optimal state-contingent evolution of inflation and output. Given a solution for the desired evolution of the variables $\{\pi_t, y_t\}$, equations (1.2) and (1.3) can then be used to determine the required state-contingent evolution of the variables $\{i_t, m_t\}$ in order for monetary policy to be consistent with the desired paths of inflation and output.

Let us suppose that the goal of policy is to minimize a discounted loss function of the form

$$E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} [\pi_t^2 + \lambda(x_t - x^*)^2], \quad (1.6)$$

where $x_t \equiv y_t - y_t^n$ is the “output gap”, x^* is a target level for the output gap (positive, in the case of greatest practical relevance), and $\lambda > 0$ measures the relative importance assigned to output-gap stabilization as opposed to inflation stabilization. Here (1.6) is simply assumed as a simple representation of conventional central-bank objectives; but a welfare-theoretic foundation for an objective of precisely this form is given in section 2. It should be noted that the discount factor β in (1.6) is the same as the coefficient on inflation expectations in (1.1). This is not accidental; it is shown in section 2 that when microfoundations are provided for both the aggregate-supply tradeoff and the stabilization objective, the same factor β (indicating the rate of time preference of the representative household) appears in both expressions.⁷

Given the objective (1.6), it is convenient to write the model structural relations in terms of the same two variables (inflation and the output gap) that appear in the policymaker’s objective function. Thus we rewrite (1.1)–(1.2) as

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1} + u_t, \quad (1.7)$$

$$x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1} - r_t^n), \quad (1.8)$$

⁷If one takes (1.6) to simply represent central-bank preferences (or perhaps the bank’s legislative mandate), that need not coincide with the interests of the representative household, the discount factor in (1.6) need not be the same as the coefficient in (1.1). The consequences of assuming different discount factors in the two places are considered by Kirsanova *et al.* (2009).

where

$$r_t^n \equiv \rho_t + \sigma^{-1}[E_t y_{t+1}^n - y_t^n]$$

is the “natural rate of interest,” *i.e.*, the (generally time-varying) real rate of interest required each period in order to keep output equal to its natural rate at all times.⁸ Our problem is then to determine the state-contingent evolution of the variables $\{\pi_t, x_t, i_t\}$ consistent with structural relations (1.7)–(1.8) that will minimize the loss function (1.6).

Supposing that there is no constraint on the ability of the central bank to adjust the level of the short-term interest rate i_t as necessary to satisfy (1.7), then the optimal paths of $\{\pi_t, x_t\}$ are simply those paths that minimize (1.6) subject to the constraint (1.7). The form of this problem immediately allows some important conclusions to be reached. The solution for the optimal state-contingent paths of inflation and the output gap depends only on the evolution of the exogenous disturbance process $\{u_t\}$ and *not* on the evolution of the disturbances $\{y_t^n, \rho_t, \epsilon_t^m\}$, to the extent that disturbances of these latter types have no consequences for the path of $\{u_t\}$. One can further distinguish between shocks of the latter three types in that disturbances to the path of $\{y_t^n\}$ should affect the path of output (though not the output gap), while disturbances to the path of $\{\rho_t\}$ (again, to the extent that these are independent of the expected paths of $\{y_t^n, u_t\}$) should be allowed to affect neither inflation nor output, but only the path of (both nominal and real) interest rates and the money supply, and disturbances to the path of $\{\epsilon_t^m\}$ (if without consequences for the other disturbance terms) should not be allowed to affect inflation, output, or interest rates, but only the path of the money supply (which should be adjusted to completely accommodate these shocks). The effects of disturbances to the path of $\{y_t^n\}$ on the path of $\{y_t\}$ should also be of an especially simple form under optimal policy: actual output should respond one-for-one to variations in the natural rate of output, so that such variations have no effect on the path of the output gap.

1.2 Optimal Equilibrium Dynamics

The characterization of optimal equilibrium dynamics is simple in the case that only disturbances of the two types $\{y_t^n, \rho_t\}$ occur, given the remarks at the end of the previous section. However, the existence of “cost-push shocks” u_t creates a tension

⁸For further discussion of this concept, see Woodford (2003, chap. 4).

between the goals of inflation and output stabilization,⁹ in which case the problem is less trivial; an optimal policy must balance the two goals, neither of which can be given absolute priority. This case is of particular interest, since it also introduces dynamic considerations — a difference between optimal policy under commitment from the outcome of discretionary optimization, superiority of history-dependent policy over purely forward-looking policy — that are in fact quite pervasive in contexts where private-sector behavior is forward-looking, and can occur for reasons having nothing to do with “cost-push shocks,” even though in the present (very simple) model they arise only when we assume that the $\{u_t\}$ terms have non-zero variance.

It suffices, as discussed in the previous section, to consider the state-contingent paths $\{\pi_t, x_t\}$ that minimize (1.6) subject to the constraint that condition (1.7) be satisfied for each $t \geq t_0$. We can write a Lagrangian for this problem

$$\begin{aligned} \mathcal{L}_{t_0} &= E_{t_0} \sum_{t=0}^{\infty} \beta^{t-t_0} \left\{ \frac{1}{2} [\pi_t^2 + \lambda(x_t - x^*)^2] + \varphi_t [\pi_t - \kappa x_t - \beta E_t \pi_{t+1}] \right\} \\ &= E_{t_0} \sum_{t=0}^{\infty} \beta^{t-t_0} \left\{ \frac{1}{2} [\pi_t^2 + \lambda(x_t - x^*)^2] + \varphi_t [\pi_t - \kappa x_t - \beta \pi_{t+1}] \right\}, \end{aligned}$$

where φ_t is a Lagrange multiplier associated with constraint (1.7), and hence a function of the state of the world in period t (since there is a distinct constraint of this form for each possible state of the world at that date). The second line has been simplified using the law of iterated expectations to observe that

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

Differentiation of the Lagrangian then yields first-order conditions

$$\pi_t + \varphi_t - \varphi_{t-1} = 0, \tag{1.9}$$

$$\lambda(x_t - x^*) - \kappa \varphi_t = 0, \tag{1.10}$$

for each $t \geq t_0$, where in (1.9) for $t = t_0$ we substitute the value

$$\varphi_{t_0-1} = 0, \tag{1.11}$$

⁹The economic interpretation of this residual in the aggregate-supply relation (1.7) is discussed further in section 2.

as there is in fact no constraint required for consistency with a period $t_0 - 1$ aggregate-supply relation if the policy is being chosen after period $t_0 - 1$ private decisions have already been made.

Using (1.9) and (1.10) to substitute for π_t and x_t respectively in (1.7), we obtain a stochastic difference equation for the evolution of the multipliers,

$$E_t \left[\beta \varphi_{t+1} - \left(1 + \beta + \frac{\kappa^2}{\lambda} \right) \varphi_t + \varphi_{t-1} \right] = \kappa x^* + u_t, \quad (1.12)$$

that must hold for all $t \geq 0$, along with the initial condition (1.11).

The characteristic equation

$$\beta \mu^2 - \left(1 + \beta + \frac{\kappa^2}{\lambda} \right) \mu + 1 = 0 \quad (1.13)$$

has two real roots

$$0 < \mu_1 < 1 < \mu_2,$$

as a result of which (1.12) has a unique bounded solution in the case of any bounded process for the disturbances $\{u_t\}$. Writing (1.12) in the alternative form

$$E_t[\beta(1 - \mu_1 L)(1 - \mu_2 L)\varphi_{t+1}] = \kappa x^* + u_t,$$

standard methods easily show that the unique bounded solution is of the form

$$(1 - \mu_1 L)\varphi_t = -\beta^{-1}\mu_2^{-1}E_t[(1 - \mu_2^{-1}L^{-1})^{-1}(\kappa x^* + u_t)],$$

or alternatively,

$$\varphi_t = \mu \varphi_{t-1} - \mu \sum_{j=0}^{\infty} \beta^j \mu^j [\kappa x^* + E_t u_{t+j}], \quad (1.14)$$

where I now simply write μ for the smaller root (μ_1) and use the fact that $\mu_2 = \beta^{-1}\mu_1^{-1}$ to eliminate μ_2 from the equation.

This is an equation that can be solved each period for φ_t given the previous period's value of the multiplier and current expectations regarding current and future "cost-push" terms. Starting from the initial condition (1.11), and given a law of motion for $\{u_t\}$ that allows the conditional expectations to be computed, it is possible to solve (1.14) iteratively for the complete state-contingent evolution of the multipliers. Substitution of this solution into (1.9)–(1.10) allows one to solve for the implied

state-contingent evolution of inflation and output. Substitution of these solutions in turn into (1.8) then yields the implied evolution of the nominal interest rate, and substitution of all of these solutions into (1.3) yields the implied evolution of the money supply.

The solution for the optimal path of each variable can be decomposed into a deterministic part — representing the expected path of the variable before anything is learned about the realizations of the disturbances $\{u_t\}$, including the value of u_{t_0} — and a sum of additional terms indicating the perturbations of the variable’s value in any period t due to the shocks realized in each of periods t_0 through t . Here the relevant shocks include all events that change the expected path of the disturbances $\{u_t\}$, including “news shocks” at date t or earlier that only convey information about cost-push terms at dates later than t ; but they do not include events that change the value of our convey information about the variables $\{y_t^n, \rho_t, \epsilon_t^m\}$, without any consequences for the expected path of $\{u_t\}$.

If we assume that the unconditional (or ex ante) expected value of each of the cost-push terms is zero, then the deterministic part of the solution for $\{\varphi_t\}$ is given by

$$\bar{\varphi}_t = -\frac{\lambda}{\kappa}x^*(1 - \mu^{t-t_0+1})$$

for all $t \geq t_0$. The implied deterministic part of the solution for the path of inflation is given by

$$\bar{\pi}_t = (1 - \mu)\frac{\lambda}{\kappa}x^*\mu^{t-t_0} \tag{1.15}$$

for all $t \geq t_0$.

An interesting feature of this solution is that the optimal long-run average rate of inflation should be zero, regardless of the size of x^* and of the relative weight λ attached to output-gap stabilization. It should not surprise anyone to find that the optimal average inflation rate is zero if $x^* = 0$, so that a zero average inflation rate implies $x_t = x^*$ on average; but it might have been expected that when a zero average inflation rate implies $x_t < x^*$ on average, an inflation rate that is above zero on average forever would be preferable. This turns out not to be the case, despite the fact that the New Keynesian Phillips curve (1.7) implies that a higher average inflation rate would indeed result in at least slightly higher average output forever. The reason is that an increase in the inflation rate aimed at (and anticipated) for some period $t > t_0$ lowers average output in period $t - 1$ in addition to raising average

output in period t , as a result of the effect of the higher expected inflation on the Phillips-curve tradeoff in period $t - 1$. And even though the factor β in (1.7) implies the reduction of output in period $t - 1$ is not quite as large as the increase in output in period t (this is the reason that permanently higher average inflation would imply permanently higher average output), the discounting in the objective function (1.6) implies that the policymaker's objective is harmed as much (to first order) by the output loss in period $t - 1$ as it is helped by the output gain in period t . The first-order effects on the objective therefore cancel; the second-order effects make a departure from the path specified in (1.15) strictly worse.

Another interesting feature of our solution for the optimal state-contingent path of inflation is that the *price level* p_t should be *stationary*: while cost-push shocks are allowed to affect the inflation rate under an optimal policy, any increase in the price level as a consequence of a positive cost-push shock must subsequently be undone (through lower-than-average inflation after the shock), so that the expected long-run price level is unaffected by the occurrence of the shock. This can be seen by observing that (1.9) can alternatively be written

$$p_t + \varphi_t = p_{t-1} + \varphi_{t-1}, \quad (1.16)$$

which implies that the cumulative change in the (log) price level over any horizon must be the additive inverse of the cumulative change in the Lagrange multiplier over the same horizon. Since (1.14) implies that the expected value of the Lagrange multiplier far in the future never changes (assuming that $\{u_t\}$ is a stationary, and hence mean-reverting, process), it follows that the expected price level far in the future can never change, either. This suggests that a version of price-level targeting may be a convenient way of bringing about inflation dynamics of the desired sort, as is discussed further below in section 1.4.

As a concrete example, suppose that u_t is an i.i.d., mean-zero random variable, the value of which is learned only at date t . In this case, (1.14) reduces to

$$\tilde{\varphi}_t = \mu \tilde{\varphi}_{t-1} - \mu u_t,$$

where $\tilde{\varphi}_t \equiv \varphi_t - \bar{\varphi}_t$ is the non-deterministic component of the path of the multiplier. Hence a positive cost-push shock at some date temporarily makes φ_t more negative, after which the multiplier returns (at an exponentially decaying rate) to the path it had previously been expected to follow. This impulse response of the multiplier




Figure 1: Impulse responses to a transitory cost-push shock under an optimal policy commitment, and in the Markov-perfect equilibrium with discretionary policy.

to the shock implies impulse responses for the inflation rate, output (and similarly the output gap), and the log price level of the kind shown in Figure 1.¹⁰ (Here the solid line in each panel represents the impulse response under an optimal policy commitment.) Note that both output and the log price level return to the paths that would have been expected in the absence of the shock at the same exponential rate as does the multiplier.

¹⁰The figure reproduces Figure 7.3 from Woodford (2003), where the numerical parameter values used are discussed. The alternative assumption of discretionary policy is discussed in the next section.

1.3 The Value of Commitment

An important general observation about this characterization of the optimal equilibrium dynamics is that they do *not* correspond to the equilibrium outcome in the case of an optimizing central bank that chooses its policy each period without making any commitments about future policy decisions. Sequential decisionmaking of that sort is not equivalent to the implementation of an optimal plan chosen once and for all, even when each of the sequential policy decisions is made with a view to achievement of the same policy objective (1.6). The reason is that in the case of what is often called *discretionary* policy,¹¹ a policymaker has no reason, when making a decision at a given point in time, to take into account the consequences for her own success in achieving her objectives at an earlier time of people's having been able to anticipate a different decision at the present time. And yet, if the outcomes achieved by policy depend not only on the current policy decision but also on expectations about future policy, it will quite generally be the case that outcomes can be improved, at least to some extent, through strategic use of the tool of modifying intended later actions precisely for the sake of inducing different expectations at an earlier time. For this reason, implementation of an optimal policy requires advance *commitment* regarding policy decisions, in the sense that one must not imagine that it is proper to optimize afresh each time a choice among alternative actions must be taken. Some procedure must be adopted that internalizes the effects of predictable patterns in policy on expectations; what sort of procedure this might be in practice is discussed further in section 1.4.

The difference that can be made by a proper form of commitment can be illustrated by comparing the optimal dynamics, characterized in the previous section, with the equilibrium dynamics in the same model if policy is made through a process of discretionary (sequential) optimization. Here I shall assume that in the case of

¹¹It is worth noting that the critique of “discretion” offered here has nothing to do with what that word often means, namely, the use of judgment about the nature of a particular situation of a kind that cannot easily be reduced to a mechanical function of a small number of objectively measurable quantities. Policy can often be improved by the use of more information, including information that may not be easily quantified or agreed upon. If one thinks that such information can only be used by a policymaker that optimizes afresh at each date, then there may be a close connection between the two concepts of “discretion,” but this is not obviously true. On the use of judgment in implementing optimal policy, see Svensson (2003, 2005).

discretion, the outcome is the one that represents a Markov perfect equilibrium of the non-cooperative “game” among successive decisionmakers.¹² This means that I shall assume that equilibrium play at any date is a function only of states that are relevant for determining the decisionmakers’ success at achieving their goals from that date onward.¹³

Let s_t be a state vector that includes all information available at date t about the path $\{u_{t+j}\}$ for $j \geq 0$.¹⁴ Then since the objectives of policymakers from date t onward depend only on inflation and output-gap outcomes from date t onward, in a way that is independent of outcomes prior to date t (owing to the additive separability of the loss function (1.6)), and since the possible rational-expectations equilibrium evolutions of inflation and output from date t onward depend only on the cost-push shocks from date t onward, independently of the economy’s history prior to date t (owing to the absence of any lagged variables in the aggregate-supply relation (1.7)), in a Markov perfect equilibrium both π_t and x_t should depend only on the current state vector s_t . Moreover, since both policymakers and the public should understand that inflation and the output gap at any time are determined purely by factors independent of past monetary policy, the policymaker at date t should not believe that her period t decision has any consequences for the probability distribution of inflation or the output gap in periods later than t , and private parties should have expectations regarding inflation in periods later than t that are unaffected by policy decisions in period t .

It follows that the discretionary policymaker in period t expects her decision to

¹²In the case of optimization without commitment, one can equivalently suppose that there is not a single decisionmaker, but a sequence of decisionmakers, each of whom chooses policy for only one period. This makes it clear that even though each decision results from optimization, an individual decision may not be made in a way that takes account of the consequences of the decision for the success of the “other” decisionmakers.

¹³There can be other equilibria of this “game” as well, but I shall not seek to characterize them here. Apart from the appeal of this refinement of Nash equilibrium, I would assert that even the *possibility* of a bad equilibrium as a result of discretionary optimization is a reason to try to design a procedure that would exclude such an outcome; it is not necessary to argue that this particular equilibrium is the *inevitable* outcome.

¹⁴In the case of the i.i.d. cost-push shocks considered above, s_t consists solely of the current value of u_t . But if u_t follows an AR(k) process, s_t consists of $(u_t, u_{t-1}, \dots, u_{t-(k-1)})$, and so on.

affect only the values of the terms

$$\pi_t^2 + \lambda(x_t - x^*)^2 \quad (1.17)$$

in the loss function (1.6); all other terms are either already given by the time of the decision or expected to be determined by factors that will not be changed by the current period's decision. Inflation expectations $E_t\pi_{t+1}$ will be given by some quantity π_t^e that depends on the economy's state in period t but that can be taken as given by the policymaker. Hence the discretionary policymaker (correctly) understands that she faces a tradeoff of the form

$$\pi_t = \kappa x_t + \beta\pi_t^e + u_t \quad (1.18)$$

between the achievable values of the two variables that can be affected by current policy. The policymaker's problem in period t is therefore simply to choose values (π_t, x_t) that minimize (1.17) subject to the constraint (1.18). (The required choices for i_t or m_t in order to achieve this outcome are then implied by the other model equations.) The solution to this problem is easily seen to be

$$\pi_t = \frac{\lambda}{\kappa^2 + \lambda} [\kappa x^* + \beta\pi_t^e + u_t]. \quad (1.19)$$

A (Markov-perfect) rational-expectations equilibrium is then a pair of functions $\pi(s_t), \pi^e(s_t)$ such that (i) $\pi(s_t)$ is the solution to (1.19) if one substitutes $\pi_t^e = \pi^e(s_t)$, and (ii) $\pi^e(s_t) = E(\pi(s_{t+1})|s_t)$, given the law of motion for the exogenous state $\{s_t\}$. The solution is easily seen to be

$$\pi_t = \pi(s_t) \equiv \tilde{\mu} \sum_{j=0}^{\infty} \beta^j \tilde{\mu}^j [\kappa x^* + E_t u_{t+j}], \quad (1.20)$$

where

$$\tilde{\mu} \equiv \frac{\lambda}{\kappa^2 + \lambda}.$$

One can show that $\mu < \tilde{\mu} < 1$, where μ is the coefficient that appears in the optimal policy equation (1.14).

There are a number of important differences between the evolution of inflation chosen by the discretionary policymaker and the optimal commitment characterized in the previous section. The deterministic component of the solution (1.20) is a constant positive inflation rate (in the case that $x^* > 0$). This is not only obviously




Figure 2: The paths of inflation under discretionary policy, under unconstrained Ramsey policy (the “time-zero-optimal” policy), and under a policy that is “optimal from a timeless perspective.”

higher than the average inflation rate implied by (1.15) in the long run (which is zero); one can show that it is higher than the inflation rate that is chosen under the optimal commitment even initially. (Figure 2 illustrates the difference between the time paths of the deterministic component of inflation under the two policies, in a numerical example.¹⁵) This is the much-discussed “inflationary bias” of discretionary monetary policy.

The outcome of discretionary optimization differs from optimal policy also with regard to the response to cost-push shocks; and this second difference exists regardless of the value of x^* . Equation (1.20) implies that under discretion, inflation at any date t depends only on current and expected future cost-push shocks at that time. This

¹⁵This reproduces Figure 7.1 from Woodford (2003); the numerical assumptions are discussed there. The figure also shows the path of inflation under a third alternative, optimal policy from a “timeless perspective,” discussed in section 1.5.

means that there is no correction for the effects of past shocks on the price level — the rate of inflation at any point in time is independent of the past history of shocks (except insofar as they may be reflected in current or expected future cost-push terms) — as a consequence of which there will be a unit root in the path of the price level. For example, in the case of i.i.d. cost-push shocks, (1.20) reduces to

$$\pi_t = \bar{\pi} + \tilde{\mu}u_t,$$

where the average inflation rate is $\bar{\pi} = \tilde{\mu}\kappa x^*/(1 - \beta\tilde{\mu}) > 0$. In this case, a transitory cost-push shock immediately increases the log price level by more than under the optimal commitment (by $\tilde{\mu}u_t$ rather than only by μu_t), and the increase in the price level is *permanent*, rather than being subsequently undone. (See Figure 1 for a comparison between the responses under discretionary policy and those under optimal policy in the numerical example; the discretionary responses are shown by the dashed line.)

These differences both follow from a single principle: the discretionary policymaker does not taken into account the consequences of (predictably) choosing a higher inflation rate in the current period on expected inflation, and hence upon the location of the Phillips-curve tradeoff, in the previous period. Because the neglected effect of higher inflation on previous expected inflation is an adverse one, in the case that $x^* > 0$ (so that the policymaker would wish to shift the Phillips curve down if possible), neglect of this effect leads the discretionary policymaker to choose a higher inflation rate at all times than would be chosen under an optimal commitment. And because this neglected effect is especially strong immediately following a positive cost-push shock, the gap between the inflation rate chosen under discretion and the one that would be chosen under an optimal policy is even larger than average at such a time.

1.4 Implementing Optimal Policy through Forecast Targeting

Thus far, I have discussed the optimal policy commitment as if the policy authority should solve a problem of the kind considered above at some initial date to determine the optimal state-contingent evolution of the various endogenous variables, and then commit itself to follow those instructions forever after, simply looking up the

calculated optimal quantities for whatever state of the world it finds itself in at any later date. Such a thought experiment is useful for making clear the reason why a policy authority should wish to arrange to behave in a different way than the one that would result from discretionary optimization. But such an approach to policy is not feasible in practice.

Actual policy deliberations are conducted sequentially, rather than once and for all, for a simple reason: policymakers have a great deal of fine-grained information about the specific situation that has arisen, once it arises, without having any corresponding ability to list all of the situations that may arise very far in advance. Thus it is desirable to be able to implement the optimal policy through a procedure that only requires that the economy's current state — including the expected future paths of the relevant disturbances, conditional upon the state that has been reached — be recognized once it has been reached, that allows a correct decision about the current action to be reached based on this information. A view of the expected forward path of policy, conditional upon current information, may also be reached, and in general this will be necessary in order to determine the right current action; but this need not involve formulating a definite intention in advance about the responses to all of the *unexpected* developments that may arise at future dates. At the same time, if it is to implement the optimal policy, the sequential procedure must not be the kind of sequential optimization that has been described above as “discretionary policy.”

An example of a suitable sequential procedure is similar to *forecast targeting* as practiced by a number of central banks. In this approach, a contemplated forward path for policy is judged correct to the extent that quantitative projections for one or more economic variables, conditional on the contemplated policy, conform to a *target criterion*.¹⁶ The optimal policy computed in section 1.2 can easily be described in terms of the fulfillment of a target criterion. One easily sees that conditions (1.9)–(1.11) imply that the joint evolution of inflation and the output gap must satisfy

$$\pi_t + \phi(x_t - x_{t-1}) = 0 \tag{1.21}$$

for all $t > t_0$, and

$$\pi_{t_0} + \phi(x_{t_0} - x^*) = 0 \tag{1.22}$$

in period t_0 , where $\phi \equiv \lambda/\kappa > 0$. Conversely, in the case of any paths $\{\pi_t, x_t\}$ satisfying (1.21)–(1.22), there will exist a Lagrange multiplier process $\{\varphi_t\}$ (suitably

¹⁶See, e.g., Svensson (1997, 2005), Svensson and Woodford (2005) and Woodford (2007).

bounded if the inflation and output-gap processes are) such that the first-order conditions (1.9)–(1.11) are satisfied in all periods. Hence verification that a particular contemplated state-contingent evolution of inflation and output from period t_0 onward satisfy the target criteria (1.21)–(1.22) at all times, in addition to satisfying certain bounds and being consistent with the structural relation (1.7) at all times (and therefore representing a feasible equilibrium path for the economy), suffices to ensure that the evolution in question is the optimal one.

The target criterion can furthermore be used as the basis for a sequential procedure for policy deliberations. Suppose that at each date t at which another policy action must be taken, the policy authority verifies the state of the economy at that time — which in the present example means evaluating the state s_t that determines the set of feasible forward paths for inflation and the output gap, and the value of x_{t-1} , that is needed to evaluate the target criterion for period t — and seeks to determine forward paths for inflation and output (namely, the conditional expectations $\{E_t\pi_{t+j}, E_t x_{t+j}\}$ for all $j \geq 0$) that are feasible and that would satisfy the target criterion at all horizons. Assuming that $t > t_0$, the latter requirement would mean that

$$E_t\pi_{t+j} + \phi(E_t x_{t+j} - E_t x_{t+j-1}) = 0$$

at all horizons $j \geq 0$. One can easily show that there is a unique bounded solution for the forward paths of inflation and the output gap consistent with these requirements, for an arbitrary initial condition x_{t-1} and an arbitrary bounded forward path $\{E_t u_{t+j}\}$ for the cost-push disturbance.¹⁷ This means that a commitment to organize policy deliberations around the search for a forward path that conforms to the target criterion is both feasible, and sufficient to determine the forward path and hence the appropriate current action. (Associated with the unique forward paths for inflation and the output gap there will also be unique forward paths for the nominal interest rate and the money supply, so that the appropriate policy action will be determined, regardless of which variable is considered to be the policy instrument.)

By proceeding in this way, the policy authority's action at each date will be precisely the same as in the optimal equilibrium dynamics computed in section 1.2. Yet

¹⁷The calculation required to show this is exactly the same as the one used in section 1.2 to compute the unique bounded evolution for the Lagrange multipliers consistent with the first-order conditions. The conjunction of the target criterion with the structural equation (1.7) gives rise to a stochastic difference equation for the evolution of the output gap that is of exactly the same form as (1.12).

it is never necessary to calculate anything but the conditional expectation of the economy’s optimal forward path, looking forward from the particular state that has been reached at a given point in time. Moreover, the target criterion provides a useful way of communicating about the authority’s policy commitment, both internally and with the public, since it can be stated in a way that does not involve any reference to the economy’s state at the time of application of the rule: it simply states a relationship that the authority wishes to maintain between the paths of two endogenous variables, the form of which will remain the same regardless of the disturbances that may have affected the economy. This robustness of the optimal target criterion to alternative views of the types of disturbances that have affected the economy in the past or that are expected to affect it in the future is a particular advantage of this way of describing a policy commitment.¹⁸

The possibility of describing optimal policy in terms of the fulfillment of a target criterion is not special to the simple example treated above. Giannoni and Woodford (2010) establish for a very general class of optimal stabilization policy problems, including both backward-looking and forward-looking constraints, that it is possible to choose a target criterion — which, as here, is a linear relation between a small number of “target variables” that should be projected to hold at all future horizons — with the properties that (i) there exists a unique forward path that fulfills the target criterion, looking forward from any initial conditions (or at least from any initial conditions close enough to the economy’s steady state, in the case of a nonlinear model), and (ii) the state-contingent evolution so determined coincides with an optimal policy commitment (or at least, coincides with it up to a linear approximation, in the case of a nonlinear model). In the case that the objective of policy is given by (1.6), the optimal target criterion always involves only the projected paths of inflation and the output gap, regardless of the complexity of the structural model of inflation and output determination.¹⁹ When the model’s constraints are purely forward-looking — by which I mean that past states have no consequences for the set

¹⁸For further comparison of this way of formulating a policy rule with other possibilities, see Woodford (2007).

¹⁹More generally, if the objective of policy is a quadratic loss function, the optimal target criterion involves only the paths of the “target variables” that appear in the loss function. The results of Giannoni and Woodford (2010) also apply, however, to problems in which the objective of policy is not given by a quadratic loss function; it may correspond, for example, to expected household utility, as in the problem treated in section 2.

of possible forward paths for the variables that matter to the policymaker’s objective function, as in the case considered here — the optimal target criterion is necessarily purely backward-looking, *i.e.*, it is a linear relation between current and past values of the target variables, as in equation (1.21). If, instead (as is more generally the case), lagged variables enter the structural equations, the optimal target criterion involves forecasts as well, for a finite number of periods into the future. (In the less relevant case that the model’s constraints are purely backward-looking — *i.e.*, they do not involve expectations — then the optimal target criterion is purely forward-looking, in the sense that it involves only the projected paths of the target variables in current and future periods.) Examples of optimal target criteria for more complex models are discussed below, and in Giannoni and Woodford (2005).

The targeting procedure described above can be viewed as a form of “flexible inflation targeting.”²⁰ It is a form of inflation targeting because the target criterion to which the policy authority commits itself, and that is to structure all policy deliberations, implies that the projected rate of inflation, looking far enough in the future, will never vary from a specific numerical value (namely, zero). This obviously follows from the requirement that (1.21) be projected to hold at all horizons, as long as the projected output gap is the same in all periods far enough in the future. Yet it is a form of *flexible* inflation targeting because the long-run inflation target is not required to hold at all times, nor is it even necessary for the central bank to do all in its power to bring the inflation rate as close as possible to the long-run target as soon as possible; instead, temporary departures of the inflation rate from the long-run target are tolerated to the extent that they are justified by projected near-term changes in the output gap. The conception of “flexible inflation targeting” advocated here differs, however, from the view that is popular at some central banks, according to which it suffices to specify a particular future horizon at which the long-run inflation target should be reached, without any need to specify what kinds of nearer-term projected paths for the economy are acceptable. The optimal target criterion derived here demands that a specific linear relation be verified both for nearer-term projections and for projections farther in the future; and it is the requirement that this linear relationship between the inflation projection and the output-gap projection be satisfied that determines how rapidly the inflation projection should converge to the long-run inflation target. (The optimal rate of convergence will not in fact be the same regard-

²⁰On the concept of flexible inflation targeting, see generally Svensson (2010).

less of the nature of the cost-push disturbance. Thus a fixed-horizon commitment to an inflation target will in general be simultaneously too *vague* a commitment to uniquely determine an appropriate forward path (and in particular to determine the appropriate current action), and too *specific* a commitment to be consistent with optimal policy.)

While the optimal target criterion has been expressed in (1.21)–(1.22) as a flexible inflation target, it can alternatively be expressed as a form of *price level target*. Note that (1.21) can alternatively be written as $\tilde{p}_t = \tilde{p}_{t-1}$, where $\tilde{p}_t \equiv p_t + \phi x_t$ is an “output-gap-adjusted price level.” Conditions (1.21)–(1.22) together can be seen to hold if and only if

$$\tilde{p}_t = p^* \tag{1.23}$$

for all $t \geq t_0$, where $p^* \equiv p_{t_0-1} + \phi x^*$. This is an example of the kind of policy rule that Hall (1984) has called an “elastic price standard.” A target criterion of this form makes it clear that the regime is one under which a rational long-run forecast of the price level never changes (it is always equal to p^*).

Which way of expressing the optimal target criterion is better? A commitment to the criterion (1.21)–(1.22) and a commitment to the criterion (1.23) are completely equivalent to one another, under the assumption that the central bank will be able to ensure that its target criterion is precisely fulfilled at all times. But this will surely not be true in practice, for a variety of reasons; and in that case, it makes a difference which criterion the central bank seeks to fulfill each time the decision process is repeated. With target misses, the criterion (1.23) incorporates a commitment to *error correction* — to aim at a lower rate of growth of the output-gap-adjusted price level following a target overshoot, or a higher rate following a target undershoot, so that over longer periods of time the cumulative growth is precisely the targeted rate despite the target misses — while the criterion (1.21) instead allows target misses to permanently shift the absolute level of prices.

A commitment to error-correction has important advantages from the standpoint of robustness to possible errors in real-time policy judgments. For example, Gorodnichenko and Shapiro (2006) note that commitment to a price-level target reduces the harm done by a poor real-time estimate of productivity (and hence of the natural rate of output) by a central bank. If the private sector expects that inflation greater than the central bank intended (owing to a failure to recognize how stimulative policy really was, on account of an overly optimistic estimate of the natural rate of output)

will cause the central bank to aim for lower inflation later, this will restrain wage and price increases during the period when policy is overly stimulative. Hence a commitment to error-correction would not only ensure that the central bank does not exceed its long-run inflation target in the same way for many years in a row; in the case of a forward-looking aggregate-supply tradeoff of the kind implied by (1.7), it would *also* result in less excess inflation in the first place, for any given magnitude of mis-estimate of the natural rate of output.²¹

Similarly, Aoki and Nikolov (2005) show that a price-level rule for monetary policy is more robust to possible errors in the central bank’s economic model. They assume that the central bank seeks to implement a target criterion — either (1.21) or (1.23) — using a quantitative model to determine the level of the short-term nominal interest rate that will result in inflation and output growth satisfying the criterion. They find that the price-level target criterion leads to much better outcomes when the central bank starts with initially incorrect coefficient estimates in the quantitative model that it uses to calculate its policy, again because the commitment to error-correction that is implied by the price-level target leads price-setters to behave in a way that ameliorates the consequences of central-bank errors in its choice of the interest rate.

Eggertsson and Woodford (2003) reach a similar conclusion (as discussed further in section 1.6) in the case that the lower bound on nominal interest rates sometimes prevents the central bank from achieving its target. A central bank that is committed to fulfill the criterion (1.21) whenever it can — and to simply keep interest rates as low as possible if the target is undershot even with interest rates at the lower bound — has very different consequences from a commitment to fulfill the criterion (1.23) whenever possible. Following a period in which the lower bound has required a central bank to undershoot its target, leading to both deflation and a negative output gap, continued pursuit of (1.23) will require a period of “reflation” in which policy is more inflationary than on average until the absolute level of the gap-adjusted price level again catches up to the target level, whereas pursuit of (1.21) would actually require policy to be more *deflationary* than average in the period just after the lower bound ceases to bind, owing to the negative lagged output gap as a legacy of the period

²¹In section 1.7, I characterize optimal policy in the case of imperfect information about the current state of the economy, including uncertainty about the current natural rate of output, and show that optimal policy does indeed involve error-correction — in fact, a somewhat stronger form of error-correction than even that implied by a simple price-level target.

in the “liquidity trap.” A commitment to reflation is in fact highly desirable, and if credible should go a long way toward mitigation of the effects of the binding lower bound. Hence while neither (1.21) nor (1.23) is a fully optimal rule in the case that the lower bound is sometimes a binding constraint, the latter rule provides a much better approximation to optimal policy in this case.

1.5 Optimality from a “Timeless Perspective”

In the previous section I have described a sequential procedure that can be used to bring about an optimal state-contingent evolution for the economy, assuming that the central bank succeeds in conducting policy so that the target criterion is perfectly fulfilled and that private agents have rational expectations. This requires, evidently, that the sequential procedure is not equivalent to the “discretionary” approach in which the policy committee seeks each period to determine the forward path for the economy that minimizes (1.6). Yet the target criterion that is the focus of policy deliberations under the recommended procedure *can* be viewed as a first-order condition for the optimality of policy, so that the search for a forward path consistent with the target criterion amounts to the solution of an optimization problem; it is simply not the *same* optimization problem as the one assumed in our account of discretionary policy in section 1.3. Instead, the target criterion (1.21) that is required to be satisfied at each horizon in the case of the decision process in any period $t > t_0$ can be viewed as a sequence of first-order conditions that characterize the solution to a problem which has been modified in order to internalize the consequences for expectations prior to date t of the systematic character of the policy decision at date t .

One way to modify the optimization problem in a way that makes the solution to an optimization problem in period t coincide with the *continuation* of the optimal state-contingent plan that would have been chosen in period t_0 (assuming that a once-and-for-all decision had been made then about the economy’s state-contingent evolution forever after) is to add an additional constraint of the form

$$\pi_t = \bar{\pi}(x_{t-1}; s_t), \tag{1.24}$$

where

$$\bar{\pi}(x_{t-1}; s_t) \equiv (1 - \mu) \frac{\lambda}{\kappa} (x_{t-1} - x^*) + \mu \sum_{j=0}^{\infty} \beta^j \mu^j [\kappa x^* + E_t u_{t+j}].$$

Note that (1.24) is a condition that holds under the optimal state-contingent evolution characterized earlier in every period $t > t_0$.²² If at date t one solves for the forward paths for inflation and output from date t onward that minimize (1.6), subject to the constraint that one can only consider paths consistent with the initial pre-commitment (1.24), then the solution to this problem will be precisely the forward paths that conform to the target criterion (1.21) from date t onward. It will also coincide with the continuation from date t onward of the state-contingent evolution that would have been chosen at date t_0 as the solution to the unconstrained Ramsey policy problem.

I have elsewhere (Woodford, 1999) referred to policies that solve this kind of modified optimization problem from some date forward as being “optimal from a timeless perspective,” rather than from the perspective of the particular time at which the policy is actually chosen. The idea is that such a policy, even if not what the policy authority would choose if optimizing afresh at date t , represents a policy that it should have been willing to commit itself to follow from date t onward if the choice had been made at some indeterminate point in the past, when its choice would have internalized the consequences of the policy for expectations prior to date t . Policies can be shown to have this property without actually solving for an optimal commitment at some earlier date, by looking for a policy that is optimal subject to an initial pre-commitment that has the property of *self-consistency*, by which I mean that the condition in question is one that a policymaker would choose to comply with each period under the constrained-optimal policy. Condition (1.24) is an example of a self-consistent initial pre-commitment, because in the solution to the constrained optimization problem stated above, the inflation rate in each period from t onward satisfies condition (1.24).²³

The study of policies that are optimal in this modified sense is of possible interest for several reasons. First of all, while the unconstrained Ramsey policy (as characterized in section 1.2 above) involves different behavior initially than the rule that the authority commits to follow later (illustrated by the difference between the target criterion (1.22) for period t_0 and the target criterion (1.21) for periods $t > t_0$), the policy that is optimal from a timeless perspective corresponds to a time-invariant

²²The condition can be derived from (1.9), using (1.14) to substitute for φ_t and then using (1.10) for period $t - 1$ to substitute for φ_{t-1} .

²³For further discussion and additional examples, see Woodford (2003, chap. 7).

policy rule (fulfillment of the target criterion (1.21) each period). This means that policies that are optimal from a timeless perspective are easier to describe.²⁴

This increase in the simplicity of the description of the optimal policy is especially great in the case of a nonlinear structural model of the kind considered in section 2. Also in an exact nonlinear model, the unconstrained Ramsey policy will involve an evolution of the kind shown in Figure 2 if every disturbance term takes its unconditional mean value: the initial inflation rate will be higher than the long-run value, in order to exploit the Phillips curve initially (given that inflation expectations prior to t_0 cannot be affected by the policy chosen), while also obtaining the benefits from a commitment to low inflation in later periods (when the consequences of expected inflation must also be taken into account). But this means that even in a local linear approximation to the optimal response of inflation and output to random disturbances, the linear approximation would have to be taken not around a deterministic steady state, but around this time-varying path, so that the derivatives that provide the coefficients of the linear approximation would be slightly different at each date. In the case of optimization subject to a self-consistent initial pre-commitment, instead, the optimal policy will involve constant values of all endogenous variables in the case that the exogenous disturbances take their mean values forever, and we can compute a local linear approximation to the optimal policy through a perturbation analysis conducted in the neighborhood of this deterministic steady state. This approach considerably simplifies the calculations involved in characterizing the optimal policy, even if now the characterization only describes the *asymptotic* nature of the unconstrained Ramsey policy, long enough after the initial date at which the optimal commitment was originally chosen. It is for the sake of this computational advantage that this approach is adopted in section 2, as in other studies of optimal policy in microfounded models such as King *et al.* (2003).

Consideration of policies that are optimal from a timeless perspective also provides a solution to an important conundrum for the theory of optimal stabilization policy. If achievement of the benefits of commitment explained in section 1.3 requires that a policy authority commit to a particular state-contingent policy for the indefinite future at the initial date t_0 , what should happen if the policy authority learns at

²⁴For example, in the deterministic case considered in Figure 2, an initial pre-commitment of the form $\pi_0 = \bar{\pi}$ is self-consistent if and only if $\bar{\pi} = 0$. In this case, the constrained-optimal policy is simply $\pi_t = 0$ for all $t \geq 0$, as shown in the figure.

some later date that the model of the economy on the basis of which it calculated the optimal policy commitment at date t_0 is no longer accurate (if, indeed, it ever was)? It is absurd to suppose that commitment should be possible because a policy authority should have complete knowledge of the true model of the economy and this truth should never change.

Yet it is also unsatisfactory to suppose that a commitment should be made that applies only as long as the authority's model of the economy does not change, with an optimal commitment to be chosen afresh as the solution to an unconstrained Ramsey problem whenever a new model is adopted. For even if it is not predictable in advance exactly *how* one's view of the truth will change, it is predictable that it *will* change, if only because additional data should allow more precise estimation of unknown structural parameters, even in a world without structural change. And if it is known that re-optimization will occur periodically, and that an initial burst of inflation will be chosen each time that it does — on the ground that in the “new” optimization at some date t , inflation expectations prior to date t are taken as given — then the inflation that occurs initially following a re-optimization should not in fact be entirely unexpected. Thus the benefits from a commitment to low inflation will not be fully achieved, nor will the assumptions made in the calculation of the original Ramsey policy be correct. (Similarly, the benefits from a commitment to subsequently reversing the price-level effects of cost-push shocks will not be fully achieved, owing to the recognition that the follow-through on this commitment will be truncated in the event that the central bank reconsiders its model.) The problem is especially severe if one recognizes that new information about model parameters will be received *continually*. If a central bank is authorized to re-optimize whenever it changes its model, it would have a motive to re-optimize each period (using as justification some small changes in model parameters) — in the absence, that is, of a commitment not to behave in this way. But a “model-contingent commitment” of this kind would be indistinguishable from discretion.

This problem can be solved if the central bank commits itself to select a new policy that is again optimal *from a timeless perspective* each time it revises its model of the economy. Under this principle, it would not matter if the central bank announces an inconsequential “revision” of its model each period: assuming no material change in the bank's model of the economy, choice of a rule that is optimal from a timeless perspective according to that model should lead it to choose a continuation of the

same policy commitment each period, so that the outcome would be the same (and should be forecasted to be the same) as if a policy commitment had simply been made at an initial date with no allowance for subsequent reconsideration. On the other hand, in the event of a genuine change in the bank’s model of the economy, a policy rule (say, a new target criterion) appropriate to the new model *could* be adopted. The expectation that this will happen from time to time should not undermine the expectations that the policy commitment chosen under the original model was trying to create, given that people should have no reason to expect the new policy rule to differ in any particular direction from the one that is expected to be followed if there is no model change.

This proposal leads us to be interested in the problem of finding a time-invariant policy that is “optimal from a timeless perspective,” in the case of any given model of the economy. Some have, however, objected to the selection of a policy rule according to this criterion, on the ground that, even one wishes to choose a time-invariant policy rule (unlike the unconstrained Ramsey policy), there will in general be *other* time-invariant policy rules that would be superior in the sense of implying a lower expected value of the loss function (1.6) at the time that the policy rule is chosen. For example, in the case of a loss function with $x^* = 0$, Blake (2001) and McCallum and Jensen (2002) argue that even if one restricts one’s attention to policies described by time-invariant linear target criteria linking π_t , x_t and x_{t-1} , one can achieve a lower expected value of (1.6) by requiring that

$$\pi_t + \phi(x_t - \beta x_{t-1}) = 0 \tag{1.25}$$

hold each period, rather than (1.21).²⁵ (Here $\phi \equiv \lambda/\kappa$, just as in (1.21).) By comparison with a policy that requires that (1.21) hold for each $t \geq t_0$, the alternative policy does not require inflation and the output gap to depart as far from their optimal values in period t_0 simply because the initial lagged output gap x_{t_0-1} happens to have been nonzero. (Recall that under the unconstrained Ramsey policy, the value of x_{t_0-1} would have no effect on policy from t_0 onward at all.) The fact that criterion (1.25) rather than (1.21) is applied in periods $t > t_0$ *increases* expected losses (that is why (1.21) holds for all $t > t_0$ under the Ramsey policy), but given the constraint

²⁵Under the criterion proposed by these authors, one would presumably also choose a long-run inflation target slightly higher than zero, in the case that $x^* > 0$; but I here consider only the case in which $x^* = 0$, following their exposition.

that one must put the same coefficient on x_{t-1} in all periods, one can nonetheless reduce the overall discounted sum of losses by using a coefficient slightly smaller than the one in (1.21).

This result does not contradict any of our analysis above; the claim that a policy under which (1.21) is required to hold in all periods is “optimal from a timeless perspective” implies only that it minimizes (1.6) among the class of policies that also conform to the additional constraint (1.24) — or alternatively, that it minimizes a modified loss function, with an additional term added to impose a penalty for violations of this initial pre-commitment — and not that it must minimize (1.6) within a class of policies that do not all satisfy the initial pre-commitment. But does the proposal of Blake, Jensen and McCallum provide a more attractive solution to the problem of making continuation of the recommended policy rule time-consistent, in the sense that a reconsideration at a later date should lead the policy authority to choose precisely the same rule again, in the absence of any change in its model of the economy?

In fact it does not. If one imagines that at any date t_0 , the authority may reconsider its policy rule, and choose a new target criterion from among the general family

$$\pi_t + \phi_1 x_t - \phi_2 x_{t-1} = 0 \tag{1.26}$$

to apply for all $t \geq t_0$ so as to minimize (1.6) looking forward from that date, then if the objective (1.6) is computed, for each candidate rule, conditional upon the current values of (x_{t_0-1}, u_{t_0}) ,²⁶ then the values of the coefficients ϕ_1, ϕ_2 that solve this problem will depend on the values of (x_{t_0-1}, u_{t_0}) . (As explained above, there will be a tradeoff between the choice of values that make policy closer to the Ramsey policy in periods $t > t_0$, and the choice of values that make policy closer to the Ramsey policy in period t_0 . But the degree to which given coefficients make policy in period t_0 different from the Ramsey policy will depend on the value of x_{t_0-1} , and so the optimal balance to strike in order to minimize (1.6) will depend on this.) This means that if one chooses a policy at one date (based on the lagged output gap at that particular time), and then reconsiders policy at some later date (when the lagged output gap will almost surely be different, given that the policy rule does not fully stabilize the output gap),

²⁶Here I follow Blake, Jensen and McCallum in assuming that $\{u_t\}$ is Markovian, so that the value of u_{t_0} contains all information available at date t_0 about the future evolution of the cost-push disturbance.

one will not choose the same coefficients on the second occasion as on the first — that is, one will not choose to continue following the policy rule chosen on the earlier occasion.

Blake, Jensen and McCallum instead argue for a specific linear target criterion (1.25), with coefficients that are independent of the initial conditions, because they do not evaluate (1.6) conditional upon the actual initial conditions at the time of the policy choice. Instead, they propose that in the case of each candidate policy rule, the *unconditional expectation* of (1.6) should be evaluated, integrating over all possible initial conditions (x_{t_0-1}, u_{t_0}) using the ergodic distribution associated with the stationary rational-expectations equilibrium implied by the time-invariant policy rule in question. This is a criterion that allows a particular policy rule to be chosen simply on the basis of one’s model of the economy (including the stochastic process for the exogenous disturbances), and independently of the actual state of the world in which the choice is made. But note that a time-independent outcome is achieved only by specifying that each time policy is reconsidered, (1.6) must be evaluated under *fictitious* initial conditions — a sort of “veil of ignorance” in the terminology of Rawls (1971) — rather than under the conditions that actually prevail at the time that the policy is reconsidered. If one is willing to posit that candidate policies should be evaluated from the standpoint of fictitious initial conditions, then the choice of (1.21) can also be justified in that way: one would choose to conform to the target criterion (1.21) for all $t \geq t_0$, if one evaluates this rule (relative to other possibilities) under the fictitious initial condition $x_{t_0-1} = 0$,²⁷ regardless of what the *actual* value of x_{t_0-1} may have been. (Note that if $x_{t_0-1} = 0$, Ramsey policy requires precisely that (1.21) hold for all $t \geq t_0$.)

Thus the preference of Blake, Jensen and McCallum for the alternative rule (1.25) depends on their preferring to evaluate the loss function under alternative (but equally fictitious) initial conditions. While they might argue that the choice of the ergodic distribution is a reasonable choice, it has unappealing aspects. In particular, the probability distribution over initial conditions that is assumed is different in the case of each of the candidate rules that are to be evaluated, since they imply different ergodic distributions for (x_{t-1}, u_t) , so that a given rule might be judged best simply because more favorable initial conditions are assumed when evaluating that rule.²⁸

²⁷More generally, if $x^* \neq 0$, the required fictitious initial condition is that $x_{t_0-1} = x^*$.

²⁸Benigno and Woodford (2008) propose a solution to the problem of choosing an optimal policy

Moreover, the criterion proposed by Blake, Jensen and McCallum leads to the choice of a different rule than does “optimality from a timeless perspective” (as defined above) only to the extent that the discount factor β is different from 1. (Note that as $\beta \rightarrow 1$, the criteria (1.21) and (1.25) become identical.) Since the empirically realistic value of β will surely be quite close to 1, it is not obvious that the alternative criterion would lead to policies that are very different, quantitatively.

1.6 Consequences of the Interest-Rate Lower Bound

In the above characterization of optimal policy, it has been taken for granted that the evolution of the nominal interest rate required in order for the joint evolution of inflation and output computed above to be consistent with equilibrium relation (1.8) involves a non-negative nominal interest rate at all times, and hence that there exists a path for the monetary base (and the interest rate paid on base money) that can implement the required path of short-term nominal interest rates. But there is no reason, in terms of the logic of New Keynesian model, why there cannot be disturbances under which the optimal commitment characterized above would require nominal interest rates to be negative. As a simple example, suppose that there are no cost-push disturbances $\{u_t\}$, but that the natural rate of interest r_t^n is negative in some periods.²⁹ In the absence of cost-push shocks, the characterization of optimal policy given above would require zero inflation (and a zero output gap) at all times. But this will require a real interest rate equal to the natural rate at all times, and hence sometimes negative; and it will also require that expected inflation be zero at all times, so that a negative real interest rate is possible only if the nominal interest rate is negative. But in any economy in which people have the option of holding currency that earns a zero nominal return, there will be no monetary policy under

rule within some class of “simple” rules, under which the same probability distribution over initial conditions is used to evaluate all rules within the candidate family of rules.

²⁹Our assumptions above imply that in the steady state the natural rate of interest is positive, so this problem can arise only in the case of disturbances that are sufficiently large. The economic plausibility of disturbances of a magnitude sufficient for this to be true is discussed by Christiano (2004). In practice, central banks have found themselves constrained by the zero lower bound only in the aftermath of serious financial crises, as during the Great Depression, in Japan beginning in the late 1990s, and during the current Great Recession. The way in which a sufficiently large financial disturbance can cause the zero lower bound to become a binding constraint is discussed in C’urdia and Woodford (2009b).

which the nominal interest rate can be negative, so the “optimal” policy characterized above would in this case be infeasible.

To treat such cases, it is necessary to add the zero bound (1.5) to the set of constraints on feasible state-contingent evolutions of the variables $\{\pi_t, x_t, i_t\}$. The constraint (1.8) also becomes a relevant (*i.e.*, sometimes binding) constraint in this case as well. The more general statement of the optimal policy problem is then to find the state-contingent evolution $\{\pi_t, x_t, i_t\}$ that minimizes (1.6) subject to the constraints that (1.5), (1.7), and (1.8) be satisfied each period. Alternatively, the problem can be stated as the choice of a state-contingent evolution $\{\pi_t, x_t\}$ each period that minimizes (1.6) subject to the constraints that (1.7) and

$$x_t \leq E_t x_{t+1} + \sigma(E_t \pi_{t+1} + r_t^n) \quad (1.27)$$

be satisfied each period. Note that (1.27) suffices for it to be possible to find a non-negative value of i_t each period that satisfies (1.8).

This problem is analyzed by Eggertsson and Woodford (2003).³⁰ One can again form a Lagrangian, and derive first-order conditions of the form

$$\pi_t + \varphi_{1,t} - \varphi_{1,t-1} - \beta^{-1} \sigma \varphi_{2,t-1} = 0 \quad (1.28)$$

$$\lambda(x_t - x^*) + \varphi_{2,t} - \beta^{-1} \varphi_{2,t-1} - \kappa \varphi_{1,t} = 0 \quad (1.29)$$

$$\varphi_{2,t} \geq 0 \quad (1.30)$$

for each $t \geq t_0$, together with the complementary slackness condition that at each point in time, at least one of the conditions (1.27) and (1.30) must hold with equality. Here $\varphi_{1,t}$ is the Lagrange multiplier associated with constraint (1.7), called simply φ_t earlier; $\varphi_{2,t}$ is the multiplier associated with constraint (1.27), which must accordingly equal zero except when this constraint binds; and in the case of Ramsey policy (*i.e.*, optimal policy under no initial pre-commitments), one substitutes the initial values $\varphi_{1,t_0-1} = \varphi_{2,t_0-1} = 0$. Optimal policy is then characterized by the state-contingent evolution for the variables $\{\pi_t, x_t, \varphi_{1,t}, \varphi_{2,t}\}$ that satisfy conditions (1.7), (1.27), (1.28)–(1.30) and the complementary slackness condition for all $t \geq t_0$.

³⁰The treatment in that paper further develops the discussion in Woodford (1999) of the way in which a binding zero lower bound changes the analysis of optimal policy in the basic New Keynesian model.

Because of the inequality constraints and the complementary slackness condition, these conditions are nonlinear, and so cannot be solved for the evolution of inflation and output as linear functions of the disturbances, as in section 1.2.³¹ Nonetheless, some general observations about the nature of optimal policy are possible. One is that, as discussed above, optimal policy will generally be *history-dependent*, and hence not implementable by any procedure which takes account only of the projected future paths of inflation, output and the nominal interest rate in a *purely forward-looking* way. In particular, the outcomes associated with Markov-perfect equilibrium under a *discretionary* approach to policy will be sub-optimal, as this is an example of a purely forward-looking procedure; instead, an optimal outcome will require *commitment*. These features of optimal policy follow directly from the presence of the lagged Lagrange multipliers associated with the forward-looking constraints in the FOCs (1.28)–(1.30). (The fact that the multipliers are not always zero implies that commitment is necessary; the fact that they are not constant, but necessarily depend on the economy’s past state rather than on its current state, implies that optimal policy must be history-dependent.) The fact that the zero lower bound can sometimes bind introduces an additional non-zero lagged Lagrange multiplier into the FOCs, relative to those discussed in section 1.2; hence there is an additional reason for commitment and history-dependence to be important for optimal policy.

Indeed, the zero lower bound can make commitment and history-dependence important, even when they otherwise would not be. This can usefully be illustrated by considering a simple case analyzed by Eggertsson and Woodford (2003). Suppose that there are no cost-push shocks ($u_t = 0$ at all times) and that the natural rate of interest $\{r_t^n\}$ evolves in accordance with a two-state Markov process. Specifically, suppose that r_t^n is always equal to one of two possible values, a “normal” level $\bar{r} > 0$ and a “crisis” level $\underline{r} < 0$.³² Suppose furthermore that when the economy is in the “normal” state, the probability of a transition to the “crisis” state is vanishingly small, but that when it is in the “crisis” state, there is only a probability $0 < \rho < 1$ each period of remaining in that state in the following period. Finally, suppose that the central bank’s objective is given by (1.6), but with $x^* = 0$. The case in which

³¹Numerical solutions of these equations for particular illustrative cases are offered by Jung *et al.* (2005), Eggertsson and Woodford (2003), Sugo and Teranishi (2005), and Adam and Billi (2006).

³²The way in which a temporary disruption of credit markets can result in a temporarily negative value of this state variable is discussed in Cúrdia and Woodford (2009b).

$x^* = 0$ is considered because in this case there would be no difference between the outcomes under discretionary policy and under an optimal commitment — and there would accordingly be no need for policy to be history-dependent — if the natural rate of interest were to evolve according to some process under which r_t^n were always non-negative, so that the zero lower bound would not bind under optimal policy. Even in this case, we shall see that discretionary policy is sub-optimal and optimal policy is history-dependent, if a state exists in which the zero lower bound is a binding constraint.

Let us first consider the Markov-perfect equilibrium under discretionary policy. When the economy is in the “normal” state, it is expected to remain there with (essentially) probability one from then on, and so discretionary policymaking leads to the choice of a policy under which inflation is zero and is expected to equal zero (essentially with probability one) from then on; this equilibrium involves a zero output gap and a nominal interest rate each period equal to $\bar{r} > 0$, so that the zero lower bound does not bind in this state. Let us now consider the discretionary policymaker’s choice in the “crisis” state (supposing that such a state occurs, despite having been considered very unlikely *ex ante*), taking as given that once reversion to the “normal” state occurs, policy will be conducted in the way just described: there will be immediate reversion to the optimal (zero-inflation) steady state. Looking forward from some date t at which the economy is in the “crisis” state, let the sequences $\{\pi_j^c, x_j^c, i_j^c\}$ indicate the anticipated values of the variables $(\pi_{t+j}, x_{t+j}, i_{t+j})$ at each future date conditional upon the economy still being in the “crisis” state at that date. Then any feasible policy during the “crisis”, consistent with rational expectations and with our above assumption about discretionary policy in the event of reversion to the “normal” state corresponds to sequences that satisfy the conditions

$$\pi_j^c = \kappa x_j^c + \beta \rho \pi_{j+1}^c, \quad (1.31)$$

$$x_j^c = \rho x_{j+1}^c - \sigma (i_j^c - \rho \pi_{j+1}^c - \underline{r}), \quad (1.32)$$

$$i_j^c \geq 0 \quad (1.33)$$

for all $j \geq 0$.

The system of difference equations (1.31)–(1.32) has a determinate solution (*i.e.*, a unique bounded solution) for the sequences $\{\pi_j^c, x_j^c\}$ corresponding to any anticipated forward path for policy specified by a bounded sequence $\{i_j^c\}$ if and only if the model

parameters satisfy the inequality

$$\frac{\rho}{(1-\rho)(1-\beta\rho)} < \kappa^{-1}\sigma^{-1}, \quad (1.34)$$

which implies that the “crisis” state is not expected to be too persistent. I shall assume that (1.34) holds in the subsequent discussion, so that there is a determinate outcome associated with any given path of the policy rate during the “crisis,” and in particular with the assumption that the policy rate remains pinned at the zero lower bound for as long as the crisis persists.

The unique bounded outcome associated with the expectation that $i_j^c = 0$ for all j will then be one under which $\pi_j^c = \pi^c, x_j^c = x^c$ for all j , where the constant values are given by

$$\pi^c = \left[\frac{(1-\beta\rho)(1-\rho)}{\kappa\sigma} - \rho \right]^{-1} \underline{r} < 0,$$

$$x^c = \frac{1-\beta\rho}{\kappa} \pi^c < 0.$$

(The signs given here follow from assumption (1.34).) Moreover, one can show under assumption (1.34) that any forward path $\{i_j^c\}$ in which $i_j^c > 0$ for some j must involve even greater deflation and even greater negative output gap than in this solution; hence under the assumed policy objective, this outcome is the best feasible outcome, given the assumption of immediate reversion to the zero-inflation steady state upon the economy’s reversion to its “normal” state. In fact, since the solution paths $\{\pi_j^c, x_j^c\}$ are monotonic functions of each of the elements in the assumed path $\{i_j^c\}$ for policy, it follows that under any assumption about policy for $j > k$, the optimal policy for $j \leq k$ will be to choose $i_j^c = 0$ for all periods $j \leq k$. Hence under our Markovian assumption about discretionary policy after reversion to the “normal” state, there is a unique solution for discretionary policy in the “crisis” state (even without the restriction to Markovian policies), which is for the policy rate to equal zero as long as the economy remains in the “crisis” state.

In this model, discretionary policy results in both deflation and a negative output gap, that persist as long as the “crisis” state persists. Moreover, it is possible to find parameter values under which the predicted output collapse and deflation are quite severe, even if the negative value of \underline{r} is quite modest.³³ Indeed, one observes

³³Denes and Eggertsson (2009) discuss parameter values under which a two-state model similar to this one can predict an output collapse and decline in prices of the magnitudes observed in the US

that as ρ approaches the upper bound defined by (1.34), the predicted values of π^c and x^c approach $-\infty$, for fixed values of \underline{r} and the other model parameters.³⁴ This may make it plausible that the Markov-perfect discretionary outcome is not in fact the best possible outcome achievable (under rational expectations) by an appropriate monetary policy. However, the discussion above makes it clear that the achievement of any better outcome must involve an anticipation of a different approach to policy *after the economy permanently reverts to the “normal” state*. Such a commitment to a history-dependent policy can be especially welfare-enhancing in a situation like the one just considered, because the central bank is so severely constrained in what it can achieve by varying *current* policy while at the zero lower bound.

The kind of commitment that improves welfare (if credible) is one that implies that inflation will be allowed to temporarily exceed its long-run target, and that a temporary output boom will accordingly be created through monetary policy, in the period immediately following the economy’s reversion to the “normal” state of fundamentals. Both a higher expected inflation rate post-reversion and a higher expected level of output post-reversion (implying a lower marginal utility of income post-reversion) will reduce incentives for saving while the economy remains in the “crisis” state, leading to greater capacity utilization and less incentives to cut prices; less pessimism about the degree of deflation and output collapse in the case that the “crisis” state persists then becomes a further reason for less deflation and output collapse, in a “virtuous circle.” It is possible for a substantial improvement in economic conditions during the “crisis” to occur as a result of a credible commitment to even a modest boom and period of reflation following the economy’s reversion to “normal” fundamentals, as illustrated by Figure 3, reproduced from Eggertsson and Woodford (2003).³⁵

Here the equilibrium outcomes under the optimal policy commitment (the solid lines, indicating the solution to equations (1.28)–(1.30) together with the complementary slackness condition) are compared to those under the Markov-perfect discretionary policy during the Great Depression. Under their modal parameter values, the value of \underline{r} is only -4 percent per annum.

³⁴Of course, the log-linear approximations assumed in the New Keynesian model of this section will surely become highly inaccurate before that limit is reached. Hence we cannot really say exactly what should happen in equilibrium in this limit on the basis of the calculations reported in this section.

³⁵The numerical parameter values assumed in this calculation are discussed there.




Figure 3: Dynamics under an optimal policy commitment compared with the equilibrium outcome under discretionary policy (or commitment to a strict inflation target $\pi^* = 0$), for the case in which the disturbance lasts for exactly 15 quarters.

tionary equilibrium (the dashed lines, corresponding to a policy that can equivalently be described as a forward-looking inflation targeting regime with a zero inflation target); the outcomes are plotted for the case in which the random realization of the length of the crisis is 15 quarters (though this is not assumed to have been known *ex ante* by either the private sector or the central bank). Under the optimal commitment, the policy rate is kept at zero for another year, even though it would have been possible to return immediately to the zero-inflation steady state in quarter 15, as occurs under the discretionary policy. This results in a brief period in which inflation exceeds its long-run value and output exceeds the natural rate, but inflation is stabilized at zero fairly soon; nonetheless, the commitment not to return immediately to the zero-inflation steady state as soon as this is possible dramatically reduces the degree to which either prices or output fall during the years of the crisis. The dra-

matic results in Figure 3 depend on parameter values, of course: in particular, they depend on assuming that the persistence of the “crisis” state is not too far below the upper bound specified in (1.34). But it is worth noting that according to this analysis, the efficacy of such a commitment to reflationary policy is greatest precisely in those cases in which the risk of a severe crisis is greatest, that is, in which even a modest decline in r_t^n can trigger a severe output collapse and deflation.

Once again, the optimal policy commitment can usefully be formulated in terms of a target criterion. Eggertsson and Woodford (2003) show that under quite general assumptions about the exogenous disturbance processes $\{r_t^n, u_t\}$ — and not just in the highly specific example discussed above — the optimal policy commitment can be expressed in the following way. The central bank commits itself to use its interest-rate instrument each period so as to cause the output-gap-adjusted price level \tilde{p}_t (again defined as in (1.23)) to achieve a target p_t^* that can be announced in advance on the basis of outcomes in period $t - 1$, if this is achievable with a non-negative level of the policy rate; if the zero lower bound makes achievement of the target infeasible, the interest rate should at any rate be set to zero. The price-level target for the next period is then adjusted based on the degree to which the target has been achieved in the current period, in accordance with the formula

$$p_{t+1}^* = p_t^* + \beta^{-1}(1 + \kappa\sigma)\Delta_t - \beta^{-1}\Delta_{t-1}, \quad (1.35)$$

where $\Delta_t \equiv p_t^* - \tilde{p}_t$ is the (absolute value of the) period t *target shortfall*. (If the central bank behaves in this way each period, then it can be shown that there exist Lagrange multipliers $\{\varphi_{1,t}, \varphi_{2,t}\}$ such that the FOCs (1.28)–(1.30) and the complementary slackness condition are satisfied in all periods.) Under this rule, a period in which the zero lower bound prevents achievement of the target for a time results in the price-level target (and hence the economy’s long-run price level) being ratcheted up to a permanently higher level, to an extent that is greater the greater the target shortfall and the longer the period for which the target shortfalls persist; but over a period in which the zero lower bound never binds, the target is not adjusted. In particular, if the zero lower bound never binds, as assumed in section 1.2, then the optimal target criterion reduces once again to the simple requirement (1.23), as derived above.

I have noted in section 1.4 that in the case that the zero lower bound never binds, the optimal policy commitment in the basic New Keynesian model can equivalently

be expressed *either* by a “flexible inflation target” (1.21) or by a “price-level target” (1.23). In the case that the target in question can always be achieved, these two commitments have identical implications. But if the zero lower bound sometimes requires undershooting of the target, they are *not* equivalent, and in this case their welfare consequences are quite different. Neither, of course, is precisely optimal in the more general case. But the pursuit of an inflation target of the form (1.21) whenever it is possible to hit the target, with no corrections for past target misses, is a much worse policy than the pursuit of a fixed price-level target (1.23) whenever it is possible to hit the target, again without any correction for past target misses. In the two-state Markov chain example discussed above, the simple inflation targeting policy would be *even worse* than the discretionary policy discussed above (which is equivalent to the pursuit of a strict inflation target, with no adjustment for the change in the output gap); for at the time of the reversion to the “normal” state, this policy would require *even tighter* policy than the “ $\pi^* = 0$ ” policy shown in Figure 3, as the fact that the output gap has been negative during the crisis would require the central bank to aim for negative inflation and/or a negative output gap immediately following reversion as well. A commitment to this kind of “gradualism” would create exactly the kind of expectations that make the crisis even worse.³⁶

Instead, while the simple price-level target criterion (1.23) is not fully optimal in the more general case considered in this section, it represents a great improvement upon discretionary policy, since it incorporates the *type* of history-dependence that is desirable: a commitment to compensate for any target undershooting required by the zero lower bound by subsequently pursuing more inflation than would otherwise be desirable, in order to return the gap-adjusted price level to its previous target. Ideally, the price-level target would even be ratcheted up slightly owing to the target shortfall; but what it is most important is that it not be *reduced* simply as a consequence of having undershot the target in the past. Such accommodation of past target

³⁶A commitment to satisfy the target criterion (1.21) is optimal in the case discussed in section 1.4 because a negative past output gap will have occurred in equilibrium only because of an adverse cost-push shock in the recent past, and it will have been beneficial in that situation to have had people expect that subsequent expenditure growth will be restrained, in order to reduce the extent to which output had to be contracted in order to contain inflation. But if the negative output gap in the past occurred owing to a target shortfall required by the zero lower bound, the kind of expectations regarding subsequent policy that one would have wished to create in order to mitigate the distortion are exactly the opposite.

shortfalls creates expectations of the type that make the distortions resulting from the zero lower bound much more severe. At least in the numerical example considered by Eggertsson and Woodford (2003), a simple price-level target of the form (1.23) achieves nearly all of the welfare gains that are possible in principle under policy commitment.³⁷ Thus while it is not true in the more general case considered in this section that the long-run price level is constant under optimal policy, it remains true that a shift from a forward-looking inflation target to a price-level target introduces history-dependence of a highly desirable kind. Indeed, the advantages of a price-level target are particularly great in the case that the possibility of an occasionally binding zero lower bound is a serious concern.

1.7 Optimal Policy Under Imperfect Information

In the characterization of optimal policy given above, I have assumed that the central bank is fully aware of the state of the economy at each time that it makes a policy decision. This has simplified the analysis, as it was possible simply to choose the best among all possible rational-expectations equilibria, taking it for granted that the central bank possesses the information necessary to adjust its instrument in the way required in order to implement any given equilibrium. But in practice, while central banks devote considerable resources to obtaining precise information about current economic conditions, they cannot be assumed to be perfectly informed about the values of all of the state variables that are relevant according to theory of optimal policy — for example, not only about how much the aggregate-supply curve has been shifted by a given disturbance, but about how persistent people are expecting the shift to be — and the necessity to make policy decisions before the correct values of state variables can be known with certainty is an important consideration in the choice of a desirable approach to the conduct of policy. However, the methods illustrated

³⁷Levin *et al.* (2010) consider an alternative parameterization — in particular, an alternative parameterization of the disturbance process — under which a simple price-level target is not as close an approximation to fully optimal policy. Note that their parameterization is again one in which there are substantial welfare gains from commitment to a history-dependent policy, and the kind of commitment that is needed is commitment to a temporary period of reflation following a crisis in which the zero lower bound constrains policy. But in their example, optimal policy must permanently raise the price level in compensation for the earlier target shortfalls to a greater extent than is true in the case shown in Figure 3.

above have direct extensions to the case of imperfect information, at the cost of some additional complexity.

Here I shall illustrate how the form of an optimal policy rule is affected by assumptions about the central bank's information set by reconsidering the policy problem treated in section 1.2 under an alternative information assumption. Let us suppose that all private agents share a common information set, that we shall consider to represent "full information" (since the central bank's information set at each point in time is a subset of this private information set), and for any random variable z_τ , let $E_t z_\tau$ denote the conditional expectation of this variable with respect to the private information set in period t . Let us further suppose that the central bank must choose i_t , the period t value of its policy instrument, on the basis of an information set that includes full awareness of the economy's state in period $t - 1$, but only partial information (if any) about those random shocks that are realized in period t ; and let $z_{\tau|t}$ be the conditional expectation of the variable z_τ with respect to the central bank's information when choosing the value of i_t . Thus if we let I_t denote the private sector's information set in period t and I_t^{cb} the central bank's information set when setting i_t , then there is a strict nesting of the sequence of information sets

$$\dots \subset I_{t-1}^{cb} \subset I_{t-1} \subset I_t^{cb} \subset I_t \subset I_{t+1}^{cb} \subset I_{t+1} \subset \dots$$

which implies that $[E_t z_\tau]_t = z_{\tau|t}$, $E_t[z_{\tau|t+1}] = E_t z_\tau$, and so on.

Let us consider the problem of adjusting the path of $\{i_t\}$ on the basis of the central bank's partial information³⁸ so as to minimize the objective (1.6),³⁹ given that the paths of inflation and output will be determined by the structural relations (1.7)–(1.8). As discussed by Svensson and Woodford (2003, 2004) in the context of a more general class of problems of this kind, we can compute FOCs for optimal policy using a Lagrangian of the form

$$\begin{aligned} \mathcal{L}_{t_0} = E_{t_0} \sum_{t=0}^{\infty} \beta^{t-t_0} & \left\{ \frac{1}{2} [\pi_t^2 + \lambda(x_t - x^*)^2] + \varphi_{1,t} [\pi_t - \kappa x_t - \beta \pi_{t+1}] \right. \\ & \left. + \varphi_{2,t} [x_t - x_{t+1} + \sigma(i_t - \pi_{t+1})] \right\}, \end{aligned}$$

³⁸Contrary to our conclusion in the full-information case, it now matters what we assume the central bank's policy instrument to be; in general, we will not obtain the same optimal equilibrium evolution for the economy if the central bank must set the nominal interest rate on the basis of its information set as if it must set the money supply on the basis of its information set.

³⁹The expectation in the objective must now be understood to be conditional upon the central bank's information set at the time of choice of the policy commitment.

where $\varphi_{1,t}, \varphi_{2,t}$ are now Lagrange multipliers associated with the two constraints, as in the previous section.

The FOCs (obtained by differentiating the Lagrangian with respect to π_t, x_t and i_t respectively) are again of the form (1.28)–(1.29), but with (1.30) replaced in this case by

$$\varphi_{2,t|t} = 0. \tag{1.36}$$

Here (1.36) need hold only conditional upon the central bank’s period t information set, because the central bank can only adjust the value of i_t separately across states that it is able to distinguish using that information. Note that if the central bank has full information, condition (1.36) becomes simply $\varphi_{2,t} = 0$, and the FOCs reduce to the system (1.9)–(1.10) obtained in section 1.2. But in the case of imperfect information, the multiplier $\varphi_{2,t}$ is in general not equal to zero, as the constraint associated with the “intertemporal IS relation” owing to the constraints on the central bank’s ability to adjust its instrument as flexibly as it would under the full-information optimal policy.

The state-contingent evolution of the endogenous variables (including the central bank’s instrument) can be determined by solving for processes $\{\pi_t, x_t, i_t, \varphi_{1,t}, \varphi_{2,t}\}$ that satisfy conditions (1.7)–(1.8), (1.28)–(1.29) and (1.36) each period, subject to the requirements that $(\pi_t, x_t, \varphi_{1,t}, \varphi_{2,t})$ depend only on I_t , and that i_t depends only on I_t^{cb} , for some specification of the exogenous disturbance processes and of the way in which the indicator variables observed by the central bank (that may include noisy observations of current-period endogenous or exogenous variables) depend on the economy’s state. Svensson and Woodford (2004) present a general method that can be used to calculate such an equilibrium, if the central bank’s indicators are linear functions of the state variables plus Gaussian measurement error, so that the Kalman filter can be used to calculate the central bank’s conditional expectations as linear functions of the indicators; Aoki (2006) illustrates its application to the model described in this section, under a particular assumption about the central bank’s information set.⁴⁰

Rather than discuss these calculations further, I shall simply observe that once again, it is possible to describe optimal policy in terms of a target criterion, and once

⁴⁰In Aoki’s model, the central bank’s information set when choosing i_t consists of complete knowledge of the period $t - 1$ state of the economy, plus noisy observations of π_t and x_t . The central bank is assumed not to directly observe any of the period t exogenous disturbances, even imprecisely.

again the form of the optimal target criterion does not depend on the specification of the disturbance processes. In the present case, the optimal target criterion is also independent of the specification of the central bank’s information set. The targeting rule can be expressed in the following way: in any period t , the central bank should choose the value of i_t so as to ensure that

$$\tilde{p}_{t|t} = p_t^* \tag{1.37}$$

conditional upon its own expectations, where \tilde{p}_t is again the same output-gap-adjusted price level as in (1.23). The target p_t^* is a function solely of the economy’s state at $t - 1$, and evolves according the same law of motion (1.35) as in the previous section, where Δ_t is again the period t target shortfall ($p_t^* - \tilde{p}_t$), observed by the central bank by the time that it chooses the value of i_{t+1} . If the variables $\{\tilde{p}_t, \tilde{p}_{t|t}, p_t^*\}$ evolve over time in accordance with (1.35) and (1.37), then it is possible to define multipliers $\{\varphi_{1,t}, \varphi_{2,t}\}$ such that the FOCs (1.28)–(1.29) and (1.36) are satisfied each period; hence non-explosive dynamics consistent with the target criterion will necessarily correspond to the optimal equilibrium.

It is worth noting that the fact that the central bank will generally fail to precisely achieve its target for the gap-adjusted price level owing to the incompleteness of its information about the current state is *not* a reason for a central bank to choose a forward-looking inflation target (and “let bygones be bygones”) rather than a price-level target. In fact, the optimal response to this problem is for the central bank to commit not only to subsequently correct past target misses (by continuing to aim at the same price-level target as before), but actually to *over-correct* for them — permanently reducing its price-level target as a result of having allowed the gap-adjusted price level to overshoot the target, and permanently increasing it as a result of allowing it to be undershot.

It is interesting to note that the optimal target criterion has exactly the same form in the case that the central bank can fail to hit its target owing to imperfect information as in the case where it can fail to hit its target due to the zero lower bound. Thus we have a single target criterion that is optimal in both cases, and that also reduces (in the case of full information and shocks small enough for the zero lower bound never to be a problem) to the simpler target criterion discussed in section 1.4. Hence description of policy in terms of a target criterion allows a unified characterization of optimal policy in all of these cases. In section 3, it is shown that

the same target criterion is optimal in an even broader class of cases.

2 Stabilization and Welfare

In the previous section, I have simply assumed a quadratic objective for stabilization policy, that incorporates concerns that are clearly at the forefront of many policy deliberations inside central banks. In this section, I instead consider how the normative theory of monetary stabilization can be developed if one takes the objective to be the maximization of the average expected utility of households — that is, the private objectives that are assumed in deriving the behavioral relations that determine the effects of alternative monetary policies — as is done in the modern theory of public finance. This discussion requires a more explicit treatment of the microfoundations of the basic New Keynesian model, which then provide the basis for a welfare-theoretic treatment of the optimal policy problem as well.

2.1 Microfoundations of the Basic New Keynesian Model

I shall begin by deriving exact structural relations for the basic New Keynesian model. (As discussed further below, the structural relations assumed in section 1 represent a log-linearization of these relations, around a zero-inflation steady state; but the log-linearized equations alone do not suffice for welfare analysis of alternative stabilization policies, which requires at least a second-order approximation.) The exposition here follows Benigno and Woodford (2005a), who write the exact structural relations in a recursive form (with only one-period-ahead expectations of a finite number of sufficient statistics mattering for equilibrium determination each period) that facilitates the definition of optimal policy from a timeless perspective, and perturbation analysis of the system of equations that characterize optimal policy.⁴¹

The economy is made up of identical infinite-lived households, each of which seeks to maximize

$$U_{t_0} \equiv E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \left[\tilde{u}(C_t; \xi_t) - \int_0^1 \tilde{v}(H_t(j); \xi_t) dj \right], \quad (2.1)$$

⁴¹The model presented here is a variant of the monetary DSGE model originally proposed by Yun (1996). Goodfriend and King (1997) is an important early discussion of optimal policy in the context of this model.

where C_t is a Dixit-Stiglitz aggregate of consumption of each of a continuum of differentiated goods,

$$C_t \equiv \left[\int_0^1 c_t(i)^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}, \quad (2.2)$$

with an elasticity of substitution equal to $\theta > 1$, $H_t(j)$ is the quantity supplied of labor of type j , and ξ_t is a vector of exogenous disturbances, which may include random shifts of either of the functions \tilde{u} or \tilde{v} .

Each differentiated good is supplied by a single monopolistically competitive producer. There are assumed to be many goods in each of an infinite number of “industries”; the goods in each industry j are produced using a type of labor that is specific to that industry, and suppliers in the same industry also change their prices at the same time.⁴² The representative household supplies all types of labor as well as consuming all types of goods. To simplify the algebraic form of the results, it is convenient to assume isoelastic functional forms,⁴³

$$\tilde{u}(C_t; \xi_t) \equiv \frac{C_t^{1-\tilde{\sigma}^{-1}} \bar{C}_t^{\tilde{\sigma}^{-1}}}{1 - \tilde{\sigma}^{-1}}, \quad (2.3)$$

$$\tilde{v}(H_t; \xi_t) \equiv \frac{\lambda}{1 + \nu} H_t^{1+\nu} \bar{H}_t^{-\nu}, \quad (2.4)$$

where $\tilde{\sigma}, \nu > 0$, and $\{\bar{C}_t, \bar{H}_t\}$ are bounded exogenous disturbance processes. (Here \bar{C}_t and \bar{H}_t are both among the exogenous disturbances included in the vector ξ_t .)

I assume a common technology for the production of all goods, in which (industry-specific) labor is the only variable input,

$$y_t(i) = A_t f(h_t(i)) = A_t h_t(i)^{1/\phi}, \quad (2.5)$$

where A_t is an exogenously varying technology factor, and $\phi > 1$. The Dixit-Stiglitz preferences (2.2)⁴⁴ imply that the quantity demanded of each individual good i will

⁴²The assumption of segmented factor markets for different “industries” is inessential to the results obtained here, but allows a numerical calibration of the model that implies a speed of adjustment of the general price level more in line with aggregate time series evidence. For further discussion, see Woodford (2003, chap. 3).

⁴³Benigno and Woodford (2004) extend the results of this section to the case of more general preferences and production technologies.

⁴⁴In addition to assuming that household utility depends only on the quantity obtained of C_t , I assume that the government also cares only about the quantity obtained of the composite good defined by (2.2), and that it seeks to obtain this good through a minimum-cost combination of purchases of individual goods.

equal

$$y_t(i) = Y_t \left(\frac{p_t(i)}{P_t} \right)^{-\theta}, \quad (2.6)$$

where Y_t is the total demand for the composite good defined in (2.2), $p_t(i)$ is the (money) price of the individual good, and P_t is the price index,

$$P_t \equiv \left[\int_0^1 p_t(i)^{1-\theta} di \right]^{\frac{1}{1-\theta}}, \quad (2.7)$$

corresponding to the minimum cost for which a unit of the composite good can be purchased in period t . Total demand is given by

$$Y_t = C_t + G_t, \quad (2.8)$$

where G_t is the quantity of the composite good purchased by the government, treated here as an exogenous disturbance process.

The producers in each industry fix the prices of their goods in monetary units for a random interval of time, as in the model of staggered pricing introduced by Calvo (1983). Let $0 \leq \alpha < 1$ be the fraction of prices that remain unchanged in any period. A supplier that changes its price in period t chooses its new price $p_t(i)$ to maximize

$$E_t \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} \Pi(p_t(i), p_T^j, P_T; Y_T, \xi_T), \quad (2.9)$$

where $Q_{t,T}$ is the stochastic discount factor by which financial markets discount random nominal income in period T to determine the nominal value of a claim to such income in period t , and α^{T-t} is the probability that a price chosen in period t will not have been revised by period T . In equilibrium, this discount factor is given by

$$Q_{t,T} = \beta^{T-t} \frac{\tilde{u}_c(C_T; \xi_T)}{\tilde{u}_c(C_t; \xi_t)} \frac{P_t}{P_T}. \quad (2.10)$$

Profits are equal to after-tax sales revenues net of the wage bill. Sales revenues are determined by the demand function (2.6), so that (nominal) after-tax revenue equals

$$(1 - \tau_t) p_t(i) Y_t \left(\frac{p_t(i)}{P_t} \right)^{-\theta}.$$

Here τ_t is a proportional tax on sales revenues in period t ; $\{\tau_t\}$ is treated as an exogenous disturbance process, taken as given by the monetary policymaker.⁴⁵ I assume that τ_t fluctuates over a small interval around a non-zero steady-state level $\bar{\tau}$. The real wage demanded for labor of type j is assumed to be given by

$$w_t(j) = \mu_t^w \frac{\tilde{v}_h(H_t(j); \xi_t)}{\tilde{u}_c(C_t; \xi_t)}, \quad (2.11)$$

where $\mu_t^w \geq 1$ is an exogenous markup factor in the labor market (allowed to vary over time, but assumed common to all labor markets),⁴⁶ and firms are assumed to be wage-takers. I allow for exogenous variations in both the tax rate and the wage markup in order to include the possibility of “pure cost-push shocks” that affect equilibrium pricing behavior while implying no change in the efficient allocation of resources.⁴⁷

Substituting the assumed functional forms for preferences and technology, the function

$$\begin{aligned} \Pi(p, p^j, P; Y, \xi) \equiv & (1 - \tau)pY(p/P)^{-\theta} \\ & - \lambda \mu^w P \left(\frac{p^j}{P}\right)^{-\theta(1+\omega)} \left(\frac{p}{p^j}\right)^{-\theta\phi} \bar{H}^{-\nu} \left(\frac{Y}{A}\right)^{1+\omega} \left(\frac{Y-G}{C}\right)^{\bar{\sigma}-1} \end{aligned} \quad (2.12)$$

then describes the after-tax nominal profits of a supplier with price p , in an industry with common price p^j , when the aggregate price index is equal to P and aggregate demand is equal to Y . Here $\omega \equiv \phi(1+\nu) - 1 > 0$ is the elasticity of real marginal cost in an industry with respect to industry output. The vector of exogenous disturbances ξ_t now includes A_t, G_t, τ_t and μ_t^w , in addition to the preference shocks.

Each of the suppliers that revise their prices in period t chooses the same new price p_t^* , that maximizes (2.9). Note that supplier i 's profits are a concave function of the quantity sold $y_t(i)$, since revenues are proportional to $y_t(i)^{\frac{\theta-1}{\theta}}$ and hence concave in $y_t(i)$, while costs are convex in $y_t(i)$. Moreover, since $y_t(i)$ is proportional to

⁴⁵The case in which the tax rate is also chosen optimally in response to other shocks is treated in Benigno and Woodford (2003). See also Canzoneri *et al.* (2010).

⁴⁶In the case that we assume that $\mu_t^w = 1$ at all times, our model is one in which both households and firms are wage-takers, or there is efficient contracting between them. Note that apart from the markup factor, the right-hand side of (2.11) represents the representative household's marginal rate of substitution between labor of type j and consumption of the composite good.

⁴⁷It is shown below, however, that these two disturbances are not, in general, the only reasons for the existence of a “cost-push” term in the aggregate-supply relation (1.7).

$p_t(i)^{-\theta}$, the profit function is also concave in $p_t(i)^{-\theta}$. The first-order condition for the optimal choice of the price $p_t(i)$ is the same as the one with respect to $p_t(i)^{-\theta}$; hence the first-order condition with respect to $p_t(i)$,

$$E_t \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} \Pi_1(p_t(i), p_T^j, P_T; Y_T, \xi_T) = 0,$$

is both necessary and sufficient for an optimum. The equilibrium choice p_t^* (which is the same for each firm in industry j) is the solution to the equation obtained by substituting $p_t(i) = p_t^*$ and $p_T^j = p_t^*$ for all $T \geq t$ into the above first-order condition.

Under the assumed isoelastic functional forms, the optimal choice has a closed-form solution

$$\frac{p_t^*}{P_t} = \left(\frac{K_t}{F_t} \right)^{\frac{1}{1+\omega\theta}}, \quad (2.13)$$

where F_t and K_t are functions of current aggregate output Y_t , the current exogenous state ξ_t , and the expected future evolution of inflation, output, and disturbances, defined by

$$F_t \equiv E_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} f(Y_T; \xi_T) \left(\frac{P_T}{P_t} \right)^{\theta-1}, \quad (2.14)$$

$$K_t \equiv E_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} k(Y_T; \xi_T) \left(\frac{P_T}{P_t} \right)^{\theta(1+\omega)}, \quad (2.15)$$

in which expressions⁴⁸

$$f(Y; \xi) \equiv (1 - \tau) \bar{C}^{\bar{\sigma}-1} (Y - G)^{-\bar{\sigma}-1} Y, \quad (2.16)$$

$$k(Y; \xi) \equiv \frac{\theta}{\theta - 1} \lambda \phi \frac{\mu^w}{A^{1+\omega} \bar{H}^\nu} Y^{1+\omega}. \quad (2.17)$$

Relations (2.14)–(2.15) can instead be written in the recursive form

$$F_t = f(Y_t; \xi_t) + \alpha\beta E_t[\Pi_{t+1}^{\theta-1} F_{t+1}], \quad (2.18)$$

$$K_t = k(Y_t; \xi_t) + \alpha\beta E_t[\Pi_{t+1}^{\theta(1+\omega)} K_{t+1}], \quad (2.19)$$

⁴⁸Note that the definition of the function $f(Y; \xi)$ here differs from that in Benigno and Woodford (2005a). There, the function here called $f(Y; \xi)$ is written as $(1 - \tau)f(Y; \xi)$, following the notation in Benigno and Woodford (2003), where the tax rate τ_t is a policy choice rather than an exogenous disturbance. Here τ_t is included in the vector ξ_t .

where $\Pi_t \equiv P_t/P_{t-1}$. It is evident that (2.14) implies (2.18); but one can also show that processes that satisfy (2.18) each period, together with certain bounds, must satisfy (2.14). Since we are interested below only in the characterization of bounded equilibria, we can omit the statement of the bounds that are implied by the existence of well-behaved expressions on the right-hand sides of (2.14) and (2.15), and treat (2.18)–(2.19) as necessary and sufficient for processes $\{F_t, K_t\}$ to measure the relevant marginal conditions for optimal price-setting.

The price index then evolves according to a law of motion

$$P_t = [(1 - \alpha)p_t^{*1-\theta} + \alpha P_{t-1}^{1-\theta}]^{\frac{1}{1-\theta}}, \quad (2.20)$$

as a consequence of (2.7). Substitution of (2.13) into (2.20) implies that equilibrium inflation in any period is given by

$$\frac{1 - \alpha \Pi_t^{\theta-1}}{1 - \alpha} = \left(\frac{F_t}{K_t} \right)^{\frac{\theta-1}{1+\omega\theta}}, \quad (2.21)$$

where $\Pi_t \equiv P_t/P_{t-1}$. This defines a short-run aggregate supply relation between inflation and output, given the current disturbances ξ_t , and expectations regarding future inflation, output, and disturbances. Condition (2.21), together with (2.18)–(2.19), represents a nonlinear version of the relation (1.1) in the log-linear New Keynesian model of section 1; and indeed, it reduces to (1.1) when log-linearized, as discussed further below.

It remains to explain the connection between monetary policy and private-sector decisions. I abstract here from any monetary frictions that would account for a demand for central-bank liabilities that earn a substandard rate of return; but I nonetheless assume that the central bank can control the riskless short-term nominal interest rate i_t ,⁴⁹ which is in turn related to other financial asset prices through the arbitrage relation

$$1 + i_t = [E_t Q_{t,t+1}]^{-1}. \quad (2.22)$$

If we use (2.10) to substitute for the stochastic discount factor in this expression, we obtain an equilibrium relation between i_t and the path of expenditure by the representative household. Using (2.8) to substitute for C_t , we obtain a relation that

⁴⁹For discussion of how this is possible even in a “cashless” economy of the kind assumed here, see Woodford (2003, chapter 2).

links i_t , Y_t , expected future output, expected inflation, and exogenous disturbances; this is a nonlinear version of the “intertemporal IS relation” (1.2) assumed in the log-linear New Keynesian model of section 1, and it reduces precisely to (1.2) when log-linearized.

I shall assume that the zero lower bound on nominal interest rates never binds in the policy problem considered in this section,⁵⁰ so that one need not introduce any additional constraint on the possible paths of output and prices associated with a need for the chosen evolution of prices to be consistent with a non-negative nominal interest rate. In this case, one can determine the optimal state-contingent evolution of inflation and real activity without any reference to the constraint implied by the “IS relation,” and without having to solve explicitly for the implied path of interest rates, as in the treatment of optimal policy in section 1. Once one has solved for the optimal state-contingent paths for inflation and output, these solutions can be substituted into (2.22) to determine the implied state-contingent evolution of the policy. (The implied equilibrium paths of other asset prices can similarly be solved for.)

Finally, I shall assume the existence of a lump-sum source of government revenue (in addition to the fixed tax rate τ), and assume that the fiscal authority ensures intertemporal government solvency regardless of what monetary policy may be chosen by the monetary authority.⁵¹ This allows us to abstract from the fiscal consequences of alternative monetary policies in our consideration of optimal monetary stabilization policy, as is implicitly done in Clarida *et al.* (1999), and much of the literature on monetary policy rules. An extension of the analysis to the case in which only distorting taxes exist is given in Benigno and Woodford (2003).

2.2 Welfare and the Optimal Policy Problem

The goal of policy is assumed to be the maximization of the level of expected utility of a representative household, given by (2.1). Inverting the production function (2.5)

⁵⁰This can be shown to be true in the case of small enough disturbances, given that the nominal interest rate is equal to $\bar{r} = \beta^{-1} - 1 > 0$ under the optimal policy in the absence of disturbances.

⁵¹Thus I assume that fiscal policy is “Ricardian,” in the terminology of Woodford (2001). A non-Ricardian fiscal policy would imply the existence of an additional constraint on the set of equilibria that could be achieved through monetary policy. The consequences of such a constraint for the character of optimal monetary policy are discussed in Benigno and Woodford (2007).

to write the demand for each type of labor as a function of the quantities produced of the various differentiated goods, and using the identity (2.8) to substitute for C_t , where G_t is treated as exogenous, it is possible to write the utility of the representative household as a function of the expected production plan $\{y_t(i)\}$. One obtains

$$U_{t_0} \equiv E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \left[u(Y_t; \xi_t) - \int_0^1 v(y_t^j; \xi_t) dj \right], \quad (2.23)$$

where

$$\begin{aligned} u(Y_t; \xi_t) &\equiv \tilde{u}(Y_t - G_t; \xi_t), \\ v(y_t^j; \xi_t) &\equiv \tilde{v}(f^{-1}(y_t^j/A_t); \xi_t). \end{aligned}$$

In this last expression I make use of the fact that the quantity produced of each good in industry j will be the same, and hence can be denoted y_t^j ; and that the quantity of labor hired by each of these firms will also be the same, so that the total demand for labor of type j is proportional to the demand of any one of these firms.

One can furthermore express the relative quantities demanded of the differentiated goods each period as a function of their relative prices, using (2.6). This allows us to write the utility flow to the representative household in the form

$$U(Y_t, \Delta_t; \xi_t) \equiv u(Y_t; \xi_t) - v(Y_t; \xi_t) \Delta_t,$$

where

$$\Delta_t \equiv \int_0^1 \left(\frac{p_t(i)}{P_t} \right)^{-\theta(1+\omega)} di \geq 1 \quad (2.24)$$

is a measure of price dispersion at date t , and the vector ξ_t now includes the exogenous disturbances G_t and A_t as well as the preference shocks. Hence we can write our objective (2.23) as

$$U_{t_0} = E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} U(Y_t, \Delta_t; \xi_t). \quad (2.25)$$

Here $U(Y, \Delta; \xi)$ is a strictly concave function of Y for given Δ and ξ , and a monotonically decreasing function of Δ given Y and ξ .

Because the relative prices of the industries that do not change their prices in period t remain the same, one can use (2.20) to derive a law of motion of the form

$$\Delta_t = h(\Delta_{t-1}, \Pi_t) \quad (2.26)$$

for the dispersion measure defined in (2.24), where

$$h(\Delta, \Pi) \equiv \alpha \Delta \Pi^{\theta(1+\omega)} + (1 - \alpha) \left(\frac{1 - \alpha \Pi^{\theta-1}}{1 - \alpha} \right)^{\frac{\theta(1+\omega)}{\theta-1}}.$$

This is the source of welfare losses from inflation or deflation.

The only relevant constraint on the monetary authority's ability to simultaneously stabilize inflation and output in this model is the aggregate-supply relation defined by (2.21), together with the definitions (2.14)–(2.17).⁵² The ability of the central bank to control i_t in each period gives it one degree of freedom each period (in each possible state of the world) with which to determine equilibrium outcomes. Because of the existence of the aggregate-supply relation (2.21) as a necessary constraint on the joint evolution of inflation and output, there is exactly one degree of freedom to be determined each period, in order to determine particular stochastic processes $\{\Pi_t, Y_t\}$ from among the set of possible rational-expectations equilibria. Hence I shall suppose that the monetary authority can choose from among the possible processes $\{\Pi_t, Y_t\}$ that constitute rational-expectations equilibria, and consider which equilibrium it is optimal to bring about; the detail that policy is implemented through the control of a short-term nominal interest rate will not actually matter to our calculations.

The Ramsey policy problem can then be defined as the choice of processes $\{Y_t, \Pi_t, F_t, K_t, \Delta_t\}$ for dates $t \geq t_0$ that satisfy conditions (2.18)–(2.19), (2.21) and (2.26) for all $t \geq t_0$ given the initial condition Δ_{t_0-1} , so as to maximize (2.25). Because the conditions (2.18)–(2.19) are forward-looking, however, the solution to this problem will not involve constant values of the endogenous variables (*i.e.*, a steady state) for any value of the initial price dispersion Δ_{t_0-1} , even in the absence of any random variation in the exogenous variables. This would prevent us from linearizing around a steady state solution — and from obtaining a solution for optimal stabilization policy that can be described by time-invariant coefficients, even if we are content with an approximate solution that is linear in the (small) disturbances. We can instead use local analysis in the neighborhood of a steady state, and obtain policy prescriptions with a time-invariant form, if we focus on the *asymptotic* character of optimal policy, once an optimal commitment (chosen at some earlier date) has converged to the neighborhood of a steady state. As in section 1, this amounts to analyzing a particular kind

⁵²This statement assumes that the zero lower bound on nominal interest rates never binds, as discussed above.

of constrained optimization problem, where policy from date t_0 onward is taken to be subject to a set of initial precommitments, chosen so that the policy that is optimal from t_0 onward subject to those constraints corresponds to the continuation of an optimal commitment that could have been chosen at an earlier date.

The state space required to state this problem can be further reduced by using (2.21) to substitute for the variable Π_t in equations (2.18)–(2.19) and in (2.26). We then obtain a set of equilibrium relations of the form

$$F_t = f(Y_t; \xi_t) + \alpha\beta E_t \phi_F(K_{t+1}, F_{t+1}), \quad (2.27)$$

$$K_t = k(Y_t; \xi_t) + \alpha\beta E_t \phi_K(K_{t+1}, F_{t+1}), \quad (2.28)$$

$$\Delta_t = \tilde{h}(\Delta_{t-1}, K_t/F_t) \quad (2.29)$$

for each period $t \geq t_0$, where the functions ϕ_F, ϕ_K are both homogeneous degree 1 functions of K and F . These constraints involve only the paths of the variables $\{Y_t, F_t, K_t, \Delta_t\}$, and since the objective has also been stated in terms of these variables, we can state the optimal policy problem in terms of the evolution of these variables alone. (A solution for the paths of these variables immediately implies a solution for inflation, using (2.21)).

The kind of initial pre-commitments that are required to create a modified problem with a time-invariant solution are of the form

$$\phi_F(K_{t_0}, F_{t_0}) = \bar{\phi}_F, \quad \phi_K(K_{t_0}, F_{t_0}) = \bar{\phi}_K, \quad (2.30)$$

where the values $\bar{\phi}_F, \bar{\phi}_K$ are chosen as functions of the economy's initial state in a “self-consistent” way, *i.e.*, according to formulas that also hold in all later periods in the constrained-optimal equilibrium.⁵³ Alternatively, there must be precommitments to particular values for F_{t_0} and K_{t_0} .

The problem of maximizing (2.25) subject to the constraints (2.27)–(2.29) in each period and the initial pre-commitments (2.30) has associated with it a Lagrangian of the form

$$\mathcal{L}_{t_0} = E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} L(Y_t, Z_t, \Delta_t, \Delta_{t-1}; \theta_t, \Theta_t, \Theta_{t-1}; \xi_t). \quad (2.31)$$

⁵³See Benigno and Woodford (2005a, 2008) or Giannoni and Woodford (2010) for more precise statements of this condition.

Here θ_t is a Lagrange multiplier associated with the backward-looking constraint (2.29), Θ_t is a vector of two Lagrange multipliers associated with the two forward-looking constraints (2.27)–(2.28), for each $t \geq t_0$, Θ_{t_0-1} is a corresponding vector of multipliers associated with the two initial pre-commitments (2.30), and

$$L(Y, Z, \Delta, \Delta_-; \theta, \Theta, \Theta_-; \xi) \equiv U(Y, \Delta; \xi) + \theta[\tilde{h}(\Delta_-, K/F) - \Delta] + \Theta'[z(Y; \xi) - Z] + \alpha\Theta'_-\Phi(Z), \quad (2.32)$$

where I use the shorthand

$$Z_t \equiv \begin{bmatrix} F_t \\ K_t \end{bmatrix}, \quad z(Y; \xi) \equiv \begin{bmatrix} f(Y; \xi) \\ k(Y; \xi) \end{bmatrix}, \quad \Phi(Z) \equiv \begin{bmatrix} \phi_F(K, F) \\ \phi_K(K, F) \end{bmatrix}.$$

Note that the inclusion of the initial pre-commitments makes the Lagrangian a sum of terms of the same form for each period $t \geq t_0$, which results in a system of time-invariant first-order conditions.

The Lagrangian is the same as for the problem of maximizing the modified objective

$$U_{t_0} + \alpha\Theta'_{t_0-1}\Phi(Z_{t_0}) \quad (2.33)$$

subject only to constraints (2.27)–(2.29) for periods $t \geq t_0$. Here the vector of initial multipliers Θ_{t_0-1} is part of the definition of the problem; the solution to this problem can represent the continuation of a prior optimal commitment if these multipliers are chosen as a function of the economy's initial state in a self-consistent way. This is an equivalent formulation of what it means for policy to be optimal from a timeless perspective.⁵⁴

2.3 Local Characterization of Optimal Dynamics

Differentiating the Lagrangian (2.31) with respect to each of the four endogenous variables yields a system of nonlinear first-order necessary conditions (FOCs) for optimality,

$$U_Y(Y_t, \Delta_t; \xi_t) + \Theta'_t z_Y(Y_t; \xi_t) = 0, \quad (2.34)$$

$$-\theta_t \tilde{h}_2(\Delta_{t-1}, K_t/F_t) \frac{K_t}{F_t^2} - \Theta_{1t} + \alpha\Theta'_{t-1} D_1(K_t/F_t) = 0, \quad (2.35)$$

⁵⁴This is the approach used in the numerical analysis of optimal policy in a related New Keynesian DSGE model by Khan *et al.* (2003).

$$\theta_t \tilde{h}_2(\Delta_{t-1}, K_t/F_t) \frac{1}{F_t} - \Theta_{2t} + \alpha \Theta'_{t-1} D_2(K_t/F_t) = 0, \quad (2.36)$$

$$U_\Delta(Y_t, \Delta_t; \xi_t) - \theta_t + \beta E_t[\theta_{t+1} \tilde{h}_1(\Delta_t, K_{t+1}/F_{t+1})] = 0, \quad (2.37)$$

each of which must hold for all $t \geq t_0$, where $\tilde{h}_i(\Delta, K/F)$ denotes the partial derivative of $\tilde{h}(\Delta, K/F)$ with respect to its i th argument, and $D_i(K/F)$ is the i th column of the matrix

$$D(Z) \equiv \begin{bmatrix} \partial_F \phi_F(Z) & \partial_K \phi_F(Z) \\ \partial_F \phi_K(Z) & \partial_K \phi_K(Z) \end{bmatrix}.$$

(Note that because the elements of $\Phi(Z)$ are homogeneous degree 1 functions of Z , the elements of $D(Z)$ are all homogenous degree 0 functions of Z , and hence functions of K/F only. Thus we can alternatively write $D(K/F)$.) The functions U_Y, U_Δ , and z_Y denote partial derivatives of the corresponding functions with respect to the argument indicated by the subscript. An optimal policy involves processes for the variables $\{Y_t, Z_t, \Delta_t, \theta_t, \Theta_t\}$ that satisfy both the structural equations (2.27)–(2.29) and the FOCs (2.34)–(2.37) for all $t \geq t_0$, given the initial values Δ_{t_0-1} and Θ_{t_0-1} . (Alternatively, we may require the additional conditions (2.30) to be satisfied as well, and solve for the elements of Θ_{t_0-1} as additional endogenous variables.)

Here I shall be concerned solely with the optimal equilibria that involve small fluctuations around a deterministic steady state. (This requires, of course, that the exogenous disturbances be small enough, and that the initial conditions be near enough to consistency with the optimal steady state.) An *optimal steady state* is a set of constant values $(\bar{Y}, \bar{Z}, \bar{\Delta}, \bar{\theta}, \bar{\Theta})$ that solve all seven of the equations just listed in the case that $\xi_t = \bar{\xi}$ at all times and initial conditions consistent with the steady state are assumed. One can show (Benigno and Woodford, 2005a; Giannoni and Woodford, 2010) that an optimal steady state exists in which the inflation rate is zero ($\bar{\Pi} = 1$), which means that $\bar{F} = \bar{K}$ and $\bar{\Delta} = 1$ (zero price dispersion in the steady state).

Briefly, conditions (2.27)–(2.29) are all satisfied as long as \bar{Y} is the output level implicitly defined by

$$f(\bar{Y}; \bar{\xi}) = k(\bar{Y}; \bar{\xi}),$$

and $\bar{F} = \bar{K} = (1 - \alpha\beta)^{-1} k(\bar{Y}; \bar{\xi})$. Because $\tilde{h}_2(1, 1) = 0$ (the effects of a small non-zero inflation rate on the measure of price dispersion are of second order, as shown by equation (2.59) below), conditions (2.35)–(2.36) reduce in the steady state to the eigenvector condition

$$\bar{\Theta}' = \alpha \bar{\Theta}' D(1). \quad (2.38)$$

Moreover, since when evaluated at a point where $F = K$,

$$\frac{\partial \log(\phi_K/\phi_F)}{\partial \log K} = -\frac{\partial \log(\phi_K/\phi_F)}{\partial \log F} = \frac{1}{\alpha},$$

we observe that $D(1)$ has a left eigenvector $[1 \ -1]$, with eigenvalue $1/\alpha$; hence (2.38) is satisfied if and only if $\bar{\Theta}_2 = -\bar{\Theta}_1$. Condition (2.34) provides one additional condition to determine the magnitude of the elements of $\bar{\Theta}$, and condition (2.37) provides one condition to determine the value of $\bar{\theta}$. In this way, one computes a steady-state solution to the FOCs.⁵⁵

I wish now to compute a local linear approximation to the optimal dynamics, in the case of equilibria in which all variables remain forever near these steady-state values. This can be obtained by linearizing both the structural relations (2.27)–(2.29) and the FOCs (2.34)–(2.37) around the steady-state values of all variables, and finding a bounded solution of the resulting system of linear equations.⁵⁶ Let us begin with the implications of the linearized structural relations.

Log-linearizing (2.29) around the steady-state values $\bar{\Delta} = 1, \bar{K}/\bar{F} = 1$, we obtain

$$\hat{\Delta}_t = \alpha \hat{\Delta}_{t-1}, \quad (2.39)$$

where $\hat{\Delta}_t \equiv \log \Delta_t$. Thus (to first order) the price dispersion evolves deterministically, regardless of monetary policy, and converges asymptotically to zero.

Log-linearizing (2.27)–(2.28), we obtain

$$\begin{aligned} \hat{F}_t &= (1 - \alpha\beta)[f_y \hat{Y}_t + f'_\xi \tilde{\xi}_t] + \alpha\beta E_t[(\theta - 1)\pi_{t+1} + \hat{F}_{t+1}], \\ \hat{K}_t &= (1 - \alpha\beta)[k_y \hat{Y}_t + k'_\xi \tilde{\xi}_t] + \alpha\beta E_t[\theta(1 + \omega)\pi_{t+1} + \hat{K}_{t+1}], \end{aligned}$$

using the notation

$$\hat{F}_t \equiv \log(F_t/\bar{F}), \quad f_y \equiv \frac{\partial \log f}{\partial \log Y}, \quad f'_\xi \equiv \frac{\partial \log f}{\partial \xi}$$

⁵⁵The second-order conditions that must be satisfied in order for the steady state to represent a local maximum of the Lagrangian rather than some other kind of critical point are discussed in section 2.5 below. Benigno and Woodford (2005a) show that these are satisfied as long as the model parameters satisfy a certain inequality, discussed further below.

⁵⁶Essentially, this amounts to using the implicit function theorem to compute a local linear approximation to the solution that is implicitly defined by the FOCs and structural relations. For further discussion, see Woodford, (2003, appendix A.3).

and corresponding definitions when K replaces F ; $\tilde{\xi}_t$ for $\xi_t - \bar{\xi}$; and $\pi_t \equiv \log \Pi_t$ for the rate of inflation. Subtracting the first of these equations from the second, one obtains an equation that involves only the variables $\hat{K}_t - \hat{F}_t$, π_t , \hat{Y}_t , and the vector of disturbances ξ_t . Log-linearization of (2.21) yields

$$\pi_t = \frac{1 - \alpha}{\alpha} \frac{1}{1 + \omega\theta} (\hat{K}_t - \hat{F}_t); \quad (2.40)$$

using this to substitute for $\hat{K}_t - \hat{F}_t$ in the relation just mentioned, we obtain

$$\pi_t = \kappa [\hat{Y}_t + u'_\xi \tilde{\xi}_t] + \beta E_t \pi_{t+1} \quad (2.41)$$

as an implication of the log-linearized structural equations, where

$$\kappa \equiv \frac{(1 - \alpha)(1 - \alpha\beta)\omega + \sigma^{-1}}{\alpha} > 0, \quad \sigma \equiv \tilde{\sigma} \frac{\bar{C}}{\bar{Y}} > 0, \quad (2.42)$$

and

$$u'_\xi \equiv \frac{k'_\xi - f'_\xi}{k_y - f_y}.$$

(This last expression is well-defined, since $k_y - f_y = \omega + \sigma^{-1} > 0$.)

Equation (2.41), which must hold for each $t \geq t_0$, is an important restriction upon the joint paths of inflation and output that can be achieved by monetary policy; note that it has precisely the form of the aggregate supply relation assumed in section 1. The composite exogenous disturbance term $u'_\xi \tilde{\xi}_t$ includes both the disturbances represented by the “cost-push” term in (1.1) and time variation in the level of output (“natural” or “potential” output) y_t^n relative to which the “output gap” is measured in (1.1); for the moment, it is not necessary to choose how to decompose the term into those two parts. (The distinction between the two types of terms only becomes meaningful when one considers the conditions for optimal stabilization.) Note that equation (2.41) is the *only* constraint on the bounded paths for inflation and aggregate output that can be achieved by an appropriate monetary policy;⁵⁷ for in the case of any bounded processes $\{\pi_t, \hat{Y}_t\}$, one can solve the log-linear equations stated above for bounded processes $\{\hat{F}_t, \hat{K}_t\}$ consistent with the model structural equations. One can similarly solve for the implied evolution of nominal interest rates and so on.

⁵⁷It is important that here I am considering only fluctuations within a sufficiently small neighborhood of the steady-state values for variables such as Π_t and Y_t ; hence the constraint associated with the zero lower bound for nominal interest rates is not an issue.

Let us next log-linearize the FOCs (2.34)–(2.37) around the steady-state values. Log-linearizing (2.35)–(2.36) yields the vector equation

$$-\frac{\bar{\theta}}{\bar{K}} \frac{1 - \alpha \theta(1 + \omega)}{\alpha} \frac{1}{1 + \omega \theta} [(\hat{K}_t - \hat{F}_t) + \alpha \hat{\Delta}_{t-1}] \begin{bmatrix} 1 \\ -1 \end{bmatrix} - \tilde{\Theta}_t + \alpha D(1)' \tilde{\Theta}_{t-1} + \alpha M \hat{Z}_t = 0, \quad (2.43)$$

where $\tilde{\Theta}_t \equiv \Theta_t - \bar{\Theta}$, $\hat{Z}_t' \equiv [\hat{F}_t \ \hat{K}_t]'$, and M is \bar{K} times the Hessian matrix of second partial derivatives of the function $\bar{\Phi}(Z) \equiv \bar{\Theta}'\Phi(Z)$. The fact that $\bar{\Phi}(Z)$ is homogeneous of degree 1 implies that its derivatives are homogeneous of degree 0, and hence functions only of K/F ; it follows that the matrix M is of the form

$$M = m \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}, \quad (2.44)$$

where m is a scalar. Similarly, the fact that each element of $\Phi(Z)$ is homogeneous of degree 1 implies that

$$D(1)e = e,$$

where $e' \equiv [1 \ 1]$.

Pre-multiplying (2.43) by e' therefore yields

$$e' \tilde{\Theta}_t = \alpha e' \tilde{\Theta}_{t-1} \quad (2.45)$$

for all $t \geq t_0$. This implies that $e' \tilde{\Theta}_t$ converges to zero with probability 1, regardless of the realizations of the disturbances; hence under the optimal dynamics, the asymptotic fluctuations in the endogenous variables are such that

$$\tilde{\Theta}_{2,t} = -\tilde{\Theta}_{1,t} \quad (2.46)$$

at all times. And if we assume initial Lagrange multipliers such that (2.46) is satisfied for $t = t_0 - 1$ (or initial pre-commitments with that implication), then (2.46) will hold for all $t \geq t_0$. In fact, I shall assume initial pre-commitments of this kind; note that this is an example of a “self-consistent” principle for selection of the initial pre-commitments, since under the constrained-optimal policy (2.46) will indeed hold in all subsequent periods.⁵⁸ Hence the optimal dynamics satisfy (2.46) at all times.

⁵⁸It is also worth noting that if at some date t_{orig} in the past, the policymaker has made a commitment to an unconstrained Ramsey policy, then at that initial date the lagged Lagrange multipliers would have satisfied (2.46), because both elements of $\Theta_{t_{orig}-1}$ would have been equal to zero.

There must also exist a vector v such that $v_2 \neq v_1$ and such that $D(1)v = \alpha^{-1}v$, since we have already observed above that $1/\alpha$ is one of the eigenvalues of the matrix. (The vector v must also not be a multiple of e , as e is the other right eigenvector, with associated eigenvalue 1.) Pre-multiplying (2.43) by v' then yields

$$-\frac{\bar{\theta}}{\bar{K}} \frac{1 - \alpha \theta(1 + \omega)}{\alpha(1 + \omega\theta)} [(\hat{K}_t - \hat{F}_t) + \alpha \hat{\Delta}_{t-1}] - \tilde{\Theta}_{1,t} + \tilde{\Theta}_{1,t-1} - \alpha m (\hat{K}_t - \hat{F}_t) = 0. \quad (2.47)$$

Here the common factor $v_1 - v_2 \neq 0$ has been divided out from all terms, and $\tilde{\Theta}_{2,t}$ has been eliminated using (2.46). Note that conditions (2.45) and (2.47) exhaust the implications of (2.43), and hence of conditions (2.35)–(2.36). We can again use (2.40) to substitute for $\hat{K}_t - \hat{F}_t$ in condition (2.47), in order to express the condition in terms of its implications for optimal inflation dynamics. We thus obtain a relation of the form

$$\xi_\pi \pi_t + \xi_\Delta \hat{\Delta}_{t-1} = \tilde{\Theta}_{1,t} - \tilde{\Theta}_{1,t-1}. \quad (2.48)$$

Condition (2.34) can similarly be log-linearized to yield

$$\bar{Y}[U_{YY} + \bar{\Theta}' z_{YY}] \hat{Y}_t + [U'_{Y\xi} + \bar{\Theta}' z_{Y\xi}] \tilde{\xi}_t + U_{Y\Delta} \hat{\Delta}_t - \frac{\bar{K}}{\bar{Y}} (k_y - f_y) \tilde{\Theta}_{1,t} = 0,$$

again using (2.46) to eliminate $\tilde{\Theta}_{2,t}$. We can equivalently write this as

$$\bar{Y}[U_{YY} + \bar{\Theta}' z_{YY}] (\hat{Y}_t - \hat{Y}_t^*) + U_{Y\Delta} \hat{\Delta}_t - \frac{\bar{K}}{\bar{Y}} (k_y - f_y) \tilde{\Theta}_{1,t} = 0, \quad (2.49)$$

where $\hat{Y}_t^* \equiv \log(Y_t^*/\bar{Y})$, and Y_t^* is a function of the exogenous disturbances, implicitly defined by the equation

$$U_Y(Y_t^*, 1; \xi_t) + \bar{\Theta}' z_Y(Y_t^*; \xi_t) = 0. \quad (2.50)$$

This “target level of output” (introduced by Benigno and Woodford, 2005a) is related to, but not the same as, the efficient level of output Y_t^e (*i.e.*, the quantity that would be produced of each good in order to maximize expected utility, subject only to the constraints imposed by technology), implicitly defined by the equation

$$U_Y(Y_t^e, 1; \xi_t) = 0. \quad (2.51)$$

One observes that in the case that the zero-inflation steady-state level of output \bar{Y} (which would also be the steady-state level of output under flexible prices) is efficient

(in the case that $\xi_t = \bar{\xi}$ at all times), so that $U_Y(\bar{Y}, 1; \bar{\xi}) = 0$, we have $\bar{\Theta} = 0$, and Y_t^* and Y_t^e will coincide. More generally, when $\bar{\Theta} \neq 0$, the target level Y_t^* differs from Y_t^e in that it is equal on average (to a first-order approximation) to \bar{Y} , which differs from the average level of Y_t^e : in the case of empirical interest, Y_t^* is lower than Y_t^e on average, because keeping Y_t as high as Y_t^e on average is not consistent with stable prices (even on average), if it is possible at all. The way in which Y_t^* responds to shocks can also be different from the way that Y_t^e responds, again to mitigate the degree to which the output variations would otherwise require instability of prices.

Solving (2.49) for $\tilde{\Theta}_{1,t}$, and using this to substitute for $\tilde{\Theta}_{1,t}$ in (2.48), we obtain a relation of the form

$$\xi_\pi \pi_t + \lambda_x (x_t - x_{t-1}) + \lambda_\Delta (\hat{\Delta}_t - \hat{\Delta}_{t-1}) + \xi_\Delta \hat{\Delta}_{t-1} = 0, \quad (2.52)$$

that must hold for all $t > t_0$, where the “output gap” $x_t \equiv \hat{Y}_t - \hat{Y}_t^*$. If we select $\tilde{\Theta}_{1,t_0-1}$ to be the value required in order for (2.49) to hold for $t = t_0 - 1$ as well,⁵⁹ then (2.52) must hold for $t = t_0$ as well. Condition (2.52) represents a restriction on the path of the endogenous variables that is required for consistency with the FOCs. Moreover, it is the *only* restriction required for consistency with the FOCs. For in the case of any bounded processes $\{\pi_t, \hat{Y}_t, \hat{\Delta}_t\}$ consistent with (2.52) for all $t \geq t_0$, one can construct an implied process $\{\tilde{\Theta}_t\}$ using (2.49) to solve for $\tilde{\Theta}_{1,t}$ and (2.46) to solve for $\tilde{\Theta}_{2,t}$. The linearized version of (2.37), which is of the form

$$\tilde{\theta}_t = \alpha\beta E_t \tilde{\theta}_{t+1} + E_t [g(\hat{Y}_t, \hat{\Delta}_t, \pi_{t+1})]$$

for a certain linear function $g(\cdot)$, can then be “solved forward” to obtain a bounded process $\{\tilde{\theta}\}$. Thus one can construct bounded processes for the Lagrange multipliers that satisfy each of the linearized FOCs by construction.

We thus conclude that a state-contingent evolution of the economy remaining forever near enough to the optimal steady state is both feasible and an optimal plan (in the case of initial Lagrange multipliers selected as described above) if and only if the bounded processes $\{\pi_t, \hat{Y}_t, \hat{\Delta}_t\}$ satisfy (2.39), (2.41), and (2.52) for all $t \geq t_0$. It is easily seen that these equations determine unique bounded processes for these variables, given initial conditions $(\hat{Y}_{t_0-1}, \hat{\Delta}_{t_0-1})$ and a bounded process for

⁵⁹Note that this would be a self-consistent principle for selecting the initial Lagrange multipliers, since under the optimal plan for this modified problem, (2.49) will indeed hold in all periods $t \geq t_0$.

the exogenous disturbances $\{\tilde{\xi}_t\}$. We can express all three equations in terms of the variables $\{\pi_t, x_t, \hat{\Delta}_t\}$ if we rewrite (2.41) as

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1} + u_t, \quad (2.53)$$

where

$$u_t \equiv \kappa[\hat{Y}_t^* + u'_\xi \tilde{\xi}_t]. \quad (2.54)$$

Moreover, (2.39) obviously has a unique bounded solution for $\{\hat{\Delta}_t\}$, given the initial condition; treating this sequence $\{\hat{\Delta}_t\}$ as exogenously given, there remain two stochastic difference equations per period to determine the two “endogenous” variables $\{\pi_t, x_t\}$. Moreover, equations (2.52) and (2.53) are of exactly the same form as equations (1.7) and (1.21) in the model of section 1, except that (2.52) contains additional (bounded) “exogenous disturbance” terms. The presence of these additional terms does not affect the conditions for determinacy of the solution, and so one can show that there exists a unique bounded solution for any given initial conditions, using the same argument as in section 1.

This allows us to characterize (to a linear approximation) the equilibrium dynamics of all endogenous variables under an optimal policy. In the case of initial conditions such that $\hat{\Delta}_{t_0-1} = 0$ (to first order), as assumed in Benigno and Woodford (2005a), the optimal equilibrium dynamics are of exactly the sort calculated in section 1, except that the micro-founded model gives specific answers about two important issues: (i) it explains to what extent various types of “fundamental” economic disturbances (changes in technology, preferences, or fiscal policy) should change the target level of output (and hence one’s measure of the “output gap”), contribute to the “cost-push” term u_t , or both; and (ii) it gives a specific value for the coefficient ϕ in the optimal target criterion (1.21), as a function of underlying model parameters, rather than making it a function of an arbitrary weight λ in the policymaker’s objective. (The answers to these questions are discussed further in sections 2.4 and 2.6 below.)

The welfare-based analysis yields another result not obtained in section 1: it explains how the optimal dynamics of inflation and output should be affected by a substantial initial level of price dispersion. (Under an optimal policy, of course, the policymaker should eventually *never* face a situation in which the existing level of price dispersion is large; but one might wish to consider the transitional dynamics that should be chosen under a newly chosen optimal policy commitment, when actual policy in the recent past has been quite bad.) Equation (2.52) indicates that the

central bank’s target for growth of the output-gap-adjusted price level should be different depending on the inherited level of price dispersion; a larger initial price dispersion reduces the optimal target rate of growth in the gap-adjusted price level, as first shown by Yun (2005) in a special case that allowed an analytical solution.

We again find that optimal policy can be described by fulfillment of a target criterion that can be described as a “flexible inflation target”; in the case that the initial level of price dispersion is zero to first order (as it will then continue to be under optimal policy, and indeed under any policy under which the departures of the inflation rate from zero are of only first order), the optimal target criterion is again of precisely the form (1.21) derived in section 1. Moreover, the result that optimal policy in a microfounded model can be characterized by a target criterion of this general form does not depend on the multitude of special assumptions made in this example. Giannoni and Woodford (2010) show, for a very broad class of stabilization problems in which welfare is measured by the expected discounted value of some function of a vector of endogenous variables, and the paths for those variables that are consistent with equilibrium (under some suitable choice of policy) are defined by a system of nonlinear stochastic difference equations that include both backward-looking elements (like the dependence of (2.29) on Δ_{t-1}) and forward-looking elements (like the dependence of (2.27)–(2.28) on the expectations of variables in period $t+1$), that it is possible to find a linear target criterion the fulfillment of which is necessary and sufficient for policy to coincide (at least to a linear approximation) with an optimal policy commitment.

The particular endogenous variables that the target criterion involves depends, of course, on the structure of one’s model. However, in a broad range of models with some basic features in common with the one just analyzed, some measure of inflation and some measure of the output gap will again be key variables in the optimal target criterion. This can be made clearer through a further discussion of the reason why the optimal target criterion can be expressed in terms of those two variables in the case just treated.⁶⁰

⁶⁰Examples of optimal target criteria for models that generalize the one assumed in this section are presented below in section 3.

2.4 A Welfare-Based Quadratic Objective

It may be considered surprising that the FOCs that characterize optimal policy in the microfounded model end up being equivalent to the same form of target criterion as in the case of the linear-quadratic policy problem discussed in section 1. As shown in the previous section, a log-linear approximation to the structural equations of the microfounded model implies precisely the same restriction upon the joint paths of inflation and output as the “New Keynesian Phillips curve” assumed in section 1. But even so, the assumed *objective* of policy in the welfare-based analysis — the maximization of expected utility, that depends on consumption and labor effort, rather than output and inflation — might seem quite different than in the earlier analysis. This section seeks to provide insight into the source of the result, by showing that one can write a quadratic approximation to the expected utility objective assumed above — a degree of approximation that suffices for a derivation of a linear approximation to the optimal dynamics, of the kind discussed in the previous subsection — that takes exactly the form (1.6) assumed in section 1, under a suitable definition of the “output gap” in that objective and for a suitable specification of the relative weight λ assigned to the output-gap stabilization objective. (The analysis follows that in Woodford, 2003, chap. 6, and Benigno and Woodford, 2005a.)

I have already shown that in the above model, it is possible to write the utility of the representative household as a function of the evolution of two endogenous variables $\{Y_t, \Delta_t\}$. Let us consider a disturbance process under which ξ_t remains in a bounded neighborhood of $\bar{\xi}$ for all t , and plans in which Y_t remains in a bounded neighborhood of \bar{Y} and Δ_t remains in a bounded neighborhood of 1 for all t ,⁶¹ and compute a second-order Taylor series expansion of (2.25) in terms of \hat{Y}_t , $\hat{\Delta}_t$, and $\tilde{\xi}_t$. For the contribution to utility in any period, we obtain

$$\begin{aligned}
 U(Y_t, \Delta_t; \xi_t) &= \bar{Y}U_Y \hat{Y}_t + U_\Delta \hat{\Delta}_t + \frac{1}{2}(\bar{Y}U_Y + \bar{Y}^2U_{YY}) \hat{Y}_t^2 \\
 &\quad + \bar{Y}U_{Y\Delta} \hat{Y}_t \hat{\Delta}_t + \bar{Y}U'_{Y\xi} \tilde{\xi}_t \hat{Y}_t + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3), \quad (2.55)
 \end{aligned}$$

where all derivatives are evaluated at the steady state; “t.i.p.” refers to terms the value of which is independent of policy (*i.e.*, that do not involve endogenous variables), and that can therefore be ignored for purposes of the welfare ranking of alternative

⁶¹Essentially, this requires that we restrict attention to policies in which inflation never deviates too far from zero.

policies; and $\|\xi\|$ is a bound on the amplitude of the exogenous disturbances (*i.e.*, on the elements of $\tilde{\xi}_t$). Here it is assumed that the only policies considered are ones in which \hat{Y}_t and $\hat{\Delta}_t$ are of order $\mathcal{O}(\|\xi\|)$ as well (so that, for example, a term proportional to $\hat{Y}_t^2 \tilde{\xi}_t$ must be of order $\mathcal{O}(\|\xi\|^3)$). Furthermore, I have used the fact that the evolution of $\hat{\Delta}_t$ is independent of policy to first order (*i.e.*, up to a residual of order $\mathcal{O}(\|\xi\|^2)$), because of equation (2.39), to show that terms proportional to $\hat{\Delta}_t \tilde{\xi}_t$ or to $\hat{\Delta}_t^2$ are independent of policy, to second order (*i.e.*, up to a residual of order $\mathcal{O}(\|\xi\|^3)$), allowing these terms to be included in the final two catch-all terms.

While the substitution of (2.55) into (2.25) would yield a discounted quadratic objective for policy, it is not necessarily true that the use of this quadratic objective together with a log-linear approximation to the model structural equations, as in section 1, would yield a correct log-linear approximation to the dynamics under optimal policy. A correct welfare comparison among rules, even to this order of accuracy, would depend on evaluation of the objective to second order under each of the different contemplated policies, and a term such as $\bar{Y} U_Y \hat{Y}_t$ cannot (in general) be evaluated to second order in accuracy using an approximate solution for the path of \hat{Y}_t under a given policy that is only accurate to first order.⁶² This issue can be dealt with in various ways, some more generally applicable than others.

2.4.1 The Case of an Efficient Steady State

The analysis is simplified if we assume that the steady state level of output \bar{Y} is optimal in the case that $\xi_t = \bar{\xi}$ at all times; here I mean not just among the outcomes achievable by monetary policy (which it is, as discussed above), but among all allocations that are technologically possible (*i.e.*, that $Y_t^e = \bar{Y}$). This will be true if and only if $\Psi = 0$, where the steady-state inefficiency measure Ψ is defined by

$$1 - \Psi \equiv \frac{1 - \bar{\tau}}{\bar{\mu}^p \bar{\mu}^w} = \frac{1 - \bar{\tau} \theta - 1}{\bar{\mu}^w \theta} > 0.$$

Since we assume that $\theta > 1$ (implying that the desired markup of prices relative to marginal cost $\mu^p > 1$) and $\mu^w \geq 1$, this requires that $\bar{\tau} < 0$, so that there is at least a mild subsidy to production and/or sales to offset the distortion resulting from firms'

⁶²For further discussion of this issue, see Woodford (2003, chap. 6), Kim and Kim (2003), or Benigno and Woodford (2008).

market power.⁶³

In this case, $U_Y = 0$ (evaluated at the efficient steady state), eliminating two of the terms in (2.55). Most crucially for our discussion, there is no longer a linear term in \hat{Y}_t , which was problematic for the reason just discussed. The term that is proportional to $\tilde{\xi}_t \hat{Y}_t$ can also be given a simple interpretation in this case. Recall that the efficient rate of output Y_t^e is implicitly defined by the equation (2.51). Total differentiation of this equation yields

$$\bar{Y}U_{YY}\hat{Y}_t^e + U'_{Y\xi}\tilde{\xi}_t = 0, \quad (2.56)$$

where $\hat{Y}_t^e \equiv \log(Y_t^e/\bar{Y})$. Using this to substitute for the factor $U'_{Y\xi}\tilde{\xi}_t$ in (2.55), and completing the square, we obtain

$$\begin{aligned} U(Y_t, \Delta_t; \xi_t) &= \frac{1}{2}(\bar{Y}^2 U_{YY})(\hat{Y}_t - \hat{Y}_t^e)^2 + U_\Delta \hat{\Delta}_t + \bar{Y}U_{Y\Delta}\hat{Y}_t\hat{\Delta}_t \\ &\quad + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \end{aligned} \quad (2.57)$$

If we further assume an initial condition under which $\hat{\Delta}_{t_0-1} = \mathcal{O}(\|\xi\|^2)$ (a condition that will hold, at least asymptotically, under any “near-steady-state” policy in the class that we are considering), then we will have $\hat{\Delta}_t = \mathcal{O}(\|\xi\|^2)$ for all t as a consequence of (2.39).⁶⁴ We can then include the term proportional to $\hat{Y}_t\hat{\Delta}_t$ among the terms of order $\mathcal{O}(\|\xi\|^3)$, and write

$$\begin{aligned} U(Y_t, \Delta_t; \xi_t) &= \frac{1}{2}(\bar{Y}^2 U_{YY})(\hat{Y}_t - \hat{Y}_t^e)^2 - \bar{v} \hat{\Delta}_t \\ &\quad + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3), \end{aligned} \quad (2.58)$$

where $\bar{v} \equiv v(\bar{Y}; \bar{\xi}) > 0$.

⁶³This is obviously a special case and counter-factual as well, but there are various reasons to consider it. One is that it provides insight into the kind of results that will also be obtained (at least approximately) in economies where steady-state distortions are not too large (Ψ is small), and has the advantage of making the calculations simple. Another might be that, as suggested by Rotemberg and Woodford (1997), it should be considered the task of other aspects of policy to cure structural distortions that make the steady-state level of economic activity inefficient, so that one might wish to design monetary policy for an environment in which this problem has been solved by other means, rather than assuming that monetary policy rules should be judged on the basis of their ability to mitigate distortions in the average level of output.

⁶⁴Recall that (2.39) is an equation that holds up to a residual of order $\mathcal{O}(\|\xi\|^2)$. A second-order approximation is instead given by equation (2.59).

This is still not an objective that can be evaluated to second order using only a first-order solution for the paths of the endogenous variables $\{\hat{Y}_t, \hat{\Delta}_t\}$, because of the presence of the terms that is linear in $\hat{\Delta}_t$. This can be cured, however, by using a second-order approximation to (2.26) to replace the $\hat{\Delta}_t$ terms by purely quadratic terms. A Taylor expansion of the function $h(\Delta, \Pi)$ yields

$$\hat{\Delta}_t = \alpha \hat{\Delta}_{t-1} + \frac{1}{2} h_{\pi\pi} \pi_t^2 + \mathcal{O}(\|\xi\|^3). \quad (2.59)$$

where

$$h_{\pi\pi} = \frac{\alpha}{1-\alpha} \theta(1+\omega)(1+\omega\theta) > 0,$$

and again I have specialized to the case in which Δ_{t_0-1} is of order $\mathcal{O}(\|\xi\|^2)$. It then follows that

$$\sum_{t=t_0}^{\infty} \beta^{t-t_0} \hat{\Delta}_t = \frac{1}{2} \frac{h_{\pi\pi}}{1-\alpha\beta} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \pi_t^2 + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3), \quad (2.60)$$

where the terms proportional to $\hat{\Delta}_{t_0-1}$ are included in the term ‘‘t.i.p.’’ Substituting the approximation (2.58) into (2.25), and using (2.60) to substitute for the sum of $\hat{\Delta}_t$ terms, we obtain

$$U_{t_0} = \frac{1}{2} E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \left[\bar{Y}^2 U_{YY} (\hat{Y}_t - \hat{Y}_t^e)^2 - (1-\alpha\beta)^{-1} \bar{v} h_{\pi\pi} \pi_t^2 \right] + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \quad (2.61)$$

This is equal to a negative constant times a discounted loss function of the form (1.6), where the welfare-relevant output gap is defined as $x_t \equiv \hat{Y}_t - \hat{Y}_t^e$, $x^* = 0$, and the relative weight on the output-gap stabilization objective is

$$\lambda = -\frac{(1-\alpha\beta)\bar{Y}^2 U_{YY}}{\bar{v} h_{\pi\pi}} = \frac{\kappa}{\theta} > 0, \quad (2.62)$$

where κ is the same coefficient as in (2.41).

The log-linearized aggregate supply relation (2.41) can also be written in the form (1.7) assumed in section 1, where x_t is the welfare-relevant output gap just defined, for an appropriate definition of the ‘‘cost-push’’ term u_t . Note that since

$$k(Y; \xi) = \mu^p \mu^w v_y(Y; \xi) Y, \quad f(Y; \xi) = (1-\tau) u_y(Y; \xi) Y,$$

we have

$$(k'_\xi - f'_\xi) \tilde{\xi}_t = u_y^{-1} (v'_{y\xi} - u'_{y\xi}) \tilde{\xi}_t + \hat{\mu}_t^w + \hat{\tau}^t,$$

where

$$\hat{\mu}_t^w \equiv \log(\mu_t^w / \bar{\mu}^w), \quad \hat{\tau}_t \equiv -\log(1 - \tau_t / 1 - \bar{\tau}).$$

From this it follows that

$$u'_\xi \tilde{\xi}_t = -\hat{Y}_t^e + (\omega + \sigma^{-1})^{-1}(\hat{\mu}_t^w + \hat{\tau}_t). \quad (2.63)$$

Substitution into (2.41) yields an aggregate-supply relation of the form (1.7), where $x_t \equiv \hat{Y}_t - \hat{Y}_t^e$, and u_t is a positive multiple of $\hat{\mu}_t^w + \hat{\tau}_t$.

Hence in this case we obtain a linear-quadratic policy problem of exactly the form considered in section 1, except that proceeding from explicit microfoundations provides a precise interpretation for the output gap x_t , a precise value for the target x^* (here equal to zero, because the steady state is efficient), a precise value for the relative weight λ in the loss function (1.6) as a function of the model structural parameters, and a precise interpretation of the “cost-push” term u_t . An implication of these identifications is that the optimal target criterion (1.21) takes the more specific form

$$\pi_t + \theta^{-1}(x_t - x_{t-1}) = 0. \quad (2.64)$$

As in section 1, the target criterion can equivalently be expressed in the form (1.23), where the “output-gap-adjusted price level” is now defined as

$$\tilde{p}_t \equiv p_t + \theta^{-1}x_t, \quad (2.65)$$

where $p_t \equiv \log P_t$.

Of course, these results are obtained under a number of simplifying assumptions. If we do not assume that the initial dispersion of prices is small (so that $\hat{\Delta}_{t_0-1}$ is non-zero to first order), then several terms omitted in the derivation of (2.61) must be restored. However, as long as $\hat{\Delta}_{t_0-1} = \mathcal{O}(\|\xi\|)$, we can write

$$\hat{\Delta}_t \pi_t = \bar{\Delta}_t \pi_t + \mathcal{O}(\|\xi\|^3),$$

$$\hat{\Delta}_t \hat{Y}_t = \bar{\Delta}_t \hat{Y}_t + \mathcal{O}(\|\xi\|^3),$$

$$\hat{\Delta}_t^2 = \bar{\Delta}_t^2 + \mathcal{O}(\|\xi\|^3),$$

where

$$\bar{\Delta}_t = \hat{\Delta}_{t_0-1} \alpha^{t-t_0+1}$$

is a deterministic sequence (depending only on the initial condition). (Here I again rely on the fact that (2.39) holds up to a residual of order $\mathcal{O}(\|\xi\|^2)$.)

With these substitutions, (2.61) becomes takes the more general form

$$\begin{aligned}
U_{t_0} = \frac{1}{2} E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \left\{ \bar{Y}^2 U_{YY} (\hat{Y}_t - \hat{Y}_t^e)^2 - (1 - \alpha\beta)^{-1} \bar{v} h_{\pi\pi} \pi_t^2 \right. \\
\left. + 2\bar{Y} U_{Y\Delta} \bar{\Delta}_t (\hat{Y}_t - \hat{Y}_t^n) - 2(1 - \alpha\beta)^{-1} \bar{v} h_{\pi\Delta} \bar{\Delta}_t \pi_t \right\} \\
+ \text{i.p.} + \mathcal{O}(\|\xi\|^3). \tag{2.66}
\end{aligned}$$

This is equal to a negative multiple of a discounted quadratic loss function of the form

$$E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} [\pi_t^2 + \lambda x_t^2 + \gamma_\pi \bar{\Delta}_t \pi_t + \gamma_x \bar{\Delta}_t x_t], \tag{2.67}$$

where x_t and λ are defined as above, and

$$\begin{aligned}
\gamma_\pi &\equiv \frac{2(1 - \alpha)}{1 + \omega\theta} > 0, \\
\gamma_x &\equiv \frac{2(1 - \alpha)(1 - \alpha\beta)}{\alpha} \frac{1}{\theta(1 + \omega\theta)} > 0.
\end{aligned}$$

As a consequence of the additional terms in the loss function, the optimal target criterion (generalizing (1.21)) contains additional terms proportional to $\bar{\Delta}_t$ and $\bar{\Delta}_{t-1}$, or equivalently (to a linear approximation) $\hat{\Delta}_t$ and $\hat{\Delta}_{t-1}$, as shown in section 2.3.

Because $\bar{\Delta}_t$ is exponentially decreasing in t , the additional terms can alternatively be combined into a single term proportional to $\bar{\Delta}_t$, $\bar{\Delta}_{t-1}$, or $\bar{\Delta}_t - \bar{\Delta}_{t-1}$. Similarly, because (2.39) holds to first order, the additional terms in the target criterion can be reduced to a single term proportional to $\hat{\Delta}_t$, $\hat{\Delta}_{t-1}$, or $\hat{\Delta}_t - \hat{\Delta}_{t-1}$. In this last representation, the modified target criterion is of the form

$$\pi_t + \theta^{-1}(x_t - x_{t-1}) + \gamma \log(\Delta_t/\Delta_{t-1}) = 0,$$

which is equivalent to requiring that

$$p_t + \gamma \log \Delta_t + \theta^{-1} x_t = p^*$$

for some constant p^* . This last criterion is again an output-gap adjusted price level target. It differs from (1.23) only in that the price index that is used is not p_t but

rather $p_t + \gamma \log \Delta_t$; this latter quantity is just another price index (*i.e.*, the log of another homogeneous degree 1 function of the individual prices).⁶⁵

Another aspect in which the analysis in this section is special is in the assumption that the zero-inflation steady state is efficient. I next consider the consequences of relaxing this assumption.

2.4.2 The Case of Small Steady-State Distortions

The previous analysis requires only a small modification in the case that Ψ (the measure of the degree of inefficiency of the steady-state level of output) is non-zero, if Ψ is only of the order $\mathcal{O}(\|\xi\|)$, as assumed in Woodford (2003, chap. 6).⁶⁶ In this case, computing a solution that is accurate “to first order” means that the optimal equilibrium dynamics

$$\pi_t = \pi(\xi_t, \dots; \Psi)$$

of a variable such as π_t , in the case of any given small enough value of Ψ , can be approximated by a linear function

$$\pi_t = \alpha_\Psi \Psi + \alpha'_{\xi,0} \tilde{\xi}_t + \dots$$

up to an error that is of order $\mathcal{O}(\|\xi\|^2)$.

Here the coefficients $\alpha_{\xi,0}$ represent partial derivatives of the function $\pi(\cdot)$ with respect to the elements of ξ_t , evaluated at the steady state with $\xi_t = \bar{\xi}$ for all t and $\Psi = 0$, while the coefficient α_Ψ represents the partial derivative of the function with respect to Ψ , evaluated at the same steady state. It follows that the coefficients $\alpha_{\xi,0}$ are the same as in the calculation where it was assumed that $\Psi = 0$. Thus the extension proposed here does not consider the consequences of a distorted steady state for the optimal responses to shocks (this would depend on terms higher than “first order”), but only the effects of the distortions on the average values of variables such as the inflation rate. However, the latter question is of interest, for example, in

⁶⁵In the case that $\hat{\Delta}_{t_0-1} = \mathcal{O}(\|\xi\|^2)$, all price indices are the same, to first order, and it does not matter which one is used in the target criterion. In general, when there are non-trivial relative-price differences, it *will* matter which price index is targeted; this issue is discussed further in section 3.2 below.

⁶⁶Technically, this means that in order to ensure a certain degree of accuracy in our approximate characterization of the optimal dynamics, it is necessary to make the value of Ψ sufficiently small, in addition to making the amplitude of the exogenous disturbances sufficiently small.

considering whether inefficiency of the zero-inflation steady state is a reason for the optimal steady-state inflation rate to differ from zero.

In this case, $U_Y = \Psi u_y = \mathcal{O}(\|\xi\|)$, and as a consequence, the $\bar{Y}U_Y\hat{Y}_t$ term in (2.55) can still be evaluated to second order using a solution for \hat{Y}_t that is accurate only to first order.⁶⁷ The method used above to substitute for the $U_\Delta\hat{\Delta}_t$ terms then suffices once again to yield a quadratic welfare objective that can be evaluated to second order using only the log-linearized structural relations to solve for the paths of the endogenous variables. We can no longer neglect the $\bar{Y}U_Y\hat{Y}_t$ term in (2.55), but the $\bar{Y}U_Y\hat{Y}_t^2$ can still be neglected, as it is of order $\mathcal{O}(\|\xi\|^3)$. Finally, in substituting for the $\bar{Y}U'_{Y\xi}\tilde{\xi}_t\hat{Y}_t$ term in (2.55), it is important to note that (2.56) now takes the more general form

$$\bar{Y}U_{YY}\hat{Y}_t^e + U'_{Y\xi}\tilde{\xi}_t = -U_Y = -\Psi u_y. \quad (2.68)$$

Making these substitutions, we once again obtain (2.57), even though $U_Y \neq 0$. If we again simplify by restricting attention to initial conditions under which $\hat{\Delta}_{t_0-1} = \mathcal{O}(\|\xi\|^2)$, we again obtain (2.61) as an approximate quadratic loss function.

However, it is no longer appropriate to define the “output gap” as $\hat{Y}_t - \hat{Y}_t^e$, if we want the gap to be a variable that is equal to zero in the zero-inflation steady state (as in the analysis of Clarida *et al.*, 1999). If we define⁶⁸ the “natural rate of output” Y_t^n as the flexible-price equilibrium level of output (common equilibrium quantity produced of each good) in the case that τ_t and μ_t^w take their steady-state values, but all other disturbances take the (time-varying) values specified by the vector ξ_t , then Y_t^n is implicitly defined by

$$(1 - \bar{\tau})u_y(Y_t^n; \xi_t) = \mu^p \bar{\mu}^w v_y(Y_t^n; \xi_t). \quad (2.69)$$

We observe that $Y_t^n = \bar{Y}$ in the zero-inflation steady state, so that the output gap definition

$$x_t \equiv \hat{Y}_t - \hat{Y}_t^n \quad (2.70)$$

⁶⁷It is important to note that in all of the Taylor expansions discussed in this section, expansions are around the zero-inflation steady state, *not* the efficient steady-state allocation; \bar{Y} refers to the zero-inflation steady-state output level, not the efficient steady-state output level; and variables such as \hat{Y}_t and \hat{Y}_t^e are defined relative to \bar{Y} , not relative to the steady-state value of Y_t^e . When $\Psi = 0$, it is not necessary to distinguish between the two possible definitions of the steady-state level of output.

⁶⁸See Woodford (2003, chap. 6), for further discussion.

has the desired property. Moreover, total differentiation of (2.69) and comparison with (2.68) indicates that

$$\hat{Y}_t^e = \hat{Y}_t^n + (-\bar{Y}U_{YY})^{-1}\Psi u_y + \mathcal{O}(\|\xi\|^2).$$

Hence

$$\hat{Y}_t - \hat{Y}_t^e = x_t - x^*,$$

up to an error of order $\mathcal{O}(\|\xi\|^2)$, where

$$x^* \equiv -\frac{U_Y}{\bar{Y}U_{YY}} = \frac{\Psi}{\omega + \sigma^{-1}} + \mathcal{O}(\|\xi\|^2). \quad (2.71)$$

Using this to substitute for $\hat{Y}_t - \hat{Y}_t^e$ in (2.61), we obtain a quadratic objective that is a negative multiple of a loss function of the form (1.6), where now x_t is defined by (2.70) and x^* is defined by (2.71). Note that x^* has the same sign as Ψ (positive, in the case of empirical relevance), and is larger the larger is Ψ . Moreover, repeating the derivation of (2.63), we find that when $\Psi = \mathcal{O}(\|\xi\|)$, equation (2.63) continues to hold, but with \hat{Y}_t^e replaced by \hat{Y}_t^n . Hence (2.41) again reduces to an aggregate-supply relation of the form (1.7), with the output gap x_t now defined by (2.70) and the cost-push term u_t again defined as in section 2.4.1.

In this case, as shown in section 1, the optimal long-run inflation rate is zero, and optimal policy is again characterized by a target criterion of the form (1.21), the coefficients of which do not depend on x^* . Hence the optimal target criterion is again given by (2.64), regardless of the (small) value of Ψ . The value of Ψ does matter, instead, for a calculation of the inflationary bias resulting from discretionary policy; it follows from our results in section 1 that the average inflation bias is (to first order) proportional to Ψ and with the same sign as Ψ .

2.4.3 The Case of Large Steady-State Distortions

If the degree of inefficiency of the zero-inflation steady-state level of output (measured by Ψ) is instead substantial, the analysis of the previous section cannot be used. In order to obtain a quadratic objective that can be evaluated to second order using a solution for the endogenous variables that need only be accurate to first order, it is necessary to replace the terms of the form $\bar{Y}U_Y\hat{Y}_t$ with purely quadratic functions of the endogenous variables (plus a residual that may include terms independent of

policy and/or terms of order $\mathcal{O}(\|\xi\|^3)$, using second-order approximations to one or more of the model's structural relations — just as was done above to eliminate the linear terms of the form $U_\Delta \hat{\Delta}_t$. This can in fact be done, not just in the present model but quite generally (as shown by Benigno and Woodford, 2008), by computing a second-order Taylor expansion of the Lagrangian (2.31) rather than of the expected utility of the representative household. Specifically, for an arbitrary evolution of the endogenous variables $\{Y_t, Z_t, \Delta_t\}$ in which these variables remain forever close enough to their steady-state values, we wish to compute a second-order approximation to (2.31) under the assumption that the Lagrange multipliers are at all times equal to their steady-state values $(\bar{\theta}, \bar{\Theta})$.

Note that in the case of any evolution of the endogenous variables that is consistent with the structural relations, the Lagrangian is equal to expected utility. Hence a quadratic approximation to the Lagrangian represents a quadratic function of the endogenous variables that will equal expected utility, up to an error of order $\mathcal{O}(\|\xi\|^3)$, in the case of any feasible policy. It is thus an equally suitable quadratic objective as is the one obtained above from a Taylor expansion of the expected utility objective; but it will have the advantage that there are no non-zero linear terms, precisely because (as discussed above) the zero-inflation steady state satisfies the steady-state version of the FOCs obtained by differentiating (2.31). This approach simultaneously solves the problem of the $U_\Delta \hat{\Delta}_t$ terms (in a way that is equivalent to the one used above) and the problem of the $\bar{Y} U_Y \hat{Y}_t$ terms.

Under this approach, the quadratic objective that we seek is an expected discounted sum of quadratic terms, where the contribution each period is given by the quadratic terms in a second-order Taylor series expansion of the function

$$L(Y_t, Z_t, \Delta_t, \Delta_{t-1}; \bar{\theta}, \bar{\Theta}, \bar{\Theta}; \xi_t).$$

Moreover, in the case that we assume that $\hat{\Delta}_{t-1} = \mathcal{O}(\|\xi\|^2)$, all of the quadratic terms that involve $\hat{\Delta}_t$ or $\hat{\Delta}_{t-1}$ will be of order $\mathcal{O}(\|\xi\|^3)$, and can thus be neglected. Hence in this case it suffices to compute the quadratic terms in a Taylor series expansion of the function

$$\hat{L}(Y_t, Z_t; \xi_t) \equiv L(Y_t, Z_t, 1, 1; \bar{\theta}, \bar{\Theta}, \bar{\Theta}; \xi_t),$$

where $L(\cdot)$ is the function defined in (2.32), and $(\bar{\theta}, \bar{\Theta})$ are the steady-state multipliers characterized in section 2.3 above.

Note that we can write

$$\hat{L}(Y, Z; \xi) = \hat{L}_1(Y; \xi) + \hat{L}_2(Z),$$

where

$$\begin{aligned}\hat{L}_1(Y; \xi) &\equiv U(Y, 1; \xi) + \bar{\Theta}'z(Y; \xi), \\ \hat{L}_2(Z) &\equiv \bar{\theta}[\tilde{h}(1, K/F) - 1] - \bar{\Theta}'Z + \alpha\bar{\Theta}'\Phi(Z).\end{aligned}$$

It follows from our previous definition (2.50) that for any vector of disturbances ξ_t , the function $\hat{L}_1(Y; \xi_t)$ has a critical point at $Y = Y_t^*$. Hence the quadratic terms in a Taylor expansion of $\hat{L}_1(Y_t; \xi_t)$ are equal to

$$\frac{1}{2}(\bar{Y})^2[U_{YY} + \bar{\Theta}'z_{YY}](\hat{Y}_t - \hat{Y}_t^*)^2.$$

The quadratic terms in a Taylor expansion of $\hat{L}_2(Z_t)$ are instead given by

$$\begin{aligned}&\frac{1}{2}\bar{\theta}\tilde{h}_{22}(\hat{K}_t - \hat{F}_t)^2 + \frac{\alpha}{2\bar{K}}\hat{Z}'_t M \hat{Z}_t \\ &= \frac{1}{2}\bar{\theta}\tilde{h}_{22}(\hat{K}_t - \hat{F}_t)^2 + \frac{\alpha m}{2\bar{K}}(\hat{K}_t - \hat{F}_t)^2,\end{aligned}$$

where \tilde{h}_{22} is the second partial derivative of \tilde{h} with respect to its second argument (evaluated at the steady-state values $(1, 1)$), M is the same matrix as in (2.43), and the second line uses (2.44). Using (2.40), the above expression can be written (up to an error of order $\mathcal{O}(\|\xi\|^3)$) as a negative multiple⁶⁹ of π_t^2 . Combining the results from our Taylor expansions of $\hat{L}_1(Y_t; \xi_t)$ and $\hat{L}_2(Z_t)$, we obtain a quadratic objective equal to a negative multiple of (1.6), where the output gap is now defined as $x_t \equiv \hat{Y}_t - \hat{Y}_t^*$, and $x^* = 0$.⁷⁰

Benigno and Woodford (2005a) show that the relative weight on the output-gap stabilization objective is given by

$$\lambda \equiv \frac{\kappa}{\bar{\theta}} \left\{ 1 - \frac{\Psi\sigma^{-1}(s_G/1 - s_G)}{(\omega + \sigma^{-1})[\omega + \sigma^{-1} + \Psi(1 - \sigma^{-1})]} \right\}, \quad (2.72)$$

⁶⁹See Benigno and Woodford (2005a) for demonstration that this coefficient is negative.

⁷⁰It may be wondered why $x^* = 0$, even though the steady-state level of output is inefficient. Since $Y_t^* = \bar{Y}$ in the zero-inflation steady state, there is no need to correct the definition of the output gap by a constant in order to obtain an output gap that is zero on average if an average inflation rate of zero is maintained, as in our treatment of the small- Ψ case. Because Y_t^* maximizes the Lagrangian (for a given vector ξ_t) rather than the utility function, the target level of output is lower, on average, than the efficient level of output Y_t^e .

where $s_G \equiv \bar{G}/\bar{Y}$ is the fraction of steady-state output that is consumed by the government. Note that this reduces to the same value (2.62) found above, in the case that either $\Psi = 0$ (the steady-state level of output is efficient) or $s_G = 0$ (no government purchases in steady state). Also, in the case that $\Psi = \mathcal{O}(\|\xi\|)$, λ is equal to the value given in (2.62) to first order (which is all that is relevant for a welfare ranking of equilibria that is accurate to second order), in accordance with our conclusions in section 2.4.2.

With this definition of the output gap, (2.41) again takes the form (2.53), where the cost-push term u_t is defined in (2.54). Thus both the quadratic loss function and the log-linear aggregate supply relation take the forms assumed in section 1. It follows that optimal policy is characterized by a target criterion of the form (1.21), where $x_t \equiv \hat{Y}_t - \hat{Y}_t^*$, and $\phi = \lambda/\kappa$, where κ is defined in (2.42) and λ is defined in (2.72). This is again the form of target criterion shown to characterize optimal policy in section 2.3 above. Note that $\phi = \theta^{-1}$, as concluded in section 2.4.1 above, only in the case that either $\Psi = 0$ (as assumed earlier) or $s_G = 0$. If $0 < \Psi < 1$ and $0 < s_G < 1$, then it follows from (2.72) that $\phi < \theta^{-1}$ in the optimal target criterion; and in fact, it is even possible for ϕ to be negative (as discussed in the next section), though this requires parameter values that are not too realistic.

2.5 Second-Order Conditions for Optimality

In the analysis above, it has been assumed that a solution to the FOCs corresponds to an optimal evolution for the economy. In order for such a solution to represent an optimum, it must maximize the Lagrangian, which requires that the Lagrangian be locally concave — more precisely, it must be locally concave on the set of paths for the endogenous variables that are consistent with the structural equations (though not necessarily concave outside this set). This can be checked using a second-order Taylor series expansion of the Lagrangian, which involves precisely the same coefficients as already appear in our linear approximation to the FOCs.⁷¹

Benigno and Woodford (2005a) show that for the model considered here, the Lagrangian (2.31) is locally concave (on the set of paths consistent with the model

⁷¹Algebraic conditions for local concavity of the Lagrangian for a more general class of optimal policy problems, to which the problem considered here belongs, are presented in Benigno and Woodford (2008) and in Giannoni and Woodford (2010).

structural relations) near the optimal steady state if and only if the model parameters are such that

$$\lambda > -\frac{\kappa^2}{(1 + \beta^{1/2})^2}, \quad (2.73)$$

where λ is defined by (2.72). In the case that $\lambda > 0$, the quadratic approximation to the Lagrangian derived in the previous section is obviously concave (up to a constant, it is a negative multiple of a function that is convex because it is a sum of squares); but even if $\lambda < 0$, the Lagrangian continues to be concave as long as λ is not *too large* a negative quantity. This is because it is not possible to vary the path of the output gap without also varying the path of inflation (if we consider only paths consistent with the aggregate-supply relation); if λ is only modestly negative, the convexity of the loss function (1.6) in inflation will suffice to insure that the entire function is convex (so that the quadratic approximation to the Lagrangian is concave, and the Lagrangian itself is locally concave) on the set of paths consistent with the aggregate-supply relation.

Since (2.72) implies that λ is positive unless *both* Ψ and s_G are substantial fractions of 1, and the second-order condition (2.73) is not violated unless λ is sufficiently negative, the Lagrangian will be at least locally concave, except in the case of relatively extreme parameter values. (For example, as long as $s_G < 1/2$, one can show that $\lambda > 0$, for any values of the other parameters.) Hence failure of the second-order conditions is unlikely to arise in the case of an empirically realistic calibration of this particular model. Nonetheless, it is worth noting that such a failure can occur under parameter values consistent with our general assumptions.⁷²

When (2.73) is not satisfied, the solution to the FOCs is *not* actually the optimal equilibrium evolution; for example, the steady state is not the optimal equilibrium, even in the absence of stochastic disturbances. It is not possible using the purely local methods illustrated here to say what the optimal equilibrium is like; but local analysis suffices to show, for example, that arbitrary randomization (of small enough amplitude) of the paths of output and inflation can be introduced in a way that *increases* expected utility, as first shown by Dupor (2003) in a simpler New Keynesian model. This occurs because under certain circumstances, firms facing a more uncertain demand for their products will prefer to set prices that are lower, relative to

⁷²Fixing the values of $\alpha, \beta, \omega, \sigma, \theta$ and any value $0 < \Psi < 1$, one can show that any value of s_G close enough to 1 will imply a value of λ sufficiently negative to violate (2.73).

their expected marginal cost of supplying their customers, than they would if their sales were more predictable.⁷³ This in turn leads to a higher average level of output in equilibrium; and in the case that the steady state level of output is sufficiently inefficient (Ψ is sufficiently large), increasing the average level of output can matter more for welfare than the losses resulting from more variable output and greater price dispersion. Nonetheless, while technically possible, this case seems unlikely to be of practical relevance.

2.6 When is Price Stability Optimal?

In general, the model presented in section 1 implies that there is a tradeoff between inflation stabilization and output-gap stabilization, as a consequence of which an optimal policy will not completely stabilize the rate of inflation; instead, modest (and relatively transitory) variations in the rate of inflation should be accepted for the sake of increased stability of the output gap. However, the degree to which actual economic disturbances give rise to a tension between the goals of inflation stabilization and output-gap stabilization depends on the nature of the disturbance. In the notation used in section 1 (following Clarida *et al.*, 1999), an exogenous disturbance should be allowed to affect the rate of inflation only to the extent that it represents a “cost-push disturbance” of the kind denoted by the term u_t in (1.7). It is therefore of some importance to have a theory of the extent to which actual disturbances should affect the value of this term.

In the case of our analysis under the assumption that the steady-state level of output is efficient, we were able to obtain a strong conclusion: the term u_t in (1.7) is a positive multiple of $\hat{\mu}_t^w + \hat{\tau}_t$. In the absence of fluctuations in either the wage markup or the tax rate, zero inflation is optimal, as a policy that achieves zero inflation at all times would also achieve $\hat{Y}_t = \hat{Y}_t^e$ at all times, and hence a zero output gap (in the welfare-relevant sense) at all times. This is only exactly true, however, on the

⁷³This requires that firms care more about selling too little in low-demand states than about selling too much in high-demand states. This will be the case if s_G is sufficiently close to 1, since in this case, there will be a large elasticity of private consumption $Y_t - G_t$ with respect to variations in aggregate demand Y_t , and as a consequence a large elasticity of the representative household’s marginal utility of income with respect to variations in aggregate demand. The fact that the firm’s shareholders value additional income so much in the lowest-demand states motivates firms to be take care not to set prices that are too high relative to wages and other prices in the economy.

assumption that $\Psi = 0$. In the more realistic case where we assume $\Psi > 0$, most real disturbances have some non-zero “cost-push” effect, as shown by Benigno and Woodford (2005a).

Even when $\Psi > 0$, there is a special case in which complete price stability continues to be optimal. Suppose that there are no government purchases ($\bar{G} = 0$, in addition to no variation in government purchases), and that the distortion factors μ_t^w and τ_t remain forever at their steady-state values. In this case,⁷⁴

$$f_Y(Y; \xi) = (1 - \sigma^{-1})(1 - \bar{\tau})u_y(Y; \xi), \quad k_Y(Y; \xi) = (1 + \omega)\mu^p \bar{\mu}^w v_y(Y; \xi),$$

and condition (2.50) reduces to

$$[1 + \bar{\Theta}_1(1 - \sigma^{-1})(1 - \bar{\tau})]u_y(Y_t^*; \xi_t) = [1 + \bar{\Theta}_1(1 + \omega)\mu^p \bar{\mu}^w]v_y(Y_t^*; \xi_t). \quad (2.74)$$

Then since

$$\bar{\Theta}_1 = \frac{\Psi}{(\sigma^{-1} + \omega)(1 - \bar{\tau})} = -\bar{\Theta}_2,$$

the condition defining Y_t^* simplifies to

$$(1 - \Psi)u_c(Y_t^*; \xi_t) = v_y(Y_t^*; \xi_t). \quad (2.75)$$

A comparison of this equation with (2.69) indicates that $Y_t^* = Y_t^n$. (Both Y_t^* and Y_t^n move exactly in proportion to the variations in Y_t^e , and both are smaller than Y_t^e by precisely the same percentage.) It follows that $u_t = 0$, and again complete price stability will be optimal. This explains the numerical results of Khan *et al.* (2003), according to which it is optimal to use monetary policy to prevent technology shocks from having any effect on the path of the price level.⁷⁵

However, once we allow for non-zero government purchases, $f_Y(Y; \xi)$ is no longer a constant multiple of $u_y(Y; \xi)$. We must then write condition (2.50) as

$$u_y(Y_t^*; \xi_t) + \bar{\Theta}_1 f_Y(Y_t^*; \xi_t) = [1 + \bar{\Theta}_1(1 + \omega)\mu^p \bar{\mu}^w]v_y(Y_t^*; \xi_t), \quad (2.76)$$

⁷⁴Note that this derivation relies on the special isoelastic functional forms assumed for the utility and production functions. More generally, the equivalence between Y_t^* and Y_t^n derived here will not hold when $\Psi > 0$, even if $G_t = 0$ at all times, and all of the real disturbances will have “cost-push” effects, making strict price stability sub-optimal.

⁷⁵The result is derived here under the assumption of Calvo-style staggered price adjustment, but can be shown to hold under more general assumptions about the way in which the probability of price review varies with the duration since the last review (Benigno and Woodford, 2004), of the kind made by Khan *et al.* (2003). See also section 3.1.1 below for discussion of optimal policy when the Calvo assumption is relaxed.

which no longer reduces to (2.74) and hence to (2.75). An exogenous increase in G_t raises f_Y by a smaller proportion than the increase in u_y . As a consequence, Y_t^* increases less with government purchases than does Y_t^n , and $u_t \equiv \kappa(\hat{Y}_t^* - \hat{Y}_t^n)$ falls when government purchases rise. Government purchases have a negative (favorable) “cost-push” effect in this case, as a result of which it is optimal for inflation to be *reduced* slightly in response to the shock, in order to keep output from expanding as much as it would under a policy consistent with price stability. (This again explains the numerical results of Khan *et al.*, (2003).)

If we start from a steady state in which $\bar{G} > 0$ (certainly the more realistic case), then other real disturbances have “cost-push” effects as well. If $\bar{G} > 0$, an increase in Y_t does not reduce f_Y in proportion to the decline in u_y . Hence the left-hand side of (2.76) is not as sharply decreasing a function of Y as is the left-hand side of (2.74). It follows that in the case of a productivity disturbance A_t , which shifts $v_y(Y)$ without affecting $u_y(Y)$ or $f_Y(Y)$, the solution to (2.76), which is to say Y_t^* , shifts by more than does the solution to (2.74), which is equal to Y_t^n as explained above. Thus a positive technology shock increases Y_t^* more than it does Y_t^n , and consequently such a shock has a positive “cost-push” effect, making a transitory increase in inflation in response to the shock desirable. A similar conclusion holds in the case of shocks to the preference factors \bar{H}_t or \bar{C}_t . Of course, these effects can only be substantial, as a quantitative matter, in the case that government purchases are a substantial share of output (so that the elasticities of f_Y and u_y differ substantially) *and* steady-state distortions are substantial (so that $\bar{\Theta}_1$ is significantly non-zero and hence the difference between the left-hand sides of (2.74) and (2.76) is non-trivial).

Of course, there are many other reasons why complete stabilization of a general price index is unlikely to represent an optimal policy, that have been abstracted from in the model treated in this section. In particular, once we allow wages as well as prices to be sticky, or allow for asymmetries between different sectors, it is unlikely to be optimal to fully stabilize an index that involves only prices (rather than wages) and that weights prices in all sectors equally. Some of the consequences of complications of these kinds are discussed in section 3.

3 Generalizations of the Basic Model

There are many special features of the basic New Keynesian model, used in the previous section to illustrate some basic methods of analysis and to introduce certain themes of broader importance. This section considers the extent to which the specific results obtained for the basic model extend to more general classes of models.

3.1 Alternative Models of Price Adjustment

Among the special features of the model treated in section 2 is the Calvo-Yun model of staggered price adjustment. There are two aspects of this model that one may wish to generalize: first, it assumes that the probability of a firm's reconsidering its pricing strategy in any period is independent of the time since the pricing of that good was last reviewed; and second, it assumes that each supplier charges a fixed nominal price between those occasions on which the pricing policy is reviewed, rather than choosing some more complex strategy (such as a non-constant price path, or an indexation rule) that is periodically revised in the light of market conditions. Here I review available results on the consequences of relaxing both of these assumptions.

I shall focus on the question of how a *welfare-based* analysis of optimal policy changes as a result of an alternative specification of the mechanism of price adjustment, continuing to assume that the goal of policy is to maximize the expected utility of the representative household, rather than assuming an *ad hoc* stabilization objective such as (1.7) that stays the same when one changes one's model of the Phillips-curve tradeoff.⁷⁶ One reason for taking up the topic is precisely to show that the welfare-theoretic justification for a loss function of the form (1.6) provided in section 2 does *not* equally extend to variant aggregate-supply specifications that may be more empirically realistic. But another important theme of this section is that some conclusions about the character of optimal policy *can* be robust to changes in the specification of the dynamics of price adjustment. In several cases discussed below, the form of the *optimal target criterion* — including the precise numerical coefficients involved as well as the relevant definition of the “output gap” — remains invariant under changes in the parameterization of the model of price adjustment.

⁷⁶Exercises of the latter kind are fairly common, especially in work at policy institutions, but do not raise new issues of method, and I shall not attempt to survey the various results that may be obtained from varying combinations of the possible specifications of objectives and constraints.

This provides a further argument for the desirability of formulating a central bank's policy commitment in terms of a target criterion, rather than through some other possible description of intended future policy.

In the present section, I simplify the analysis by considering only the case in which the steady-state level of output is efficient, as in section 2.4.1., and in which $\hat{\Delta}_{t_0-1} = \mathcal{O}(\|\xi\|^2)$, so that (2.58) holds. If one further assumes that each relative price $\log(p_t(i)/P_t)$ is of order $\mathcal{O}(\|\xi\|)$,⁷⁷ then we can use the approximation

$$\hat{\Delta}_t = \frac{1}{2}\theta(1 + \omega)(1 + \omega\theta)\text{var}_i \log p_t(i) + \mathcal{O}(\|\xi\|^3)$$

to show that

$$U(Y_t, \Delta_t; \xi_t) = \frac{1}{2}(1 + \omega\theta)\bar{Y}u_y[\zeta x_t^2 + \theta\text{var}_i \log p_t(i)] \\ + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3),$$

where $x_t \equiv \hat{Y}_t - \hat{Y}_t^e$ and

$$\zeta \equiv \frac{\sigma^{-1} + \omega}{1 + \omega\theta} > 0$$

is a measure of the degree of “real rigidities.” Note that the price p_t^{opt} implicitly defined by the relation⁷⁸

$$\Pi_1(p_t^{opt}, p_t^{opt}, P_t; Y_t, \xi_t) = 0$$

or alternatively

$$\frac{p_t^{opt}}{P_t} = \left(\frac{k(Y_t; \xi_t)}{f(Y_t; \xi_t)} \right)^{\frac{1}{1+\omega\theta}}$$

can be approximated to first order by

$$\log p_t^{opt} = p_t + \zeta x_t + \hat{\mu}_t, \tag{3.1}$$

where

$$\hat{\mu}_t \equiv (1 + \omega\theta)^{-1}(\hat{\mu}_t^w + \hat{\tau}_t)$$

⁷⁷If we consider only policies under which $\pi_t = \mathcal{O}(\|\xi\|)$ for all t , relative prices will indeed be of order $\mathcal{O}(\|\xi\|)$ for all t , assuming that one starts from an initial condition in which this is true. In each of the exercises considered below, the optimal steady-state inflation rate continues to be zero, so that under the optimal policy, $\pi_t = \mathcal{O}(\|\xi\|)$ for all t .

⁷⁸This relation defines the industry equilibrium price at time t for a (hypothetical) sector with perfectly flexible prices. See Woodford (2003, chap. 3) for further discussion of the interpretation and significance of the parameter ζ . Note that it is a feature of the economy that is independent of any assumptions about the degree or nature of nominal rigidities.

is a composite distortion factor. This explains the significance of the coefficient ζ .

It follows that welfare is maximized (to a second-order approximation) by minimizing a quadratic loss function of the form

$$E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} [\zeta x_t^2 + \theta \text{var}_i \log p_t(i)]. \quad (3.2)$$

In the case of Calvo pricing, this loss function is proportional to (1.6); but under alternative assumptions about the dynamics of price adjustment, the connection between price dispersion and inflation is different, and so the way in which the welfare-based loss function will depend on inflation is different.

3.1.1 Structural Inflation Inertia

The Calvo-Yun model of price adjustment makes the model dynamics in section 2 highly tractable, but has some implications that are arguably unappealing. In particular, it results in a log-linear aggregate supply relation (2.41) that is purely forward-looking: neither past inflation nor past real activity have any consequences for the inflation-output tradeoff that exists at a given point in time. Empirical aggregate supply relations often instead involve some degree of *structural inflation inertia*, in the sense that a higher level of inflation in the recent past makes the inflation rate associated with a given path for real activity from now on higher.⁷⁹

In fact, as Wolman (1999), Dotsey (2002) and Sheedy (2007) note, a model of optimal price-setting of the kind considered above can imply inflation inertia, if one abandons the Calvo assumption of duration-independence of the probability of price review. If, as is arguably more plausible,⁸⁰ one instead assumes that prices are more likely to be reviewed the older they are, then when inflation has been higher than average in the recent past, old prices will be especially low relative to prices on average, and as a consequence the average percentage increase in the prices that are adjusted will be greater. This mechanism makes the overall rate of inflation higher when past inflation has been higher, for any given assumption about where newly revised prices will be set relative to the average level of current prices (which depends

⁷⁹See Fuhrer (2010) for a review of the literature on this issue.

⁸⁰Wolman (1999) argues for this kind of model as an approximation to the dynamics implied by a state-dependent pricing model of the kind analyzed by Dotsey *et al.* (1999).

on real marginal costs — and hence on the output gap — and on expected inflation from now on).

As an example (taken from Sheedy, 2007) in which the state space required to describe aggregate dynamics remains relatively small, consider a generalization of the Calvo model in which at each point in time, the fraction θ_j of all prices that were chosen j periods ago is given by

$$\theta_j = \frac{(1 - \alpha_1)(1 - \alpha_2)}{\alpha_1 - \alpha_2} [\alpha_1^{j+1} - \alpha_2^{j+1}] \quad (3.3)$$

for any integer $j \geq 0$, where

$$0 \leq \alpha_2 < \min(\alpha_1, 1 - \alpha_1) < 1.$$

Conditional on a price having already been charged for j periods, the probability that it will continue to be charged for another period, θ_j/θ_{j-1} , is less than 1, and non-increasing in j . The Calvo case is nested within this family as the case in which $\alpha_2 = 0$, in which case the probability of non-review each period, $\theta_j/\theta_{j-1} = \alpha_1$, is independent of j . When $\alpha_2 > 0$, instead, the probability of a price review each period is an increasing function of j . As in the model of section 2, let us suppose that the same price is charged until the random date at which the price of that good is again reviewed.

If we continue to maintain all of the other assumptions of section 2, each firm that reviews its price in period t faces the same optimization problem and chooses the same price p_t^* . The optimal choice is again given by (2.13), where we now define

$$F_t \equiv E_t \sum_{T=t}^{\infty} \beta^{T-t} \theta_{T-t} f(Y_T; \xi_T) \left(\frac{P_T}{P_t} \right)^{\theta-1},$$

$$K_t \equiv E_t \sum_{T=t}^{\infty} \beta^{T-t} \theta_{T-t} k(Y_T; \xi_T) \left(\frac{P_T}{P_t} \right)^{\theta(1+\omega)},$$

generalizing the previous forms (2.14)–(2.15). Log-linearizing this relation around the zero-inflation steady state (that continues to be the optimal steady state, regardless of the values of α_1, α_2), we obtain

$$\log p_t^* = \sum_{j=0}^{\infty} \omega_j E_t [\log p_{t+j}^{opt}], \quad (3.4)$$

where

$$\omega_j \equiv \frac{\beta^j \theta_j}{\sum_{i=0}^{\infty} \beta^i \theta_i}$$

for each $j \geq 0$. The log of the Dixit-Stiglitz price index is in turn given (to first order) by

$$p_t = \sum_{j=0}^{\infty} \theta_j \log p_{t-j}^*. \quad (3.5)$$

When the sequence $\{\theta_j\}$ is given by (3.3), (3.5) implies that the price index must satisfy a difference equation of the form

$$(1 - \alpha_1 L)(1 - \alpha_2 L)p_t = (1 - \alpha_1)(1 - \alpha_2) \log p_t^*, \quad (3.6)$$

and (3.4) implies that $\{p_t^*\}$ must satisfy the expectational difference equation

$$E_t[(1 - \alpha_1 \beta L^{-1})(1 - \alpha_2 \beta L^{-1}) \log p_t^*] = (1 - \alpha_1 \beta)(1 - \alpha_2 \beta) \log p_t^{opt}. \quad (3.7)$$

Substituting (3.6) for $\log p_t^*$ and (3.1) for $\log p_t^{opt}$ in (3.7), one can show that the inflation rate must satisfy an aggregate-supply relation of the form

$$\pi_t - \gamma_{-1} \pi_{t-1} = \kappa x_t + \gamma_1 E_t \pi_{t+1} + \gamma_2 E_t \pi_{t+2} + u_t, \quad (3.8)$$

where the exogenous disturbance u_t is a positive multiple of $\hat{\mu}_t$, and the coefficients satisfy $\kappa > 0, \gamma_{-1} + \gamma_1 + \gamma_2 = \beta$ as in the model of section 2, but now

$$\gamma_{-1} = \frac{\alpha_1 \alpha_2}{(1 + \alpha_1 \alpha_2 \beta)(\alpha_1 + \alpha_2) - \alpha_1 \alpha_2} \geq 0$$

is a positive coefficient (indicating structural inflation inertia) if $\alpha_2 > 0$ (so that the probability of price adjustment is increasing in duration).⁸¹

Let us now consider the consequences of this generalization for optimal policy. Modification of the model of price adjustment implies that the welfare-based stabilization objective also changes, if written as a function of the evolution of the general price index (rather than in terms of the dispersion of individual prices). The quadratic approximation $\delta_t \equiv \text{var}_i \log p_t(i)$ to the index of price dispersion evolves according to a law of motion

$$(1 - \alpha_1 L)(1 - \alpha_2 L)\delta_t = (1 - \alpha_1)(1 - \alpha_2)(\log p_t^*)^2 - (1 - \alpha_1 L)(1 - \alpha_2 L) p_t^2$$

⁸¹Sheedy (2007) finds that estimation of the model using U.S. data yields a significantly positive coefficient.

as a consequence of our assumption about the probability of revision of prices of differing durations.

Multiplying by β^{t-t_0} and summing, one finds that

$$\sum_{t=t_0}^{\infty} \beta^{t-t_0} \delta_t = \Gamma \sum_{t=t_0}^{\infty} \beta^{t-t_0} (\log p_t^*)^2 - \sum_{t=t_0}^{\infty} \beta^{t-t_0} p_t^2 + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3), \quad (3.9)$$

where

$$\Gamma \equiv \frac{(1 - \alpha_1)(1 - \alpha_2)}{(1 - \alpha_1\beta)(1 - \alpha_2\beta)}.$$

(Note that $0 < \Gamma < 1$.) This result can be used to substitute for the discounted sum of δ_t terms in (3.2), yielding a quadratic objective of the form

$$E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} [\zeta x_t^2 + \theta \Gamma (\log p_t^*)^2 - \theta p_t^2]. \quad (3.10)$$

This involves only the paths of the variables $\{x_t, p_t^*, p_t\}$. Moreover, the evolution of $\{p_t^*\}$ depends purely on the paths of the variables $\{x_t, p_t\}$ and the exogenous disturbances, because of equations (3.1) and (3.4). Thus the stabilization objective can again be expressed as a quadratic function of the paths of the variables $\{x_t, p_t\}$ and no other endogenous variables.

If we write a Lagrangian for the problem of minimizing (3.10) subject to the constraints (3.6) and (3.7), we obtain a system of FOCs

$$x_t - (1 - \alpha_1\beta)(1 - \alpha_2\beta)\varphi_t = 0, \quad (3.11)$$

$$\theta p_t + (1 - \alpha_1\beta)(1 - \alpha_2\beta)\varphi_t + E_t[\gamma(\beta L^{-1})\psi_t] = 0, \quad (3.12)$$

$$\theta \Gamma \log p_t^* + (1 - \alpha_1)(1 - \alpha_2)\psi_t + \gamma(L)\varphi_t = 0, \quad (3.13)$$

where ψ_t is a Lagrange multiplier associated with constraint (3.6), φ_t is a multiplier associated with constraint (3.7), and

$$\gamma(L) \equiv (1 - \alpha_1 L)(1 - \alpha_2 L).$$

Each of constraints (3.11)–(3.13) must hold for each $t \geq t_0$, if we adjoin to this system the initial conditions

$$\varphi_{t_0-1} = \varphi_{t_0-2} = 0. \quad (3.14)$$

If we solve (3.13) for ψ_t , and use (3.7) to substitute for p_t^* in this expression, we obtain

$$\psi_t = -\frac{\theta\gamma(L)}{\gamma(1)\gamma(\beta)}\hat{p}_t,$$

where

$$\hat{p}_t \equiv p_t - \theta^{-1}\gamma(1)\varphi_t.$$

Using this to substitute for ψ_t in (3.12), we obtain

$$E_t[A(L)\hat{p}_{t+2}] = 0, \tag{3.15}$$

where $A(L)$ is a quartic polynomial. Because the factors of $A(L)$ are of the form

$$A(L) = (1 - L)(1 - \beta^{-1}L)(1 - \lambda_1L)(1 - \lambda_2L),$$

where $0 < \lambda_1 < 1 < \lambda_2$, it follows that in any non-explosive solution (for example, any solution in which the inflation rate and output gap are forever bounded), we must have

$$(1 - \lambda_1L)(1 - L)\hat{p}_t = 0$$

for each $t \geq t_0$. Because this equation is purely backward-looking, the path $\{\hat{p}_t\}$ is uniquely determined by the initial conditions (p_{t_0-1}, p_{t_0-2}) and conditions (3.14). Note that $\{\hat{p}_t\}$ is a deterministic sequence that converges asymptotically to some constant value p^* (a homogeneous degree 1 function of (p_{t_0-1}, p_{t_0-2})).

Finally, (3.11) implies that $\tilde{p}_t = \hat{p}_t$ for all $t \geq t_0$, where \tilde{p}_t is again the output-gap-adjusted price level defined in (2.65). Hence we have also solved for a uniquely determined optimal path $\{\tilde{p}_t\}$, with the properties just discussed. Thus, as in the basic New Keynesian model, it is possible to verify whether the economy's projected evolution is consistent with the optimal equilibrium simply by verifying that the projected paths for $\{p_t\}$ and $\{x_t\}$ satisfy a certain linear relationship. Moreover, optimal policy again requires that the path of the output-gap-adjusted price level be *completely unaffected* by any random shocks that occur from date t_0 onward. As an example, Figure 4 shows the impulse responses of inflation, output and the price level to a transitory positive “cost-push” shock under an optimal policy commitment, using again the format of Figure 1. However, whereas Figure 1 corresponds to a




Figure 4: Impulse responses to a transitory cost-push shock under an optimal policy commitment, in the case of structural inflation inertia.

parameterization in which $\alpha_1 = 0.66$, $\alpha_2 = 0$, it is now assumed that $\alpha_1 = \alpha_2 = 0.5$.⁸² As a consequence, $\gamma_{-1} = 0.251$, and there is structural inflation inertia. Even though the shock has a “cost-push” effect in period zero only, optimal policy now allows the price level to continue to increase (by a small amount) in period one as well; owing to the structural inflation inertia, it would be too costly to reduce the inflation rate more suddenly. Nonetheless, under the optimal policy commitment, the impulse response of the price level is the mirror image of the impulse response of the output gap, just as in Figure 1. The difference is simply that it is necessary for the deviations of both

⁸²As in Figure 1, the shock is a one-period increase in $\hat{\mu}_0$ by 5.61 (which under the parameterization used in Figure 1 corresponds to a cost-push shock that would raise the log price level by 1 in the absence of any change in the output gap or in expected inflation). Also as in Figure 1, it is assumed that $\beta = 0.99$ and $\zeta = 0.134$. The value of θ used here is 6, a slightly smaller value than the one (taken from Rotemberg and Woodford, 1997) used in Figure 1. This is because the lower value of θ increases the degree of structural inflation inertia, making more visible the contrast between the two figures.

the price level and of the output gap from their long-run values to be more persistent when inflation is inertial.

One difference from the results obtained earlier is that in the basic New Keynesian model, optimal policy required that the target for \tilde{p}_t be the same for all $t \geq t_0$, while in the more general case, the targets for \tilde{p}_t form a deterministic sequence, but equal the constant p^* only asymptotically. When $\alpha_2 > 0$, the optimal sequence $\{\tilde{p}_t\}$ is monotonically increasing if $\pi_{t_0-1} > 0$ and monotonically decreasing if $\pi_{t_0-1} < 0$; and both the initial rate of increase of \tilde{p}_t and the cumulative increase in \tilde{p}_t over the long run should be proportional to the initial inflation rate π_{t_0-1} . Thus in the presence of structural inflation inertia, the fact that the economy starts out from a positive inflation rate has implications for the rate of inflation that should initially be targeted after the optimal policy is adopted. However, shocks that occur after the adoption of the optimal policy commitment should never be allowed to alter the targeted path for the output-gap-adjusted price level, which should eventually be held constant regardless of the disturbances to which the economy may recently have been subject.

Sheedy (2008) shows that this result holds not only for the particular parametric family of sequences $\{\theta_j\}$ defined in (3.3), but for any sequence $\{\theta_j\}$ that corresponds to the solution to a linear difference equation of arbitrary (finite) order. In any such case, under an optimal policy commitment, the path of the output-gap-adjusted price level must evolve deterministically, and converge asymptotically to a constant. Moreover, it must satisfy a time-invariant target criterion of the form

$$\delta(L)\tilde{p}_t = p^* \tag{3.16}$$

for all t after some date T , where $\delta(L)$ is a finite-order lag polynomial all of the roots of which lie outside the unit circle (so that the sequence $\{\tilde{p}_t\}$ must converge).⁸³ Hence the result that optimal policy requires the central bank to target a deterministic path for the output-gap adjusted price level is independent of the assumed duration-dependence of the probability of price review; the definition of the output-gap-adjusted price level in this target criterion is also independent of those details. This provides a further example of how the description of optimal policy in terms of a target criterion that must be fulfilled is more robust to alternative model parameterizations than other levels of description, that would be equivalent in the context of a particular quantitative specification.

⁸³In the class of examples defined by (3.3), $\delta(L) \equiv (1 - \lambda_1 L)$, and (3.16) must hold for all $t \geq t_0 + 1$.

3.1.2 Sticky Information

Another model of price adjustment treats the delays in adjustment of prices to current market conditions as resulting from infrequent updating of price-setters' information, rather than from any delays in adjustment of prices to what price-setters currently understand to be optimal, owing to costs of changing prices from what they have been in the past. In the “sticky information” model of Mankiw and Reis (2002), price-setters update their information sets only at certain dates, though they obtain full information about the current state of the world each time that they obtain any new information at all; and they continually adjust the price that they charge, to reflect what they believe to be the currently optimal price, on the basis of the most recent information available.

Suppose, following Mankiw and Reis, that the probability that a firm (the monopoly producer of a single differentiated product) updates its information is a function only of the time that has elapsed since it has last done so.⁸⁴ For any integer $j \geq 0$, let θ_j be the fraction of the population of firms at any date that have last updated their information j periods earlier, where $\{\theta_j\}$ is a non-increasing sequence of non-negative quantities that sum to 1, as in the previous section. In Mankiw and Reis (2002) and in Ball *et al.* (2005), it is assumed that the probability of updating is independent of the time since the firm last updated its information, so that $\theta_j = (1 - \alpha)\alpha^j$ for some $0 < \alpha < 1$; but here I shall allow for duration-dependence of a fairly general sort. I shall let J denote the largest integer j such that $\theta_j > 0$; this may be infinite (as in the case assumed by Mankiw and Reis, 2002), but I shall also allow for the possibility that there is a finite maximum duration between dates at which information is updated (as, for example, in Koenig, 2004).

In any period t , a firm that last updated its information in period $t - j$ chooses its price p to maximize $E_{t-j}\Pi(p, p_t^j, P_t; Y_t, \xi_t)$. Let the solution to this problem be denoted $p_{t,t-j}^*$; to a log-linear approximation, it is given by

$$\log p_{t,t-j}^* = E_{t-j} \log p_t^{opt}, \quad (3.17)$$

where p_t^{opt} will again be given (to a log-linear approximation) by (3.1). To a similar

⁸⁴This probability is taken here, as in the original paper of Mankiw and Reis (2002), to be exogenously given, and to be the same for all firms. Reis (2006) considers instead the endogenous determination of the time interval between information acquisitions. The consequences of such endogeneity for optimal policy have not yet been addressed.

log-linear approximation, the log general price index p_t will be given by

$$p_t = \sum_{j=0}^{\infty} \theta_j \log p_{t,t-j}^*. \quad (3.18)$$

Combining these equations, we find that the model implies an aggregate-supply relation of the form

$$\sum_{j=0}^{\infty} \theta_j [p_t - E_{t-j} p_t] = \sum_{j=0}^{\infty} \theta_j E_{t-j} [\zeta x_t + \hat{\mu}_t].$$

This is a form of expectations-augmented Phillips curve, which provides a possible explanation for apparently inertial inflation (even if it is not true structural inflation inertia): a higher than usual inflation rate is associated with a given output gap if the inflation rate was *expected* at some past date to be higher than usual, which tends to have been the case when actual inflation was higher than usual at that past date.

This model of pricing similarly implies that

$$\begin{aligned} \text{var}_i \log p_t(i) &= \sum_{j=0}^{\infty} \theta_j (\log p_{t,t-j}^*)^2 - \left(\sum_{j=0}^{\infty} \theta_j \log p_{t,t-j}^* \right)^2 \\ &= \sum_{j=0}^{\infty} \theta_j (\log p_{t,t-j}^*)^2 - p_t^2. \end{aligned} \quad (3.19)$$

Our problem is then to find a state-contingent evolution for the variables $\{p_t, x_t, p_{t,t-j}^*\}$ for all $t \geq t_0$ (and for each such t , all $0 \leq j \leq t - t_0$) so as to minimize (3.2), in which we substitute (3.19) for the price dispersion terms, subject to the constraints that (3.17) must hold for each $(t, t - j)$ and (3.18) must hold for each t .

We can write a Lagrangian for this problem of the form

$$\begin{aligned} \mathcal{L}_{t_0} &= E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \left\{ \frac{\zeta}{2} x_t^2 + \frac{\theta}{2} \left[\sum_{j=0}^{\infty} \theta_j (\log p_{t,t-j}^*)^2 - p_t^2 \right] \right. \\ &\quad + \sum_{j=0}^{t-t_0} \psi_{t,t-j} [\log p_{t,t-j}^* - p_t - \zeta x_t - \hat{\mu}_t] \\ &\quad \left. + \varphi_t \left[p_t - \sum_{j=0}^{\infty} \theta_j \log p_{t,t-j}^* \right] \right\}, \end{aligned} \quad (3.20)$$

where $\psi_{t,t-j}$ is a Lagrange multiplier associated with constraint (3.17) and φ_t is a multiplier associated with constraint (3.18). Differentiation yields FOCs

$$\theta_j \theta \log p_{t,t-j}^* + \psi_{t,t-j} - \theta_j \varphi_t = 0 \quad (3.21)$$

for each $0 \leq j \leq t - t_0$ and each $t \geq t_0$; and

$$\theta p_t + \sum_{j=0}^{t-t_0} \psi_{t,t-j} - \varphi_t = 0, \quad (3.22)$$

$$\zeta x_t - \zeta \sum_{j=0}^{t-t_0} \psi_{t,t-j} = 0 \quad (3.23)$$

for all $t \geq t_0$.

Since $\psi_{t,t-j}$ must be measurable with respect to period $t - j$ information (note that there is only one constraint (3.17) for each possible state of the world at date $t - j$, and not a separate constraint for each state that may be reached at date t), condition (3.21) implies that φ_t must be measurable with respect to period $t - j$ information, so that

$$\varphi_t = E_{t-j} \varphi_t \quad (3.24)$$

for all j such that $\theta_j > 0$, which is to say, for all $j \leq \min J, t - t_0$.

Solving (3.21) for $\psi_{t,t-j}$ and substituting for these multipliers in (3.22) and (3.23) yields

$$\sigma_{t-t_0} \varphi_t = \theta \hat{p}_{t,t_0-1}, \quad (3.25)$$

$$x_t = (1 - \sigma_{t-t_0}) \varphi_t + \theta [\hat{p}_{t,t_0-1} - p_t], \quad (3.26)$$

where

$$\sigma_j \equiv \sum_{i>j} \theta_i$$

is the probability of a firm's having information more than j periods old, and

$$\hat{p}_{t,t_0-1} \equiv \sum_{i=1}^{\infty} \theta_{t-t_0+i} \log p_{t,t_0-i}^*$$

is the contribution to p_t from prices set on the basis of information dating from prior to date t_0 .

In the case of any $t < t_0 + J$, one has $\sigma_{t-t_0} > 0$, and (3.25) can be solved for φ_t . Substituting this value into (3.26), we find that the FOCs require that

$$\tilde{p}_t = \sigma_{t-t_0}^{-1} \hat{p}_{t,t_0-1} \quad (3.27)$$

for all $t < t_0 + J$. If $J = \infty$, as in the case assumed by Mankiw and Reis (2002) and Ball *et al.* (2005), we once again find that optimal policy requires that the output-gap-adjusted price level follow a deterministic path for all $t \geq t_0$. (This path is not, however, necessarily constant even in the long run: if prior to date t_0 , the public has expected prices to increase over the long run at a steady rate of two percent per year, then under the optimal Ramsey policy from date t_0 onward, \tilde{p}_t should increase at two percent per year in the long run.)

This result differs from the one obtained by Ball *et al.*, because they assume that monetary policy must be determined on the basis of the economy's state in the previous period, rather than its current state. Under this information constraint on the central bank, they find that optimal policy requires that $E_{t-1}\tilde{p}_t$ must evolve according to a deterministic path, though shocks in period t can still cause surprise variations in \tilde{p}_t relative to what was expected a period in advance. In a sense, this result provides an even stronger argument for price-level targeting: for while the optimal target criterion for the full-information case can equivalently be stated as a requirement that $\tilde{\pi}_t \equiv \pi_t + \theta^{-1}(x_t - x_{t-1})$ follow a deterministic path, in the case that the central bank must make its policy decision a period in advance, it is *not* optimal for the projection $E_{t-1}\tilde{\pi}_t$ to evolve deterministically. (It should instead depend on the error that was made in period $t - 2$ in forecasting $\tilde{\pi}_{t-1}$.)

Returning to the case of full information on the part of the central bank, it is possible to generalize our result to specifications under which $J < \infty$. In this case, (3.25) places no restriction on φ_t for any $t \geq t_0 + J$. However, in such a case, (3.26) implies that $\tilde{p}_t = \varphi_t$. It then follows from (3.24) that

$$\tilde{p}_t = E_{t-J}\tilde{p}_t$$

for all $t \geq t_0 + J$.⁸⁵ Hence the more general result is that under an optimal policy, \tilde{p}_t must follow a deterministic path for the first J periods after the adoption of the optimal policy, and thereafter must always be perfectly predictable J periods in advance. However, the path of $\{E_{t-J}\tilde{p}_t\}$ can be completely arbitrary for $t \geq t_0 + J$, and can depend in an arbitrary way upon shocks occurring between period t_0 and period $t - J$. Thus while a policy rule that ensures that $\{\tilde{p}_t\}$ follows a deterministic target path is *among* the optimal policies even when J is finite, this strong requirement is

⁸⁵See Koenig (2004) for a similar result in a model where nominal wages are set on the basis of sticky information.

no longer necessary for optimality.

The finding that it is optimal for the output-gap adjusted price level defined in (2.65) to follow a deterministic target path in this general class of sticky-information models as well as in the general class of sticky-price models considered in the previous section suggests that the result holds under quite weak assumptions about the timing of price adjustments and the information upon which they are based. Indeed, Kitamura (2008) analyzes optimal policy in a model of price adjustment that combines stickiness of prices with stickiness of information, and finds that it is again optimal for the output-gap-adjusted price level to evolve deterministically, regardless of the values assigned the parameters specifying either the degree of price stickiness or the degree of stickiness of information.

Of course, one cannot conclude that it is universally true that optimal policy requires that the output-gap-adjusted price level defined in (2.65) evolve deterministically. While we have obtained this result under a variety of assumptions about price adjustment, each of the models considered shares a large number of common features — we have in each case assumed the same specification of the demand side of the model, the same structure of production costs, and full information on the part of the central bank. Varying these assumptions can change the form of the optimal target criterion. Nonetheless, it does seem that a description of optimal policy in terms of a target criterion is more robust than other levels of description.

3.2 Which Price Index to Stabilize?

In the basic New Keynesian model of section 2 (as well as the generalizations just considered), the differentiated goods enter the model in a completely symmetric way, and furthermore only aggregate disturbances are considered, that affect the supply and demand for each good in an identical way. In such a model, there is a single obvious way to measure “the general level” of prices, an index in which each goods price enters with an identical weight. We have written the model structural equations in terms of their implications for the evolution of this symmetric price index, and derived optimal target criterion that involve the path of inflation, measured by changes in the log of this price index.

In actual economies, however, there are many reasons for the prices of different goods not to perfectly co-move with one another, apart from the differences in the

timing of price reviews or differences in the information sets of price-setters considered above, and the question as to which measure of inflation (or of the price level) should be targeted by a central bank is an important practical question for the theory of monetary policy. It would not be correct to say that the theory expounded above implies that an equally-weighted price index (or alternatively, one in which all goods are weighted by their long-run average expenditure shares) should be used in the target criterion; in fact, the models considered above are ones in which *there is no relevant difference* among alternative possible price indices. (Any price index that averages a large enough number of different prices for sampling error to be minimal — assuming that the criterion for selection of prices to include in the index is uncorrelated with any systematic differences in the timing of price reviews or updating of information sets by suppliers of the particular goods — will evolve in essentially the same way in response to aggregate disturbances.)

It is therefore important to extend the theory developed above to deal with environments in which the factors that determine prices can differ across sectors of the economy. It is important to consider the consequences of disturbances that have asymmetric effects on different sectors of the economy, and also to allow for different structural parameters in the case of different sectors. Below, I give particular attention to heterogeneity in the degree to which prices are sticky in different sectors of the economy.⁸⁶ The model sketched here is still a highly stylized one, with only two sectors and only a few types of heterogeneity. But it will illustrate how the methods introduced above can also be applied to a multi-sector model, and may also provide some insight into which of the conclusions obtained above in the case of the basic New Keynesian model are more likely to be apply to more general settings.

3.2.1 Sectoral Heterogeneity and Asymmetric Disturbances

Let us consider a slightly more general version of the two-sector model proposed by Aoki (2001).⁸⁷ Instead of assuming that the consumption index C_t that enters the utility function of the representative household is defined by the CES index (2.2), let

⁸⁶For evidence both on the degree of heterogeneity across sectors of the U.S. economy, and on the degree to which this heterogeneity matters for quantitative predictions about aggregate dynamics, see, e.g., Carvalho (2006) and Nakamura and Steinsson (2010).

⁸⁷The two-sector model presented here also closely resembles the treatment of a two-country monetary union in Benigno (2004).

us suppose that it is a CES aggregate of two sub-indices,

$$C_t \equiv \left[(n_1 \varphi_{1t})^{\frac{1}{\eta}} C_{1t}^{\frac{\eta-1}{\eta}} + (n_2 \varphi_{2t})^{\frac{1}{\eta}} C_{2t}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \quad (3.28)$$

for some elasticity of substitution $\eta > 0$. The sub-indices are in turn CES aggregates of the quantities purchased of the continuum of differentiated goods in each of the two sectors,

$$C_{jt} \equiv \left[\int_{N_j} c_t(i)^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}},$$

for $j = 1, 2$, where the intervals of goods belonging to the two sectors are respectively $N_1 \equiv [0, n_1]$ and $N_2 \equiv (n_1, 1]$, and once again $\theta > 1$.⁸⁸ In the aggregator (3.28), n_j is the number of goods of each type ($n_2 \equiv 1 - n_1$), and the random coefficients φ_{jt} are at all times positive and satisfy the identity $n_1 \varphi_{1t} + n_2 \varphi_{2t} = 1$. (The variation in the φ_{jt} thus represents a single disturbance each period, a shift in the relative demand for the two sectors' products.)

It follows from this specification of preferences that the minimum cost of obtaining a unit of the sectoral composite good C_{jt} will be given by the sectoral price index

$$P_{jt} \equiv \left[\int_{N_j} p_t(i)^{1-\theta} di \right]^{\frac{1}{1-\theta}} \quad (3.29)$$

for $j = 1, 2$, and that the minimum cost of obtaining a unit of C_t will correspondingly be given by the overall price index

$$P_t \equiv \left[n_1 \varphi_{1t} P_{1t}^{1-\eta} + n_2 \varphi_{2t} P_{2t}^{1-\eta} \right]^{\frac{1}{1-\eta}}.$$

Assuming that both households and the government care only about achieving as large as possible a number of units of the aggregate composite good at minimum cost, the demand function for any individual good i in sector j will be of the form

$$y_t(i) = Y_{jt} (p_t(i)/P_{jt})^{-\theta},$$

where the demand for the sectoral composite good is given by

$$Y_{jt} = n_j \varphi_{jt} Y_t (P_{jt}/P_t)^{-\eta}$$

⁸⁸We need not assume that $\eta > 1$ in order for there to be a well-behaved equilibrium of the two-sector model under monopolistic competition. In fact, the limiting case in which $\eta \rightarrow 1$ and the aggregator (3.28) becomes Cobb-Douglas is frequently assumed; see, e.g., Benigno (2004).

for each sector j . Note that the random factors φ_{jt} appear as multiplicative disturbances in the sectoral demand functions; this is one of the types of asymmetric disturbance that we wish to consider.

A common production technology of the form (2.5) is again assumed for each good, with the exception that the multiplicative productivity factor A_t is now allowed to be sector-specific. (That is, there is an exogenous factor A_{1t} for each of the firms in sector 1, and another factor A_{2t} for each of the firms in sector 2.) This allowance for sector-specific productivity variation is another form of asymmetric disturbance. There is similarly assumed as above to be a common disutility of labor function for each of the types of labor, except that now the preference shock \bar{H}_t is allowed to be sector-specific as well. The model thus allows for three types of asymmetric disturbances: variation in relative demands for the goods produced in the two sectors; variation in the relative productivity of labor in the two sectors; and variation in the relative willingness of households to supply labor to the two sectors. Each of these three types of asymmetric disturbances would result in variation in the relative quantity supplied of goods in the two sectors, and in the relative price of goods in the two sectors, even in the case of complete price flexibility (with full information). If we allow for time variation in a wage markup μ_t^w or in a proportional tax rate τ_t on sales revenues, these can be sector-specific as well. The latter types of asymmetric disturbances do not result in any asymmetry of the efficient allocation of resources, but would again be sources of asymmetry in both prices and quantities in a flexible-price equilibrium.

As in the basic New Keynesian model, I shall assume Calvo pricing by each firm; but the probability α that a firm fails to reconsider its price in a given period is now allowed to depend on the sector j . Again it is useful to derive a log-linear approximation to the model dynamics, near a long-run steady state with zero inflation. The calculations are also simplest if I assume that in this steady state, all sector-specific disturbances have common values in the two sectors: $\bar{\varphi}_1 = \bar{\varphi}_2$, $\bar{A}_1 = \bar{A}_2$, and so on. In this case, the prices of all goods are the same in the steady state (as before), and the steady-state allocation of resources is the same as in the one-sector model. I allow, however, for (small) asymmetric departures from the symmetric steady state.

Using the same methods as in section 2, one can show⁸⁹ that to a log-linear approximation, the dynamics of the two sectoral price indices are given by a pair of

⁸⁹See Woodford (2003, chap. 3, sec. 2.5) for details of the calculation.

sector-specific Phillips curves

$$\pi_{jt} = \kappa_j(\hat{Y}_t - \hat{Y}_t^n) + \gamma_j(p_{Rt} - p_{Rt}^n) + \beta E_t \pi_{j,t+1} + u_{jt}, \quad (3.30)$$

for $j = 1, 2$. Here $\pi_{jt} \equiv \Delta \log P_{jt}$ is the sectoral inflation rate for sector j ; \hat{Y}_t is the percentage deviation of production of the aggregate composite good (*not* the sectoral composite good!) from its steady-state level, as before; \hat{Y}_t^n is the flexible-price equilibrium level of production of the aggregate composite good, when the wage markups and the tax rates are held fixed at their (common) steady-state levels, as in section 2.4.2; $p_{Rt} \equiv \log(P_{2t}/P_{1t})$ is a measure of the relative price of the two sectoral composite goods; p_{Rt}^n is the flexible-price equilibrium relative price, again with the wage markups and tax rates fixed at their steady-state levels; and u_{jt} is a sector-specific “cost-push” disturbance that depends only on the deviations of μ_{jt}^w and τ_{jt} from their steady-state levels. To this log-linear approximation, \hat{Y}_t^n depends only on the “aggregate” disturbances, defined as $a_t \equiv \sum_j n_j a_{jt}$ (where $a_{jt} \equiv \log A_{jt}$ for $j = 1, 2$), and so on, and is the same function of these disturbances as in the one-sector model;⁹⁰ while \hat{p}_{Rt}^n depends only on the “relative” disturbances, defined as $a_{Rt} \equiv a_{2t} - a_{1t}$, and so on.

The coefficients of equation (3.30) are given by

$$\kappa_j \equiv \frac{(1 - \alpha_j)(1 - \alpha_j \beta) \omega + \sigma^{-1}}{\alpha_j} \frac{1}{1 + \omega \theta} > 0$$

for $j = 1, 2$, and

$$\gamma_1 \equiv n_2 \frac{(1 - \alpha_1)(1 - \alpha_1 \beta) \frac{1 + \omega \eta}{1 + \omega \theta}}{\alpha_1} > 0, \quad \gamma_2 \equiv -n_1 \frac{(1 - \alpha_2)(1 - \alpha_2 \beta) \frac{1 + \omega \eta}{1 + \omega \theta}}{\alpha_2} < 0,$$

where $0 < n_j < 1$ is the number of goods in sector j . Because the only kind of “structural” asymmetry allowed for here is heterogeneity in the degree of price stickiness, the slope coefficients κ_j differ across the two sectors if and only if the α_j are different: $\kappa_1 < \kappa_2$ if and only if $\alpha_2 < \alpha_1$. The coefficients γ_j instead have opposite signs, owing to the asymmetry in the way each sector’s price index enters the definition of p_{Rt} .

Optimal policy is particularly easy to characterize using the method of linear-quadratic approximation explained in section 2.4, and if we restrict ourselves to the case of “small steady-state distortions,” as in section 2.4.2. In this case one can show⁹¹ that the expected utility of the representative household varies inversely (to

⁹⁰This is one of the simplifications that results from log-linearizing around a symmetric steady state.

⁹¹For details of the calculations, see Woodford (2003, chap. 6, sec. 4.3).

a second-order approximation) with a discounted loss function of the form

$$E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \left[\sum_j w_j \pi_{jt}^2 + \lambda_x (x_t - x^*)^2 + \lambda_R (p_{Rt} - p_{Rt}^n)^2 \right], \quad (3.31)$$

where $x_t \equiv \hat{Y}_t - \hat{Y}_t^n$ denotes the output gap; the relative weights on the two inflation objectives (normalized to sum to 1) are given by

$$w_j \equiv \frac{n_j \kappa}{\kappa_j} > 0,$$

in which expression the “average” Phillips curve slope is defined as

$$\kappa \equiv (n_1 \kappa_1^{-1} + n_2 \kappa_2^{-1})^{-1} > 0;$$

and the other two relative weights are given by

$$\lambda_x \equiv \frac{\kappa}{\theta} > 0, \quad \lambda_R \equiv n_1 n_2 \frac{\eta(1 + \omega\eta)}{\omega + \sigma^{-1}} \lambda_x > 0.$$

The optimal level of the output gap x^* is again the same function of the steady-state distortions as in the one-sector model.

Each of the terms in (3.31) has a simple interpretation. As usual, deviations of aggregate output from the efficient level (given the “aggregate” technology and preference shocks) — or equivalently, deviations of the output gap from the optimal level x^* — lower welfare. But even given an efficient level of production of the aggregate composite good, a non-zero “relative price gap” $p_{Rt} - p_{Rt}^n$ will imply an inefficient relative level of production of the two sectoral composite goods, so variations in the relative price gap lower welfare as well. And finally, instability of *either* sectoral price level leads to relative-price distortions within that sector, and hence an inefficient composition of sectoral output, even if the quantity supplied of the sectoral composite good is efficient. Elimination of all of these sources of inefficiency in the equilibrium allocation of resources would require simultaneous stabilization of all four of the variables appearing in separate quadratic terms of (3.31). But in general, monetary policy cannot simultaneously stabilize all four variables, even in the absence of variation in the cost-push terms, if there is exogenous variation in the “natural relative price” p_{Rt}^n , as there almost inevitably will be if there are asymmetric disturbances to technology and/or preferences. This means that the conditions

required for complete price stability to be fully optimal are even more stringent (and implausible) in a multi-sector economy.⁹²

As in section 2.4.2, a log-linear approximation to the optimal evolution of the endogenous variables can be obtained by finding the state-contingent paths of the variables $\{P_{1t}, P_{2t}, Y_t\}$ that minimize (3.31) subject to the constraints (3.30). The FOCs for this problem can be written as

$$w_j \pi_{jt} + \varphi_{jt} - \varphi_{j,t-1} + (-1)^{j-1} \psi_t = 0, \quad (3.32)$$

$$\lambda_x(x_t - x^*) - \sum_j \kappa_j \varphi_{jt} = 0, \quad (3.33)$$

$$\lambda_R(p_{Rt} - p_{Rt}^n) - \sum_j \gamma_j \varphi_{jt} + \psi_t - \beta E_t \psi_{t+1} = 0, \quad (3.34)$$

where (3.32) must hold for $j = 1, 2$. Here φ_{jt} is the Lagrange multiplier associated with constraint (3.30) (for $j = 1, 2$), and ψ_t is the Lagrange multiplier associated with the identity

$$p_{Rt} = p_{R,t-1} + \pi_{2t} - \pi_{1t}.$$

The optimal state-contingent dynamics are then obtained by solving the four FOCs (3.32)–(3.34) and the three structural equations (equations (3.30) plus the identity) each period for the paths of the seven endogenous variables $\{\pi_{jt}, p_{Rt}, x_t, \varphi_{jt}, \psi_t\}$, given stochastic processes for the composite exogenous disturbances $\{p_{Rt}^n, u_{jt}\}$.

Figure 5 illustrates the kind of solution implied by these equations in a numerical example. In this example, the two sectors are assumed to be of equal size ($n_1 = n_2 = 0.5$), but prices in sector 2 are assumed to be more flexible; specifically, while the overall frequency of price change is assumed to be the same as in the example considered in Figure 1 (where $\alpha = 0.66$ for all firms), the model is now parameterized so that prices adjust roughly twice as often in sector 2 as in sector 1 ($\alpha_1 = 0.77, \alpha_2 = 0.55$). In other respects, the model is parameterized as in Figure 1.⁹³ The disturbance

⁹²It can be shown, however, that even in the presence of asymmetric disturbances to technology and preferences, if the degree of price stickiness is the same in both sectors ($\alpha_1 = \alpha_2$) and there are no cost-push disturbances, it is optimal to completely stabilize an equally-weighted price index; and just as in the one-sector model, this policy will completely stabilize the output gap x_t . (See Woodford, 2003, chap. 6, sec. 4.3.) However, this result no longer holds if $\alpha_1 \neq \alpha_2$.

⁹³The parameter values assumed for β, ζ, ω , and θ , as well as for the average frequency of price adjustment, are taken from Woodford (2003, Table 5.1). In addition, it is assumed that $\eta = 1$, so that the expenditure shares of the two sectors remain constant over time, despite permanent shifts in the relative price p_{Rt} .




Figure 5: Impulse responses of the output gap and of the two log sectoral price indices under an optimal policy commitment, in the case of an asymmetric real disturbance that permanently increases the natural relative price of sector-2 goods.

assumed is one that immediately and permanently increases the (log) natural relative price p_{Rt}^n of the flexible-price sector by one percentage point; note that it does not matter for the calculations reported here whether this is due to a shift in relative demand or a shift in relative costs of production. (All quantities on the vertical axis are in percentage points.) The figure shows that under an optimal policy, the long-run increase in the relative price of sector-2 goods results from an increase in the sector-2 price index of about 84 basis points and a decrease in the sector-1 price index of about 16 basis points. While this means that an equally-weighted (or expenditure-weighted) price index increases in response to the shock, one observes that the output gap is temporarily reduced during the period that prices are adjusting. Hence this type of disturbance gives rise to phenomena of the sort captured by a “cost-push shock” in a one-sector model like that of section 1, though in the present case no variations in the degree of market power or in tax distortions are needed to cause such an effect

(and the terms u_{jt} are both equal to zero in this example).

While the optimal dynamics are generally more complex than in the one-sector model, one important conclusion of the analysis in sections 1 and 2 remains valid: under optimal policy, no disturbances should be allowed to permanently shift a (suitably defined) measure of the general price level. In the case that the exogenous process $\{p_{Rt}^n\}$ is stationary, so that there are no permanent shifts in the natural relative price, this will be true regardless of the price index used to measure the general level of prices. If, instead, we suppose that the process $\{p_{Rt}^n\}$ has a unit root, so that permanent shifts occur in the natural relative price, then there will be no possibility of using monetary policy to stabilize *all* prices, and one can at most maintain a constant long-run level for some particular price index; but once again, there exists a price index for which a constant long-run price level target remains optimal. This price index is defined (to a log-linear approximation) by

$$\bar{p}_t \equiv \sum_j w_j \log P_{jt}.$$

Note that in the numerical example of Figure 5, $w_1 = 0.84, w_2 = 0.16$, so that the responses shown in the figure imply no long-run change in \bar{p}_t .

The optimality of maintaining a constant long-run expected value for \bar{p}_t can be demonstrated from the form of the FOCs as follows. Let us suppose that the cost-push disturbances $\{u_{jt}\}$ are stationary processes with means equal to zero, and that while $\{p_{Rt}^n\}$ may have a unit root, its first difference Δp_{Rt}^n is stationary with mean zero, so that at each point in time, there is a well-defined long-run expected value

$$p_{Rt}^{n\infty} \equiv \lim_{j \rightarrow \infty} E_t p_{R,t+j}^n$$

for the natural relative price. One can show that there exists a solution to the FOCs together with the structural equations in which each of the endogenous variables $\{\pi_{jt}, p_{Rt}, x_t, \varphi_{jt}, \psi_t\}$ is also difference-stationary (if not actually stationary), and so has a well-defined long-run expected value at all times; and this is the solution corresponding to the optimal equilibrium. Here I shall confine myself to arguing that in any difference-stationary solution, \bar{p}_t must be a stationary variable, so that its long-run expected value is some constant p^* .

It follows from the FOCs (3.32) that at any time, the long-run expected values of the two sectoral inflation rates must satisfy

$$w_1 \pi_{1t}^\infty = -\psi_t^\infty, \quad w_2 \pi_{2t}^\infty = \psi_t^\infty.$$

But in order for p_{Rt} to have a well-defined long-run expected value, the long-run expected values of the two sectoral inflation rates must be identical. It then follows that the FOCs can be satisfied only if $\pi_{jt}^\infty = 0$ for $j = 1, 2$, and that $\psi_t^\infty = 0$ as well. Thus one finds as in the one-sector model that the optimal long-run average inflation rate is zero, and that this is equally true in both sectors, so that it is true regardless of the price index used to measure an overall inflation rate.

It then follows from the sectoral Phillips curves (3.30) that in order for both of the long-run expected sectoral inflation rates to be zero the long-run expected values of the output gap and of the relative price gap must satisfy

$$\kappa_j x_t^\infty + \gamma_j (p_{Rt}^\infty - p_{Rt}^{n\infty}) = 0 \quad (3.35)$$

at all times, for $j = 1, 2$. But it is not possible for (3.35) to be simultaneously satisfied for both j unless

$$x_t^\infty = 0, \quad p_{Rt}^\infty = p_{Rt}^{n\infty}$$

at all times. And if these conditions hold at all times, the FOCs (3.33)–(3.34) respectively require that

$$\sum_j \kappa_j \varphi_{jt}^\infty = \lambda_x x^*, \quad \sum_j \gamma_j \varphi_{jt}^\infty = 0$$

at all times. But these conditions cannot be jointly satisfied unless the φ_{jt}^∞ take certain constant values φ_j^* at all times.

Finally, summing the two FOCs (3.30), one obtains

$$\sum_j w_j \pi_{jt} + \sum_j \varphi_{jt} - \sum_j \varphi_{j,t-1} = 0,$$

which can be alternatively written in the form

$$\Delta \bar{p}_t + \sum_j \varphi_{jt} - \sum_j \varphi_{j,t-1} = 0.$$

This implies that the quantity $\bar{p}_t + \sum_j \varphi_{jt}$ must remain constant over time, regardless of the disturbances affecting the economy. If we let the amount by which the constant equilibrium value of this quantity exceeds $\sum_j \varphi_j^*$ be denoted p^* , then it follows that

$$\lim_{j \rightarrow \infty} E_t \bar{p}_{t+j} = p^*$$

at all times. Hence, as noted above, optimal policy requires complete stabilization of the expected long-run value of the log price index \bar{p}_t .

The occurrence of real disturbances that permanently shift equilibrium relative prices is often thought to provide an important argument against the desirability of price-level targeting. It is commonly argued that it is appropriate to allow a one-time (permanent) shift in the general level of prices in response to such a shock, though this should not be allowed to give rise to expectations of ongoing inflation; hence a constant long-run target for inflation is appropriate, but not a constant long-run target for the price level. But I have shown in the present model that, while it is true that under an optimal policy, the long-run expected value of *all* measures of the inflation rate should remain constant, and the long-run expected values of *most* measures of the general price level should *not* remain constant, in the case of a shock to the natural relative price, there is nonetheless a particular price index that long-run value of which *should* remain constant, even in the case of real disturbances of this kind. Moreover, a description of optimal policy in terms of the long-run price-level target is superior to a description in terms of the long-run inflation target alone (or even long-run targets for each of the sectoral inflation rates); for the long-run inflation target alone would not tell the public what to expect about the cumulative increase in prices in each of the two sectors during the period of adjustment to the new long-run relative price. A commitment to a fixed long-run value for \bar{p}_t instead suffices to clarify what long-run values should be expected for each of the sectoral price indices at any point in time, given current long-run relative-price expectations. It therefore specifies the precise extent to which a given increase in the relative price of sector 2 goods should occur through inflation in sector 2 as opposed to deflation in sector 1.⁹⁴

One observes from the definition of the coefficients w_j that, for any given degree of price stickiness in the two sectors, the coefficient w_j is proportional to the size n_j (or the expenditure share) of each sector. One may also observe that, for any given relative sizes of the two sectors, and fixing the degree of price flexibility in the other

⁹⁴It is also sometimes argued that an increase in the general level of prices is desirable in response to a relative-price shock, in order to ensure that there is never deflation in any sector. But the reason for avoiding deflation is that expected declines in *all* prices can easily create a situation in which the zero lower bound on nominal interest rates becomes a binding constraint. Deflation in one sector only, when coupled with higher than average inflation in other sectors so that there is no decline in the expected overall inflation rate, does not imply that unusually low nominal interest rates will be required to achieve the desired path of prices. Hence there is no reason to regard temporary sectoral deflation as particularly problematic.

sector (at some value $0 < \alpha_{-j} < 1$), w_j is a monotonically increasing function of α_j , ranging from 0 when $\alpha_j = 0$ (the case of completely flexible prices in sector j , to precisely the fraction n_j when α_j is equal to α_{-j} , to a limiting value of 1 as α_j approaches 1. Thus the long-run price level target should be defined in terms of a price index that is weighted by expenditure only in the case that the degree of price stickiness in both sectors is the same.⁹⁵

In the case that prices are sticky only in one sector, and completely flexible in the other, the price-level target should be defined purely in terms of an index of prices in the sticky-price sector, as shown by Aoki (2001). This provides a theoretical justification for a long-run target for a “core” price index, which omits extremely flexible prices such as those of food and energy. However, in general, the optimal price-level target will involve an index that puts weights on prices in different sectors that differ from their expenditure shares, even among those prices that are not excluded from the index altogether. In the present model, if $0 < \alpha_j < 1$ in both sectors, the optimal price index will put some weight on prices in each sector; but the relative weights will not generally equal the relative expenditure shares. In particular, $w_1/w_2 > n_1/n_2$ if and only if $\alpha_1 > \alpha_2$, as in the example considered in Figure 5; the sector in which prices are more flexible should receive a lower weight, relative to its share in total expenditure.

While the price index \bar{p}_t should have a constant long-run level, it is not generally optimal in a multi-sector model for even this measure of the price level to be held constant at all times; and (contrary to the result obtained in sections 1 and 2 for the one-sector model) this is true even when there are no “cost-push” disturbances $\{u_{jt}\}$.⁹⁶ One can, however, fully characterize optimal policy, even in the short run, by

⁹⁵This agrees with the conclusion of Benigno (2004) regarding the optimal weights on regional inflation rates in the inflation target for a monetary union. Benigno assumes, however, that some price index is to be stabilized at all times, rather than only in the long run, and optimizes over policies in that restricted class.

⁹⁶Complete stability of the price index \bar{p}_t is optimal in two special cases: if there are no cost-push disturbances, and (i) one of the sectors has completely flexible prices, or (ii) prices are equally flexible in the two sectors. In the first case, it is optimal to completely stabilize the price index for the sticky-price sector, as shown by Aoki (2001), as this achieves the flexible-price allocation of resources. In the second case, it is optimal to completely stabilize the expenditure-weighted price index, as shown in Woodford (2003, chap. 6, sec. 4.3); in this case, the evolution of the relative price p_{Rt} is independent of monetary policy, and an analysis similar to that for the one-sector model continues to apply.

a time-invariant target criterion. One can show that there exist Lagrange multipliers such that all of the FOCs (3.32)–(3.34) are satisfied each period, if and only if the target criterion

$$\Delta\tilde{p}_t - \beta E_t \Delta\tilde{p}_{t+1} = -\Gamma[(\bar{p}_t - p^*) + \phi_x x_t + \phi_R(p_{Rt} - p_{Rt}^n)] \quad (3.36)$$

is satisfied each period, where p^* is the long-run price-level target discussed above; \tilde{p}_t is the output-gap-adjusted price level defined in (2.65), again using p_t to denote $\log P_t$;⁹⁷ and the coefficients are given by

$$\Gamma \equiv \frac{\kappa_1 \kappa_2}{\kappa} \frac{1 + \omega \eta}{\omega + \sigma^{-1}} > 0, \quad \phi_x \equiv \theta^{-1}, \quad \phi_R \equiv n_1 n_2 \frac{\eta}{\theta} (\kappa_1^{-1} - \kappa_2^{-1}) \kappa.$$

Note that ϕ_R is positive if and only if prices in sector 2 are more flexible than those in sector 1 ($\alpha_1 > \alpha_2$). Hence the term in the square brackets in (3.36) is a greater positive quantity the greater the extent to which the price index \bar{p}_t exceeds its long-run target value p^* , output exceeds the natural rate, or the relative price of the goods with more flexible prices exceeds its natural value.

Since each of the terms in (3.36) other than the price-level gap term $(\bar{p}_t - p^*)$ is necessarily stationary (under the maintained assumption of a difference-stationary solution), it follows that a policy that conforms to this target criterion will make the price-level gap stationary as well. This implies a long-run average inflation rate of zero, and this, as explained above, requires that the output gap and relative-price gap each have long-run average values of zero as well. Hence each of the terms in (3.36) other than the price-level gap term has a long-run average value of zero. It then follows that the long-run average value of the price-level gap must be zero as well, so that conformity to the target criterion (3.36) implies that the long-run average value of \bar{p}_t will equal p^* . Thus this target criterion does indeed guarantee consistency with the long-run price-level target. At the same time, it specifies the precise rate of short-term adjustment of prices that should be targeted under an optimal policy.

Figure 6 illustrates how the optimal responses to a permanent relative-price shock shown in Figure 5 conform to this target criterion. The figure plots the responses under optimal policy of the variables \bar{p}_t , x_t and $p_{Rt} - p_{Rt}^n$, the projections of which appear on the right-hand side of (3.36). (Note that the path of $\{\bar{p}_t\}$ converges asymptotically to the same level, denoted zero on the vertical axis, as it would have had in

⁹⁷In deriving (3.36) from the FOCs, I use the fact that to a log-linear approximation, $p_t = \sum_j n_j \log P_{jt}$.




Figure 6: Impulse responses of the variables referred to in the target criterion (3.36), for the same numerical example as in Figure 5. Here the “price level” plotted is the asymmetrically weighted price index \bar{p}_t .

the absence of the shock, while the relative price converges asymptotically to the natural relative price, and output converges asymptotically to the natural rate of output. The path of the output gap shown here is the same as in Figure 5.) The line labeled “target” plots the response of the composite target variable (a linear combination of the three variables just mentioned) that appears inside the square brackets on the right-hand side of (3.36). The line labeled “adjustment” is equal to Γ^{-1} times the price-adjustment terms on the left-hand side of (3.36). The fact that the “adjustment” response and the “target” response are mirror images of one another shows that the criterion (3.36) is fulfilled at all horizons.

Integrating forward in time, the optimal target criterion can alternatively be written in the form

$$\Delta \tilde{p}_t = -\Gamma \sum_{j=0}^{\infty} \beta^j E_t g_{t+j}, \quad (3.37)$$

where g_t is the composite target variable plotted in Figure 6. This version of the

criterion suggests an approach to the implementation of optimal policy through a forecast-targeting procedure. At each decision point, the central bank would compute projections of the forward paths of the price level gap, the output gap, and the relative-price gap, under a contemplated forward path for policy, and also a projection for path of the rate of growth of the output-gap-adjusted price level. It would then judge whether the contemplated path for policy is appropriate by checking whether the growth in the gap-adjusted price level is justified by the projected levels (specifically, by a forward-looking moving average of the levels) of the three gap variables, in the way specified by equation (3.37).

A simpler target criterion can be proposed that, while not precisely optimal in general, captures the main features of optimal policy. This is the simple proposal that policy be used to ensure that the composite gap g_t follow a deterministic path, converging asymptotically to zero. This simpler target criterion approximates optimal policy for the following reason. If the degree of price flexibility in the two sectors is quite asymmetric, then Γ is a large positive quantity, and (3.36) essentially requires that the value of g_t be near zero at all times. On the other hand, if the degree of price flexibility in the two sectors is nearly the same, then ϕ_R is near zero, and \bar{p}_t is nearly the same price index as p_t , so that g_t is approximately equal to $\tilde{p}_t - p^*$.⁹⁸ The criterion (3.36) can then be approximated by the criterion

$$E_t[A(L)g_{t+1}] = 0, \quad (3.38)$$

where $A(L) \equiv \beta - (1 + \beta + \Gamma)L + L^2$. Factoring $A(L)$ as $\beta(1 - \mu L)(1 - \beta^{-1}\mu^{-1}L)$, where $0 < \mu < 1$, one can show that (3.36) holds if and only if

$$(1 - \mu L)g_t = 0. \quad (3.39)$$

Condition (3.39) is equivalent to (3.36) in the case that $\alpha_1 = \alpha_2$, and will have similar implications as long as α_1 and α_2 are not too dissimilar. But (3.39) is also equivalent to (3.36) in the case that $\alpha_j = 0$ in one sector (but not both); note that in this case $\mu = 0$, and (3.39) reduces to the requirement that $g_t = 0$. And the implications of (3.39) will also be approximately the same as those of (3.36) in any case in which α_j is small enough in one sector. Hence (3.39) has implications somewhat similar to those of (3.36) over the entire range of possible assumptions

⁹⁸These two quantities are exactly equal in the case that $\alpha_1 = \alpha_2$ exactly.

about the relative degree of price stickiness in the two sectors.⁹⁹ And (3.39) implies that $\{g_t\}$ must evolve deterministically, converging asymptotically to zero. Note that this approximate target criterion is a very close cousin to the one shown to be optimal in a variety of one-sector models. Again it can be viewed as stating that a gap-adjusted price level must evolve deterministically, and be asymptotically constant; the only difference is that now the price level is not necessarily an expenditure-weighted index of prices, and the gap adjustment includes an adjustment for the relative-price gap.

3.2.2 Sticky Wages as Well as Prices

Similar complications arise if we assume that wages are sticky, and not just the prices of produced goods. In the models considered above, wages are assumed to be perfectly flexible (or equivalently, efficient contracting in the labor market is assumed); this makes it possible for a policy that stabilizes the general price level to eliminate all distortions resulting from nominal rigidities, at least in one-sector models. If wages are also sticky, this will not be case. Moreover, if *both* wages and prices are sticky, there will in general be *no* monetary policy that eliminates all distortions resulting from nominal rigidities. The existence of random shifts in the “natural real wage” (the one that would result in equilibrium with complete flexibility of both wages and prices, with the distortion factors held at their steady-state values) requires adjustments in wages, prices, or both to occur, which will necessarily create inefficiencies owing to the mis-alignment of wages or prices that are set at different times, if the adjustments of both wages and prices are staggered rather than synchronous.

Erceg *et al.* (2000) introduce wage stickiness in a manner closely analogous to the Calvo model of price adjustment.¹⁰⁰ They assume that firms each hire labor of a large number of distinct types, with a production technology that makes output an increasing concave function of a CES aggregate of the distinct labor inputs; this results in a downward-sloping demand curve for each type of labor, the location of

⁹⁹The case shown in Figures 5 and 6 is not one in which it is optimal for g_t to be held precisely constant in response to the shock; nonetheless, the optimal change in the path of g_t is not large and is quite smooth.

¹⁰⁰More complex versions of their model of wage and price adjustment are at the heart of most current-vintage empirical DSGE models, such as the models of Christiano *et al.* (2005) and Smets and Wouters (2007).

which is independent of the wage demands of the suppliers of that type of labor. Wages are assumed to be set for each type of labor by a single (monopolistically competitive) representative of the suppliers of that type of labor, acting in their joint interest, and to be fixed in terms of money for a random time interval. The probability that the wage for a given type of labor is reconsidered in any given period is assumed to be independent of both the time since the last reconsideration of the wage and of the relation between the existing wage and current market conditions.

Under these assumptions, the joint dynamics of wage and price adjustment satisfy (to a log-linear approximation) the following pair of coupled equations:¹⁰¹

$$\pi_t = \kappa_p(\hat{Y}_t - \hat{Y}_t^n) + \xi_p(w_t - w_t^n) + \beta E_t \pi_{t+1} + u_{pt}, \quad (3.40)$$

$$\pi_{wt} = \kappa_w(\hat{Y}_t - \hat{Y}_t^n) - \xi_w(w_t - w_t^n) + \beta E_t \pi_{w,t+1} + u_{wt}, \quad (3.41)$$

where $\pi_{wt} \equiv \Delta \log W_t$ is the rate of wage inflation (rate of change of the Dixit-Stiglitz index of wages W_t); $w_t \equiv \log(W_t/P_t)$ is the log real wage; w_t^n (the “natural real wage”) is a function of exogenous disturbances that indicates the equilibrium real wage under flexible wages and prices, in the case that all distortion factors are fixed at their steady-state values; u_{pt} is an exogenous cost-push factor for price dynamics given wages (reflecting variations in VAT or payroll tax rates paid by firms, or in the market power of the suppliers of individual goods); and u_{wt} is an exogenous cost-push factor for wage dynamics given prices (reflecting variations in a wage income tax rate or a sales tax rate paid by consumers in addition to the sticky goods price, or in the market power of the suppliers of individual types of labor). The coefficient ξ_p is a positive factor that is larger the more frequently prices are adjusted, and ξ_w is correspondingly a positive factor that is larger the more frequently wages are adjusted. The output-gap response coefficients are defined as

$$\kappa_p \equiv \xi_p \epsilon_{mc,p} > 0, \quad \kappa_w \equiv \xi_w \epsilon_{mc,w} > 0,$$

where the elasticity of average real marginal cost with respect to increases in aggregate output, $\epsilon_{mc} \equiv \omega + \sigma^{-1}$, has been decomposed into the sum of two parts: the part $\epsilon_{mc,w} \equiv \nu\phi + \sigma^{-1}$ due to the increase in the average real wage when output increases,

¹⁰¹See Erceg *et al.* (2000) or the exposition in Woodford (2003, chap. 3, sec. 4.1) for the derivation. The notation used here follows Woodford (2003), except for the inclusion here of the possibility of the cost-push terms.

and the part $\epsilon_{mc,p} \equiv \phi - 1$ due to the increase in real marginal cost relative to the real wage (owing to diminishing returns to labor).

Again the analysis of optimal policy is simplest if we restrict ourselves to the case of “small steady-state distortions,” as in section 2.4.2.¹⁰² In this case one can show¹⁰³ that the expected utility of the representative household varies inversely (to a second-order approximation) with a discounted loss function of the form

$$E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} [\lambda_p \pi_t^2 + \lambda_w \pi_{wt}^2 + \lambda_x (x_t - x^*)^2], \quad (3.42)$$

where $x_t \equiv \hat{Y}_t - \hat{Y}_t^n$ again denotes the output gap; the weights on the two inflation measures are two positive coefficients (normalized to sum to 1) with relative magnitude

$$\frac{\lambda_w}{\lambda_p} = \frac{\theta_w}{\theta_p \phi} \frac{\xi_p}{\xi_w},$$

where θ_w, θ_p are the elasticities of substitution among different types of labor and among different goods, respectively; and the relative weight on the output-gap objective is

$$\lambda_x \equiv \frac{\epsilon_{mc}}{\theta_p \xi_p^{-1} + \theta_w \phi^{-1} \xi_w^{-1}} > 0.$$

Thus when wages and prices are both sticky, variability of either wage or price inflation distorts the allocation of resources and reduces welfare; the relative weight on wage inflation in the quadratic loss function is greater the stickier are wages (or the more flexible are prices), and greater the more substitutable are the different types of labor (or the less substitutable are the different goods).

One observes that there is a precise analogy between the form of this linear-quadratic policy problem and the one considered in the previous section, if we identify goods price inflation here with “sector 1 inflation” in the previous model, wage inflation with “sector 2 inflation,” and the real wage with the sector 2 relative price.

¹⁰²Benigno and Woodford (2005b) generalize the analysis to the case of large steady-state distortions, using the method explained in section 2.4.3. They derive a quadratic loss function of the same form as in (3.42) for this more general case, except that the output gap must be defined relative to a more complex function of the exogenous disturbances, and the coefficients $\lambda_w, \lambda_p, \lambda_x$ are more complex functions of the model parameters, involving in particular the degree of inefficiency of the steady-state level of output.

¹⁰³For details of the calculations, see Woodford (2003, chap. 6, sec. 4.4).

(The only difference is that in the case of the Erceg *et al.* model, there is no term proportional to the squared “relative price gap” in the quadratic loss function; the present model corresponds to the special case $\lambda_R = 0$ of the model considered earlier.) Hence the calculations discussed in the previous section have immediate implications for this model as well.

It follows that under optimal policy, there is a “price index” \bar{p}_t the long-run value of which should be unaffected by disturbances, even when those disturbances have permanent effects on output or the real wage; the only difference between this case and the ones discussed previously is that the “price index” in question is an index of both goods prices and wages, specifically

$$\bar{p}_t \equiv \lambda_p \log P_t + \lambda_w \log W_t.$$

Of course, it follows from this that if there are disturbances that permanently shift the natural real wage w_t^n , then there will exist no index of goods prices *alone* that is stationary under optimal policy; but the principle that it is desirable to maintain long-run stability of the price level remains valid, under the understanding that the correct definition of “price stability” should be stability of \bar{p}_t .¹⁰⁴

Similarly, it follows from the same arguments as in the previous section that optimal policy can be characterized by a target criterion of the form

$$\Delta[\hat{p}_t + \hat{\theta}^{-1} x_t] = -\Gamma \sum_{j=0}^{\infty} \beta^j E_t[(\bar{p}_{t+j} - p^*) + \bar{\theta}^{-1} x_{t+j}], \quad (3.43)$$

where \hat{p}_t is another (differently weighted) index of both prices and wages,

$$\hat{p}_t \equiv \frac{\kappa_p \lambda_p \log P_t + \kappa_w \lambda_w \log W_t}{\kappa_p \lambda_p + \kappa_w \lambda_w};$$

the coefficients $\bar{\theta}$ and $\hat{\theta}$ are two different weighted averages of θ_p and $\phi^{-1}\theta_w$,

$$\bar{\theta} \equiv \frac{\xi_p^{-1} \theta_p + \xi_w^{-1} \phi^{-1} \theta_w}{\xi_p^{-1} + \xi_w^{-1}},$$

$$\hat{\theta} \equiv \frac{\epsilon_{mc,p} \theta_p + \epsilon_{mc,w} \phi^{-1} \theta_w}{\epsilon_{mc}};$$

¹⁰⁴If the natural real wage is a stationary random variable, then the result just mentioned implies that the long-run expected value of $\log P_t$ should also be constant. However, if there is a unit root in productivity, as is often argued, then the natural real wage should possess a unit root as well.

and

$$\Gamma \equiv \frac{\xi_p \xi_w \epsilon_{mc}}{\kappa_p \lambda_p + \kappa_w \lambda_w} > 0.$$

Hence optimal policy can be implemented through a forecast-targeting procedure under which it is necessary at each decision point to compute projected future paths of the general level of prices, the general level of wages, and the output gap, in order to verify that (3.43) is satisfied under the intended forward path of policy.

4 Research Agenda

This chapter has shown how it is possible to analyze monetary stabilization policy using techniques similar to those used in the modern theory of public finance, in particular the Ramsey theory of dynamic optimal taxation. The methods and some characteristic issues that arise in this project have been illustrated using a particular class of relative simple models of the monetary transmission mechanism. Several key themes that have emerged are nonetheless likely to be of broad applicability to more complex (and more realistic) models. These include the advantages of a suitably chosen policy commitment (assuming that commitment is possible and can be made credible to the public) over the outcome associated with discretionary policymaking, and the convenience of formulating a desirable policy commitment in terms of a target criterion that the central bank should seek to fulfill through adjustment of its policy instrument or instruments.

The degree to which other, more specific results generalize to more realistic settings deserves further investigation. In a range of different models, I have shown that optimal policy requires not merely that there be a well-defined long-run inflation rate that remains invariant in the face of economic disturbances, but that there be a well-defined long-run *price level* that is unaffected by shocks, if the price level is measured by a suitably defined index of the prices of different goods. I have given examples (in sections 1.6 and 1.7) where it is not quite true that the long-run forecast of the price level should remain unchanged by all shocks; but even in these cases, optimal policy is characterized by *error correction*, in the sense that when a disturbance deflects the output-gap-adjusted price level from the path that it would otherwise have been expected to take, the gap-adjusted price level should subsequently be brought back to that path *and even somewhat beyond it* — this over-correction being the reason

why the price level is not actually stationary under the optimal commitment. Thus it is quite generally desirable, in the settings considered here, for a central bank to commit itself to error-correction of the sort implied by a price-level target.

Another recurrent theme has been the desirability, in the shorter run, of maintaining a deterministic path for an *output-gap-adjusted* price level, rather than for a measure of the price level itself. A range of models has been discussed in each of which this very simple target criterion represents an optimal commitment, and the appropriate relative weight on the output gap has been the same (*i.e.*, equal to the reciprocal of the elasticity of substitution between differentiated goods) in many of these cases.¹⁰⁵ In the more complex models considered in section 3.2, the optimal target criterion is in general no longer so simple, yet it continues to be the case that temporary departures of the (appropriately defined) price level from its long-run target should be proportional to certain measures of temporary real distortions, with a gap between the level of aggregate output and a time-varying “natural rate” appearing as at least one important aspect of the real distortions that justify such temporary variations in the price level.

While we have thus obtained quite consistent results across a range of specifications, that incorporate (at least in simple ways) a number of key elements of empirical models of the monetary transmission mechanism, it must nonetheless be admitted that all of the models considered in this chapter are simple in some of the same ways. Not only are they all representative-household models, but they all assume that all goods are final goods produced using labor as the only variable factor of production, and they treat all private expenditure as indistinguishable from non-durable consumer expenditure (that is, there is no allowance for endogenous variation in rate of growth of productive capacity).¹⁰⁶ They are also all closed economies, and I have

¹⁰⁵Giannoni and Woodford (2005) show that the same result obtains in yet another case, a model that incorporates habit-formation in private expenditure, as do many empirical New Keynesian DSGE models.

¹⁰⁶The literature that evaluates particular parametric families of simple policy rules in the context of a particular quantitative model frequently does allow for more complex technologies and endogenous capital accumulation (*e.g.*, Schmitt-Grohé and Uribe, 2004, 2007). (This literature is reviewed in another chapter of this *Handbook* (Taylor and Williams, 2010), and so is not reviewed here.) Often studies of this kind find that optimal simple rules are fairly similar to those that would be optimal (within the same parametric family) in the case of a simpler model, without endogenous capital accumulation. But it is not clear how dependent these results may be on other restrictive

tacitly assumed throughout that lump-sum taxes exist and that the fiscal authority can be expected to adjust them so as to ensure intertemporal solvency of the government, regardless of the monetary policy chosen by the central bank, so that it has been possible to consider alternative monetary policies while abstracting from any consequences for the government's budget.¹⁰⁷ Finally, the models in this chapter all abstract from the kinds of labor-market frictions that have been important not only in real models of unemployment dynamics, but in some of the more recent monetary DSGE models.¹⁰⁸ Analysis of the form of optimal policy commitments in settings that are more complex in these respects is highly desirable, and these represent important directions for further development of the literature. Such developments are clearly possible in principle, since the general methods used to characterize optimal policy commitments in this chapter have been shown to be applicable to quite general classes of nonlinear policy problems, with state spaces of arbitrary (finite) size, by Benigno and Woodford (2008) and Giannoni and Woodford (2010).

Among the respects in which the models considered here omit the complexity of actual economies is their complete neglect of the role of financial intermediaries in the monetary transmission mechanism. This means that the analyses of optimal policy above have abstracted entirely from a set of considerations that have played a very large role in monetary policy deliberations in the recent past (most notably in 2008), namely, the degree to which monetary policy should take account of variations in financial conditions, such as changes in spreads between the interest rates paid by different borrowers. Extension of the theory of monetary stabilization policy to deal with questions of this kind is of particular importance at present.

While work of this kind remains relatively preliminary at the time of writing, it should be possible to apply the general methods explained in this chapter to models that incorporate both a non-trivial role for financial intermediation and the possibility of disturbances to the efficiency of private intermediation. Cúrdia and Woodford (2009a) provide one example of how this can be done. They consider a model

aspects of the specifications within which the welfare comparisons are made, such as the assumption that the only disturbances that ever occur are of a few simple kinds.

¹⁰⁷Extensions of the theory of optimal monetary stabilization policy to deal with these latter two issues are fairly well developed, but I omit discussion of them here because these are the topics of two other chapters of this *Handbook*; see Corsetti *et al.* (2010) for open-economy issues and Canzoneri *et al.* (2010) for the interaction between monetary and fiscal policies.

¹⁰⁸See, for example, the chapters in this *Handbook* by Gali (2010) and Christiano *et al.* (2010).

in which infinite-lived households differ in their opportunities for productive (*i.e.*, utility-producing, since again the model is one that abstracts from effects of private expenditure on productive capacity) expenditure, so that financial intermediation can improve the allocation of resources; and they allow for two reasons why a positive spread between the interest rate at which intermediaries lend to borrowers and the rate at which they themselves are financed by savers can persist in equilibrium. (On the one hand, loan origination may require the consumption of real resources that increase with the bank’s scale of operation; and on the other hand, banks may be unable to discriminate between borrowers who can be forced to repay their debts and others who will be able to avoid repayment, so that all borrowers will have to be charged an interest rate higher than the bank’s cost of funds to reflect expected losses on bad loans.) Random variation in either of these aspects of the lending “technology” can cause equilibrium credit spreads to vary for reasons that originate in the financial sector. Cúrdia and Woodford also allow for endogenous variation in credit spreads due to changes in the volume of lending, in response to disturbances to preferences, technology or fiscal policy.

Methods similar to those expounded above can be used to characterize an optimal policy commitment, if we take the average expected utility of the households of the different types (weighting the utility of each type by its population fraction) as the objective of policy. Cúrdia and Woodford (2009a) obtain an especially simple characterization of optimal interest-rate policy in the special case that (i) no resources are consumed by intermediaries, (ii) the fraction of loans that are bad is an exogenously varying quantity that is independent of an intermediary’s scale of operation, and (iii) the steady state is undistorted, as in section 2.4.1. In this case (in which intermediation is still essential, owing to the heterogeneity, and credit spreads can be non-zero, as a result of shocks to the fraction of loans that are bad), a linear approximation to the optimal policy commitment is again obtained by committing to the fulfillment at all times of a target criterion of the form (1.21) — or alternatively, of the form (1.23).¹⁰⁹

Thus it continues in this case to be possible to characterize optimal policy purely in terms of the projected evolution of inflation (or the price level) and of the output gap.

¹⁰⁹For the reasons discussed in section 1, the latter formulation of the optimal target criterion is once again more robust. In particular, the problem of a sometimes binding zero lower bound on the policy rate is more likely to arise as a result of disturbances to the size of equilibrium credit spreads.

Financial conditions are relevant to the central bank’s deliberations, but because they must be monitored in order to determine the path of the policy rate that is required in order to achieve paths for the price level and for the output gap consistent with the target criterion, and not (at least in this special case) because they influence the form of the target criterion itself. This result is obtained as an exact analytical result only in a fairly special case; but Cúrdia and Woodford also find that under a variety of calibrations of the model that are intended to be more realistic, a commitment to the “flexible inflation targeting” criterion continues to provide a reasonably close approximation to optimal interest-rate policy, even if it is no longer precisely the optimal policy.¹¹⁰

This chapter has discussed the character of fully optimal policy in a variety of fairly simple models, for which it is possible to obtain analytical solutions for the optimal equilibrium dynamics and for the target criteria that can implement these equilibria. While the methods illustrated here can be applied much more generally, the resulting characterization of optimal policy can rapidly become more complex, as the discussion in section 3 has already indicated. In particular, even in the case of a fairly small macroeconomic model (that necessarily abstracts from a great deal of the richness of available economic data), the optimal target criterion may be much too complex to be useful as a public expression of a central bank’s policy commitment — at least, to the extent that the point of such a commitment is to allow public understanding of what it should expect future policy to be like. As a practical matter, then, it is important to formulate recommendations for relatively simple target criteria, that, while not expected to be fully optimal, nonetheless approximate optimal policy to a reasonable extent.

Analysis of the properties of simple policy rules — both calculations of the optimal rule within some parametric family of simple rules, and analysis of the robustness of particular rules to alternative model specifications — has been extensive in the case of simple interest-rate *reaction functions* such as the “Taylor rule.”¹¹¹ A similar anal-

¹¹⁰The introduction of credit frictions also allows, at least in principle, for additional dimensions of central-bank policy in addition to the traditional tool of influencing the level of money-market interest rates; it now becomes relevant whether the central bank purchases only Treasury securities for its balance sheet, or instead extends credit in various forms to the private sector as well. Methods similar to those discussed above can be used to analyze optimal policy along these additional dimensions as well, as discussed in Cúrdia and Woodford (2009b).

¹¹¹This literature is reviewed in this *Handbook* by Taylor and Williams (2010).

ysis of the performance of simple target criteria in quantitative models with some claim to empirical realism still needs to be undertaken, and this is an important area for future research. While any target criterion must be implemented by adjustment of a policy instrument (that for most central banks, under current institutional arrangements, will be an operating target for an overnight interest rate such as the federal funds rate), it is far from obvious that a description of the central bank's policy commitment in terms of an interest-rate reaction function is more desirable than a "higher-level" description of the policy commitment in terms of a target criterion that the central bank seeks to fulfill.¹¹² In particular, while there has thus far been a larger literature assessing the robustness of simple reaction functions to alternative model specifications, it is far from obvious that policy rules specified at that level are more robust than equally simple target criteria. The results of this chapter have shown that in the case of models with simple structural equations, but that may be subject to many different types of stochastic disturbances (with potentially complex dynamics), simple target criteria can be found that are fully optimal across a wide range of specifications of the stochastic disturbance processes, whereas no interest-rate rule can be formulated in these examples that is equally robust to alternative disturbance processes. While this is only one of the kinds of robustness with which central banks must be concerned, and while the "robustly optimal" target criteria that are shown quite generally to exist in Giannoni and Woodford (2010) are only simple in the case of models with simple structural equations, these results suggest that further exploration of the robustness of simple target criteria are well worth undertaking.

Moreover, while the results of this chapter pertain only to fully optimal target criteria for simple models, I believe that the theory of optimal target criteria is likely to prove useful in the design of simple (and only approximately optimal) target criteria for more complex economies. A theoretical understanding of which types of target criteria are superior, at least in cases that are simple enough to be fully understood, is likely to provide guidance as to which simple criteria are plausible candidates to be approximately optimal in a broader range of circumstances. (This is illustrated by the result of Cúrdia and Woodford, 2009a, discussed above, in which the target criterion that would be optimal in a simpler case was found still to be approximately optimal under a range of alternative parameterizations of a more complex model.)

¹¹²See Svensson (2003) and Woodford (2007) for further discussion of this issue.

And knowing the form of a fully optimal target criterion in a model of interest, even when that criterion is too complex to be proposed as a practical policy commitment, may be useful in suggesting simpler target criteria, that continue to capture key features of the optimal criterion, and that may provide useful approximations to an optimal policy. (This is illustrated by the discussion in section 3.2.1 above of a simpler version of the optimal target criterion.)

Another important limitation of all of the analyses of optimal policy in this chapter has been the assumption that the policy rule that is adopted will be fully credible to the private sector, and that the outcome that is realized will be a rational-expectations equilibrium consistent with the central bank's policy commitment. Given that the target criteria discussed in this chapter each determine a unique non-explosive equilibrium, the central bank is assumed to be able to confidently predict the equilibrium implied by its policy commitment, and the choice of a policy commitment has accordingly been treated as equivalent to the choice of the most preferred among all possible rational-expectations equilibria of the model in question. But in practice, it is an important question whether a central bank can assume that a policy commitment about which it is quite serious will necessarily be fully credible with or correctly understood by the private sector; and even granting that the commitment itself is understood and believed, it is an important question whether private decision-makers will all understand why the economy should evolve in the way required by the rational-expectations equilibrium, and will necessarily have the expectations required for that evolution to be the one that actually occurs. And to the extent that it is unclear whether the outcome actually realized must be precisely the one predicted by rational-expectations equilibrium analysis, it is unclear whether a policy commitment that would be optimal in that case should actually be considered desirable. For example, a rule that would be sub-optimal under the assumption of rational expectations might be preferable on the ground that performance under this rule will not deteriorate as greatly under plausible departures from rational expectations as would occur in the case of a rule that would be better under the hypothesis of rational expectations.

This is another aspect of the broader question of the robustness of policy rules — just as one should be concerned about whether a rule that might be predicted to lead to a good outcome under a particular model of the economy will also lead to outcomes that are reasonably good if the correct model of the economy is somewhat different than had been assumed, one should be concerned about whether a rule that

is predicted to lead to a good outcome under the assumption of rational expectations will also lead to outcomes that are reasonably good in the case of expectations that fail to precisely conform to this assumption. One approach to this question is to model expectations as being formed in accordance with some explicit statistical model of learning from the data observed up to some point in time, and continuing to evolve as new data are observed.

One can analyze the optimal conduct of policy under a particular model of adaptive learning (assumed to be understood by the policymaker), and one can also analyze the robustness of particular policy proposals to alternative specifications of the learning process. Among other questions, one can consider the degree to which policy recommendations that would be optimal under rational expectations continue to lead to at least nearly-optimal outcomes under learning processes that are not too different from rational expectations, in the sense that the learning algorithm implies that forecasts should eventually converge to the rational-expectations forecasts once enough data have been observed, or that they should fluctuate around the rational-expectations forecasts without departing too far from them most of the time. While the literature addressing issues of this kind is still fairly new, some suggestive results exist, as summarized by Evans and Honkapohja (2009) and Gaspar *et al.* (2010). The results that exist suggest that some of the important themes of the literature on optimal policy under rational expectations continue to apply when adaptive learning is assumed instead; for example, a number of studies have found that it continues to be important to pursue a policy that stabilizes inflation expectations in response to cost-push disturbances to a greater extent than would occur under a discretionary policy that takes inflation expectations to be independent of policy — though if a mechanical model of adaptive learning is taken to be strictly correct, the stability of expectations must be maintained entirely through constraining the variability of the observed inflation rate, and not through any public announcements of policy targets or commitments.

Woodford (2010) illustrates an alternative approach to the problem of robustness of policy to departures from rational expectations. Rather than assuming a particular model of expectation formation that is known to the policymaker, it is assumed that private-sector expectations may differ in arbitrary ways from the forecasts that would represent rational expectations according to the central bank's model, but it is assumed that private-sector expectations will not be *too far* from being correct,

where the distance between private-sector expectations and those that the central bank regards as correct is measured using a relative-entropy criterion.¹¹³ A policy is then sought that will be ensure as good as possible an outcome in the case of any private-sector forecasts that do not depart from correct expectations (according to the central bank’s model) by more than a certain amount. In the case of the baseline policy problem considered above in section 1, “robustly optimal” policy is characterized as a function of the size of possible departures from rational expectations that are contemplated; and while the “robustly optimal” policy is not precisely like the optimal policy commitment characterized in section 1 (except when the contemplated departures are of size zero), it has many of the qualitative features of that policy. For example, it continues to be the case that commitment can achieve a much better worst-case outcome than discretionary policymaking (assuming either that $x^* > 0$ or a positive variance for cost-push shocks); that the robustly optimal long-run inflation target is zero even when $x^* > 0$; that the robustly optimal commitment allows inflation to respond less to cost-push shocks than would occur under discretionary policy; and that the robustly optimal commitment implies that following an increase in prices due to a cost-push shock, the central bank should plan to undo the increase in the level of prices by keeping inflation lower than its long-run value for a period of time. The reasons for the desirability of each of these elements of the robustly optimal policy are essentially the same as in section 1.

Further work on optimal (and robust) policy design when plausible departures from rational expectations are allowed for should be a high priority. Among the goals of such an inquiry should be the clarification not only of the appropriate targets for monetary policy, but of the way in which it makes sense for a central bank to respond to observed private-sector expectations that differ from its own forecasts. The latter issue, which is one of great practical importance for central bankers, is plainly one that cannot be analyzed within a framework that simply assumes rational expectations.¹¹⁴ While the results of such analysis cannot be anticipated before they

¹¹³A similar criterion has been extensively used in the literature on the design of policies that are robust to model mis-specification, as discussed by Hansen and Sargent (2010).

¹¹⁴It is unlikely that the most robust approach to policy will be one under which the central bank simply ignores any evidence of private-sector forecasts that depart from its own. See, e.g., Evans and Honkapohja (2006) and Preston (2008) for analyses of policy under adaptive learning dynamics in which policies that respond to observed private-sector forecasts are more robust to alternative specifications of the learning dynamics than policies (that would be optimal under rational expectations)

have been obtained, a clear understanding of the policy that would be optimal in the case that it were correctly understood is likely to be a useful starting point for the analysis of the subtler problems raised by diversity of opinions.

that ignore private-sector forecasts.

References

- [1] Adam, Klaus, and Roberto Billi, “Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates,” *Journal of Money, Credit and Banking* 38: 1877-1905 (2006).
- [2] Aoki, Kosuke, “Optimal Monetary Policy Responses to Relative Price Changes,” *Journal of Monetary Economics* 48: 55-80 (2001).
- [3] Aoki, Kosuke, “Optimal Commitment Policy Under Noisy Information,” *Journal of Economic Dynamics and Control* 30: 81-109 (2006).
- [4] Aoki, Kosuke, and Kalin Nikolov, “Rule-Based Monetary Policy under Central Bank Learning,” CEPR discussion paper no. 5056, May 2005.
- [5] Ball, Laurence, N. Gregory Mankiw, and Ricardo Reis, “Monetary Policy for Inattentive Economies,” *Quarterly Journal of Economics* 52: 703-725 (2005).
- [6] Benigno, Pierpaolo, “Optimal Monetary Policy in a Currency Area,” *Journal of International Economics* 63: 293-320 (2004).
- [7] Benigno, Pierpaolo, and Michael Woodford, “Optimal Monetary and Fiscal Policy: A Linear-Quadratic Approach,” in M. Gertler and K. Rogoff, eds., *NBER Macroeconomics Annual 2003*, Cambridge, MA: MIT Press.
- [8] Benigno, Pierpaolo, and Michael Woodford, “Inflation Stabilization and Welfare: The Case of a Distorted Steady State,” unpublished draft, New York University, May 2004.
- [9] Benigno, Pierpaolo, and Michael Woodford, “Inflation Stabilization And Welfare: The Case Of a Distorted Steady State,” *Journal of the European Economic Association* 3: 1185-1236 (2005a).
- [10] Benigno, Pierpaolo, and Michael Woodford, “Optimal Monetary Policy when Wages and Prices are Sticky: The Case of a Distorted Steady State,” in J. Faust, A. Orphanides, and D. Reifschneider, eds., *Models and Monetary Policy*, Washington, D.C.: Federal Reserve Board, 2005b.

- [11] Benigno, Pierpaolo, and Michael Woodford, “Optimal Inflation Targeting under Alternative Fiscal Regimes,” in F.S. Mishkin and K. Schmidt-Hebbel, eds., *Monetary Policy under Inflation Targeting*, Santiago, Chile: Central Bank of Chile, 2007.
- [12] Benigno, Pierpaolo, and Michael Woodford, “Linear-Quadratic Approximation of Optimal Policy Problems,” NBER working paper no. 12672, revised August 2008.
- [13] Blake, Andrew P. (2001), “A ‘Timeless Perspective’ on Optimality in Forward-Looking Rational Expectations Models,” NIESR Discussion Papers 188, National Institute of Economic and Social Research.
- [14] Canzoneri, Matthew, Robert Cumby, and Behzad Diba, “The Interaction Between Monetary and Fiscal Policy,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, volume 3, Amsterdam: North-Holland, forthcoming 2010.
- [15] Carvalho, Carlos, “Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks,” *The BE Journal of Macroeconomics* 2(1: Frontiers) (2006).
- [16] Christiano, Lawrence J., “The Zero Bound, Zero Inflation Targeting, and Output Collapse,” unpublished, Northwestern University, March 2004.
- [17] Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy* 113: 1-45 (2005).
- [18] Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin, “DSGE Models for Monetary Policy,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, volume 3, Amsterdam: North-Holland, forthcoming 2010.
- [19] Clarida, Richard, Jordi Gali and Mark Gertler, “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature* 37: 1661-1707 (1999).
- [20] Corsetti, Giancarlo, Luca Dedola, and Sylvain LeDuc, “Optimal Monetary Policy in Open Economies,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, volume 3, Amsterdam: North-Holland, forthcoming 2010.

- [21] Cúrdia, Vasco, and Michael Woodford, “Credit Frictions and Optimal Monetary Policy,” unpublished, Federal Reserve Bank of New York, August 2009a.
- [22] Cúrdia, Vasco, and Michael Woodford, “Conventional and Unconventional Monetary Policy,” CEPR Discussion Paper no. 7514, October 2009b.
- [23] Denes, Matthew, and Gauti B. Eggertsson, “A Bayesian Approach to Estimating Tax and Spending Multipliers,” Federal Reserve Bank of New York Staff Report no. 403, November 2009.
- [24] Dotsey, Michael, “Pitfalls in Interpreting Tests of Backward-Looking Pricing in New Keynesian Models,” Federal Reserve Bank of Richmond *Economic Quarterly* 88(1): 37-50 (2002).
- [25] Dotsey, Michael, Robert G. King, and Alexander L. Wolman, “State-Dependent Pricing and the General-Equilibrium Dynamics of Money and Output,” *Quarterly Journal of Economics* 114: 655-690 (1999).
- [26] Dupor, William , “Optimal Random Monetary Policy with Nominal Rigidity,” *Journal of Economic Theory* 112: 66-78 (2003).
- [27] Eggertsson, Gauti B., and Michael Woodford, “The Zero Bound on Interest Rates and Optimal Monetary Policy,” *Brookings Papers on Economic Activity* 2003:1, pp. 139-211.
- [28] Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin, “Optimal Monetary Policy with Staggered Wage and Price Contracts,” *Journal of Monetary Economics* 46: 281-313 (2000).
- [29] Evans, George W., and Seppo Honkapohja, “Monetary Policy, Expectations and Commitment,” *Scandinavian Journal of Economics* 108: 15-38 (2006).
- [30] Evans, George W., and Seppo Honkapohja, “Expectations, Learning and Monetary Policy: An Overview of Recent Research,” K. Schmitt-Hebbel and C.E. Walsh, *Monetary Policy Under Uncertainty and Learning*, Santiago, Chile: Central Bank of Chile, 2009.

- [31] Fuhrer, Jeffrey C., “Inflation Persistence,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, volume 3, Amsterdam: North-Holland, forthcoming 2010.
- [32] Gali, Jordi, “Monetary Policy and Unemployment,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, volume 3, Amsterdam: North-Holland, forthcoming 2010.
- [33] Gaspar, Vitor, Frank R. Smets, and David Vestin, “Inflation Expectations, Adaptive Learning, and Optimal Monetary Policy,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, volume 3, Amsterdam: North-Holland, forthcoming 2010.
- [34] Giannoni, Marc P., and Michael Woodford, “Optimal Inflation Targeting Rules,” in B.S. Bernanke and M. Woodford, eds., *The Inflation Targeting Debate*, Chicago: University of Chicago Press, 2005.
- [35] Giannoni, Marc P., and Michael Woodford, “Optimal Target Criteria for Stabilization Policy,” NBER Working Paper no. 15757, February 2010.
- [36] Goodfriend, Marvin, and Robert G. King, “The New Neoclassical Synthesis and the Role of Monetary Policy,” *NBER Macroeconomics Annual* 12: 231-283 (1997).
- [37] Gorodnichenko, Yuriy, and Matthew D. Shapiro, “Monetary Policy when Potential Output is Uncertain: Understanding the Growth Gamble of the 1990s,” NBER working paper no. 12268, May 2006.
- [38] 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.
- [39] Hansen, Lars Peter, and Thomas J. Sargent, “Wanting Robustness in Macroeconomics,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, volume 3, Amsterdam: North-Holland, forthcoming 2010.
- [40] Jensen, Christian, and Bennett C. McCallum, “The Non-Optimality of Proposed Monetary Policy Rules Under Timeless-Perspective Commitment,” *Economics Letters* 77(2): 163-168 (2002).

- [41] Jung, Taehun, Yuki Teranishi, and Tsutomu Watanabe, “Optimal Monetary Policy at the Zero-Interest-Rate Bound,” *Journal of Money, Credit and Banking* 37: 813-835 (2005).
- [42] Khan, Aubhik, Robert G. King, and Alexander L. Wolman, “Optimal Monetary Policy,” *Review of Economic Studies* 70: 825-860 (2003).
- [43] Kim, Jinill, and Sunghyun Kim , “Spurious Welfare Reversal in International Business Cycle Models”, *Journal of International Economics* 60: 471-500 (2003).
- [44] Kirsanova, Tatiana, David Vines, and Simon Wren-Lewis, “Inflation Bias with Dynamic Phillips Curves and Impatient Policy Makers,” *The B.E. Journal of Macroeconomics* 9(1: Topics), Article 32 (2009).
- [45] Kitamura, Tomiyuki, “Optimal Monetary Policy under Sticky Prices and Sticky Information,” unpublished, Ohio State University, March 2008.
- [46] Koenig, Evan, “Optimal Monetary Policy in Economies with ‘Sticky-Information’ Wages,” Federal Reserve Bank of Dallas Research Department working paper no. 0405, December 2004.
- [47] Levin, Andrew, David López-Salido, Edward Nelson, and Tack Yun, “Limitations on the Effectiveness of Forward Guidance at the Zero Lower Bound,” *International Journal of Central Banking* 6: 143-189 (2010).
- [48] Mankiw, N. Gregory, and Ricardo Reis, “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics* 117: 1295-1328 (2002).
- [49] Nakamura, Emi, and Jón Steinsson, “Monetary Non-Neutrality in a Multi-Sector Menu Cost Model,” *Quarterly Journal of Economics*, forthcoming 2010.
- [50] Preston, Bruce, “Adaptive Learning and the Use of Forecasts in Monetary Policy,” *Journal of Economic Dynamics and Control* 32: 3661-3681 (2008).
- [51] Rawls, John, *A Theory of Justice*, Cambridge, MA: Harvard University Press, 1971.

- [52] Reis, Ricardo, “Inattentive Producers,” *Review of Economic Studies* 73: 793-821 (2006).
- [53] Rotemberg, Julio J., and Michael Woodford, “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy,” *NBER Macroeconomics Annual 1997*, pp. 297-346.
- [54] Schmitt-Grohé, Stephanie, and Martin Uribe, “Optimal Operational Monetary Policy in the Christiano-Eichenbaum-Evans Model of the U.S. Business Cycle,” unpublished, Duke University, August 2004.
- [55] Schmitt-Grohé, Stephanie, and Martin Uribe, “Optimal Inflation Stabilization in a Medium-Scale Macroeconomic Model,” in F.S. Mishkin and K. Schmidt-Hebbel, eds., *Monetary Policy under Inflation Targeting*, Santiago, Chile: Central Bank of Chile, 2007.
- [56] Schmitt-Grohé, Stephanie, and Martin Uribe, “The Optimal Rate of Inflation,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, volume 3, Amsterdam: North-Holland, forthcoming 2010.
- [57] Sheedy, Kevin D., “Intrinsic Inflation Persistence,” London School of Economics, Centre for Economic Performance discussion paper no. 837, November 2007.
- [58] Sheedy, Kevin D., “Robustly Optimal Monetary Policy,” unpublished, London School of Economics, May 2008.
- [59] Smets, Frank, and Raf Wouters, “Shocks and Frictions in US Business Cycles,” *American Economic Review* 97: 586-606 (2007).
- [60] Sugo, Tomohiro, and Yuki Teranishi, “Optimal Monetary Policy Rule under the Non-Negativity Constraint on Nominal Interest Rates,” *Economics Letters* 89: 95-100 (2005).
- [61] Svensson, Lars E.O., “Inflation Forecast Targeting: Implementing and Monitoring Inflation Targeting,” *European Economic Review* 41: 1111-1146 (1997).
- [62] Svensson, Lars E.O., “What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules,” *Journal of Economic Literature* 41: 426-477 (2003).

- [63] Svensson, Lars E.O., “Monetary Policy with Judgment: Forecast Targeting,” *International Journal of Central Banking* 1: 1-54 (2005).
- [64] Svensson, Lars E.O., “Inflation Targeting,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, volume 3, Amsterdam: North-Holland, forthcoming 2010.
- [65] Svensson, Lars E.O., and Michael Woodford, “Optimal Policy with Partial Information in a Forward-Looking Model: Certainty-Equivalence Redux,” NBER Working Paper no. 9430, January 2003.
- [66] Svensson, Lars E.O., and Michael Woodford, “Indicator Variables for Optimal Policy Under Asymmetric Information,” *Journal of Economic Dynamics and Control* 28: 661-690 (2004).
- [67] Svensson, Lars E.O., and Michael Woodford, “Implementing Optimal Policy through Inflation-Forecast Targeting,” in B.S. Bernanke and M. Woodford, eds., *The Inflation Targeting Debate*, Chicago: University of Chicago Press, 2005.
- [68] Taylor, John B., and John C. Williams, “Monetary Policy Rules,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, volume 3, Amsterdam: North-Holland, forthcoming 2010.
- [69] Wolman, Alexander L., “Sticky Prices, Marginal Cost, and the Behavior of Inflation,” Federal Reserve Bank of Richmond *Economic Quarterly* 85(4): 29-48 (1999).
- [70] Woodford, Michael, “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, 1999.
- [71] Woodford, Michael, “Fiscal Requirements for Price Stability,” *Journal of Money, Credit and Banking* 33: 669-728 (2001).
- [72] Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, 2003.
- [73] Woodford, Michael, “Forecast Targeting as a Monetary Policy Strategy: Policy Rules in Practice,” NBER working paper no. 13716, December 2007.

- [74] Woodford, Michael, “Robustly Optimal Monetary Policy under Near-Rational Expectations,” *American Economic Review* 100: 274-303 (2010).
- [75] Yun, Tack, “Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles,” *Journal of Monetary Economics* 37: 345-370 (1996).
- [76] Yun, Tack, “Optimal Monetary Policy with Relative Price Distortions,” *American Economic Review* 95: 89-109 (2005).