12 DYNAMIC MODELS WITH STICKY PRICES

Chapter 12 Contents

A. Topics and Tools

In Romer's Chapter 6, we studied a firm's decision to change prices vs. keeping prices sticky as though the price change were an isolated event that would happen only once. Firms in the Chapter 6 model have a pre-set menu price of ambiguous origin, then decide whether or not to change it taking into account the current period's profits at the pre-existing price vs. the optimal price.

A more complete model would consider the implications of today's price setting for future profits as well as current profits. The price that the firm sets today—whether it be the pre-existing price or a newly changed one—becomes the pre-set menu price for the next period, so it has an effect that extends beyond the current period. Chapter

7 explores dynamic models of price-setting using the tools that we developed in Chapter 6.

Section 7.1 develops a general framework for optimal price-setting in a dynamic model. This framework is then applied to alternative situations in subsequent sections. Section 7.2 examines a "predetermined-price" model in which firms make pricing decisions for two periods at a time, though they may set a different price for the first and second period. There are two groups of firms that set prices at different times with one group making two-period pricing decisions in even periods and the other in odd periods, so in each period half of the prices are newly set and half were set one period before.

Section 7.3 considers a "fixed-price" model that is similar to the predetermined-price model except that firms set the same price for the first and second periods on their price "contract." Section 7.4 examines a workhorse model of the literature, the Calvo model, in which a fixed fraction α of randomly chosen firms re-set their prices each period. For example, if $\alpha = 25\%$, then a firm would have a 25% chance of resetting its price in any given quarter, so on average the firm sets its price once per year.

The models of Sections 7.2 through 7.4 all have "time-dependent pricing," in which the decision to change price does not depend on economic conditions. In the real world, it is likely that a large shock would cause firms to change their pricing strategies regardless of how long their existing prices had been in effect. In other words, the length of time over which prices are fixed (or the probability that a firm resets its price) is endogenous. Section 7.5 considers two important models with "state-dependent pricing."

Price stickiness alone cannot explain a widely observed phenomenon of modern economies: inflation inertia. Empirical evidence suggests that *inflation* is sometimes sticky, which cannot be caused by menu costs. (For example, menu costs present no impediment to firms in reducing inflation from a positive value to zero, but in fact firms seem to continue raising prices even after aggregate-demand growth slows down.) After a digression on empirical studies of price stickiness (which are covered in Coursebook Chapter 13), Section 7.7 covers several models that aim to explain inflation inertia. Sections 7.8 and 7.9 then conclude the chapter by summarizing a "canonical" new Keynesian model and variations on it.

B. Price Stickiness vs. Inflation Inertia

As noted above, there is an important difference between the stickiness of prices and the stickiness of inflation, which we call inflation inertia. Chapter 6 of Romer taught us that menu costs can lead to price stickiness, especially if they are combined with real rigidities. But if that is the only source of rigidity in pricing, it should be trivially easy to end inflation. There are no costs to *not changing prices*, which means that firms would have no reluctance to keeping their prices constant (zero inflation) if aggregate demand were to stop growing. Thus, ending inflation should be costless, with no increase in unemployment or recession in output and no carryover inflation.

That does not seem to be the case. To explain why disinflations (reductions in the inflation rate) would be associated with temporary output declines, we need some rigidity in *inflation*, not just in *prices*. With such ***inflation inertia***, disinflation (if not perfectly anticipated) lowers inflation below what people expected and have built into their plans, leading to prices that are "too high" and reductions in the demand for output relative to the natural level: recessions.

Menu costs provide a simple intuitive story leading to price stickiness. Is there a similarly simple and compelling argument for inflation inertia? Why would firms continue to raise prices (and incur menu costs) when aggregate demand is no longer rising? This requires a different and complementary theory of nominal stickiness—one that applies to adjusting the rate of change of prices rather than their level.

To see where such stickiness might occur, consider the steps that a modern firm must undertake to change its pricing policy. First, it must gain the information necessary to decide on the optimal price strategy. Then it must process that information and decide on an optimal strategy, which often involves a meeting of central and regional management executives. Once it has decided on a strategy, the firm must implement it. Note that it is only at the last stage that menu costs may be important.

If a large part of the cost of setting a pricing strategy is information gathering and decision making, then firms will want to economize on the frequency that they undertake these tasks, just as the economize on how frequently they change physical price tags in the presence of menu costs. Suppose that the firm chooses to do this only once per year. Further suppose that the firm has just chosen a plan to increase prices by 1% each quarter for the next year, but now aggregate demand growth shudders to a halt and a 0% price change would have been optimal. It might be better for the firm to continue with its 1% per quarter increases than to collect new information and re-do the strategy meeting to change its pricing policy. In such a case, *inflation* would have inertia—prices would continue to rise based on inflation inertia even after the aggregate-demand stimulus to inflation ceased and despite menu costs that would save the firm money if it immediately stopped raising inflation.

Romer presents two models that incorporate inflation inertia, one by Christiano, Eichenbaum, and Evans and one by Mankiw and Reis. In the former, firms that are between pricing-decision meetings keep raising prices at the former inflation rate, which causes prices of such firms to keep rising as in the example above. In the Mankiw and Reis model, firms have "sticky information"—it is at the information-gathering stage of pricing decisions that they react only with a lag. Both models have similar implications: the current inflation rate in the Phillips curve depends both on expected future inflation *and on past inflation*. This leads to inflation inertia in the models.

C. Understanding Romer's Chapter 7

Part B of Romer's Chapter 6 examined the incentives of each individual firm in deciding whether to change its price or keep it fixed. In Chapter 7, we embed these firms into a macroeconomic model and consider the macroeconomic implications of price stickiness.

Romer's "building blocks"

Romer begins in Section 7.1 by developing a dynamic version of the imperfect competition model of Section 6.5. This model is based on utility maximization by households and profit maximization by firms, so its microfoundations are quite completely developed. Most of the elements of this model are familiar from the imperfect competition model of Chapter 6, but some take slightly new forms.

For example, the utility function (7.1) is a discrete-time lifetime utility function similar to ones we used in the Diamond growth model, the real-business-cycle model, and the new Keynesian model in Chapter 6. Utility is an "additively separable" sum of utility from consumption and disutility from labor. The additivity of the utility function simplifies the analysis by making the marginal utility of consumption independent of labor and vice versa. The condition $V' > 0$ means that more work leads to more disutility (working is disliked), and $V'' > 0$ implies that the more you work, the greater is the marginal disutility of work. These conditions are the flip-side of an assumption that leisure has positive but diminishing marginal utility.

The discount factor in equation (7.1) is written simply as $\beta \in (0,1)$. As noted in the discussion of the imperfect competition model of Chapter 6, you can think of β as equal to $1/(1 + \rho)$ if you wish, with ρ being the marginal rate of time preference; it is just a more compact notation.

Equation (7.4) is the first-order condition relating to the trade-off between consumption at time t and labor at time t. It says that the marginal disutility of working (the left-hand side) must equal the marginal utility of the goods that can be bought with an additional unit of work (the right-hand side).

The new Keynesian *IS* curve in equation (7.7) is the same one we derived in Romer's equation (6.8). As with traditional *IS* curves, it slopes downward in (Y, r) space.

The theory of the firm in the discussion on pages 313 through 316 is a little tricky. We usually simply assume that each firm maximizes the present value of its stream of profits. Here, the firm is assumed to maximize the *utility* of its stream of profits to the shareholders. With a competitive credit market, these assumptions are equivalent. To see this, consider Romer's equation (2.48) in the discussion of the Diamond model on page 79. This equation applies to the equilibrium between consumption in periods 1 and 2. Solving it for $1 + r_{t+1}$ yields

$$1+r_{t+1}=(1+\rho)\frac{C_{2,t+1}^{\theta}}{C_{1,t}^{\theta}}=(1+\rho)\frac{C_{1,t}^{-\theta}}{C_{2,t+1}^{-\theta}}=(1+\rho)\frac{U'(C_{1,t})}{U'(C_{2,t+1})}. \tag{1}$$

The right-hand equality in equation (1) follows directly from the definition of the utility function. In the Diamond model, individuals live for only two periods, so the only relevant comparison is between t and $t + 1$.

The owners of firms in the new Keynesian model are longer-lived, so we must also consider consumption tradeoffs between more distant points in time. If we were to generalize equation (1) to reflect the tradeoff between consumption at time zero and time t, the corresponding equation would be

$$1+\overline{r}_t\equiv\prod_{s=1}^{t}(1+r_s)=(1+\rho)^t\frac{U'(C_0)}{U'(C_t)}. \tag{2}$$

Taking the reciprocals of both sides of equation (2) yields

$$\frac{1}{1+\overline{r}_t}=\left(\frac{1}{1+\rho}\right)^t\frac{U'(C_t)}{U'(C_0)}. \tag{3}$$

If, as we suggested above, the discount factor β can be thought of as $1/(1 + \rho)$, then we can rewrite (3) as

$$\frac{1}{1+\overline{r}_t}=\beta^t\frac{U'(C_t)}{U'(C_0)}=\lambda_t. \tag{4}$$

This λ_t term is defined by Romer in text in the paragraph after equation (7.8). From the derivation above, we can see that it serves the same role as the usual discount factor involving the interest rate. In particular, if the interest rate were constant between time 0 and time t, equation (4) would simplify to

$$\left(\frac{1}{1+r}\right)^t = \beta^t \frac{U'(C_t)}{U'(C_0)} = \lambda_t. \tag{5}$$

Thus, the λ_t term in equation (7.9) can be interpreted as a discount factor in which the equilibrium interest rate from the consumption side of the model has been substituted in.

Another potentially confusing component of equation (7.9) is q_t, which denotes the probability that a price set today has not been changed t periods later. This probability depends on the firm's future decisions about whether or not to change price—the decisions we analyzed in the previous section. In the remaining sections of Chapter 7, Romer looks at several alternative models for q, including time-dependent models in which the pattern of price-changing is exogenous (as with fixed-length contracts) and state-dependent models in which the decision to change prices depends on economic conditions. For now, we simply treat q as a parameter, leaving its determination unspecified.

This leads us to Romer's maximand shown in equation (7.9), which is more complex than it appears because of uncertainty. By making a couple of reasonable simplifying assumptions and using a second-order Taylor series approximation to the effect of prices on profits, he arrives at equation (7.14), which has a useful intuitive interpretation.

To understand this equation you need to be very clear about what p_i and p_t^* represent.

- p_i is the price that the firm sets now, knowing that it will be in place both in the current period and (probably) in some future periods.
- p_t^* is the price that would be ideal for the firm to set for period t if it were to set the price independently for each period. Firms would set $p_t = p_t^*$ in every period if there were no costs of price adjustment.

Equation (7.13) shows that the firm should set a price that equals the average of the ideal prices in each future period, weighted in proportion to the probability that the current price will still be in effect during that future period. For example, if it is known that the newly set price will be in effect for two periods, then the optimal price for the firm to set is the (unweighted) average between the desired price in the first period and the desired price in the second period. If the new price will be in effect for the first period and there is a 50% chance it will be in effect for the second period (but not any longer), then the firm should set the price at a weighted average of the two ideal prices with a 2/3 weight given to the current period and a 1/3 weight to the second.

Equation (7.14) and its discounted form (7.15) are central to the dynamic new Keynesian model. They describe the solution to a basic problem: how to set a price that will carry over into future time periods. The solution is a logical one: set a price that is the average of the prices you'll want in the future.

The final piece of the puzzle in this section is the somewhat cryptic paragraph starting at the bottom of page 315. He asserts that the "profit-maximizing real price is [proportional to] the real wage." If we can set a distinct price for period t, then we would want to maximize R_t in equation (7.8). We can derive Romer's result easily by maximizing equation (7.8) with respect to (P_{it}/P) and setting the result equal to zero:

$$\frac{dR_t}{d\left(P_{it}/P\right)} = Y_t\left[(1-\eta)\left(\frac{P_{it}}{P}\right)^{-\eta} + \eta\frac{W_t}{P_t}\left(\frac{P_{it}}{P}\right)^{-\eta-1}\right]. \tag{6}$$

The derivative in (6) can equal zero only if the bracketed expression is zero, which implies

$$(1-\eta)\left(\frac{P_{it}}{P}\right)^{-\eta} = -\eta\frac{W_t}{P_t}\left(\frac{P_{it}}{P}\right)^{-\eta-1}$$

or

$$\frac{P_{it}}{P} = \frac{\eta}{\eta-1}\frac{W_t}{P_t}. \tag{7}$$

Equation (7), which is identical to Romer's equation (6.57), demonstrates the assertion that the desired price is proportional to the real wage.

From Romer's equation (7.6),

$$\frac{W_t}{P_t} = BY_t^{\gamma+\theta-1}. \tag{8}$$

Plugging (8) into (7) yields

$$\frac{P_{it}}{P} = B\frac{\eta}{\eta-1}Y_t^{\gamma+\theta-1},$$

or in log terms,

$$p_{it} - p_t = b + \ln\left(\frac{\eta}{\eta - 1}\right) + (\gamma + \theta - 1)\,y_t. \tag{9}$$

This corresponds to Romer's equation (7.16) for p^* with $b = \ln(B)$. Obviously, the first two terms on the right side of (9) are not zero in general. However, Romer is correct in saying that setting them to zero does not change the fundamental result and it keeps the algebra simple.

Romer's equations (7.17) and (7.18) give us the "building blocks" we need in order to proceed with the analysis of macro models with sticky prices. All that remains is to specify the pattern of price stickiness. Romer considers the basic patterns of price stickiness shown in Table 1. Romer has adapted each of these models in simplified form using common notation. This means that Romer's versions of these models do not correspond exactly to the versions in the original sources. However, the basic conclusions of the models are representative of those of the more widely varying models in the literature.

The predetermined-price model is a cousin of the wage-contract model developed in a seminal paper by Stanley Fischer (1977). In this model, prices are set for two periods at a time, with half of the firms in the economy setting their prices in even periods and the other half in odd periods. The price a firm sets for the first of the two periods is not necessarily the same as the price set for the second. Romer calls this model the "predetermined-price" model. He shows that monetary policy can have a positive countercyclical role under these assumptions (as in Fischer's original wage-contract model). Monetary shocks have real effects that last two periods.

Model number two is based on another wage-contract model originally due to John Taylor (1979). Romer calls this the "fixed-price" model. This model also has overlapping price setting for two periods at a time. It differs from the predetermined-price model in that firms are constrained to set the *same* price for the two periods rather than a different price for the first and second periods of the "contract." This model has similar implications for monetary policy, but leads to quite a different dynamic response to a monetary shock. In the predetermined-price model, monetary effects lasted only as long as the longest price contract (two periods). In the fixed-price case, the real effects of the monetary shock are longer lasting, damping out to zero only asymptotically.

The Calvo model allows the duration of any particular price to vary rather than being fixed at two periods. In each period, the firm changes its price with probability α and keeps it constant with probability $1 - \alpha$. This model has properties that are similar to the Taylor model, but it allows derivation of the new Keynesian Phillips curve.

Table 1. Classification of price-setting regimes under imperfect competition.

Model	Section	Description
Predetermined prices (Fischer)	7.2	Prices set for two periods at a time. May set a different price for first and second periods.
Fixed prices (Taylor)	7.3	Prices set for two periods. Same price must be set for both periods of contract.
Calvo	7.4	Constant (exogenous) probability of firm re-setting price in any period.
Caplin-Spulber	7.5	Models how firms decide when to change prices in simple, constant inflation setting.
Danziger-Golosov-Lucas	7.5	Extends the Caplin-Spulber model to allow for differences across firms and idiosyncratic demand shocks.
Christiano-Eichenbaum-Evans	7.7	Adds indexation of price changes into the Calvo model so that inflation between price reviews the firm raises prices at the previous period's inflation rate.
Mankiw-Reis	7.7	Like predetermined prices except firms set prices when they receive new macroeconomic information, which happens randomly with probability α per period.

The next group of models allow the frequency of price change to be determined by the agent's need to change prices, rather than according to a strict and exogenous schedule. In the Caplin-Spulber model, agents adjust prices when the gap between their existing price and their optimal price becomes large enough. While Romer does not present all of the underlying logic to justify this behavior, he does use this model to show that money can be neutral under sticky prices. The Danziger-Golosov-Lucas model builds on the idea of state-dependent pricing to consider a setting in which firms experience monetary shocks and shocks to their individual demand curves, as in the Lucas model of the previous chapter.

One shortcoming of the Calvo and other time-dependent models is that they cannot explain "inflation inertia," the empirically observed tendency of inflation to persist when monetary growth slows or stops. Menu costs make it very easy to stop changing prices, so dropping the rate of inflation to zero should be costless.

The final two models attempt to explain "inflation stickiness." Christiano, Eichenbaum, and Evans allow firms in the Calvo model to "index" their prices to last period's inflation rate when in periods where they do not make an explicit price change. In the Mankiw and Reis (2002) model, firms set a pricing policy for the indeterminate future based on their current information. As in the predetermined price model, they may set a different level of price for each period, so for example they may set a policy of increasing price by 2% each year indefinitely. A randomly selected share α of firms receives new macroeconomic information each period and, when they receive new information, reformulates strategies. In this model, the adjustment costs do not result from changing prices but rather from acquiring detailed macroeconomic information and reformulating a dynamic pricing policy.

Macroeconomic equilibrium with predetermined prices

As described in Table 1, agents in the predetermined-price model set prices in advance for each of the next two periods. The reasons for this price stickiness are not addressed until later, but one can most easily think of this price stickiness as fixed-term contracts that are established every two periods.

Romer adopts the (potentially confusing) notation that p_t^1 is the price set for period t by the half of the people who set prices at the end of period $t - 1$ and p_t^2 is the period t price set by the other half of the economy who established their prices at the end of period $t - 2$. The superscripts here *are not exponents*, so do not think of the p^2 term as a square. Once you get the notation down, the algebra on pages 318 and 319 should be pretty easy to understand.

The ***law of iterated projections*** on page 318 warrants some discussion. What this law says is that your current (2019) expectation of the price that will prevail in the year 2021 cannot be different than your current (2019) expectation of the price that you will expect in 2020 to prevail in 2021. If you have rational expectations, then your expectation of the 2021 price will only change from 2019 to 2020 due to *new* information that becomes available in 2020. Since you do not, by definition, have that information now, you cannot anticipate how your expectation will change and $E_{2019}[E_{2020}[P_{2021}]] = E_{2019}[P_{2021}]$.[1]

The solution of the model is summarized by equations (7.27) and (7.28). The former shows that the price level in period t depends on the expectations of the period t money supply during the two periods in which the prices for the current period were set ($t - 2$ and $t - 1$). The latter shows that the deviation of output from its steady-state value (by normalization, this value is one, or zero in log terms) is due to two "money-surprise" terms. The first of these measures the change in the prediction about the current money supply based on information that arrived last period; the latter is the deviation of the current money supply from what was expected last period.

Romer describes two key implications of these results. The first is that, in contrast to Lucas's model, countercyclical monetary policy has a positive role here. Shocks that are one period old still affect real output (through $E_{t-1}m_t - E_{t-2}m_t$). The central bank can observe these shocks and respond by changing the money supply in period t. This monetary-policy reaction is a change in $m_t - E_{t-1}m_t$ that brings y_t back to its steady-state level (zero).

[1] For a detailed mathematical discussion of the law of iterated projections, see Sargent (1987), Chapter X, Section 3.

The second key result is that a modified version of monetary neutrality continues to hold in this model. Changes in the money supply that people know about more than two periods in advance have a proportional effect on prices and no effect on output. Consider the effect of an unexpected, one-time, permanent increase in the money supply happening at date t. This will affect output in t and $t + 1$, but from periods $t + 2$ onward, prices will be proportionately higher and output will be unaffected by the shock. Thus, the non-neutrality of money in this model has a finite lifetime equal to the length of the "contract"—the longest amount of time in advance that prices are set.

Macroeconomic equilibrium with "fixed" prices

Romer's third model of this section is based on Taylor's overlapping wage-contract model. As in the case of the Fischer model, Romer adapts the model to look at overlapping price setting rather than wage-setting, and to do so must take explicit account of imperfect competition in the product market.

It is important to keep in mind the difference between "fixed" and "predetermined" prices as Romer uses the terms. In each case, firms set prices for two-period intervals, with half of the firms setting prices in even-numbered periods and half in odd-numbered periods. However, in the fixed-price model, the firms must decide on a single level of price to prevail in both of the upcoming periods. With predetermined prices they are able to specify a different price for the first and second periods. This seemingly minor alteration of the price-setting structure has a substantial effect on the dynamic behavior of the model.

Equation (7.32) shows that firms setting prices today (for the next two periods) set a price that is based on the average of the price set last period and their expectations of the price to be set next period, along with the current money supply. To understand the rationale for this, think about the markets in which the currently set price x_t will prevail. During the first period, firms setting x_t will be competing with firms who set prices last period at x_{t-1}. Since the other half of the market has a preset price of x_{t-1}, firms will not want to deviate too much from this price lest they lose too much of the market (if $x_t > x_{t-1}$) or fail to take advantage of profit opportunities afforded by their competitors' overpricing (if $x_t < x_{t-1}$). Similarly, during the second period that the price currently being set will be in effect, it will be competing against the price to be set next period, about which our current expectation is $E_t [x_{t+1}]$. For similar reasons, they would like to keep x_t fairly close to $E_t [x_{t+1}]$.

Thus, over the two periods, the average price against which we expect the currently set price to compete is the average of x_{t-1} and $E_t [x_{t+1}]$, which is the first part of (7.32). The second part shows the effect of the optimal long-run price, which under our simple normalization is m_t. In long-run equilibrium with no monetary shocks, $x_t = x_{t-1} = E_t [x_{t+1}] = m_t$, which shows that money is neutral in the long run in this model.

The method of undetermined coefficients, which Romer uses to solve the fixed-price model, may be familiar to you from the real-business-cycle chapter. We posit a hypothetical solution such as (7.33), then use the properties of the model to demonstrate the correspondence between the parameters of the solution (μ, λ, and ν) and the parameters of the original model (in this case, just ϕ). The mathematical derivation carried through to equation (7.44) executes this procedure.

The only difficult aspect of this derivation is the fact that the equation for λ in terms of ϕ is quadratic, which means that there are two different values of λ that solve the model. Romer notes that one value is greater than one and the other is less than one in absolute value, and that only the value that is less than one leads to a stable equilibrium. The use of stability to choose which of two possible roots to choose is an extension of ***Samuelson's correspondence principle***, which argues that because equilibrium in the actual economy is stable rather than explosive, we are justified in ignoring possible parameter values that lead to explosive behavior.

The fixed-price model implies that monetary shocks will have long-lived effects on output. Instead of the truncated impact of the predetermined-price model, where output is affected for a finite number of periods, the effects of a monetary shock in this model will die away gradually, converging on neutrality only asymptotically. The relationship between the mean lag in the effect of monetary policy and the length of the contract is called the "contract multiplier." Since inflation and output fluctuations are more persistent than our models often predict, the contract multiplier in a Taylor-type model has been proposed as a potential explanation.[2]

While lag operators are extremely useful for many tasks in time-series analysis, they are not crucial here. It does not seem worth your time to read and learn the material about them on pages 324–326, so you may skip this section.

Evaluation of the Fischer and Taylor models

The Fischer and Taylor models were criticized on several grounds. One was the ad-hoc nature of the length of the price (wage) contract. Why should agents set prices or wages for two periods rather than one, when utility would be higher if they set them in each period? Another case of $50 bills lying on the sidewalk?

A considerable literature on the optimal length of contracts ensued.[3] Models of contract length commonly assume that there are fixed negotiation costs that must be incurred each time a contract is negotiated (or the price is set, in our context). Because of these fixed costs, agents will fix prices/wages for a finite contract period. The length of the contract period is determined by striking a balance between the disequilibrium

[2] For example, see Chari, Kehoe, and McGrattan (2000).

[3] An excellent and quite readable paper is Gray (1978).

costs of being away from the optimal price or wage (which rises with longer contracts) and the per-period negotiation cost (which rises with shorter contracts).

Another common criticism was that the models exclude the possibility of ***indexed contracts,*** which would allow prices or wages to respond fully to monetary changes and thus eliminate the source of non-neutrality in the model. An indexed contract could make the price or wage a function of the money supply, similar to cost-of-living adjustments based on the actual CPI. With appropriately indexed contracts, markets would always clear. This means that indexed contracts would lead to welfare gains, so agents would have a strong incentive to use them rather than the predetermined-price or fixed-price arrangements that are assumed by Fischer and Taylor. Although this was an important criticism, other economists argued that in the labor-market context, indexed contracts might have to be unreasonably complex in order to fully offset changes in the money supply. It might be difficult to design indexing rules that would allow for both monetary shocks and also changes in productivity or in the relative demand for the product.

Ultimately, the most damning evidence against the Fischer and Taylor wage-contract models (which does not apply to the price-based versions in Romer) was that both models rely on strongly countercyclical real wages to produce their basic results. In both cases, a contract wage that turns out to be too high, given the price level that ends up prevailing in the period, leads to a high real wage that causes firms to reduce employment and output and leads to a recession. Since real wages seem to be mildly procyclical rather than strongly countercyclical, these models have lost much of their initial popularity.

Calvo model

The Fischer and Taylor models are both adaptations from early wage-contract models. Most of the price-stickiness literature has used variations on Guillermo Calvo's model. This model is very tractable and, as a first approximation, seems reasonably realistic.

Whereas the Fischer and Taylor models feature a fixed schedule of dates at which firms adjust prices, the Calvo model instead assumes a fixed probability α that each firm's price will be adjusted in any period. This means that the probability that the currently-set price is still in effect *j* periods later is $q_j = (1-\alpha)^j$, since for the price to stay in force for *j* period the firm must be in the $1 - \alpha$ share of the population that does not adjust its price for *j* periods in a row.

If x_t is the price set in period *t* by the share α of firms that set their price anew, then the average price is $p_t = \alpha x_t + (1-\alpha) p_{t-1}$, as in Romer's equation (7.53). The remaining analysis on pages 327 and 328 uses a time-series trick to derive the new Keynesian Phillips curve (7.60). This equation is similar to the Friedman-Phelps Phillips curve and the Lucas supply curve in that it shows a short-run positive relationship between the difference between actual and expected inflation and the difference between output and natural output.

Caplin-Spulber model

The Caplin-Spulber model differs from the Fischer and Taylor models in that there is no fixed frequency at which prices are set. Instead, agents change prices when their desired price gets far enough from their existing price. This is an example of ***state-dependent*** rather than ***time-dependent*** pricing. Under certain quite restrictive assumptions, one can show that the *Ss* pricing rule of the Caplin-Spulber model is optimal. (See the citations in Romer.) In particular, the model is usually applied to a situation of steady inflation of the general (average) price level.[4]

Under an *Ss* rule, firms keep prices constant until the difference between actual and desired price reaches some trigger threshold *s*, then they reset the price. However, they do not reset the price to the *current* desired price because that price will be out of date an instant later. Instead, they overshoot their current desired price and set the price *S* units above it. As inflation continues, their fixed price gradually falls in relation to the desired price, which rises with the general price level in the economy. Eventually, their relative price falls to *s*, at which point another price increase is triggered.

You can think of the dynamic behavior of a single firm's price in the *Ss* model as being a step function of time that moves above and below an upward-sloping line, which represents the steady increase in the desired price. When the desired-price line gets far enough above the previous step, the agent raises price and leaps above the line, staying there until the desired-price line catches up and again moves far enough above.

Interestingly, money is neutral in the Caplin-Spulber model because an unusually large increase in the money supply pushes more firms above the price-adjustment threshold. Although each firm's price is fixed for a finite period of time, as in Romer's variants of the Fischer and Taylor models, the *aggregate* price level is perfectly flexible. This comes about because the length of the price contract is endogenous. A monetary shock makes it more or less desirable to change prices, thus causing the average price level to respond directly to the shock.

Danziger-Golosov-Lucas model

We won't spend much time talking about this model. The central idea here is that if firms face both aggregate shocks and relative-price shocks, then a monetary shock will induce some those firms whose prices are furthest from the optimal price to adjust.

---

[4] Subsequent papers have applied state-dependent models using simulation techniques to get around the complexity that precludes analytic solution in all but simple cases. See Dotsey, King, and Wolman (1999), Dotsey and King (2005), and Caballero and Engel (2007).

This means that the initial price adjustment will be faster than if the firms were randomly selected as in the Calvo model.

Mankiw and Reis: sticky information vs. sticky prices

Mankiw and Reis (2002) present a model that attempts to "correct" several implications of the new Keynesian Phillips curve that are both counterintuitive and counterfactual. Romer discusses the problem of *inflation inertia* at the end of Section 7.6. However, this is only one of three problems that Mankiw and Reis discuss and try to solve in their paper. They also consider the counterfactual implication of some sticky-price models that a credible disinflation (reduction in inflation) can be expansionary and the existence of a discrepancy between the speed of adjustment of inflation to changes in monetary policy between the implications of sticky-price models and the empirical evidence.

Laurence Ball (1994) looks carefully at the implications of a simple sticky-price model for sudden disinflations. Working with the Calvo model, he finds that while a credible *deflation*, in which the central bank lowers the money supply over time and prices must fall, leads to a recession, a credible *disinflation*, in which the central bank reduces the *rate of growth of the money supply* from a positive value to zero, leads not only to no recession but actually causes output to be above full employment.

Ball explains the intuition of the difference between the effects of deflation and disinflation as follows:

> To understand the difference between deflation and disinflation, recall why the former reduces output: prices set before deflation is announced are too high once money begins to fall. In the case of disinflation, the overhang of preset prices is a less serious problem. Prices set before an announcement of disinflation are set higher than the current money stock in anticipation of further increases in money. … However, the overhang of high prices is easily overcome if money growth, while falling, remains positive for some time. The level of money quickly catches up to the highest preset price and can then be stabilized costlessly. (Ball 1994, 286–287)

Essentially, *price* stickiness does not imply *inflation* stickiness, which means that a reduction in money *growth* has different effects than a change in the *level* of the money supply.

Fuhrer and Moore (1995) argue that the sticky price model implies that there should be no persistence in inflation, whereas the data indicate that the autocorrelations of the inflation rate are quite high. Again, it is prices that are sticky in the theoretical model, whereas the data seem to indicate stickiness of inflation.

Finally, Mankiw (2001) compares the theoretical time-path of the response of inflation to changes in money growth to econometrically estimated responses. He finds that sticky-price theory suggests that inflation should adjust quickly to changes in money growth, but the evidence suggests that adjustment is slow.

Mankiw and Reis replace the assumption of sticky prices with one of "sticky information." Because it is information that is sticky rather than prices, this model introduces stickiness or persistence into *inflation*. They do this by restricting the frequency with which firms can adopt new pricing strategies.

In each period, a fraction α of firms receives current macroeconomic information and updates its pricing strategy and the remaining fraction $(1 - \alpha)$ keeps its old pricing policy intact. Thus, α share of firms are setting their price at the level that is currently optimal and $(1 - \alpha)$ are using older information. Similarly, in the previous period, α share had the opportunity to reset pricing policy and $(1 - \alpha)$ did not, so the share $(1 - \alpha)^2$ are using information more than one period old. Following the same logic, $(1 - \alpha)^i$ is the share of firms whose information and pricing policies are at least i periods old and $1 - (1 - \alpha)^i = \lambda_i$ is the fraction that has updated less than i periods ago.

The algebra on page 347 is involved, but the logic is straightforward. By obtaining the expression for a in equation (7.81), Romer solves for the equilibrium output expression using (7.78). This expression shows that output is a function of monetary policy shocks extending back into the indefinite past.

The Mankiw-Reis model exhibits considerable inflation persistence due to the interaction of nominal and real rigidity. The persistence is in inflation rates rather than price levels because of the predetermined-price nature of the individual pricing policies that firms adopt. Moreover, simulations in the Mankiw and Reis (2002) show that the model can lead to a delayed reaction to monetary policy in which the maximum effect of policy changes on output do not occur immediately.

In a later follow-up to their 2002 paper, Mankiw and Reis (2006) examine the ability of macro models to reproduce three stylized facts about modern business cycles: (1) inflation rises when output is above its trend level, (2) real wages are smoother than labor productivity, and (3) most real variables have gradual, hump-shaped responses to shocks.

They find that in order to explain these three phenomena, more than simply sticky information of price setters is required. They introduce additional information stickiness (or "inattentiveness") in household consumption planning in workers' labor supply decisions. Only if all three forms of inattentiveness are present in the model can the three stylized facts all be explained.

D. Suggestions for Further Reading

Original papers on wage-contracting models

Fischer, Stanley, "Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule," *Journal of Political Economy* 85(1), February 1977, 191-205.

Taylor, John B., "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy* 88(1), February 1980, 1-23.

Barro, Robert J., "Long-Term Contracting, Sticky Prices, and Monetary Policy," *Journal of Monetary Economics* 3, July 1977, 305-316.

Papers on optimal indexing

Gray, Jo Anna, "Wage Indexation: A Macroeconomic Approach," *Journal of Monetary Economics* 2(2), April 1976, 221-35.

Gray, Jo Anna, "On Indexation and Contract Length," *Journal of Political Economy* 86(1), February 1978, 1-18.

Karni, Edi, "On Optimal Wage Indexation," *Journal of Political Economy* 91(2), April 1983, 282-92.

Blanchard, Olivier, "Wage Indexing Rules and the Behavior of the Economy," *Journal of Political Economy* 87(4), August 1979, 798-815.

Vroman, Susan B., "Inflation Uncertainty and Contract Duration," *Review of Economics and Statistics* 71(4), November 1989, 677-81.

Seminal papers on sticky prices

Mankiw, N. Gregory, "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly," *Quarterly Journal of Economics* 100(2), May 1986, 455–470.

Akerlof, George A., and Janet L. Yellen, "A Near-Rational Model of the Business Cycle, with Wage and Price Inertia," *Quarterly Journal of Economics* 100(Supplement), 1985, 823–838.

Other theoretical approaches to price adjustment

Barro, Robert J., "A Theory of Monopolistic Price Adjustment," *Review of Economic Studies* 34(1), January 1972, 17–26.

Rotemberg, Julio J., "Monopolistic Price Adjustment and Aggregate Output," *Review of Economic Studies* 44(4), October 1982, 517–531.

E. Works Cited in Text

Ball, Laurence. 1994. Credible Disinflation with Staggered Price Setting. *American Economic Review* 84 (1):282-289.

Caballero, Ricardo J, and Eduardo M. R. A. Engel. 2007. Price Stickiness in Ss Models: New Interpretations of Old Results. *Journal of Monetary Economics* 54S:S100-S121.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. 2000. Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem? *Econometrica* 68 (5):1151-79.

Dotsey, Michael, and Robert G. King. 2005. Implications of State-Dependent Pricing for Dynamic Macroeconomic Models. *Journal of Monetary Economics* 52 (1):213-242.

Dotsey, Michael, Robert G. King, and Alexander L. Wolman. 1999. State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output. *Quarterly Journal of Economics* 114 (2):655-690.

Fischer, Stanley. 1977. Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule. *Journal of Political Economy* 85 (1):191-205.

Fuhrer, Jeffrey, and George Moore. 1995. Inflation Persistence. *Quarterly Journal of Economics* 110 (1):127-160.

Gray, JoAnna. 1978. On Indexation and Contract Length. *Journal of Political Economy* 86 (1):1-18.

Mankiw, N. Gregory. 2001. The Inexorable and Mysterious Tradeoff Between Inflation and Unemployment. *Economic Journal* 111 (471):C45-C61.

Mankiw, N. Gregory, and Ricardo Reis. 2002. Sticky Information and Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics* 117 (4):1295-1328.

———. 2006. Pervasive Stickiness. *American Economic Review. March* 96 (2):164-169.

Sargent, Thomas J. 1987. *Macroeconomic Theory*. Second ed. Boston: Academic Press.

Taylor, John B. 1979. Staggered Wage Setting in a Macro Model. *American Economic Review* 69 (2):108-13.