

Chapter 7

DYNAMIC STOCHASTIC GENERAL-EQUILIBRIUM MODELS OF FLUCTUATIONS

Our analysis of macroeconomic fluctuations in the previous two chapters has developed two very incomplete pieces. In Chapter 5, we considered a full intertemporal macroeconomic model built from microeconomic foundations with explicit assumptions about the behavior of the underlying shocks. The model generated quantitative predictions about fluctuations, and is therefore an example of a quantitative *dynamic stochastic general-equilibrium*, or *DSGE*, model. The problem is that, as we saw in Section 5.10, the model appears to be an empirical failure. For example, it implies that monetary disturbances do not have real effects; it rests on large aggregate technology shocks for which there is little evidence; and its predictions about the effects of technology shocks and about business-cycle dynamics appear to be far from what we observe.

To address the real effects of monetary shocks, Chapter 6 introduced nominal rigidity. It established that barriers to price adjustment and other nominal frictions can cause monetary changes to have real effects, analyzed some of the determinants of the magnitude of those effects, and showed how nominal rigidity has important implications for the impacts of other disturbances. But it did so at the cost of abandoning most of the richness of the model of Chapter 5. Its models are largely static models with one-time shocks; and to the extent their focus is on quantitative predictions at all, it is only on addressing broad questions, notably whether plausibly small barriers to price adjustment can lead to plausibly large effects of monetary disturbances.

Researchers' ultimate goal is to build a model of fluctuations that combines the strengths of the models of the previous two chapters. This chapter will not take us all the way to that goal, however. There are two reasons. First, there is no consensus about the ingredients that are critical to include in such a model. Second, the state-of-the-art models in this effort (for example, Erceg, Henderson, and Levin, 2000, Smets and Wouters, 2003, and Christiano, Eichenbaum, and Evans, 2005) are quite complicated. If there

were strong evidence that one of these models captured the essence of modern macroeconomic fluctuations, it would be worth covering in detail. But in the absence of such evidence, the models are best left for more specialized treatments.

Instead, the chapter moves us partway toward constructing a realistic DSGE model of fluctuations. The bulk of the chapter extends the analysis of the microeconomic foundations of incomplete nominal flexibility to dynamic settings. This material vividly illustrates the lack of consensus about how best to build a realistic dynamic model of fluctuations: counting generously, we will consider seven distinct models of dynamic price adjustment. As we will see, the models often have sharply different implications for the macroeconomic consequences of microeconomic frictions in price adjustment. This analysis shows the main issues in moving to dynamic models of price-setting and illustrates the list of ingredients to choose from, but it does not identify a specific “best practice” model.

The main nominal friction we considered in Chapter 6 was a fixed cost of changing prices, or menu cost. In considering dynamic models of price adjustment, it is therefore tempting to assume that the only nominal imperfection is that firms must pay a fixed cost each time they change their price. There are two reasons not to make this the only case we consider, however. First, it is complicated: analyzing models of dynamic optimization with fixed adjustment costs is technically challenging and only rarely leads to closed-form solutions. Second, the vision of price-setters constantly monitoring their prices and standing ready to change them at any moment subject only to an unchanging fixed cost may be missing something important. Many prices are reviewed on a schedule and are only rarely changed at other times. For example, many wages are reviewed annually; some union contracts specify wages over a three-year period; and many companies issue catalogues with prices that are in effect for six months or a year. Thus price changes are not purely *state dependent* (that is, triggered by developments within the economy, regardless of the time over which the developments have occurred); they are partly *time dependent* (that is, triggered by the passage of time).

Because time-dependent models are easier, we will start with them. Section 7.1 presents a common framework for all the models of this part of the chapter. Sections 7.2 through 7.4 then consider three baseline models of time-dependent price adjustment: the Fischer, or Fischer-Phelps-Taylor, model (Fischer, 1977; Phelps and Taylor, 1977); the Taylor model (Taylor, 1979); and the Calvo model (Calvo, 1983). All three models posit that prices (or wages) are set by multiperiod contracts or commitments. In each period, the contracts governing some fraction of prices expire and must be renewed; expiration is determined by the passage of time, not economic developments. The central result of the models is that multiperiod contracts lead to gradual adjustment of the price level to nominal disturbances. As a result, aggregate demand disturbances have persistent real effects.

The Taylor and Calvo models differ from the Fischer model in one important respect. The Fischer model assumes that prices are *predetermined* but not *fixed*. That is, when a multiperiod contract sets prices for several periods, it can specify a different price for each period. In the Taylor and Calvo models, in contrast, prices are fixed: a contract must specify the same price each period it is in effect.

The difference between the Taylor and Calvo models is smaller. In the Taylor model, opportunities to change prices arrive deterministically, and each price is in effect for the same number of periods. In the Calvo model, opportunities to change prices arrive randomly, and so the number of periods a price is in effect is stochastic. In keeping with the assumption of time-dependence rather than state-dependence, the stochastic process governing price changes operates independently of other factors affecting the economy. The qualitative implications of the Calvo model are the same as those of the Taylor model. Its appeal is that it yields simpler inflation dynamics than the Taylor model, and so is easier to embed in larger models.

Section 7.5 then turns to two baseline models of state-dependent price adjustment, the Caplin-Spulber and Danziger-Golosov-Lucas models (Caplin and Spulber, 1987; Danziger, 1999; Golosov and Lucas, 2007). In both, the only barrier to price adjustment is a constant fixed cost. There are two differences between the models. First, money growth is always positive in the Caplin-Spulber model, while the version of the Danziger-Golosov-Lucas model we will consider assumes no trend money growth. Second, the Caplin-Spulber model assumes no firm-specific shocks, while the Danziger-Golosov-Lucas model includes them. Both models deliver strong results about the effects of monetary disturbances, but for very different reasons.

After Section 7.6 examines some empirical evidence, Section 7.7 considers two more models of dynamic price adjustment: the Calvo-with-indexation model and the Mankiw-Reis model (Christiano, Eichenbaum, and Evans, 2005; Mankiw and Reis, 2002). These models are more complicated than the models of the earlier sections, but appear to have more hope of fitting key facts about inflation dynamics.

The final two sections begin to consider how dynamic models of price adjustment can be embedded in models of the business cycle. Section 7.8 presents an example of a complete DSGE model with nominal rigidity. The model is the canonical three-equation new Keynesian model of Clarida, Galí, and Gertler (2000). Unfortunately, in many ways this model is closer to the baseline real-business-cycle model than to our ultimate objective: much of the model's appeal is tractability and elegance, not realism. Section 7.9 therefore discusses elements of other DSGE models with monetary non-neutrality. Because of the models' complexity and the lack of agreement about their key ingredients, however, it stops short of analyzing other fully specified models.

Before proceeding, it is important to emphasize that the issue we are interested in is incomplete adjustment of *nominal* prices and wages. There are

many reasons—involving uncertainty, information and renegotiation costs, incentives, and so on—that prices and wages may not adjust freely to equate supply and demand, or that firms may not change their prices and wages completely and immediately in response to shocks. But simply introducing some departure from perfect markets is not enough to imply that nominal disturbances matter. All the models of unemployment in Chapter 10, for example, are real models. If one appends a monetary sector to those models without any further complications, the classical dichotomy continues to hold: monetary disturbances cause all nominal prices and wages to change, leaving the real equilibrium (with whatever non-Walrasian features it involves) unchanged. Any microeconomic basis for failure of the classical dichotomy requires some kind of *nominal* imperfection.

7.1 Building Blocks of Dynamic New Keynesian Models

Overview

We will analyze the various models of dynamic price adjustment in a common framework. The framework draws heavily on the model of exogenous nominal rigidity in Section 6.1 and the model of imperfect competition in Section 6.5.

Time is discrete. Each period, imperfectly competitive firms produce output using labor as their only input. As in Section 6.5, the production function is one-for-one; thus aggregate output and aggregate labor input are equal. The model omits the government and international trade; thus, as in the models of Chapter 6, aggregate consumption and aggregate output are equal.

For simplicity, for the most part we will neglect uncertainty. Households maximize utility, taking the paths of the real wage and the real interest rate as given. Firms, which are owned by the households, maximize the present discounted value of their profits, subject to constraints on their price-setting (which vary across the models we will consider). Finally, a central bank determines the path of the real interest rate through its conduct of monetary policy.

Households

There is a fixed number of infinitely lived households that obtain utility from consumption and disutility from working. The representative household's objective function is

$$\sum_{t=0}^{\infty} \beta^t [U(C_t) - V(L_t)], \quad 0 < \beta < 1. \quad (7.1)$$

As in Section 6.5, C is a consumption index that is a constant-elasticity-of-substitution combination of the household's consumption of the individual goods, with elasticity of substitution $\eta > 1$. We make our usual assumptions about the functional forms of $U(\bullet)$ and $V(\bullet)$:¹

$$U(C_t) = \frac{C_t^{1-\theta}}{1-\theta}, \quad \theta > 0, \quad (7.2)$$

$$V(L_t) = \frac{B}{\gamma} L_t^\gamma, \quad B > 0, \quad \gamma > 1. \quad (7.3)$$

Let W denote the nominal wage and P denote the price level. Formally, P is the price index corresponding to the consumption index, as in Section 6.5. Throughout this chapter, however, we use the approximation we used in the Lucas model in Section 6.9 that the log of the price index, which we will denote p , is simply the average of firms' log prices.

An increase in labor supply in period t of amount dL increases the household's real income by $(W_t/P_t) dL$. The first-order condition for labor supply in period t is therefore

$$V'(L_t) = U'(C_t) \frac{W_t}{P_t}. \quad (7.4)$$

Because the production function is one-for-one and the only possible use of output is for consumption, in equilibrium C_t and L_t must both equal Y_t . Combining this fact with (7.4) tells us what the real wage must be given the level of output:

$$\frac{W_t}{P_t} = \frac{V'(Y_t)}{U'(Y_t)}. \quad (7.5)$$

Substituting the functional forms in (7.2)–(7.3) into (7.5) and solving for the real wage yields

$$\frac{W_t}{P_t} = B Y_t^{\theta+\gamma-1}. \quad (7.6)$$

Equation (7.6) is similar to equation (6.56) in the model of Section 6.5.

Since we are making the same assumptions about consumption as before, the new Keynesian IS curve holds in this model (see equation [6.8]):

$$\ln Y_t = \ln Y_{t+1} - \frac{1}{\theta} r_t. \quad (7.7)$$

Firms

Firm i produces output in period t according to the production function $Y_{it} = L_{it}$, and, as in Section 6.5, faces demand function $Y_{it} = Y_t (P_{it}/P_t)^{-\eta}$. The

¹ The reason for introducing B in (7.3) will be apparent below.

firm's real profits in period t , R_t , are revenues minus costs:

$$\begin{aligned} R_t &= \left(\frac{P_{it}}{P_t}\right) Y_{it} - \left(\frac{W_t}{P_t}\right) Y_{it} \\ &= Y_t \left[\left(\frac{P_{it}}{P_t}\right)^{1-\eta} - \left(\frac{W_t}{P_t}\right) \left(\frac{P_{it}}{P_t}\right)^{-\eta} \right]. \end{aligned} \quad (7.8)$$

Consider the problem of the firm setting its price in some period, which we normalize to period 0. As emphasized above, we will consider various assumptions about price-setting, including ones that imply that the length of time a given price is in effect is random. Thus, let q_t denote the probability that the price the firm sets in period zero is in effect in period t . Since the firm's profits accrue to the households, it values the profits according to the utility they provide to households. The marginal utility of the representative household's consumption in period t relative to period 0 is $\beta^t U'(C_t)/U'(C_0)$; denote this quantity λ_t .

The firm therefore chooses its price in period 0, P_i , to maximize $\sum_{t=0}^{\infty} q_t \lambda_t R_t \equiv A$, where R_t is the firm's profits in period t if P_i is still in effect. Using equation (7.8) for R_t , we can write A as

$$A = \sum_{t=0}^{\infty} q_t \lambda_t Y_t \left[\left(\frac{P_i}{P_t}\right)^{1-\eta} - \left(\frac{W_t}{P_t}\right) \left(\frac{P_i}{P_t}\right)^{-\eta} \right]. \quad (7.9)$$

One can say relatively little about the P_i that maximizes A in the general case. Two assumptions allow us to make progress, however. The first, and most important, is that inflation is low and that the economy is always close to its flexible-price equilibrium. The other is that households' discount factor, β , is close to 1.

To see the usefulness of these assumptions, rewrite (7.9) as

$$A = \sum_{t=0}^{\infty} q_t \lambda_t Y_t P_t^{\eta-1} (P_i^{1-\eta} - W_t P_i^{-\eta}). \quad (7.10)$$

The production function implies that marginal cost is constant and equal to W_t , and the elasticity of demand for the firm's good is constant. Thus the price that maximizes profits in period t , which we denote P_t^* , is a constant times W_t (see equation [6.55]). Equivalently, W_t is a constant times P_t^* . Thus we can write the expression in parentheses in (7.10) as a function of just P_i and P_t^* . As before, we will end up working with variables expressed in logs rather than levels. Thus, rewrite (7.10) as

$$A = \sum_{t=0}^{\infty} q_t \lambda_t Y_t P_t^{\eta-1} F(p_i, p_t^*), \quad (7.11)$$

where p_i and p_t^* denote the logs of P_i and P_t^* .

Our simplifying assumptions have two important implications about (7.11). The first is that the variation in $\lambda_t Y_t P_t^{\eta-1}$ is negligible relative to the

variation in q_t and p_t^* . The second is that $F(\bullet)$ can be well approximated by a second-order approximation around $p_i = p_t^*$.² Period- t profits are maximized at $p_i = p_t^*$; thus at $p_i = p_t^*$, $\partial F(p_i, p_t^*)/\partial p_i$ is zero and $\partial^2 F(p_i, p_t^*)/\partial p_i^2$ is negative. It follows that

$$F(p_i, p_t^*) \simeq F(p_t^*, p_t^*) - K(p_i - p_t^*)^2, \quad K > 0. \quad (7.12)$$

This analysis implies that the problem of choosing P_i to maximize A can be simplified to the problem,

$$\min_{p_i} \sum_{t=0}^{\infty} q_t (p_i - p_t^*)^2. \quad (7.13)$$

Finding the first-order condition for p_i and rearranging gives us

$$p_i = \sum_{t=0}^{\infty} \omega_t p_t^*, \quad (7.14)$$

where $\omega_t \equiv q_t / \sum_{\tau=0}^{\infty} q_\tau$. ω_t is the probability that the price the firm sets in period 0 will be in effect in period t divided by the expected number of periods the price will be in effect. Thus it measures the importance of period t to the choice of p_i . Equation (7.14) states that the price firm i sets is a weighted average of the profit-maximizing prices during the time the price will be in effect.

Finally, paralleling our assumption of certainty equivalence in the Lucas model in Section 6.9, we assume that when there is uncertainty, firms base their prices on expectations of the p_t^* 's:

$$p_i = \sum_{t=0}^{\infty} \omega_t E_0[p_t^*], \quad (7.15)$$

where $E_0[\bullet]$ denotes expectations as of period 0. Again, (7.15) is a legitimate approximation under appropriate assumptions.

A firm's profit-maximizing real price, P^*/P , is $\eta/(\eta - 1)$ times the real wage, W/P . And we know from equation (7.6) that w_t equals $p_t + b + (\theta + \gamma - 1)y_t$ (where $b \equiv \ln B$, $w_t \equiv \ln W_t$, and $y_t \equiv \ln Y_t$). Thus, the profit-maximizing price is

$$p^* = p + \ln[\eta/(\eta - 1)] + b + (\theta + \gamma - 1)y. \quad (7.16)$$

Note that (7.16) is of the form $p^* = p + c + \phi y$, $\phi > 0$, of the static model of Section 6.5 (see [6.58]). To simplify this, let m denote log nominal GDP, $p + y$, define $\phi \equiv \theta + \gamma - 1$, and assume $\ln[\eta/(\eta - 1)] + b = 0$ for simplicity.³ This yields

$$p_t^* = \phi m_t + (1 - \phi)p_t. \quad (7.17)$$

² These claims can be made precise with appropriate formalizations of the statements that inflation is small, the economy is near its flexible-price equilibrium, and β is close to 1.

³ It was for this reason that we introduced B in (7.3).

Substituting this expression into (7.15) gives us

$$p_i = \sum_{t=0}^{\infty} \omega_t E_0[\phi m_t + (1 - \phi)p_t]. \quad (7.18)$$

The Central Bank

Equation (7.18) is the key equation of the aggregate supply side of the model, and equation (7.7) describes aggregate demand for a given real interest rate. It remains to describe the determination of the real interest rate. To do this, we need to bring monetary policy into the model.

One approach, along the lines of Section 6.4, is to assume that the central bank follows some rule for how it sets the real interest rate as a function of macroeconomic conditions. This is the approach we will use in Section 7.8 and in much of Chapter 11. Our interest here, however, is in the aggregate supply side of the economy. Thus, along the lines of what we did in Part B of Chapter 6, we will follow the simpler approach of taking the path of nominal GDP (that is, the path of m_t) as given. We will then examine the behavior of the economy in response to various paths of nominal GDP, such as a one-time, permanent increase in its level or a permanent increase in its growth rate. As described in Section 6.5, a simple interpretation of the assumption that the path of nominal GDP is given is that the central bank has a target path of nominal GDP and conducts monetary policy to achieve it. This approach allows us to suppress not only the money market, but also the *IS* equation, (7.7).

7.2 Predetermined Prices: The Fischer Model

Framework and Assumptions

We now turn to the Fischer model of staggered price adjustment.⁴ The model follows the framework of the previous section. Price-setting is assumed to take a particular form, however: each price-setter sets prices every other period for the next two periods. And as emphasized above, the model assumes that the price-setter can set different prices for the two periods. That is, a

⁴ The original versions of the Fischer and Taylor models focused on staggered adjustment of wages; prices were in principle flexible but were determined as markups over wages. For simplicity, we assume instead that staggered adjustment applies directly to prices. Staggered wage adjustment has qualitatively similar implications. The key difference is that the microeconomic determinants of the parameter ϕ in the equation for desired prices, (7.17), are different under staggered wage adjustment (Huang and Liu, 2002).

firm setting its price in period 0 sets one price for period 1 and one price for period 2. Since each price will be in effect for only one period, equation (7.15) implies that each price (in logs) equals the expectation as of period 0 of the profit-maximizing price for that period. In any given period, half of price-setters are setting their prices for the next two periods. Thus at any point, half of the prices in effect are those set the previous period, and half are those set two periods ago.

No specific assumptions are made about the process followed by aggregate demand. For example, information about m_t may be revealed gradually in the periods leading up to t ; the expectation of m_t as of period $t-1$, $E_{t-1}m_t$, may therefore differ from the expectation of m_t the period before, $E_{t-2}m_t$.

Solving the Model

In any period, half of prices are ones set in the previous period, and half are ones set two periods ago. Thus the average price is

$$p_t = \frac{1}{2}(p_t^1 + p_t^2), \quad (7.19)$$

where p_t^1 denotes the price set for t by firms that set their prices in $t-1$, and p_t^2 the price set for t by firms that set their prices in $t-2$. Our assumptions about pricing from the previous section imply that p_t^1 equals the expectation as of period $t-1$ of p_{it}^* , and p_t^2 equals the expectation as of $t-2$ of p_{it}^* . Equation (7.17) therefore implies

$$\begin{aligned} p_t^1 &= E_{t-1}[\phi m_t + (1-\phi)p_t] \\ &= \phi E_{t-1}m_t + (1-\phi)\frac{1}{2}(p_t^1 + p_t^2), \end{aligned} \quad (7.20)$$

$$\begin{aligned} p_t^2 &= E_{t-2}[\phi m_t + (1-\phi)p_t] \\ &= \phi E_{t-2}m_t + (1-\phi)\frac{1}{2}(E_{t-2}p_t^1 + p_t^2), \end{aligned} \quad (7.21)$$

where $E_{t-\tau}$ denotes expectations conditional on information available through period $t-\tau$. Equation (7.20) uses the fact that p_t^2 is already determined when p_t^1 is set, and thus is not uncertain.

Our goal is to find how the price level and output evolve over time, given the behavior of m . To do this, we begin by solving (7.20) for p_t^1 ; this yields

$$p_t^1 = \frac{2\phi}{1+\phi}E_{t-1}m_t + \frac{1-\phi}{1+\phi}p_t^2. \quad (7.22)$$

Since the left- and right-hand sides of (7.22) are equal, the expectation as of $t-2$ of the two sides must be equal. Thus,

$$E_{t-2}p_t^1 = \frac{2\phi}{1+\phi}E_{t-2}m_t + \frac{1-\phi}{1+\phi}p_t^2, \quad (7.23)$$

where we have used the law of iterated projections to substitute $E_{t-2}m_t$ for $E_{t-2}E_{t-1}m_t$.

We can substitute (7.23) into (7.21) to obtain

$$p_t^2 = \phi E_{t-2}m_t + (1 - \phi) \frac{1}{2} \left(\frac{2\phi}{1 + \phi} E_{t-2}m_t + \frac{1 - \phi}{1 + \phi} p_t^2 + p_t^2 \right). \quad (7.24)$$

Solving this expression for p_t^2 yields simply

$$p_t^2 = E_{t-2}m_t. \quad (7.25)$$

We can now combine the results and describe the equilibrium. Substituting (7.25) into (7.22) and simplifying gives

$$p_t^1 = E_{t-2}m_t + \frac{2\phi}{1 + \phi} (E_{t-1}m_t - E_{t-2}m_t). \quad (7.26)$$

Finally, substituting (7.25) and (7.26) into the expressions for the price level and output, $p_t = (p_t^1 + p_t^2)/2$ and $y_t = m_t - p_t$, implies

$$p_t = E_{t-2}m_t + \frac{\phi}{1 + \phi} (E_{t-1}m_t - E_{t-2}m_t), \quad (7.27)$$

$$y_t = \frac{1}{1 + \phi} (E_{t-1}m_t - E_{t-2}m_t) + (m_t - E_{t-1}m_t). \quad (7.28)$$

Implications

Equation (7.28) shows the model's main implications. First, unanticipated aggregate demand shifts have real effects; this is shown by the $m_t - E_{t-1}m_t$ term. Because price-setters are assumed not to know m_t when they set their prices, these shocks are passed one-for-one into output.

Second, aggregate demand shifts that become anticipated after the first prices are set affect output. Consider information about aggregate demand in t that becomes available between period $t - 2$ and period $t - 1$. In practice, this might correspond to the release of survey results or other leading indicators of future economic activity, or to indications of likely shifts in monetary policy. As (7.27) and (7.28) show, proportion $1/(1 + \phi)$ of information about m_t that arrives between $t - 2$ and $t - 1$ is passed into output, and the remainder goes into prices. The reason that the change is not neutral is straightforward: not all prices are completely flexible in the short run.

One implication of these results is that interactions among price-setters can either increase or decrease the effects of microeconomic price stickiness. One might expect that since half of prices are already set and the other half are free to adjust, half of the information about m_t that arrives between $t - 2$ and $t - 1$ is passed into prices and half into output. But in general this is not correct. The key parameter is ϕ : the proportion of the shift that is passed into output is not $\frac{1}{2}$ but $1/(1 + \phi)$ (see [7.28]).

Recall that ϕ measures the degree of real rigidity: ϕ is the responsiveness of price-setters' desired real prices to aggregate real output, and so a smaller value of ϕ corresponds to greater real rigidity. When real rigidity is large, price-setters are reluctant to allow variations in their relative prices. As a result, the price-setters that are free to adjust their prices do not allow their prices to differ greatly from the ones already set, and so the real effects of a monetary shock are large. If ϕ exceeds 1, in contrast, the later price-setters make large price changes, and the aggregate real effects of changes in m are small.⁵

Finally, and importantly, the model implies that output does not depend on $E_{t-2}m_t$ (given the values of $E_{t-1}m_t - E_{t-2}m_t$ and $m_t - E_{t-1}m_t$). That is, any information about aggregate demand that all price-setters have had a chance to respond to has no effect on output. Thus the model does not provide an explanation of persistent effects of movements in aggregate demand. We will return to this issue in Section 7.7.

7.3 Fixed Prices: The Taylor Model

The Model

We now change the model of the previous section by assuming that when a firm sets prices for two periods, it must set the same price for both periods. In the terminology introduced earlier, prices are not just predetermined, but fixed.

We make two other, less significant changes to the model. First, a firm setting a price in period t now does so for periods t and $t + 1$ rather than for periods $t + 1$ and $t + 2$. This change simplifies the model without affecting the main results. Second, the model is much easier to solve if we posit a specific process for m . A simple assumption is that m is a random walk:

$$m_t = m_{t-1} + u_t, \quad (7.29)$$

where u is white noise. The key feature of this process is that an innovation to m (the u term) has a long-lasting effect on its level.

Let x_t denote the price chosen by firms that set their prices in period t . Here equation (7.18) for price-setting implies

$$\begin{aligned} x_t &= \frac{1}{2}(p_{it}^* + E_t p_{it+1}^*) \\ &= \frac{1}{2}\{[\phi m_t + (1 - \phi)p_t] + [\phi E_t m_{t+1} + (1 - \phi)E_t p_{t+1}]\}, \end{aligned} \quad (7.30)$$

where the second line uses the fact that $p^* = \phi m + (1 - \phi)p$.

Since half of prices are set each period, p_t is the average of x_t and x_{t-1} . In addition, since m is a random walk, $E_t m_{t+1}$ equals m_t . Substituting these

⁵ Haltiwanger and Waldman (1989) show more generally how a small fraction of agents who do not respond to shocks can have a disproportionate effect on the economy.

facts into (7.30) gives us

$$x_t = \phi m_t + \frac{1}{4}(1 - \phi)(x_{t-1} + 2x_t + E_t x_{t+1}). \quad (7.31)$$

Solving for x_t yields

$$x_t = A(x_{t-1} + E_t x_{t+1}) + (1 - 2A)m_t, \quad A \equiv \frac{1}{2} \frac{1 - \phi}{1 + \phi}. \quad (7.32)$$

Equation (7.32) is the key equation of the model.

Equation (7.32) expresses x_t in terms of m_t , x_{t-1} , and the expectation of x_{t+1} . To solve the model, we need to eliminate the expectation of x_{t+1} from this expression. We will solve the model in two different ways, first using the method of undetermined coefficients and then using *lag operators*. The method of undetermined coefficients is simpler. But there are cases where it is cumbersome or intractable; in those cases the use of lag operators is often fruitful.

The Method of Undetermined Coefficients

As described in Section 5.6, the idea of the method of undetermined coefficients is to guess the general functional form of the solution and then to use the model to determine the precise coefficients. In the model we are considering, in period t two variables are given: the money stock, m_t , and the prices set the previous period, x_{t-1} . In addition, the model is linear. It is therefore reasonable to guess that x_t is a linear function of x_{t-1} and m_t :

$$x_t = \mu + \lambda x_{t-1} + \nu m_t. \quad (7.33)$$

Our goal is to determine whether there are values of μ , λ , and ν that yield a solution of the model.

Although we could now proceed to find μ , λ , and ν , it simplifies the algebra if we first use our knowledge of the model to restrict (7.33). We have normalized the constant in the expression for firms' desired prices to zero, so that $p_{it}^* = p_t + \phi y_t$. As a result, the equilibrium with flexible prices is for y to equal zero and for each price to equal m . In light of this, consider a situation where x_{t-1} and m_t are equal. If period- t price-setters also set their prices to m_t , the economy is at its flexible-price equilibrium. In addition, since m follows a random walk, the period- t price-setters have no reason to expect m_{t+1} to be on average either more or less than m_t , and hence no reason to expect x_{t+1} to depart on average from m_t . Thus in this situation p_{it}^* and $E_t p_{it+1}^*$ are both equal to m_t , and so price-setters will choose $x_t = m_t$. In sum, it is reasonable to guess that if $x_{t-1} = m_t$, then $x_t = m_t$. In terms of (7.33), this condition is

$$\mu + \lambda m_t + \nu m_t = m_t \quad (7.34)$$

for all m_t .

Two conditions are needed for (7.34) to hold. The first is $\lambda + \nu = 1$; otherwise (7.34) cannot be satisfied for all values of m_t . Second, when we impose $\lambda + \nu = 1$, (7.34) implies $\mu = 0$. Substituting these conditions into (7.33) yields

$$x_t = \lambda x_{t-1} + (1 - \lambda)m_t. \quad (7.35)$$

Our goal is now to find a value of λ that solves the model.

Since (7.35) holds each period, it implies $x_{t+1} = \lambda x_t + (1 - \lambda)m_{t+1}$. Thus the expectation as of period t of x_{t+1} is $\lambda x_t + (1 - \lambda)E_t m_{t+1}$, which equals $\lambda x_t + (1 - \lambda)m_t$. Using (7.35) to substitute for x_t then gives us

$$\begin{aligned} E_t x_{t+1} &= \lambda[\lambda x_{t-1} + (1 - \lambda)m_t] + (1 - \lambda)m_t \\ &= \lambda^2 x_{t-1} + (1 - \lambda^2)m_t. \end{aligned} \quad (7.36)$$

Substituting this expression into (7.32) yields

$$\begin{aligned} x_t &= A[x_{t-1} + \lambda^2 x_{t-1} + (1 - \lambda^2)m_t] + (1 - 2A)m_t \\ &= (A + A\lambda^2)x_{t-1} + [A(1 - \lambda^2) + (1 - 2A)]m_t. \end{aligned} \quad (7.37)$$

Thus, if price-setters believe that x_t is a linear function of x_{t-1} and m_t of the form assumed in (7.35), then, acting to maximize their profits, they will indeed set their prices as a linear function of these variables. If we have found a solution of the model, these two linear equations must be the same. Comparison of (7.35) and (7.37) shows that this requires

$$A + A\lambda^2 = \lambda \quad (7.38)$$

and

$$A(1 - \lambda^2) + (1 - 2A) = 1 - \lambda. \quad (7.39)$$

It is easy to show that (7.39) simplifies to (7.38). Thus we only need to consider (7.38). This is a quadratic in λ . The solution is

$$\lambda = \frac{1 \pm \sqrt{1 - 4A^2}}{2A}. \quad (7.40)$$

Using the definition of A in equation (7.32), one can show that the two values of λ are

$$\lambda_1 = \frac{1 - \sqrt{\phi}}{1 + \sqrt{\phi}}, \quad (7.41)$$

$$\lambda_2 = \frac{1 + \sqrt{\phi}}{1 - \sqrt{\phi}}. \quad (7.42)$$

Of the two values, only $\lambda = \lambda_1$ gives reasonable results. When $\lambda = \lambda_1$, $|\lambda| < 1$, and so the economy is stable. When $\lambda = \lambda_2$, in contrast, $|\lambda| > 1$,

and thus the economy is unstable: the slightest disturbance sends output off toward plus or minus infinity. As a result, the assumptions underlying the model—for example, that sellers do not ration buyers—break down. For that reason, we focus on $\lambda = \lambda_1$.

Thus equation (7.35) with $\lambda = \lambda_1$ solves the model: if price-setters believe that others are using that rule to set their prices, they find it in their own interests to use that same rule.

We can now describe the behavior of output. y_t equals $m_t - p_t$, which in turn equals $m_t - (x_{t-1} + x_t)/2$. With the behavior of x given by (7.35), this implies

$$\begin{aligned} y_t &= m_t - \frac{1}{2} \{ [\lambda x_{t-2} + (1 - \lambda)m_{t-1}] + [\lambda x_{t-1} + (1 - \lambda)m_t] \} \\ &= m_t - \left[\lambda \frac{1}{2} (x_{t-2} + x_{t-1}) + (1 - \lambda) \frac{1}{2} (m_{t-1} + m_t) \right]. \end{aligned} \quad (7.43)$$

Using the facts that $m_t = m_{t-1} + u_t$ and $(x_{t-1} + x_{t-2})/2 = p_{t-1}$, we can simplify this to

$$\begin{aligned} y_t &= m_{t-1} + u_t - \left[\lambda p_{t-1} + (1 - \lambda)m_{t-1} + (1 - \lambda) \frac{1}{2} u_t \right] \\ &= \lambda (m_{t-1} - p_{t-1}) + \frac{1 + \lambda}{2} u_t \\ &= \lambda y_{t-1} + \frac{1 + \lambda}{2} u_t. \end{aligned} \quad (7.44)$$

Implications

Equation (7.44) is the key result of the model. As long as λ_1 is positive (which is true if $\phi < 1$), (7.44) implies that shocks to aggregate demand have long-lasting effects on output—effects that persist even after all firms have changed their prices. Suppose the economy is initially at the equilibrium with flexible prices (so y is steady at 0), and consider the effects of a positive shock of size u^0 in some period. In the period of the shock, not all firms adjust their prices, and so not surprisingly, y rises; from (7.44), $y = [(1 + \lambda)/2]u^0$. In the following period, even though the remaining firms are able to adjust their prices, y does not return to normal even in the absence of a further shock: from (7.44), y is $\lambda[(1 + \lambda)/2]u^0$. Thereafter output returns slowly to normal, with $y_t = \lambda y_{t-1}$ each period.

The response of the price level to the shock is the flip side of the response of output. The price level rises by $[1 - (1 + \lambda)/2]u^0$ in the initial period, and then fraction $1 - \lambda$ of the remaining distance from u^0 in each subsequent period. Thus the economy exhibits price-level inertia.

The source of the long-lasting real effects of monetary shocks is again price-setters' reluctance to allow variations in their relative prices. Recall that $p_{it}^* = \phi m_t + (1 - \phi)p_t$, and that $\lambda_1 > 0$ only if $\phi < 1$. Thus there is gradual adjustment only if desired prices are an increasing function of the price level. Suppose each price-setter adjusted fully to the shock at the first

opportunity. In this case, the price-setters who adjusted their prices in the period of the shock would adjust by the full amount of the shock, and the remainder would do the same in the next period. Thus γ would rise by $u^0/2$ in the initial period and return to normal in the next.

To see why this rapid adjustment cannot be the equilibrium if ϕ is less than 1, consider the firms that adjust their prices immediately. By assumption, all prices have been adjusted by the second period, and so in that period each firm is charging its profit-maximizing price. But since $\phi < 1$, the profit-maximizing price is lower when the price level is lower, and so the price that is profit-maximizing in the period of the shock, when not all prices have been adjusted, is less than the profit-maximizing price in the next period. Thus these firms should not adjust their prices fully in the period of the shock. This in turn implies that it is not optimal for the remaining firms to adjust their prices fully in the subsequent period. And the knowledge that they will not do this further dampens the initial response of the firms that adjust their prices in the period of the shock. The end result of these forward- and backward-looking interactions is the gradual adjustment shown in equation (7.35).

Thus, as in the model with prices that are predetermined but not fixed, the extent of incomplete price adjustment in the aggregate can be larger than one might expect simply from the knowledge that not all prices are adjusted every period. Indeed, the extent of aggregate price sluggishness is even larger in this case, since it persists even after every price has changed. And again a low value of ϕ —that is, a high degree of real rigidity—is critical to this result. If ϕ is 1, then λ is 0, and so each price-setter adjusts his or her price fully to changes in m at the earliest opportunity. If ϕ exceeds 1, λ is negative, and so p moves by more than m in the period after the shock, and thereafter the adjustment toward the long-run equilibrium is oscillatory.

Lag Operators

A different, more general approach to solving the model is to use lag operators. The lag operator, which we denote by L , is a function that lags variables. That is, the lag operator applied to any variable gives the previous period's value of the variable: $LZ_t = Z_{t-1}$.

To see the usefulness of lag operators, consider our model without the restriction that m follows a random walk. Equation (7.30) continues to hold. If we proceed analogously to the derivation of (7.32), but without imposing $E_t m_{t+1} = m_t$, straightforward algebra yields

$$x_t = A(x_{t-1} + E_t x_{t+1}) + \frac{1-2A}{2} m_t + \frac{1-2A}{2} E_t m_{t+1}, \quad (7.45)$$

where A is as before. Note that (7.45) simplifies to (7.32) if $E_t m_{t+1} = m_t$.

The first step is to rewrite this expression using lag operators. x_{t-1} is the lag of x_t : $x_{t-1} = Lx_t$. In addition, if we adopt the rule that when L is applied to

an expression involving expectations, it lags the date of the variables but not the date of the expectations, then x_t is the lag of $E_t x_{t+1}$: $LE_t x_{t+1} = E_t x_t = x_t$.⁶ Equivalently, using L^{-1} to denote the inverse lag function, $E_t x_{t+1} = L^{-1} x_t$. Similarly, $E_t m_{t+1} = L^{-1} m_t$. Thus we can rewrite (7.45) as

$$x_t = A(Lx_t + L^{-1}x_t) + \frac{1-2A}{2}m_t + \frac{1-2A}{2}L^{-1}m_t, \quad (7.46)$$

or

$$(I - AL - AL^{-1})x_t = \frac{1-2A}{2}(I + L^{-1})m_t. \quad (7.47)$$

Here I is the identity operator (so $Iz_t = z_t$ for any z). Thus $(I + L^{-1})m_t$ is shorthand for $m_t + L^{-1}m_t$, and $(I - AL - AL^{-1})x_t$ is shorthand for $x_t - Ax_{t-1} - AE_t x_{t+1}$.

Now observe that we can “factor” $I - AL - AL^{-1}$ as $(I - \lambda L^{-1})(I - \lambda L)(A/\lambda)$, where λ is again given by (7.40). Thus we have

$$(I - \lambda L^{-1})(I - \lambda L)x_t = \frac{\lambda}{A} \frac{1-2A}{2}(I + L^{-1})m_t. \quad (7.48)$$

This formulation of “multiplying” expressions involving the lag operator should be interpreted in the natural way: $(I - \lambda L^{-1})(I - \lambda L)x_t$ is shorthand for $(I - \lambda L)x_t$ minus λ times the inverse lag operator applied to $(I - \lambda L)x_t$, and thus equals $(x_t - \lambda Lx_t) - (\lambda L^{-1}x_t - \lambda^2 x_t)$. Simple algebra and the definition of λ can be used to verify that (7.48) and (7.47) are equivalent.

As before, to solve the model we need to eliminate the term involving the expectation of the future value of an endogenous variable. In (7.48), $E_t x_{t+1}$ appears (implicitly) on the left-hand side because of the $I - \lambda L^{-1}$ term. It is thus natural to “divide” both sides by $I - \lambda L^{-1}$. That is, consider applying the operator $I + \lambda L^{-1} + \lambda^2 L^{-2} + \lambda^3 L^{-3} + \dots$ to both sides of (7.48). $I + \lambda L^{-1} + \lambda^2 L^{-2} + \dots$ times $I - \lambda L^{-1}$ is simply I ; thus the left-hand side is $(I - \lambda L)x_t$. And $I + \lambda L^{-1} + \lambda^2 L^{-2} + \dots$ times $I + L^{-1}$ is $I + (1 + \lambda)L^{-1} + (1 + \lambda)\lambda L^{-2} + (1 + \lambda)\lambda^2 L^{-3} + \dots$.⁷ Thus (7.48) becomes

$$\begin{aligned} (I - \lambda L)x_t \\ = \frac{\lambda}{A} \frac{1-2A}{2} [I + (1 + \lambda)L^{-1} + (1 + \lambda)\lambda L^{-2} + (1 + \lambda)\lambda^2 L^{-3} + \dots] m_t. \end{aligned} \quad (7.49)$$

⁶ Since $E_t x_{t-1} = x_{t-1}$ and $E_t m_t = m_t$, we can think of all the variables in (7.45) as being expectations as of t . Thus in the analysis that follows, the lag operator should always be interpreted as keeping all variables as expectations as of t . The *backshift operator*, B , lags both the date of the variable and the date of the expectations. Thus, for example, $BE_t x_{t+1} = E_{t-1} x_t$. Whether the lag operator or the backshift operator is more useful depends on the application.

⁷ Since the operator $I + \lambda L^{-1} + \lambda^2 L^{-2} + \dots$ is an infinite sum, this requires that $\lim_{n \rightarrow \infty} (I + \lambda L^{-1} + \lambda^2 L^{-2} + \dots + \lambda^n L^{-n})(I + L^{-1})m_t$ exists. This requires that $\lambda^n L^{-(n+1)}m_t$ (which equals $\lambda^n E_t m_{t+n+1}$) converges to 0. For the case where $\lambda = \lambda_1$ (so $|\lambda| < 1$) and where m is a random walk, this condition is satisfied.

Rewriting this expression without lag operators yields

$$x_t = \lambda x_{t-1} + \frac{\lambda}{A} \frac{1-2A}{2} [m_t + (1+\lambda)(E_t m_{t+1} + \lambda E_t m_{t+2} + \lambda^2 E_t m_{t+3} + \dots)]. \quad (7.50)$$

Expression (7.50) characterizes the behavior of newly set prices in terms of the exogenous money supply process. To find the behavior of the aggregate price level and output, we only have to substitute this expression into the expressions for p ($p_t = (x_t + x_{t-1})/2$) and y ($y_t = m_t - p_t$).

In the special case when m is a random walk, all the $E_t m_{t+i}$'s are equal to m_t . In this case, (7.50) simplifies to

$$x_t = \lambda x_{t-1} + \frac{\lambda}{A} \frac{1-2A}{2} \left(1 + \frac{1+\lambda}{1-\lambda}\right) m_t. \quad (7.51)$$

It is straightforward to show that expression (7.38), $A + A\lambda^2 = \lambda$, implies that equation (7.51) reduces to equation (7.35), $x_t = \lambda x_{t-1} + (1-\lambda)m_t$. Thus when m is a random walk, we obtain the same result as before. But we have also solved the model for a general process for m .

Although this use of lag operators may seem mysterious, in fact it is no more than a compact way of carrying out perfectly standard manipulations. We could have first derived (7.45) (expressed without using lag operators) by simple algebra. We could then have noted that since (7.45) holds at each date, it must be the case that

$$E_t x_{t+k} - A E_t x_{t+k-1} - A E_t x_{t+k+1} = \frac{1-2A}{2} (E_t m_{t+k} + E_t m_{t+k+1}) \quad (7.52)$$

for all $k \geq 0$.⁸ Since the left- and right-hand sides of (7.52) are equal, it must be the case that the left-hand side for $k = 0$ plus λ times the left-hand side for $k = 1$ plus λ^2 times the left-hand side for $k = 2$ and so on equals the right-hand side for $k = 0$ plus λ times the right-hand side for $k = 1$ plus λ^2 times the right-hand side for $k = 2$ and so on. Computing these two expressions yields (7.50). Thus lag operators are not essential; they serve merely to simplify the notation and to suggest ways of proceeding that might otherwise be missed.⁹

⁸ The reason that we cannot assume that (7.52) holds for $k < 0$ is that the law of iterated projections does not apply backward: the expectation today of the expectation at some date *in the past* of a variable need not equal the expectation today of the variable.

⁹ For a more thorough introduction to lag operators and their uses, see Sargent (1987, Chapter 9).

7.4 The Calvo Model and the New Keynesian Phillips Curve

Overview

In the Taylor model, each price is in effect for the same number of periods. One consequence is that moving beyond the two-period case quickly becomes intractable. The Calvo model (Calvo, 1983) is an elegant variation on the model that avoids this problem. Calvo assumes that price changes, rather than arriving deterministically, arrive stochastically. Specifically, he assumes that opportunities to change prices follow a *Poisson process*: the probability that a firm is able to change its price is the same each period, regardless of when it was last able to change its price. As in the Taylor model, prices are not just predetermined but fixed between the times they are adjusted.

This model's qualitative implications are similar to those of the Taylor model. Suppose, for example, the economy starts with all prices equal to the money stock, m , and that in period 1 there is a one-time, permanent increase in m . Firms that can adjust their prices will want to raise them in response to the rise in m . But if ϕ in the expression for the profit-maximizing price ($p_t^* = \phi m_t + (1 - \phi)p_t$) is less than 1, they put some weight on the overall price level, and so the fact that not all firms are able to adjust their prices mutes their adjustment. And the smaller is ϕ , the larger is this effect. Thus, just as in the Taylor model, nominal rigidity (the fact that not all prices adjust every period) leads to gradual adjustment of the price level, and real rigidity (a low value of ϕ) magnifies the effects of nominal rigidity.¹⁰

The importance of the Calvo model, then, is not in its qualitative predictions. Rather, it is twofold. First, the model can easily accommodate any degree of price stickiness; all one needs to do is change the parameter determining the probability that a firm is able to change its price each period. Second, it leads to a simple expression for the dynamics of inflation. That expression is known as the *new Keynesian Phillips curve*.

Deriving the New Keynesian Phillips Curve

Each period, fraction α ($0 < \alpha \leq 1$) of firms set new prices, with the firms chosen at random. The average price in period t therefore equals α times the price set by firms that set new prices in t , x_t , plus $1 - \alpha$ times the average price charged in t by firms that do not change their prices. Because the firms that change their prices are chosen at random (and because the number of

¹⁰ See Problem 7.6.

firms is large), the average price charged by the firms that do not change their prices equals the average price charged by all firms the previous period. Thus we have

$$p_t = \alpha x_t + (1 - \alpha) p_{t-1}, \quad (7.53)$$

where p is the average price and x is the price set by firms that are able to change their prices. Subtracting p_{t-1} from both sides gives us an expression for inflation:

$$\pi_t = \alpha(x_t - p_{t-1}). \quad (7.54)$$

That is, inflation is determined by the fraction of firms that change their prices and the relative price they set.

In deriving the rule in Section 7.1 for how a firm sets its price as a weighted average of the expected profit-maximizing prices while the price is in effect (equation [7.14]), we assumed the discount factor was approximately 1. For the Fisher and Taylor models, where prices are only in effect for two periods, this assumption simplified the analysis at little cost. But here, where firms need to look indefinitely into the future, it is not innocuous. Extending expression (7.14) to the case of a general discount factor implies

$$x_t = \sum_{j=0}^{\infty} \frac{\beta^j q_j}{\sum_{k=0}^{\infty} \beta^k q_k} E_t p_{t+j}^*, \quad (7.55)$$

where β is the discount factor and, as before, q_j is the probability the price will still be in effect in period $t+j$. Calvo's Poisson assumption implies that q_j is $(1 - \alpha)^j$. Thus (7.55) becomes

$$x_t = [1 - \beta(1 - \alpha)] \sum_{j=0}^{\infty} \beta^j (1 - \alpha)^j E_t p_{t+j}^*. \quad (7.56)$$

Firms that can set their prices in period $t+1$ face a very similar problem. Period t is no longer relevant, and all other periods get a proportionally higher weight. It therefore turns out to be helpful to express x_t in terms of p_t^* and $E_t x_{t+1}$. To do this, rewrite (7.56) as

$$\begin{aligned} x_t &= [1 - \beta(1 - \alpha)] E_t p_t^* + \beta(1 - \alpha) [1 - \beta(1 - \alpha)] \left[\sum_{j=0}^{\infty} \beta^j (1 - \alpha)^j E_t p_{t+1+j}^* \right] \\ &= [1 - \beta(1 - \alpha)] p_t^* + \beta(1 - \alpha) E_t x_{t+1}, \end{aligned} \quad (7.57)$$

where the second line uses the fact that p_t^* is known at time t and expression (7.56) shifted forward one period. To relate (7.57) to (7.54), subtract p_t from both sides of (7.57), and rewrite $x_t - p_t$ as $(x_t - p_{t-1}) - (p_t - p_{t-1})$.

This gives us

$$(x_t - p_{t-1}) - (p_t - p_{t-1}) = [1 - \beta(1 - \alpha)](p_t^* - p_t) + \beta(1 - \alpha)(E_t x_{t+1} - p_t). \quad (7.58)$$

We can now use (7.54): $x_t - p_{t-1}$ is π_t/α , and $E_t x_{t+1} - p_t$ is $E_t \pi_{t+1}/\alpha$. In addition, $p_t - p_{t-1}$ is just π_t , and $p_t^* - p_t$ is ϕy_t . Thus (7.58) becomes

$$(\pi_t/\alpha) - \pi_t = [1 - \beta(1 - \alpha)]\phi y_t + \beta(1 - \alpha)(E_t \pi_{t+1}/\alpha), \quad (7.59)$$

or

$$\begin{aligned} \pi_t &= \frac{\alpha}{1 - \alpha} [1 - \beta(1 - \alpha)]\phi y_t + \beta E_t \pi_{t+1} \\ &= \kappa y_t + \beta E_t \pi_{t+1}, \quad \kappa \equiv \frac{\alpha[1 - (1 - \alpha)\beta]\phi}{1 - \alpha}. \end{aligned} \quad (7.60)$$

Discussion

Equation (7.60) is the new Keynesian Phillips curve.¹¹ Like the accelerationist Phillips curve of Section 6.4 and the Lucas supply curve of Section 6.9, it states that inflation depends on a core or expected inflation term and on output. Higher output raises inflation, as does higher core or expected inflation.

There are two features of this Phillips curve that make it “new.” First, it is derived by aggregating the behavior of price-setters facing barriers to price adjustment. Second, the inflation term on the right-hand side is different from previous Phillips curves. In the accelerationist Phillips curve, it is last period’s inflation. In the Lucas supply curve, it is the expectation of current inflation. Here it is the current expectation of next period’s inflation. These differences are important—a point we will return to in Section 7.6.

Although the Calvo model leads to a particularly elegant expression for inflation, its broad implications stem from the general assumption of staggered price adjustment, not the specific Poisson assumption. For example, one can show that the basic equation for pricing-setting in the Taylor model, $x_t = (p_{it}^* + E_t p_{it+1}^*)/2$ (equation [7.30]) implies

$$\pi_t^x = E_t \pi_{t+1}^x + 2\phi(y_t + E_t y_{t+1}), \quad (7.61)$$

where π^x is the growth rate of newly set prices. Although (7.61) is not as simple as (7.60), its basic message is the same: a measure of inflation depends on a measure of expected future inflation and expectations of output.

¹¹ The new Keynesian Phillips curve was originally derived by Roberts (1995).

7.5 State-Dependent Pricing

The Fischer, Taylor, and Calvo models assume that the timing of price changes is purely time dependent. The other extreme is that it is purely state dependent. Many retail stores, for example, can adjust the timing of their price change fairly freely in response to economic developments. This section therefore considers state-dependent pricing.

The basic message of analyses of state-dependent pricing is that it leads to more rapid adjustment of the overall price level to macroeconomic disturbances for a given average frequency of price changes. There are two distinct reasons for this result. The first is the *frequency effect*: under state-dependent pricing, the number of firms that change their prices is larger when there is a larger monetary shock. The other is the *selection effect*: the composition of the firms that adjust their prices changes in response to a shock. In this section, we consider models that illustrate each effect.

The Frequency Effect: The Caplin-Spulber Model

Our first model is the Caplin-Spulber model. The model is set in continuous time. Nominal GDP is always growing; coupled with the assumption that there are no firm-specific shocks, this causes profit-maximizing prices to always be increasing. The specific state-dependent pricing rule that price-setters are assumed to follow is an *Ss policy*. That is, whenever a firm adjusts its price, it sets the price so that the difference between the actual price and the optimal price at that time, $p_i - p_i^*$, equals some target level, S . The firm then keeps its nominal price fixed until money growth has raised p_i^* sufficiently that $p_i - p_i^*$ has fallen to some trigger level, s . Then, regardless of how much time has passed since it last changed its price, the firm resets $p_i - p_i^*$ to S , and the process begins anew.

Such an *Ss policy* is optimal when inflation is steady, aggregate output is constant, and there is a fixed cost of each nominal price change (Barro, 1972; Sheshinski and Weiss, 1977). In addition, as Caplin and Spulber describe, it is also optimal in some cases where inflation or output is not constant. And even when it is not fully optimal, it provides a simple and tractable example of state-dependent pricing.

Two technical assumptions complete the model. First, to keep prices from overshooting s and to prevent bunching of the distribution of prices across price-setters, m changes continuously. Second, the initial distribution of $p_i - p_i^*$ across price-setters is uniform between s and S . We continue to use the assumptions of Section 7.1 that $p_i^* = (1 - \phi)p + \phi m$, p is the average of the p_i 's, and $y = m - p$.

Under these assumptions, shifts in aggregate demand are completely neutral in the aggregate despite the price stickiness at the level of the individual price-setters. To see this, consider an increase in m of amount

$\Delta m < S - s$ over some period of time. We want to find the resulting changes in the price level and output, Δp and Δy . Since $p_i^* = (1 - \phi)p + \phi m$, the rise in each firm's profit-maximizing price is $(1 - \phi)\Delta p + \phi\Delta m$. Firms change their prices if $p_i - p_i^*$ falls below s ; thus firms with initial values of $p_i - p_i^*$ that are less than $s + [(1 - \phi)\Delta p + \phi\Delta m]$ change their prices. Since the initial values of $p_i - p_i^*$ are distributed uniformly between s and S , this means that the fraction of firms that change their prices is $[(1 - \phi)\Delta p + \phi\Delta m]/(S - s)$. Each firm that changes its price does so at the moment when its value of $p_i - p_i^*$ reaches s ; thus each price increase is of amount $S - s$. Putting all this together gives us

$$\begin{aligned}\Delta p &= \frac{(1 - \phi)\Delta p + \phi\Delta m}{S - s}(S - s) \\ &= (1 - \phi)\Delta p + \phi\Delta m.\end{aligned}\tag{7.62}$$

Equation (7.62) implies that $\Delta p = \Delta m$, and thus that $\Delta y = 0$. Thus the change in money has no impact on aggregate output.¹²

The reason for the sharp difference between the results of this model and those of the models with time-dependent adjustment is that the number of firms changing their prices at any time is endogenous. In the Caplin-Spulber model, the number of firms changing their prices at any time is larger when aggregate demand is increasing more rapidly; given the specific assumptions that Caplin and Spulber make, this has the effect that the aggregate price level responds fully to changes in m . In the Fischer, Taylor, and Carlo models, in contrast, the number of firms changing their prices at any time is fixed; as a result, the price level does not respond fully to changes in m . Thus this model illustrates the frequency effect.

The Selection Effect: The Danziger-Golosov-Lucas Model

A key fact about price adjustment, which we will discuss in more detail in the next section, is that it varies enormously across firms and products. For example, even in environments of moderately high inflation, a substantial fraction of price changes are price cuts.

This heterogeneity introduces a second channel through which state-dependent pricing dampens the effects of nominal disturbances. With state-dependent pricing, the composition of the firms that adjust their prices responds to shocks. When there is a positive monetary shock, for example,

¹² In addition, this result helps to justify the assumption that the initial distribution of $p_i - p_i^*$ is uniform between s and S . For each firm, $p_i - p_i^*$ equals each value between s and S once during the interval between any two price changes; thus there is no reason to expect a concentration anywhere within the interval. Indeed, Caplin and Spulber show that under simple assumptions, a given firm's $p_i - p_i^*$ is equally likely to take on any value between s and S .

the firms that adjust are disproportionately ones that raise their prices. As a result, it is not just the number of firms changing their prices that responds to the shock; the average change of those that adjust responds as well.

Here we illustrate these ideas using a simple example based on Danziger (1999). However, the model is similar in spirit to the richer model of Golosov and Lucas (2007).

Each firm's profit-maximizing price in period t depends on aggregate demand, m_t , and an idiosyncratic variable, ω_{it} ; ω is independent across firms. For simplicity, ϕ in the price-setting rule is set to 1. Thus $p_{it}^* = m_t + \omega_{it}$.

To show the selection effect as starkly as possible, we make strong assumptions about the behavior of m and ω . Time is discrete. Initially, m is constant and not subject to shocks. Each firm's ω follows a random walk. The innovation to ω , denoted ε , can take on either positive or negative values and is distributed uniformly over a wide range (in a sense to be specified momentarily).

When profit-maximizing prices can either rise or fall, as is the case here, the analogue of an Ss policy is a *two-sided Ss policy*. If a shock pushes the difference between the firm's actual and profit-maximizing prices, $p_i - p_i^*$, either above some upper bound S or below some lower bound s , the firm resets $p_i - p_i^*$ to some target K . As with the one-sided Ss policy in the Caplin-Spulber model, the two-sided policy is optimal in the presence of fixed costs of price adjustment under appropriate assumptions. Again, however, here we just assume that firms follow such a policy.

The sense in which the distribution of ε is wide is that regardless of a firm's initial price, there is some chance the firm will raise its price and some chance that it will lower it. Concretely, let A and B be the lower and upper bounds of the distribution of ε . Then our assumptions are $S - B < s$ and $s - A > S$, or equivalently, $B > S - s$ and $A < -(S - s)$. To see the implications of these assumptions, consider a firm that is at the upper bound, S , and so appears to be on the verge of cutting its price. The assumption $B > S - s$ means that if it draws that largest possible realization of ε , its $p - p^*$ is pushed below the lower bound s , and so it raises its price. Thus every firm has some chance of raising its price each period. Likewise, the assumption $A < -(S - s)$ implies that every firm has some chance of cutting its price.

The steady state of the model is relatively simple. Initially, all $p_i - p_i^*/s$ must be between s and S . For any $p_i - p_i^*$ within this interval, there is a range of values of ε of width $S - s$ that leaves $p_i - p_i^*$ between s and S . Thus the probability that the firm does not adjust its price is $(S - s)/(B - A)$. Conditional on not adjusting, $p_i - p_i^*$ is distributed uniformly on $[sS]$. And with probability $1 - [(S - s)/(B - A)]$ the firm adjusts, in which case its $p_i - p_i^*$ equals the reset level, K .

This analysis implies that the distribution of $p_i - p_i^*$ consists of a uniform distribution over $[sS]$ with density $1/(B - A)$, plus a spike of mass $1 - [(S - s)/(B - A)]$ at K . This is shown in Figure 7.1. For convenience, we assume that $K = (S + s)/2$, so that the reset price is midway between s and S .




FIGURE 7.1 The steady state of the Danziger model

Now consider a one-time monetary shock. Specifically, suppose that at the end of some period, after firms have made price-adjustment decisions, there is an unexpected increase in m of amount $\Delta m < K - s$. This raises all p_i^* 's by Δm . That is, the distribution in Figure 7.1 shifts to the left by Δm . Because pricing is state-dependent, firms can change their prices at any time. The firms whose $p_i - p_i^*$'s are pushed below s therefore raise them to K . The resulting distribution is shown in Figure 7.2.

Crucially, the firms that adjust are not a random sample of firms. Instead, they are the firms whose actual prices are furthest below their optimal prices, and thus that are most inclined to make large price increases. For small values of Δm , the firms that raise their prices do so by approximately $K - s$. If instead, in the spirit of time-dependent models, we picked firms at random and allowed them to change their prices, their average price increase would be Δm .¹³ Thus there is a selection effect that sharply increases the initial price response.

Now consider the next period: there is no additional monetary shock, and the firm-specific shocks behave in their usual way. But because of the initial monetary disturbance, there are now relatively few firms near S . Thus the firms whose idiosyncratic shocks cause them to change their prices are

¹³ The result that the average increase is Δm is exactly true only because of the assumption that $K = (S + s)/2$. If this condition does not hold, there is a constant term that does not depend on the sign or magnitude of Δm .




FIGURE 7.2 The initial effects of a monetary disturbance in the Danziger model

disproportionately toward the bottom of the $[sS]$ interval, and so price changes are disproportionately price increases. Given the strong assumptions of the model, the distribution of $p_i - p_i^*$ returns to its steady state after just one period. And the distribution of $p_i - p_i^*$ being unchanged is equivalent to the distribution of p_i moving one-for-one with the distribution of p_i^* . That is, actual prices on average adjust fully to the rise in m . Note that this occurs even though the fraction of firms changing their prices in this period is exactly the same as normal (all firms change their prices with probability $1 - [(S - s)/(B - A)]$, as usual), and even though all price changes in this period are the result of firm-specific shocks.

Discussion

The assumptions of these examples are chosen to show the frequency and selection effects as starkly as possible. In the Danziger-Golosov-Lucas model, the assumption of wide, uniformly distributed firm-specific shocks is needed to deliver the strong result that a monetary shock is neutral after just one period. With a narrower distribution, for example, the effects would be more persistent. Similarly, a nonuniform distribution of the shocks generally leads to a nonuniform distribution of firms' prices, and so weakens the frequency effect. In addition, allowing for real rigidity (that is, allowing

ϕ in the expression for firms' desired prices to be less than 1) causes the behavior of the nonadjusters to influence the firms that change their prices, and so causes the effects of monetary shocks to be larger and longer lasting.

Similarly, if we introduced negative as well as positive monetary shocks to the Caplin-Spulber model, the result would be a two-sided Ss rule, and so monetary shocks would generally have real effects (see, for example, Caplin and Leahy, 1991, and Problem 7.7). In addition, the values of S and s may change in response to changes in aggregate demand. If, for example, high money growth today signals high money growth in the future, firms widen their Ss bands when there is a positive monetary shock; as a result, no firms adjust their prices in the short run (since no firms are now at the new, lower trigger point s), and so the positive shock raises output (Tsiddon, 1991).¹⁴

In short, the strong results of the simple cases considered in this section are not robust. What is robust is that state-dependent pricing gives rise naturally to the frequency and selection effects, and that those effects can be quantitatively important. For example, Golosov and Lucas show in the context of a much more carefully calibrated model that the effects of monetary shocks can be much smaller with state-dependent pricing than in a comparable economy with time-dependent pricing.

7.6 Empirical Applications

Microeconomic Evidence on Price Adjustment

The central assumption of the models we have been analyzing is that there is some kind of barrier to complete price adjustment at the level of individual firms. It is therefore natural to investigate pricing policies at the microeconomic level. By doing so, we can hope to learn whether there are barriers to price adjustment and, if so, what form they take.

The microeconomics of price adjustment have been investigated by many authors. The broadest studies of price adjustment in the United States are the survey of firms conducted by Blinder (1998), the analysis of the data underlying the Consumer Price Index by Klenow and Kryvtsov (2008), and the analysis of the data underlying the Consumer Price Index and the Producer Price Index by Nakamura and Steinsson (2008). Blinder's and Nakamura and Steinsson's analyses show that the average interval between price changes for intermediate goods is about a year. In contrast, Klenow and Kryvtsov's and Nakamura and Steinsson's analyses find that the typical period between price changes for final goods and services is only about 4 months.

The key finding of this literature, however, is not the overall statistics concerning the frequency of adjustment. Rather, it is that price adjustment

¹⁴ See Caballero and Engel (1993) for a more detailed analysis of these issues.




FIGURE 7.3 Price of a 9.5 ounce box of Triscuits (from Chevalier, Kashyap, and Rossi, 2000; used with permission)

does not follow any simple pattern. Figure 7.3, from Chevalier, Kashyap, and Rossi (2000), is a plot of the price of a 9.5 ounce box of Triscuit crackers at a particular supermarket from 1989 to 1997. The behavior of this price clearly defies any simple summary. One obvious feature, which is true for many products, is that temporary “sale” prices are common. That is, the price often falls sharply and is then quickly raised again, often to its previous level. Beyond the fact that sales are common, it is hard to detect any regular patterns. Sales occur at irregular intervals and are of irregular lengths; the sizes of the reductions during sales vary; the intervals between adjustments of the “regular” price are heterogeneous; the regular price sometimes rises and sometimes falls; and the sizes of the changes in the regular price vary. Other facts that have been documented include tremendous heterogeneity across products in the frequency of adjustment; a tendency for some prices to be adjusted at fairly regular intervals, most often once a year; the presence of a substantial fraction of price decreases (of both regular and sale prices), even in environments of moderately high inflation; and the presence for many products of a second type of sale, a price reduction that is not reversed and that is followed, perhaps after further reductions, by the disappearance of the product (a “clearance” sale).

Thus the microeconomic evidence does not show clearly what assumptions about price adjustment we should use in building a macroeconomic

model. Time-dependent models are grossly contradicted by the data, and purely state-dependent models fare only slightly better. The time-dependent models are contradicted by the overwhelming presence of irregular intervals between adjustments. Purely state-dependent models are most clearly contradicted by two facts: the frequent tendency for prices to be in effect for exactly a year, and the strong tendency for prices to revert to their original level after a sale.

In thinking about the aggregate implications of the evidence on price adjustment, a key issue is how to treat sales. At one extreme, they could be completely mechanical. Imagine, for example, that a store manager is instructed to discount goods representing 10 percent of the store's sales by an average of 20 percent each week. Then sale prices are unresponsive to macroeconomic conditions, and so should be ignored in thinking about macroeconomic issues. If we decide to exclude sales, we then encounter difficult issues of how to define them and how to treat missing observations and changes in products. Klenow and Kryvtsov's and Nakamura and Steinsson's analyses suggest, however, that across goods, the median frequency of changes in regular prices of final goods is about once every 7 months. For intermediate goods, sales are relatively unimportant, and so accounting for them has little impact on estimates of the average frequency of adjustment.

The other possibility is that sale prices respond to macroeconomic conditions; for example, they could be more frequent and larger when the economy is weak. At the extreme, sales should not be removed from the data at all in considering the macroeconomic implications of the microeconomics of price adjustment.

Another key issue for the aggregate implications of these data is heterogeneity. The usual summary statistic, and the one used above, is the median frequency of adjustment across goods. But the median masks an enormous range, from goods whose prices typically adjust more than once a month to ones whose prices usually change less than once a year. Carvalho (2006) poses the following question. Suppose the economy is described by a model with heterogeneity, but a researcher wants to match the economy's response to various types of monetary disturbances using a model with a single frequency of adjustment. What frequency should the researcher choose? Carvalho shows that in most cases, one would want to choose a frequency less than the median or average frequency. Moreover, the difference is magnified by real rigidity: as the degree of real rigidity rises, the importance of the firms with the stickiest prices increases. Carvalho shows that to best match the economy's response to shocks using a single-sector model, one would often want to use a frequency of price adjustment a third to a half of the median across heterogeneous firms. Thus heterogeneity has important effects.

Finally, Levy, Bergen, Dutta, and Venable (1997) look not at prices, but at the costs of price adjustment. Specifically, they report data on each step of

the process of changing prices at supermarkets, such as the costs of putting on new price tags or signs on the shelves, of entering the new prices into the computer system, and of checking the prices and correcting errors. This approach does not address the possibility that there may be more sophisticated, less expensive ways of adjusting prices to aggregate disturbances. For example, a store could have a prominently displayed discount factor that it used at checkout to subtract some proportion from the amount due; it could then change the discount factor rather than the shelf prices in response to aggregate shocks. The costs of changing the discount factor would be dramatically less than the cost of changing the posted price on every item in the store.

Despite this limitation, it is still interesting to know how large the costs of changing prices are. Levy et al.'s basic finding is that the costs are surprisingly high. For the average store in their sample, expenditures on changing prices amount to between 0.5 and 1 percent of revenues. To put it differently, the average cost of a price change in their stores in 1991–1992 was about 50 cents. Thus the common statement that the physical costs of nominal price changes are extremely small is not always correct: for the stores that Levy et al. consider, these costs, while not large, are far from trivial.

In short, empirical work on the microeconomics of price adjustment and its macroeconomic implications is extremely active. A few examples of recent contributions in addition to those discussed above are Dotsey, King, and Wolman (1999), Klenow and Willis (2006), Gopinath and Rigobon (2008), and Midrigan (2009).

Inflation Inertia

We have encountered three aggregate supply relationships that include an inflation term and an output term: the accelerationist Phillips curve of Section 6.4, the Lucas supply curve of Section 6.9, and the new Keynesian Phillips curve of Section 7.4. Although the three relationships look broadly similar, in fact they have sharply different implications. To see this, consider the experiment of an anticipated fall in inflation in an economy with no shocks. The accelerationist Phillips curve, $\pi_t = \pi_{t-1} + \lambda(y_t - \bar{y}_t)$ (see [6.22]–[6.23]), implies that disinflation requires below-normal output. The Lucas supply curve, $\pi_t = E_{t-1}\pi_t + \lambda(y_t - \bar{y}_t)$ (see [6.84]), implies that disinflation can be accomplished with no output cost. Finally, for the new Keynesian Phillips curve (equation [7.60]), it is helpful to rewrite it as

$$E_t[\pi_{t+1}] - \pi_t = \left(\frac{1-\beta}{\beta}\right)\pi_t - \frac{\kappa}{\beta}(y_t - \bar{y}_t). \quad (7.63)$$

With β close to 1, the $[(1-\beta)/\beta]\pi_t$ term is small. Thus the new Keynesian Phillips curve implies that anticipated disinflation is associated with an output *boom*.

The view that high inflation has a tendency to continue unless there is a period of low output is often described as the view that there is *inflation inertia*. That is, “inflation inertia” refers not to inflation being highly serially correlated, but to it being costly to reduce. Of the three Phillips curves, only the accelerationist one implies inertia. The Lucas supply curve implies that there is no inertia, while the new Keynesian Phillips curve (as well as other models of staggered price-setting) implies that there is “anti-inertia.”¹⁵

Ball (1994b) performs a straightforward test for inflation inertia. Looking at a sample of nine industrialized countries over the period 1960–1990, he identifies 28 episodes where inflation fell substantially. He reports that in all 28 cases, observers at the time attributed the decline to monetary policy. Thus the view that there is inflation inertia predicts that output was below normal in the episodes; the Lucas supply curve suggests that it need not have departed systematically from normal; and the new Keynesian Phillips curve implies that it was above normal. Ball finds that the evidence is overwhelmingly supportive of inflation inertia: in 27 of the 28 cases, output was on average below his estimate of normal output during the disinflation.

Ball’s approach of choosing episodes on the basis of ex post inflation outcomes could create bias, however. In particular, suppose the disinflations had important unanticipated components. If prices were set on the basis of expectations of higher aggregate demand than actually occurred, the low output in the episodes does not clearly contradict any of the models.

Gali and Gertler (1999) therefore take a more formal econometric approach. Their main interest is in testing between the accelerationist and new Keynesian views. They begin by positing a hybrid Phillips curve with backward-looking and forward-looking elements:

$$\pi_t = \gamma_b \pi_{t-1} + \gamma_f E_t \pi_{t+1} + \kappa(y_t - \bar{y}_t) + e_t. \quad (7.64)$$

They point out, however, that what the $\kappa(y_t - \bar{y}_t)$ term is intended to capture is the behavior of firms’ real marginal costs. When output is above normal, marginal costs are high, which increases desired relative prices. In the model of Section 7.1, for example, desired relative prices rise when output rises because the real wage increases. Gali and Gertler therefore try a more direct approach to estimating marginal costs. Real marginal cost equals the real wage divided by the marginal product of labor. If the production function is Cobb-Douglas, so that $Y = K^\alpha(AL)^{1-\alpha}$, the marginal product of labor is $(1 - \alpha)Y/L$. Thus real marginal cost is $wL/[(1 - \alpha)Y]$, where w is the real wage. That is, marginal cost is proportional to the share of income going

¹⁵ The result that models of staggered price adjustment do not imply inflation inertia is due to Fuhrer and Moore (1995) and Ball (1994a).

to labor (see also Sbordone, 2002). Galí and Gertler therefore focus on the equation:

$$\pi_t = \gamma_b \pi_{t-1} + \gamma_f E_t \pi_{t+1} + \lambda S_t + e_t, \quad (7.65)$$

where S_t is labor's share.¹⁶

Galí and Gertler estimate (7.65) using quarterly U.S. data for the period 1960-1997.¹⁷ A typical set of estimates is

$$\pi_t = 0.378 \pi_{t-1} + 0.591 E_t \pi_{t+1} + 0.015 S_t + e_t, \quad (7.66)$$

(0.020) (0.016) (0.004)

where the numbers in parentheses are standard errors. Thus their results appear to provide strong support for the importance of forward-looking expectations.

In a series of papers, however, Rudd and Whelan show that in fact the data provide little evidence for the new Keynesian Phillips curve (see especially Rudd and Whelan, 2005, 2006). They make two key points. The first concerns labor's share. Galí and Gertler's argument for including labor's share in the Phillips curve is that under appropriate assumptions, it captures the rise in firms' marginal costs when output rises. Rudd and Whelan (2005) point out, however, that in practice labor's share is low in booms and high in recessions. In Galí and Gertler's framework, this would mean that booms are times when the economy's flexible-price level of output has risen even more than actual output, and when marginal costs are therefore unusually low. A much more plausible possibility, however, is that there are forces other than those considered by Galí and Gertler moving labor's share over the business cycle, and that labor's share is therefore a poor proxy for marginal costs.

Since labor's share is countercyclical, the finding of a large coefficient on expected future inflation and a positive coefficient on the share means that inflation tends to be above future inflation in recessions and below future inflation in booms. That is, inflation tends to fall in recessions and rise in booms, consistent with the accelerationist Phillips curve and not with the new Keynesian Phillips curve.

¹⁶ How can labor's share vary if production is Cobb-Douglas? Under perfect competition (and under imperfect competition if price is a constant markup over marginal cost), it cannot. But if prices are not fully flexible, it can. For example, if a firm with a fixed price hires more labor at the prevailing wage, output rises less than proportionally than the rise in labor, and so labor's share rises.

¹⁷ For simplicity, we omit any discussion of their estimation procedure, which, among other things, must address the fact that we do not have data on $E_t \pi_{t+1}$. Section 8.3 discusses estimation when there are expectational variables.

Rudd and Whelan's second concern has to do with the information content of current inflation. Replacing y_t with a generic marginal cost variable, mc_t , and then iterating the new Keynesian Phillips curve, (7.60), forward implies

$$\begin{aligned}\pi_t &= \kappa mc_t + \beta E_t \pi_{t+1} \\ &= \kappa mc_t + \beta[\kappa E_t mc_{t+1} + \beta E_t \pi_{t+2}] \\ &= \dots \\ &= \kappa \sum_{i=0}^{\infty} \beta^i E_t mc_{t+i}.\end{aligned}\tag{7.67}$$

Thus the model implies that inflation should be a function of expectations of future marginal costs, and thus that it should help predict marginal costs. Rudd and Whelan (2005) show, however, that the evidence for this hypothesis is minimal. When marginal costs are proxied by an estimate of $y - \bar{y}$, inflation's predictive power is small and goes in the wrong direction from what the model suggests. When marginal costs are measured using labor's share (which, as Rudd and Whelan's first criticism shows, may be a poor proxy), the performance is only slightly better. In this case, inflation's predictive power for marginal costs is not robust, and almost entirely absent in Rudd and Whelan's preferred specification. They also find that the hybrid Phillips curve performs little better (Rudd and Whelan, 2006). They conclude that there is little evidence in support of the new Keynesian Phillips curve.¹⁸

The bottom line of this analysis is twofold. First, the evidence we have on the correct form of the Phillips curve is limited. The debate between Galí and Gertler and Rudd and Whelan, along with further analysis of the econometrics of the new Keynesian Phillips curve (for example, King and Plosser, 2005), does not lead to clear conclusions on the basis of formal econometric studies. This leaves us with the evidence from less formal analyses, such as Ball's, which is far from airtight. Second, although the evidence is not definitive, it points in the direction of inflation inertia and provides little support for the new Keynesian Phillips curve.

Because of this and other evidence, researchers attempting to match important features of business-cycle dynamics typically make modifications to models of price-setting (often along the lines of the ones we will encounter in the next section) that imply inertia. Nonetheless, because of its simplicity

¹⁸ This discussion does not address the question of why Galí and Gertler's estimates suggest that the new Keynesian Phillips curve fits well. Rudd and Whelan argue that this has to do with the specifics of Galí and Gertler's estimation procedure, which we are not delving into. Loosely speaking, Rudd and Whelan's argument is that because inflation is highly serially correlated, small violations of the conditions needed for the estimation procedure to be valid can generate substantial upward bias in the coefficient on $E_t \pi_{t+1}$.

and elegance, the new Keynesian Phillips curve is still often used in theoretical models. Following that pattern, we will meet it again in Section 7.8 and in Chapter 11.

7.7 Models of Staggered Price Adjustment with Inflation Inertia

The evidence in the previous section suggests that a major limitation of the micro-founded models of dynamic price adjustment we have been considering is that they do not imply inflation inertia. A central focus of recent work on price adjustment is therefore bringing inflation inertia into the models. At a general level, the most common strategy is to assume that firms' prices are not fixed between the times they review them, but adjust in some way. These adjustments are assumed to give some role to past inflation, or to past beliefs about inflation. The result is inflation inertia.

The two most prominent approaches along these lines are those of Christiano, Eichenbaum, and Evans (2005) and Mankiw and Reis (2002). Christiano, Eichenbaum, and Evans assume that between reviews, prices are adjusted for past inflation. This creates a direct role for past inflation in price behavior. But whether this reasonably captures important microeconomic phenomena is not clear. Mankiw and Reis return to Fischer's assumption of prices that are predetermined but not fixed. This causes past beliefs about what inflation would be to affect price changes, and so creates behavior similar to inflation inertia. In contrast to Fischer, however, they make assumptions that imply that some intervals between reviews of prices are quite long, which has important quantitative implications. Again, however, the strength of the microeconomic case for the realism of their key assumption is not clear.

The Christiano, Eichenbaum, and Evans Model: The New Keynesian Phillips Curve with Indexation

Christiano, Eichenbaum, and Evans begin with Calvo's assumption that opportunities for firms to review their prices follow a Poisson process. As in the basic Calvo model of Section 7.4, let α denote the fraction of firms that review their prices in a given period. Where Christiano, Eichenbaum, and Evans depart from Calvo is in their assumption about what happens between reviews. Rather than assuming that prices are fixed, they assume they are indexed to the previous period's inflation rate. This assumption captures the fact that even in the absence of a full-fledged reconsideration of their prices, firms can account for the overall inflationary environment. The assumption that the indexing is to lagged rather than current inflation

reflects the fact that firms do not continually obtain and use all available information.

Our analysis of the model is similar to the analysis of the Calvo model in Section 7.4. Since the firms that review their prices in a given period are chosen at random, the average (log) price in period t of the firms that do not review their prices is $p_{t-1} + \pi_{t-1}$. The average price in t is therefore

$$p_t = (1 - \alpha)(p_{t-1} + \pi_{t-1}) + \alpha x_t, \quad (7.68)$$

where x_t is the price set by firms that review their prices. Equation (7.68) implies

$$\begin{aligned} x_t - p_t &= x_t - [(1 - \alpha)(p_{t-1} + \pi_{t-1}) + \alpha x_t] \\ &= (1 - \alpha)x_t - (1 - \alpha)(p_{t-1} + \pi_{t-1}) \\ &= (1 - \alpha)(x_t - p_t) - (1 - \alpha)(p_{t-1} + \pi_{t-1} - p_t) \\ &= (1 - \alpha)(x_t - p_t) + (1 - \alpha)(\pi_t - \pi_{t-1}). \end{aligned} \quad (7.69)$$

Thus,

$$x_t - p_t = \frac{1 - \alpha}{\alpha}(\pi_t - \pi_{t-1}). \quad (7.70)$$

Equation (7.70) shows that to find the dynamics of inflation, we need to find $x_t - p_t$. That is, we need to determine how firms that review their prices set their relative prices in period t . As in the Calvo model, a firm wants to set its price to minimize the expected discounted sum of the squared differences between its optimal and actual prices during the period before it is next able to review its price. Suppose a firm sets a price of x_t in period t and that it does not have an opportunity to review its price before period $t + j$. Then, because of the lagged indexation, its price in $t + j$ (for $j \geq 1$) is $x_t + \sum_{\tau=0}^{j-1} \pi_{t+\tau}$. The profit-maximizing price in $t + j$ is $p_{t+j} + \phi y_{t+j}$, which equals $p_t + \sum_{\tau=1}^j \pi_{t+\tau} + \phi y_{t+j}$. Thus the difference between the profit-maximizing and actual prices in $t + j$, which we will denote $e_{t,t+j}$, is

$$e_{t,t+j} = (p_t - x_t) + (\pi_{t+j} - \pi_t) + \phi y_{t+j}. \quad (7.71)$$

Note that (7.71) holds for all $j \geq 0$. The discount factor is β , and the probability of nonadjustment each period is $1 - \alpha$. Thus, similarly to equation (7.56) in the Calvo model without indexation, the firm sets

$$x_t - p_t = [1 - \beta(1 - \alpha)] \sum_{j=0}^{\infty} \beta^j (1 - \alpha)^j [(E_t \pi_{t+j} - \pi_t) + \phi E_t y_{t+j}]. \quad (7.72)$$

As in the derivation of the new Keynesian Phillips curve, it is helpful to rewrite this expression in terms of period- t variables and the expectation of

$x_{t+1} - p_{t+1}$. Equation (7.72) implies

$$\begin{aligned} x_{t+1} - p_{t+1} & \\ &= [1 - \beta(1 - \alpha)] \sum_{j=0}^{\infty} \beta^j (1 - \alpha)^j [(E_{t+1} \pi_{t+1+j} - \pi_{t+1}) + \phi E_{t+1} y_{t+1+j}]. \end{aligned} \quad (7.73)$$

Rewriting the π_{t+1} term as $\pi_t + (\pi_{t+1} - \pi_t)$ and taking expectations as of t (and using the law of iterated projections) gives us

$$\begin{aligned} E_t[x_{t+1} - p_{t+1}] &= -E_t[\pi_{t+1} - \pi_t] \\ &+ [1 - \beta(1 - \alpha)] \sum_{j=0}^{\infty} \beta^j (1 - \alpha)^j [(E_t \pi_{t+1+j} - \pi_t) + \phi E_t y_{t+1+j}]. \end{aligned} \quad (7.74)$$

We can therefore rewrite (7.72) as

$$x_t - p_t = [1 - \beta(1 - \alpha)] \phi y_t + \beta(1 - \alpha) \{E_t[x_{t+1} - p_{t+1}] + E_t[\pi_{t+1} - \pi_t]\}. \quad (7.75)$$

The final step is to use (7.70) applied to both periods t and $t+1$: $x_t - p_t = [(1 - \alpha)/\alpha](\pi_t - \pi_{t-1})$, $E_t[x_{t+1} - p_{t+1}] = [(1 - \alpha)/\alpha](E_t[\pi_{t+1}] - \pi_t)$. Substituting these expressions into (7.75) and performing straightforward algebra yields

$$\begin{aligned} \pi_t &= \frac{1}{1 + \beta} \pi_{t-1} + \frac{\beta}{1 + \beta} E_t \pi_{t+1} + \frac{1}{1 + \beta} \frac{\alpha}{1 - \alpha} [1 - \beta(1 - \alpha)] \phi y_t \\ &\equiv \frac{1}{1 + \beta} \pi_{t-1} + \frac{\beta}{1 + \beta} E_t \pi_{t+1} + \chi y_t. \end{aligned} \quad (7.76)$$

Equation (7.76) is the new *Keynesian Phillips curve with indexation*. It resembles the new Keynesian Phillips curve except that instead of a weight of β on expected future inflation and no role for past inflation, there is a weight of $\beta/(1 + \beta)$ on expected future inflation and a weight of $1/(1 + \beta)$ on lagged inflation. If β is close to 1, the weights are both close to one-half. An obvious generalization of (7.76) is

$$\pi_t = \gamma \pi_{t-1} + (1 - \gamma) E_t \pi_{t+1} + \chi y_t, \quad 0 \leq \gamma \leq 1. \quad (7.77)$$

Equation (7.77) allows for any mix of weights on the two inflation terms.

Because they imply that past inflation has a direct impact on current inflation, and thus that there is inflation inertia, expressions like (7.76) and (7.77) often appear in modern dynamic stochastic general-equilibrium models with nominal rigidity.

The Model's Implications for the Costs of Disinflation

The fact that equation (7.76) (or [7.77]) implies inflation inertia does not mean that the model can account for the apparent output costs of disinflation. To see this, consider the case of $\beta = 1$, so that (7.76) becomes $\pi_t = (\pi_{t-1}/2) + (E_t[\pi_{t+1}]/2) + \chi y_t$. Now suppose that there is a perfectly

anticipated, gradual disinflation that occurs at a uniform rate: $\pi_t = \pi_0$ for $t \leq 0$; $\pi_t = 0$ for $t \geq T$; and $\pi_t = [(T - t)/T]\pi_0$ for $0 < t < T$. Because the disinflation proceeds linearly and is anticipated, π_t equals the average of π_{t-1} and $E_t[\pi_{t+1}]$ in all periods except $t = 0$ and $t = T$. In period 0, π_0 exceeds $(\pi_{t-1} + E_t[\pi_{t+1}])/2$, and in period T , it is less than $(\pi_{t-1} + E_t[\pi_{t+1}])/2$ by the same amount. Thus the disinflation is associated with above-normal output when it starts and an equal amount of below-normal output when it ends, and no departure of output from normal in between. That is, the model implies no systematic output cost of an anticipated disinflation.

One possible solution to this difficulty is to reintroduce the assumption that β is less than 1. This results in more weight on π_{t-1} and less on $E_t[\pi_{t+1}]$, and so creates output costs of disinflation. For reasonable values of β , however, this effect is small.

A second potential solution is to appeal to the generalization in equation (7.77) and to suppose that $\gamma > (1 - \gamma)$. But since (7.77) is not derived from microeconomic foundations, this comes at the cost of abandoning the initial goal of grounding our understanding of inflation dynamics in microeconomic behavior.

The final candidate solution is to argue that the prediction of no systematic output costs of an anticipated disinflation is reasonable. Recall that Ball's finding is that disinflations are generally associated with below-normal output. But recall also that the fact that disinflations are typically less than fully anticipated means that the output costs of actual disinflations tend to overstate the costs of perfectly anticipated disinflations. Perhaps the bias is sufficiently large that the average cost of an anticipated disinflation is zero.

The bottom line is that adding indexation to Calvo pricing introduces some inflation inertia. But whether that inertia is enough to explain actual inflation dynamics is not clear.

The other important limitation of the model is that its key microeconomic assumption appears unrealistic—we do not observe actual prices rising mechanically with lagged inflation. At the same time, however, it could be that price-setters behave in ways that cause their average prices to rise roughly with lagged inflation between the times that they seriously rethink their pricing policies in light of macroeconomic conditions, and that this average adjustment is masked by the fact that individual nominal prices are not continually adjusted.

The Mankiw-Reis Model

Mankiw and Reis take a somewhat different approach to obtaining inflation inertia. Like Christiano, Eichenbaum, and Evans, they assume some adjustment of prices between the times that firms review their pricing policies. Their assumption, however, is that each time a firm reviews its price, it sets a *path* that the price will follow until the next review. That is, they

reintroduce the idea from the Fischer model that prices are predetermined but not fixed.

Recall that a key result from our analysis in Section 7.2 is that with predetermined prices, a monetary shock ceases to have real effects once all price-setters have had an opportunity to respond. This is often taken to imply that predetermined prices cannot explain persistent real effects of monetary shocks. But recall also that when real rigidity is high, firms that do not change their prices have a disproportionate impact on the behavior of the aggregate economy. This raises the possibility that a small number of firms that are slow to change their price paths can cause monetary shocks to have important long-lasting effects with predetermined prices. This is the central idea of Mankiw and Reis's model (see also Devereux and Yetman, 2003).

Although the mechanics of the Mankiw-Reis model involve predetermined prices, their argument for predetermination differs from Fischer's. Fischer motivates his analysis in terms of labor contracts that specify a different wage for each period of the contract; prices are then determined as markups over wages. But such contracts do not appear sufficiently widespread to be a plausible source of substantial aggregate nominal rigidity. Mankiw and Reis appeal instead to what they call "sticky information." It is costly for price-setters to obtain and process information. Mankiw and Reis argue that as a result, they may choose not to continually update their prices, but to periodically choose a path for their prices that they follow until they next gather information and adjust their path.

Specifically, Mankiw and Reis begin with a model of predetermined prices like that of Section 7.2. Opportunities to adopt new price paths do not arise deterministically, as in the Fischer model, however. Instead, as in the Calvo and Christiano-Eichenbaum-Evans models, they follow a Poisson process. Paralleling those models, each period a fraction α of firms adopt a new piece path (where $0 < \alpha \leq 1$). And again $y_t = m_t - p_t$ and $p_t^* = p_t + \phi y_t$.

Our analysis of the Fischer model provides a strong indication of what the solution of the model will look like. Because a firm can set a different price for each period, the price it sets for a given period, period t , will depend only on information about y_t and p_t . It follows that the aggregate price level, p_t (and hence y_t), will depend only on information about m_t ; information about m in other periods will affect y_t and p_t only to the extent it conveys information about m_t . Further, if the value of m_t were known arbitrarily far in advance, all firms would set their prices for t equal to m_t , and so y_t would be zero. Thus, departures of y_t from zero will come only from information about m_t revealed after some firms have set their prices for period t . And given the log-linear structure of the model, its solution will be log-linear.

Consider information about m_t that arrives in period $t - i$ ($i \geq 0$); that is, consider $E_{t-i}m_t - E_{t-(i+1)}m_t$. If we let a_i denote the fraction of $E_{t-i}m_t - E_{t-(i+1)}m_t$ that is passed into the aggregate price level, then the information about m_t that arrives in period $t - i$ raises p_t by $a_i(E_{t-i}m_t - E_{t-(i+1)}m_t)$ and raises y_t by $(1 - a_i)(E_{t-i}m_t - E_{t-(i+1)}m_t)$. That is, y_t will be given by

an expression of the form

$$y_t = \sum_{i=0}^{\infty} (1 - a_i)(E_{t-i}m_t - E_{t-(i+1)}m_t). \quad (7.78)$$

To solve the model, we need to find the a_i 's. To do this, let λ_i denote the fraction of firms that have an opportunity to change their price for period t in response to information about m_t that arrives in period $t - i$ (that is, in response to $E_{t-i}m_t - E_{t-(i+1)}m_t$). A firm does *not* have an opportunity to change its price for period t in response to this information if it does not have an opportunity to set a new price path in any of periods $t - i$, $t - (i - 1), \dots, t$. The probability of this occurring is $(1 - \alpha)^{i+1}$. Thus,

$$\lambda_i = 1 - (1 - \alpha)^{i+1}. \quad (7.79)$$

Because firms can set a different price for each period, the firms that adjust their prices are able to respond freely to the new information. We know that $p_t^* = (1 - \phi)p_t + \phi m_t$ and that the change in p_t in response to the new information is $a_i(E_{t-i}m_t - E_{t-(i+1)}m_t)$. Thus, the firms that are able to respond raise their prices for period t by $(1 - \phi)a_i(E_{t-i}m_t - E_{t-(i+1)}m_t) + \phi(E_{t-i}m_t - E_{t-(i+1)}m_t)$, or $[(1 - \phi)a_i + \phi](E_{t-i}m_t - E_{t-(i+1)}m_t)$. Since fraction λ_i of firms are able to adjust their prices and the remaining firms cannot respond at all, the overall price level responds by $\lambda_i[(1 - \phi)a_i + \phi](E_{t-i}m_t - E_{t-(i+1)}m_t)$. Thus a_i must satisfy

$$\lambda_i[(1 - \phi)a_i + \phi] = a_i. \quad (7.80)$$

Solving for a_i yields

$$\begin{aligned} a_i &= \frac{\phi\lambda_i}{1 - (1 - \phi)\lambda_i} \\ &= \frac{\phi[1 - (1 - \alpha)^{i+1}]}{1 - (1 - \phi)[1 - (1 - \alpha)^{i+1}]}, \end{aligned} \quad (7.81)$$

where the second line uses (7.79) to substitute for λ_i . Finally, since $p_t + y_t = m_t$, we can write p_t as

$$p_t = m_t - y_t. \quad (7.82)$$

Implications

To understand the implications of the Mankiw-Reis model, it is helpful to start by examining the effects of a shift in the level of aggregate demand (as opposed to its growth rate).¹⁹ Specifically, consider an unexpected, one-time, permanent increase in m in period t of amount Δm . The increase raises

¹⁹ The reason for not considering this experiment for the Christiano-Eichenbaum-Evans model is that the model's implications concerning such a shift are complicated. See Problem 7.9.

$E_t m_{t+i} - E_{t-1} m_{t+i}$ by Δm for all $i \geq 0$. Thus p_{t+i} rises by $a_i \Delta m$ and y_{t+i} rises by $(1 - a_i) \Delta m$.

Equation (7.80) implies that the a_i 's are increasing in i and gradually approach 1. Thus the permanent increase in aggregate demand leads to a rise in output that gradually disappears, and to a gradual rise in the price level. If the degree of real rigidity is high, the output effects can be quite persistent even if price adjustment is frequent. Mankiw and Reis assume that a period corresponds to a quarter, and consider the case of $\lambda = 0.25$ and $\phi = 0.1$. These assumptions imply price adjustment on average every four periods and substantial real rigidity. For this case, $a_8 = 0.55$. Even though by period 8 firms have been able to adjust their price paths twice on average since the shock, there is a small fraction—7.5 percent—that have not been able to adjust at all. Because of the high degree of real rigidity, the result is that the price level has only adjusted slightly more than halfway to its long-run level.

Another implication concerns the time pattern of the response. Straightforward differentiation of (7.81) shows that if $\phi < 1$, then $d^2 a_i / d\lambda_i^2 > 0$. That is, when there is real rigidity, the impact of a given change in the number of additional firms adjusting their prices is greater when more other firms are adjusting. Thus there are two competing effects on how the a_i 's vary with i . The fact that $d^2 a_i / d\lambda_i^2 > 0$ tends to make the a_i 's rise more rapidly as i rises, but the fact that fewer additional firms are getting their first opportunity to respond to the shock as i increases tends to make them rise less rapidly. For the parameter values that Mankiw and Reis consider, the a_i 's rise first at an increasing rate and then a decreasing one, with the greatest rate of increase occurring after about eight periods. That is, the peak effect of the demand expansion on inflation occurs with a lag.²⁰

Now consider a disinflation. For concreteness, we start with the case of an immediate, unanticipated disinflation. In particular, assume that until date 0 all firms expect m to follow the path $m_t = gt$ (where $g > 0$), but that the central bank stabilizes m at 0 starting at date 0. Thus $m_t = 0$ for $t \geq 0$.

Because of the policy change, $E_0 m_t - E_{-1} m_t = -gt$ for all $t \geq 0$. This expression is always negative—that is, the actual money supply is always below what was expected by the firms that set their price paths before date 0. Since the a_i 's are always between 0 and 1, it follows that the disinflation lowers output. Specifically, equations (7.78) and (7.81) imply that the path of y is given by

$$\begin{aligned} y_t &= (1 - a_t)(-gt) \\ &= -\frac{(1 - \alpha)^{t+1}}{1 - (1 - \phi)[1 - (1 - \alpha)^{t+1}]} gt \quad \text{for } t \geq 0. \end{aligned} \tag{7.83}$$

²⁰ This is easier to see in a continuous-time version of the model (see Problem 7.11). In this case, equation (7.81) becomes $a(i) = \phi(1 - e^{-\alpha i}) / [1 - (1 - \phi)(1 - e^{-\alpha i})]$. The sign of $a'(i)$ is determined by the sign of $(1 - \phi)e^{-\alpha i} - \phi$. For Mankiw and Reis's parameter values, this is positive until $i \simeq 8.8$ and then negative.

The $(1 - a_t)$'s are falling over time, while g_t is rising. Initially the linear growth of the g_t term dominates, and so the output effect increases. Eventually, however, the fall in the $(1 - a_t)$'s dominates, and so the output effect decreases, and asymptotically it approaches zero. Thus the switch to a lower growth rate of aggregate demand produces a recession whose trough is reached with a lag. For the parameter values described above, the trough occurs after seven quarters.

For the first few periods after the policy shift, most firms still follow their old price paths. Moreover, the firms that are able to adjust do not change their prices for the first few periods very much, both because m is not yet far below its old path and because (if $\phi < 1$) they do not want to deviate far from the prices charged by others. Thus initially inflation falls little. As time passes, however, these forces all act to create greater price adjustment, and so inflation falls. In the long run, output returns to normal and inflation equals the new growth rate of aggregate demand, which is zero. Thus, consistent with what we appear to observe, a shift to a disinflationary policy first produces a recession, and then a fall in inflation.

The polar extreme from a completely anticipated disinflation is one that is anticipated arbitrarily far in advance. The model immediately implies that such a disinflation is not associated with any departure of output from normal. If all firms know the value of m_t for some period t when they set their prices, then, regardless of what they expect about m in any other period, they set $p_t = m_t$, and so we have $y_t = 0$.

For any disinflation, either instantaneous or gradual, that is not fully anticipated, there are output costs. The reason is simple: any disinflation involves a fall of aggregate demand below its prior path. Thus for sufficiently large values of τ , m_t is less than $E_{t-\tau}m_t$, and so the prices for period t that are set in period $t - \tau$ are above m_t . As a result, the average value of prices, p_t , exceeds m_t , and thus y_t (which equals $m_t - p_t$) is negative. Finally, recall that the a_i 's are increasing in i . Thus the further in advance a change in aggregate demand is anticipated, the smaller are its real effects.

At the same time, the model is not without difficulties. As with the Christiano-Eichenbaum-Evans model, its assumptions about price-setting do not match what we observe at the microeconomic level: many prices and wages are fixed for extended periods, and there is little evidence that many price-setters or wage-setters set price or wage paths of the sort that are central to the model. And some phenomena, such as the finding described in Section 6.10 that aggregate demand disturbances appear to have smaller and less persistent real effects in higher-inflation economies, seem hard to explain without fixed prices. It is possible that to fully capture the major features of fluctuations, our microeconomic model will need to incorporate important elements both of adjustments between formal reviews, as in the models of this section, and of fixed prices.

Another limitation of the Christiano-Eichenbaum-Evans and Mankiw-Reis models, like all models of pure time-dependence, is that the assumption

of an exogenous and unchanging frequency of changes in firms' pricing plans is clearly too strong. The frequency of adjustment is surely the result of some type of optimizing calculation, not an exogenous parameter. Perhaps more importantly, it could change in response to policy changes, and this in turn could alter the effects of the policy changes. That is, a successful model may need to incorporate elements of both time-dependence and state-dependence.

This leaves us in an unsatisfactory position. It appears that any model of price behavior that does not include elements of both fixed prices and mechanical price adjustments, and elements of both time-dependence and state-dependence, will fail to capture important macroeconomic phenomena. Yet the hope that a single model could incorporate all these features and still be tractable seems far-fetched. The search for a single workhorse model of pricing behavior—or for a small number of workhorse models together with an understanding of when each is appropriate—continues.

7.8 The Canonical New Keynesian Model

The next step in constructing a complete model of fluctuations is to integrate a model of dynamic price adjustment into a larger model of the economy. Given the wide range of models of pricing behavior we have seen, it is not possible to single out one approach as the obvious starting point. Moreover, dynamic general-equilibrium models with the behavior of inflation built up from microeconomic foundations quickly become complicated. In this section, we therefore consider only an illustrative, relatively simple general-equilibrium model.

Assumptions

The specific model we consider is the canonical three-equation new Keynesian model of Clarida, Galí, and Gertler (2000). The price-adjustment equation is the new Keynesian Phillips curve of Section 7.4. This treatment of price adjustment has two main strengths. The first is its strong microeconomic foundations: it comes directly from an assumption of infrequent adjustment of nominal prices. The other is its comparative simplicity: inflation depends only on expected future inflation and current output, with no role for past inflation or for more complicated dynamics. The aggregate-demand equation of the model is the new Keynesian *IS* curve of Sections 6.1 and 7.1. The final equation describes monetary policy. So far in this chapter, because our goal has been to shed light on the basic implications of various assumptions concerning price adjustment, we have considered only simple paths of the money supply (or aggregate demand). To build a model that is more

useful for analyzing actual macroeconomic fluctuations, however, we need to assume that the central bank follows a rule for the interest rate along the lines of Section 6.4. In particular, in keeping with the forward-looking character of the new Keynesian Phillips curve and the new Keynesian *IS* curve, we assume the central bank follows a *forward-looking interest-rate rule*, adjusting the interest rate in response to changes in expected future inflation and output.

The other ingredient of the model is its shocks: it includes serially correlated disturbances to all three equations. This allows us to analyze disturbances to private aggregate demand, price-setting behavior, and monetary policy. Finally, for convenience, all the equations are linear and the constant terms are set to zero. Thus the variables should be interpreted as differences from their steady-state or trend values.

The three core equations are:

$$y_t = E_t[y_{t+1}] - \frac{1}{\theta}r_t + u_t^{IS}, \quad \theta > 0, \quad (7.84)$$

$$\pi_t = \beta E_t[\pi_{t+1}] + \kappa y_t + u_t^\pi, \quad 0 < \beta < 1, \quad \kappa > 0, \quad (7.85)$$

$$r_t = \phi_\pi E_t[\pi_{t+1}] + \phi_y E_t[y_{t+1}] + u_t^{MP}, \quad \phi_\pi > 0, \quad \phi_y \geq 0. \quad (7.86)$$

Equation (7.84) is the new Keynesian *IS* curve, (7.85) is the new Keynesian Phillips curve, and (7.86) is the forward-looking interest-rate rule. The shocks follow independent AR-1 processes:

$$u_t^{IS} = \rho_{IS}u_{t-1}^{IS} + e_t^{IS}, \quad -1 < \rho_{IS} < 1, \quad (7.87)$$

$$u_t^\pi = \rho_\pi u_{t-1}^\pi + e_t^\pi, \quad -1 < \rho_\pi < 1, \quad (7.88)$$

$$u_t^{MP} = \rho_{MP}u_{t-1}^{MP} + e_t^{MP}, \quad -1 < \rho_{MP} < 1, \quad (7.89)$$

where e^{IS} , e^π , and e^{MP} are white-noise disturbances that are uncorrelated with one another.

The model is obviously extremely stylized. To give just a few examples, all behavior is forward-looking; the dynamics of inflation and aggregate demand are very simple; and the new Keynesian Phillips curve is assumed to describe inflation dynamics despite its poor empirical performance. Nonetheless, because its core ingredients are so simple and have such appealing microeconomic foundations, the model is a key reference point in modern models of fluctuations. The model and variants of it are frequently used, and it has been modified and extended in many ways.

Because of its forward-looking elements, for some parameter values the model has sunspot solutions, like those we encountered in the model of Section 6.4. Since we discussed such solutions there and will encounter them again in our discussion of monetary policy in a model similar to this one in Section 11.5, here we focus only on the fundamental, non-sunspot solution.

The Case of White-Noise Disturbances

The first step in solving the model is to express output and inflation in terms of their expected future values and the disturbances. Applying straightforward algebra to (7.84)–(7.85) gives us

$$y_t = -\frac{\phi_\pi}{\theta} E_t[\pi_{t+1}] + \left(1 - \frac{\phi_y}{\theta}\right) E_t[y_{t+1}] + u_t^{IS} - \frac{1}{\theta} u_t^{MP}, \quad (7.90)$$

$$\pi_t = \left(\beta - \frac{\phi_\pi \kappa}{\theta}\right) E_t[\pi_{t+1}] + \left(1 - \frac{\phi_y}{\theta}\right) \kappa E_t[y_{t+1}] + \kappa u_t^{IS} + u_t^\pi - \frac{\kappa}{\theta} u_t^{MP}. \quad (7.91)$$

An important and instructive special case of the model occurs when there is no serial correlation in the disturbances (so $\rho_{IS} = \rho_\pi = \rho_{MP} = 0$). In this case, because of the absence of any backward-looking elements and any information about the future values of the disturbances, there is no force causing agents to expect the economy to depart from its steady state in the future. That is, the fundamental solution has $E_t[y_{t+1}]$ and $E_t[\pi_{t+1}]$ always equal to zero. To see this, note that with $E_t[y_{t+1}] = E_t[\pi_{t+1}] = 0$, equations (7.86), (7.90), and (7.91) simplify to

$$y_t = u_t^{IS} - \frac{1}{\theta} u_t^{MP}, \quad (7.92)$$

$$\pi_t = \kappa u_t^{IS} + u_t^\pi - \frac{\kappa}{\theta} u_t^{MP}, \quad (7.93)$$

$$r_t = u_t^{MP}. \quad (7.94)$$

If (7.92)–(7.94) describe the behavior of output, inflation, and the real interest rate, then, because we are considering the case where the u 's are white noise, the expectations of future output and inflation are always zero. (7.92)–(7.94) therefore represent the fundamental solution to the model in this case.

These expressions show the effects of the various shocks. A contractionary monetary-policy shock raises the real interest rate and lowers output and inflation. A positive shock to private aggregate demand raises output and inflation and has no impact on the real interest rate. And an unfavorable inflation shock raises inflation but has no other effects. These results are largely conventional. The *IS* shock fails to affect the real interest rate because monetary policy is forward-looking, and so does not respond to the increases in current output and inflation. The fact that monetary policy is forward-looking is also the reason the inflation shock does not spill over to the other variables.

The key message of this case of the model, however, is that the model, like the baseline real-business-cycle model of Chapter 5, has no internal propagation mechanisms. Serial correlation in output, inflation, and the real interest rate can come only from serial correlation in the driving processes.

As a result, a major goal of extensions and variations of the model—such as those we will discuss in the next section—is to introduce forces that cause one-time shocks to trigger persistent changes in the macroeconomy.

The General Case

A straightforward way to solve the model in the general case is to use the method of undetermined coefficients. Given the model's linear structure and absence of backward-looking behavior, it is reasonable to guess that the endogenous variables are linear functions of the disturbances. For output and inflation, we can write this as

$$y_t = a_{IS}u_t^{IS} + a_{\pi}u_t^{\pi} + a_{MP}u_t^{MP}, \quad (7.95)$$

$$\pi_t = b_{IS}u_t^{IS} + b_{\pi}u_t^{\pi} + b_{MP}u_t^{MP}. \quad (7.96)$$

This conjecture and the assumptions about the behavior of the disturbances in (7.87)–(7.89) determine $E_t[y_{t+1}]$ and $E_t[\pi_{t+1}]$: $E_t[y_{t+1}]$ equals $a_{IS}\rho_{IS}u_t^{IS} + a_{\pi}\rho_{\pi}u_t^{\pi} + a_{MP}\rho_{MP}u_t^{MP}$, and similarly for $E_t[\pi_{t+1}]$. We can then substitute these expressions and (7.95) and (7.96) into (7.90) and (7.91). This yields:

$$\begin{aligned} a_{IS}u_t^{IS} + a_{\pi}u_t^{\pi} + a_{MP}u_t^{MP} &= -\frac{\Phi_{\pi}}{\theta}(b_{IS}\rho_{IS}u_t^{IS} + b_{\pi}\rho_{\pi}u_t^{\pi} + b_{MP}\rho_{MP}u_t^{MP}) \\ &+ \left(1 - \frac{\Phi_Y}{\theta}\right)(a_{IS}\rho_{IS}u_t^{IS} + a_{\pi}\rho_{\pi}u_t^{\pi} + a_{MP}\rho_{MP}u_t^{MP}) + u_t^{IS} - \frac{1}{\theta}u_t^{MP}, \end{aligned} \quad (7.97)$$

$$\begin{aligned} b_{IS}u_t^{IS} + b_{\pi}u_t^{\pi} + b_{MP}u_t^{MP} &= \left(\beta - \frac{\Phi_{\pi}\kappa}{\theta}\right)(b_{IS}\rho_{IS}u_t^{IS} + b_{\pi}\rho_{\pi}u_t^{\pi} + b_{MP}\rho_{MP}u_t^{MP}) \\ &+ \left(1 - \frac{\Phi_Y}{\theta}\right)\kappa(a_{IS}\rho_{IS}u_t^{IS} + a_{\pi}\rho_{\pi}u_t^{\pi} + a_{MP}\rho_{MP}u_t^{MP}) + \kappa u_t^{IS} + u_t^{\pi} - \frac{\kappa}{\theta}u_t^{MP}. \end{aligned} \quad (7.98)$$

For the equations of the model to be satisfied when output and inflation are described by equations (7.95) and (7.96), the two sides of (7.97) must be equal for all values of u_t^{IS} , u_t^{π} , and u_t^{MP} . That is, the coefficients on u_t^{IS} on the two sides must be equal, and similarly for the coefficients on u_t^{π} and u_t^{MP} . This gives us three equations—one involving a_{IS} and b_{IS} , one involving a_{π} and b_{π} , and one involving a_{MP} and b_{MP} . Equation (7.98) gives us three more equations. Once we have found the a 's and b 's, equations (7.95) and (7.96) tell us the behavior of output and inflation. We can then use (7.86) and the expressions for $E_t[\pi_{t+1}]$ and $E_t[y_{t+1}]$ to find the behavior of the real interest rate. Thus solving the model is just a matter of algebra.

Unfortunately, the equations determining the a 's and b 's are complicated, the algebra is tedious, and the resulting solutions for the a 's and b 's are complex and unintuitive. To get a sense of the model's implications, we will therefore assume values for the parameters and find their implications for how the economy responds to shocks. Specifically, following Galí (2008, Section 3.4.1), we interpret a time period as a quarter, and assume $\theta = 1$,

$\kappa = 0.1275$, $\beta = 0.99$, $\phi_\pi = 0.5$, and $\phi_y = 0.125$. For each of the disturbances, we will consider both the case of no serial correlation and a serial correlation coefficient of 0.5 to see how serial correlation affects the behavior of the economy.

Consider first a monetary-policy shock. With $\rho_{MP} = 0$, our parameter values and equations (7.92)–(7.94) imply that $y_t = -u_t^{MP}$, $\pi_t = -0.13u_t^{MP}$, and $r_t = u_t^{MP}$. With $\rho_{MP} = 0.5$, they imply that $y_t = -1.60u_t^{MP}$, $\pi_t = -0.40u_t^{MP}$, and $r_t = 0.80u_t^{MP}$. Intuitively, the fact that output and inflation will be below normal in later periods mutes the rise in the real interest rate. But because of the fall in future output, a larger fall in current output is needed for households to satisfy their Euler equation in response to the rise in the real rate. And both the greater fall in output and the decline in future inflation strengthen the response of inflation. As the economy returns to its steady state, the real rate is above normal and output is rising, consistent with the new Keynesian *IS* curve. And inflation is rising and output is below normal, consistent with the new Keynesian Phillips curve.

Next, consider an *IS* shock. When $\rho_{IS} = 0$, our parameter values imply $y_t = u_t^{IS}$, $\pi_t = 0.13u_t^{IS}$, and $r_t = 0$. When ρ_{IS} rises to 0.5, we obtain $y_t = 1.60u_t^{IS}$, $\pi_t = 0.40u_t^{IS}$, and $r_t = 0.20u_t^{IS}$. Again, the impact of the shock on future output magnifies the output response via the new Keynesian *IS* curve. In addition, the increases in future inflation strengthen the inflation response through the new Keynesian Phillips curve. And with future output and inflation affected by the shock, the current real interest rate responds through the forward-looking interest-rate rule.

Finally, consider an inflation shock. As described above, in the absence of serial correlation, the shock is translated one-for-one into inflation and has no effect on output or the real interest rate. With $\rho_\pi = 0.5$, in contrast, $y_t = -0.80u_t^\pi$, $\pi_t = 1.78u_t^\pi$, and $r_t = 0.40u_t^\pi$. The persistence of the inflation shock increases the response of current inflation (through the forward-looking term of the new Keynesian Phillips curve) and raises the real interest rate (through the inflation term of the forward-looking interest-rate rule). The increase in the real rate reduces current output through the *IS* curve; and this effect is magnified by the fact that the curve is forward-looking.

7.9 Other Elements of Modern New Keynesian DSGE Models of Fluctuations

The model of Section 7.8 is a convenient illustrative model. But it is obviously far short of being rich enough to be useful for many applications. A policymaker wanting to forecast the path of the economy or evaluate the likely macroeconomic effects of some policy intervention would certainly need a considerably more complicated model.

A large and active literature is engaged in constructing and estimating more sophisticated quantitative DSGE models that, at their core, have important resemblances to the model of the previous section. The models do not lend themselves to analytic solutions or to transparency. But they are in widespread use not just in academia, but in central banks and other policymaking institutions. This section briefly discusses some of the most important modifications and extensions of the baseline model. Many of the changes come from the models of Christiano, Eichenbaum, and Evans (2005), Erceg, Henderson, and Levin (2000), and Smets and Wouters (2003).

Aggregate Supply

The canonical new Keynesian model uses the new Keynesian Phillips curve to model the behavior of inflation. Richer models often extend this in two ways. First, recall that the evidence in favor of the distinctive predictions of the new Keynesian Phillips curve—notably its implication that an anticipated disinflation is associated with an output boom—is weak. Thus modern models often introduce inflation inertia. Because of its tractability, the usual approach is to posit a relationship along the lines suggested by the new Keynesian Phillips curve with indexation. Typically, the coefficients on lagged and expected future inflation are not constrained to equal $1/(1 + \beta)$ and $\beta/(1 + \beta)$, as in equation (7.76), but follow the more general set of possibilities allowed by equation (7.77).

Second, to better capture the behavior of prices and wages, the models often assume incomplete adjustment not just of goods prices, but also of wages. The most common approach is to assume Calvo wage adjustment (with an adjustment frequency potentially different from that for price changes). Under appropriate assumptions, the result is a new Keynesian Phillips curve for wage inflation:

$$\pi_t^w = \beta E_t[\pi_{t+1}^w] + \kappa_w \gamma_t, \quad (7.99)$$

where π^w is wage inflation. A natural alternative, paralleling the treatment of prices, is to assume indexation to lagged wage inflation between adjustments, leading to an equation for wage inflation analogous to the new Keynesian Phillips curve with indexation.

Aggregate Demand

There are two major limitations of the new Keynesian *IS* curve. First, and most obviously, it leaves out investment, government purchases, and net exports. Virtually every model intended for practical use includes investment modeled as arising from the decisions of profit-maximizing firms. Government purchases are almost always included as well; they are generally

modeled as exogenous. And there are numerous open-economy extensions. Examples include Obstfeld and Rogoff (2002); Corsetti and Pesenti (2005); Benigno and Benigno (2006); and Galí (2008, Chapter 7).

Second, the basic new Keynesian *IS* curve, even when it is extended to include other components of output, tends to imply large and rapid responses to shocks. To better match the data, the models therefore generally include ingredients that slow adjustment. With regard to consumption, the most common approach is to assume *habit formation*. That is, a consumer's utility is assumed to depend not just on the level of consumption, but also on its level relative to some reference amount, such as the consumer's or others' past consumption. Under appropriate assumptions, this slows the response of consumption to shocks. On the investment side, the most common way of slowing responses is to assume directly that there are costs of adjusting investment.

We will see in Chapter 8 that households' current income appears to have an important effect on their consumption, and we will see in Chapter 9 that firms' current cash flow may be important to their investment decisions. The new Keynesian *IS* curve, with or without the various modifications we have discussed, does not allow for these possibilities. To let current income affect the demand for goods, the usual approach is to assume that some fraction of consumption is determined by rule-of-thumb or liquidity-constrained households that devote all their current income to consumption.²¹ This assumption can magnify the economy's responses to various disturbances and can introduce a role for shocks that shift the timing of income, which would otherwise not affect behavior.

Credit-Market Imperfections

The crisis of 2008–2009 has made it clear that non-Walrasian features of credit markets have important macroeconomic consequences. Disruptions in credit markets can cause large swings in economic activity, and credit-market imperfections can have large effects on how other shocks affect the macroeconomy. As a result, introducing credit-market imperfections into new Keynesian DSGE models is an active area of research.

Three recent efforts in this area are those by Cúrdia and Woodford (2009), Gertler and Karadi (2009), and Christiano, Motto, and Rostagno (2009). In all three models, there is a financial sector that intermediates between saving and investment. Cúrdia and Woodford's model is conceptually the simplest.

²¹ The models generally do not give current cash flow a role in investment. For some purposes, the assumption of rule-of-thumb consumers has similar implications, making it unnecessary to add this complication. In addition, some models that include credit-market imperfections, along the lines of the ones we will discuss in a moment, naturally imply an impact of cash flow on investment.

They assume a costly intermediation technology. The spread between borrowing and lending rates changes because of changes both in the marginal cost of intermediation and in intermediaries' markups. These fluctuations have an endogenous component, with changes in the quantity of intermediation changing its marginal cost, and an exogenous component, with shocks to both the intermediation technology and markups.

In Gertler and Karadi's model, the spread arises from constraints on the size of the intermediation sector. Intermediaries have limited capital. Because high leverage would create harmful incentives, the limited capital restricts intermediaries' ability to attract funds from savers. The result is that they effectively earn rents on their capital, charging more to borrowers than they pay to savers. Again, the spread moves both endogenously and exogenously. Various types of shocks affect intermediaries' capital, and so change their ability to attract funds and the spread. And shocks to the value of their capital directly affect their ability to attract funds, and so again affect the spread. Both endogenous and exogenous movements in the spread are propagated to the remainder of the economy.

Christiano, Motto, and Rostagno, building on their earlier work (Christiano, Motto, and Rostagno, 2003), focus on frictions in the relationship between intermediaries and borrowers. The limited capital of borrowers and the riskiness of their investments affect their ability to borrow and the interest rates they must pay. As a result, borrowing rates and the quantity of borrowing move endogenously in response to various types of disturbances. In addition, Christiano, Motto, and Rostagno assume that loan contracts are written in nominal terms (along the lines we discussed in Section 6.9), so that any disturbance that affects the price level affects borrowers' real indebtedness, which in turn affects the rest of the economy. And, as in the other models, there are exogenous disturbances to the factors governing spreads. Christiano, Motto, and Rostagno consider not only shocks to borrowers' net worth and to the riskiness of their projects, but also the arrival of news about the riskiness of future projects.

All three papers represent early efforts to incorporate financial-market imperfections and disruptions into larger models. Recent events leave no doubt that those imperfections and disruptions are important. But the question of how to best incorporate them in larger macroeconomic models is very much open.

Policy

The policy assumptions of more sophisticated new Keynesian DSGE models of fluctuations depart from the simple interest-rate rule we considered in Section 7.8 in three main ways. The first, and most straightforward, is to consider other interest-rate rules. A seemingly infinite variety of interest-rate rules have been considered. The rules consider gradual adjustment,

responses to current values or past values of variables instead of (or in addition to) their expected future values, responses to growth rates rather than levels of variables, and the possible inclusion of many variables other than output and inflation. A common strategy in this literature is to ask how some change in the rule, such as the addition of a new variable, affects macroeconomic outcomes, such as the variability of inflation and output.

The second, larger departure is to replace the assumption of a prespecified policy rule with the assumption that policymakers maximize some objective function. The objective function may be specified directly; for example, policymakers can be assumed to have a quadratic loss function over inflation and output. Alternatively, the function may be derived from microeconomic foundations; most commonly, policymakers' goal is assumed to be to maximize the expected utility of the representative household in the model. With this approach, it is necessary to specify a model rich enough that inflation affects welfare. Once the objective is in place (either by assumption or by derivation), policymakers' decisions come from maximizing that function.

A natural way to meld the approach based on interest-rate rules and the approach based on maximization is to ask how well various simple rules approximate optimal policy. There is a widespread view that policymakers would be reluctant to follow a complicated rule or the prescriptions of one particular model. Thus it is important to ask whether there are simple rules that perform relatively well across a range of models. We will investigate both modifications of simple interest-rate rules and the derivation of optimal policy further in Chapter 11, where we examine monetary policy in more depth.

The third way that recent models extend the analysis of policy is by considering policy instruments other than the short-term interest rate. One set of additional policy instruments are those associated with fiscal policy, notably government purchases, transfers, and various tax rates. And models that incorporate imperfections in credit markets naturally allow for consideration of various government interventions in those markets.

Discussion

Assessments of this research program fall along a continuum between two extremes. Although few economists are at either extreme, they are useful reference points.

One extreme is that we are well on the way to having models of the macroeconomy that are sufficiently well grounded in microeconomic assumptions that their parameters can be thought of as structural (in the sense that they do not change when policies change), and that are sufficiently realistic that they can be used to obtain welfare-based recommendations about the conduct of policy. Advocates of this view can point to the

facts that the models are built up from microeconomic foundations; that estimated versions of the models match some important features of fluctuations reasonably well; that many policymakers value the models enough to put weight on their predictions and recommendations; that there is microeconomic evidence for many of their assumptions; and that their sophistication is advancing rapidly.

The other extreme is that the models are ad hoc constructions that are sufficiently distant from reality that their policy recommendations are unreliable and their predictions likely to fail if the macroeconomic environment changes. Advocates of this view can point to two main facts. First, despite the models' complications, there is a great deal they leave out. For example, until the recent crisis, the models' treatment of credit-market imperfections was generally minimal. Second, the microeconomic case for some important features of the models is questionable. Most notably, the models include assumptions that generate inertia in decision making: inflation indexation in price adjustment, habit formation in consumption, and adjustment costs in investment. The inclusion of these features is mainly motivated not by microeconomic evidence, but by a desire to match macroeconomic facts. For example, at the microeconomic level we see nominal prices that are fixed for extended periods, not frequently adjusted to reflect recent inflation. Similarly, as we will see in Chapter 9, standard models of investment motivated by microeconomic evidence involve costs of adjusting the capital stock, not costs of adjusting investment. The need to introduce these features, in this view, suggests that the models have significant gaps.

Almost all macroeconomists agree that the models have important strengths and weaknesses, and thus that the truth lies between the two extremes. Nonetheless, where in that range the truth is matters for how macroeconomists should conduct their research. The closer it is to the first extreme, the greater the value of extending the models and of examining new phenomena by incorporating them into the models. The closer it is to the second extreme, the greater the value of working on new issues in narrower models and of postponing efforts to construct integrative models until our understanding of the component pieces is further advanced.

Problems

- 7.1. The Fischer model with unbalanced price-setting.** Suppose the economy is described by the model of Section 7.2, except that instead of half of firms setting their prices each period, fraction f set their prices in odd periods and fraction $1 - f$ set their prices in even periods. Thus the price level is $f p_t^1 + (1 - f) p_t^2$ if t is even and $(1 - f) p_t^1 + f p_t^2$ if t is odd. Derive expressions analogous to (7.27) and (7.28) for p_t and y_t for even and odd periods.
- 7.2. The instability of staggered price-setting.** Suppose the economy is described as in Problem 7.1, and assume for simplicity that m is a random walk

(so $m_t = m_{t-1} + u_t$, where u is white noise and has a constant variance). Assume the profits a firm loses over two periods relative to always having $p_t = p_t^*$ is proportional to $(p_{it} - p_{it}^*)^2 + (p_{it+1} - p_{it+1}^*)^2$. If $f < 1/2$ and $\phi < 1$, is the expected value of this loss larger for the firms that set their prices in odd periods or for the firms that set their prices in even periods? In light of this, would you expect to see staggered price-setting if $\phi < 1$?

7.3. Synchronized price-setting. Consider the Taylor model. Suppose, however, that every other period all the firms set their prices for that period and the next. That is, in period t prices are set for t and $t + 1$; in $t + 1$, no prices are set; in $t + 2$, prices are set for $t + 2$ and $t + 3$; and so on. As in the Taylor model, prices are both predetermined and fixed, and firms set their prices according to (7.30). Finally, assume that m follows a random walk.

- What is the representative firm's price in period t , x_t , as a function of m_t , $E_t m_{t+1}$, p_t , and $E_t p_{t+1}$?
- Use the fact that synchronization implies that p_t and p_{t+1} are both equal to x_t to solve for x_t in terms of m_t and $E_t m_{t+1}$.
- What are y_t and y_{t+1} ? Does the central result of the Taylor model—that nominal disturbances continue to have real effects after all prices have been changed—still hold? Explain intuitively.

7.4. Consider the Taylor model with the money stock white noise rather than a random walk; that is, $m_t = \varepsilon_t$, where ε_t is serially uncorrelated. Solve the model using the method of undetermined coefficients. (Hint: In the equation analogous to (7.33), is it still reasonable to impose $\lambda + \nu = 1$?)

7.5. Repeat Problem 7.4 using lag operators.

7.6. Consider the experiment described at the beginning of Section 7.4. Specifically, a Calvo economy is initially in long-run equilibrium with all prices equal to m , which we normalize to zero. In period 1, there is a one-time, permanent increase in m to m_1 .

Let us conjecture that the behavior of the price level for $t \geq 1$ is described by an expression of the form $p_t = (1 - \lambda^t)m_1$.

- Explain why this conjecture is or is not reasonable.
- Find λ in terms of the primitive parameters of the model (α , β , and ϕ).
- How do increases in each of α , β , and ϕ affect λ ? Explain your answers intuitively.

7.7. State-dependent pricing with both positive and negative inflation. (Caplin and Leahy, 1991.) Consider an economy like that of the Caplin-Spulber model. Suppose, however, that m can either rise or fall, and that firms therefore follow a simple two-sided Ss policy: if $p_i - p_i^*(t)$ reaches either S or $-S$, firm i changes p_i so that $p_i - p_i^*(t)$ equals 0. As in the Caplin-Spulber model, changes in m are continuous.

Assume for simplicity that $p_i^*(t) = m(t)$. In addition, assume that $p_i - p_i^*(t)$ is initially distributed uniformly over some interval of width S ; that is, $p_i - p_i^*(t)$ is distributed uniformly on $[X, X + S]$ for some X between $-S$ and 0.

- (a) Explain why, given these assumptions, $p_t - p_t^*(t)$ continues to be distributed uniformly over some interval of width S .
- (b) Are there any values of X for which an infinitesimal increase in m of dm raises average prices by less than dm ? by more than dm ? by exactly dm ? Thus, what does this model imply about the real effects of monetary shocks?
- 7.8.** (This follows Ball, 1994a.) Consider a continuous-time version of the Taylor model, so that $p(t) = (1/T) \int_{\tau=0}^T x(t-\tau)d\tau$, where T is the interval between each individual's price changes and $x(t-\tau)$ is the price set by individuals who set their prices at time $t-\tau$. Assume that $\phi = 1$, so that $p_t^*(t) = m(t)$; thus $x(t) = (1/T) \int_{\tau=0}^T E_t m(t+\tau)d\tau$.
- (a) Suppose that initially $m(t) = gt$ ($g > 0$), and that $E_t m(t+\tau)$ is therefore $(t+\tau)g$. What are $x(t)$, $p(t)$, and $y(t) = m(t) - p(t)$?
- (b) Suppose that at time 0 the government announces that it is steadily reducing money growth to zero over the next interval T of time. Thus $m(t) = t[1 - (t/2T)]g$ for $0 < t < T$, and $m(t) = gT/2$ for $t \geq T$. The change is unexpected, so that prices set before $t = 0$ are as in part (a).
- (i) Show that if $x(t) = gT/2$ for all $t > 0$, then $p(t) = m(t)$ for all $t > 0$, and thus that output is the same as it would be without the change in policy.
- (ii) For $0 < t < T$, are the prices that firms set more than, less than, or equal to $gT/2$? What about for $T \leq t \leq 2T$? Given this, how does output during the period $(0, 2T)$ compare with what it would be without the change in policy?
- 7.9.** Consider the new Keynesian Phillips curve with indexation, equation (7.76), under the assumptions of perfect foresight and $\beta = 1$, together with our usual aggregate demand equation, $y_t = m_t - p_t$.
- (a) Express p_{t+1} in terms of its lagged values and m_t .
- (b) Consider an anticipated, permanent, one-time increase in m : $m_t = 0$ for $t < 0$, $m_t = 1$ for $t \geq 0$. Sketch how you would find the resulting path of p_t . (Hint: Use the lag operator approach from Section 7.3.)
- 7.10. The new Keynesian Phillips curve with partial indexation.** Consider the analysis of the new Keynesian Phillips curve with indexation in Section 7.7. Suppose, however, that the indexation is only partial. That is, if a firm does not have an opportunity to review its price in period t , its price in t is the previous period's price plus $\gamma\pi_{t-1}$, $0 \leq \gamma \leq 1$. Find an expression for π_t in terms of π_{t-1} , $E_t\pi_{t+1}$, y_t , and the parameters of the model. Check that your answer simplifies to the new Keynesian Phillips curve when $\gamma = 0$ and to the new Keynesian Phillips curve with indexation when $\gamma = 1$. (Hint: Start by showing that $[\alpha/(1-\alpha)](x_t - p_t) = \pi_t - \gamma\pi_{t-1}$.)
- 7.11.** Consider a continuous-time version of the Mankiw-Reis model. Opportunities to review pricing policies follow a Poisson process with arrival rate $\alpha > 0$. Thus the probability that a price path set at time t is still being followed at time $t+i$ is $e^{-\alpha i}$. The other assumptions of the model are the same as before.

(a) Show that the expression analogous to (7.81) is $a(i) = \frac{\phi(1 - e^{-\alpha i})}{[1 - (1 - \phi)(1 - e^{-\alpha i})]}$.

(b) Consider the experiment of a permanent fall in the growth rate of aggregate demand discussed in Section 7.7. That is, until $t = 0$, all firms expect $m(t) = gt$ (where $g > 0$); thereafter, they expect $m(t) = 0$.

(i) Find the expression analogous to (7.83).

(ii) Find an expression for inflation, $\dot{p}(t)$, for $t \geq 0$. Is inflation ever negative during the transition to the new steady state?

(iii) Suppose $\phi = 1$. When does output reach its lowest level? When does inflation reach its lowest level?

7.12. Consider the model of Section 7.8. Suppose, however, that monetary policy responds to current inflation and output: $r_t = \phi_\pi \pi_t + \phi_y y_t + u_t^{MP}$.

(a) For the case of white-noise disturbances, find expressions analogous to (7.92)–(7.94). What are the effects of an unfavorable inflation shock in this case?

(b) Describe how you would solve this model using the method of undetermined coefficients (but do not actually solve it).