

NBER WORKING PAPERS SERIES

ASYMMETRIC PRICE ADJUSTMENT AND ECONOMIC FLUCTUATIONS

Laurence Ball

N. Gregory Mankiw

Working Paper No. 4089

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 1992

We are grateful to Boris Simkovich for research assistance, to Ben Bernanke, Antonio Fatas, Andrew Metrick, David Romer, and Julio Rotemberg for suggestions, and to the NBER, the National Science Foundation, and the Sloan Foundation for financial support. This paper is part of NBER's research programs in Economic Fluctuation and Monetary Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

ASYMMETRIC PRICE ADJUSTMENT AND ECONOMIC FLUCTUATIONS

ABSTRACT

This paper considers a possible explanation for asymmetric adjustment of nominal prices. We present a menu-cost model in which positive trend inflation causes firms' relative prices to decline automatically between price adjustments. In this environment, shocks that raise firms' desired prices trigger larger price responses than shocks that lower desired prices. We use this model of asymmetric adjustment to address three issues in macroeconomics: the effects of aggregate demand, the effects of sectoral shocks, and the optimal rate of inflation.

Laurence Ball
Department of Economics
Fisher Hall
Princeton University
Princeton, NJ 08544
and NBER

N. Gregory Mankiw
Department of Economics
Littauer 223
Harvard University
Cambridge, MA 02138
and NBER

I. INTRODUCTION

Rigidities in nominal wages and prices are central to many theories of economic fluctuations. Economists frequently argue that these rigidities are asymmetric: prices are more flexible when going up than when going down. (See, for example, Tobin, 1972.) Some textbooks capture this idea with a convex aggregate supply curve. Most often, however, these asymmetries are simply assumed, or they are justified through irrational taboos on decreases in nominal wages or prices. By contrast, in modern theories based on costs of price adjustment, nominal rigidities are usually symmetric: prices respond similarly to positive and negative shocks. This paper explores one possible microeconomic foundation for asymmetric price adjustment.

We consider a model in which firms make regularly scheduled price changes and, by paying a menu cost, can also make special adjustments in response to shocks. In this model, asymmetries arise naturally with the addition of one feature: positive trend inflation. With trend inflation, positive shocks to firms' desired prices trigger greater adjustment than do negative shocks of the same size. The intuition is that inflation causes firms' relative prices to decline automatically between adjustments. When a firm wants a lower relative price, it need not pay the menu cost, because inflation does much of the work. By contrast, a positive shock means that the firm's desired relative price rises while its actual relative price is falling, creating a large gap between desired and actual prices. As a result, positive shocks are more likely to induce price adjustment than are negative shocks, and the

positive adjustments that occur are larger than the negative adjustments.

It is perhaps not surprising that price adjustment is asymmetric in the presence of trend inflation: Tsiddon (1991b) presents a similar result in a related model. The model developed here--which includes a combination of time-contingent and state-contingent pricing--has the advantage of substantial tractability. The simplicity of our framework allows us to apply the model to a number of issues in macroeconomics. In particular, we explore the implications of asymmetric price adjustment for the effects of aggregate demand, the effects of sectoral shocks, and the optimal rate of inflation.

Aggregate Demand: Our model implies that shifts in aggregate demand have asymmetric effects on output. Since prices are sticky downward, a fall in aggregate demand reduces output substantially. A rise in demand has a smaller absolute effect on output, because prices adjust more quickly. We thus provide a theoretical rationale for the empirical finding that monetary shocks have asymmetric effects (Cover, 1988; DeLong and Summers, 1988). The model also implies that, with a symmetric distribution of demand shocks, the distribution of output is skewed to the left. Thus the model explains the skewed distribution that Sichel (1989) finds for the postwar United States.

Despite these asymmetries, our model exhibits a strong form of the natural-rate hypothesis: the average level of (log) output is invariant to the distribution of aggregate demand. We thus provide a counterexample to the claim that asymmetric rigidity provides a

rationale for demand stabilization. Lucas (1987) argues that economic fluctuations have small welfare consequences, because average output is unaffected and the losses from variability are small. Advocates of active demand management (such as DeLong and Summers) respond that, with asymmetric rigidities, the output gains in booms are smaller than the losses in recessions. They conclude, contrary to Lucas, that eliminating fluctuations would raise average output. Perhaps surprisingly, this argument is incorrect in our model. When firms set their initial prices, they know that their response to shocks will be asymmetric: an initial price is more likely to remain in effect when demand falls than when it rises. Firms therefore reduce initial prices below the level they would set if there were no shocks. Lower initial prices raise output by exactly enough to offset the effects of asymmetric ex post adjustment.

Sectoral Shocks: Our model also has implications for the effects of sectoral shocks. In particular, it implies that shifts in relative prices are inflationary in the short run. When a shock raises some firms' desired prices and lowers others', the desired increases trigger greater price adjustment than the desired decreases. Thus, in periods with a large variance of relative-price changes, inflation rises above trend and output falls (assuming aggregate demand is held constant). In other words, large shifts in relative prices are adverse supply shocks.

Our model is therefore consistent with the evidence, documented by Fischer (1981) among others, that inflation is correlated with relative-price variability. In most explanations

for this correlation, such as that proposed by Sheshinski and Weiss (1977), inflation causes relative-price variability. As Fischer points out, however, causality can also run in the opposite direction if price rigidity is asymmetric. We argue that causality from relative-price shocks to inflation is an essential part of the story.¹

Optimal Inflation: Finally, our model has implications for the optimal rate of inflation. Classical economics suggests that optimal inflation is zero or even negative (Friedman, 1969). Tobin (1972) uses the assumption of asymmetric nominal rigidity to argue that a positive rate is optimal. With downward rigidity of prices, inflation is needed for relative prices to fall. Thus, inflation is "grease" that facilitates sectoral adjustment.

Our model of asymmetric rigidity does not support Tobin's argument. Whereas Tobin treats the asymmetry of adjustment as exogenous, in our model it is an endogenous response to inflation. The asymmetry would disappear if inflation were eliminated. Thus, even if one observes downward rigidity under inflation, one should not conclude that downward relative adjustments would be difficult at zero inflation. In our model, zero inflation is optimal. Inflation is costly because (under staggered adjustment) it creates

¹This point complements the one in Ball and Mankiw (1992). In that paper, we show that sectoral shocks influence inflation if the distribution of shocks is asymmetric. Here we assume that shocks are symmetric and focus on asymmetries in adjustment. The empirical results in our other paper suggest that both kinds of asymmetries influence inflation.

inefficient relative-price variability without any offsetting benefit.

The rest of this paper contains five sections. Section II demonstrates that rigidity is asymmetric in a simple partial-equilibrium model. Sections III, IV and V develop our three implications for macroeconomics. Section VI concludes and briefly discusses the empirical evidence.

II. ASYMMETRIC RIGIDITY

In this section we present a partial-equilibrium model in which trend inflation and costs of price adjustment produce asymmetric responses to shocks.

A. Assumptions

We consider a single firm whose desired relative price, in logs, is θ . The firm's desired nominal price is $p+\theta$, where p is the aggregate price level (again in logs). There is steady inflation at rate π , which the firm takes as exogenous. The price level in period t is $p_t = \pi^t$. We consider various assumptions about the dynamics of θ .

Price adjustment occurs as follows. Every even period, after observing the current θ , the firm sets a single price for that period and the following odd period. The firm's commitment to maintain the price for both periods, however, is not absolute. The firm can make an extra adjustment in an odd period by paying a menu cost C . As shown below, the firm does so when there is a large shock to θ .

This specification combines elements of time-contingent price adjustment (such as in Ball, Mankiw, and Romer, 1988) and state-contingent price adjustment (such as in Caplin and Spulber, 1987). As with time-contingent pricing, the firm adjusts on a regular schedule (every two periods); but firms also have the option, as with state-contingent pricing, of adjusting whenever circumstances change substantially. It is realistic to combine these two polar cases, because price setting in actual economies has elements of both. Indeed, our specification closely matches price adjustment in some settings: labor contracts and catalogs are written for fixed periods but can be rewritten in special circumstances. More important, our combination of time- and state-contingent pricing is substantially more tractable than the purely state-contingent case analyzed by many authors.²

The firm's one-period loss function is $(q-q^*)^2 + DC$, where q is the firm's actual price, $q^* = p + \theta$ is its desired price, and D is a

²It is essential for our results that price setting is partly state-contingent: with purely time-contingent adjustment, the effects of shocks are symmetric even with trend inflation (Ball, Mankiw, and Romer, 1988). It is unclear whether some time-contingency is also essential, or whether all our results carry over to a purely state-contingent model. As mentioned in the introduction, Tsiddon (1991b) presents a model with state-contingent pricing in which trend inflation yields asymmetric adjustment. Given the rich dynamics of his model, however, Tsiddon is able to analyze only large, one-time changes in the money stock, and to contrast positive and negative changes of equal size. He is not able to study, for example, a "negative" monetary shock in which the money stock grows by a positive amount but by less than expected. For other, related analyses of state-contingent adjustment, see Bertola and Caballero (1990), Caballero and Engel (1991), Caplin and Leahy (1991), and Tsiddon (1991a).

dummy that equals one if the firm pays the menu cost.³ The firm chooses even-period prices and whether to adjust in odd periods to minimize the average of its loss, with no discounting. If the firm does adjust in an odd period, it chooses the current q^* , since the price lasts only one period. Finally, for the purposes of this section, we simplify matters by assuming that the menu cost C exceeds $\pi^2/2$ (otherwise, the firm would adjust in odd periods even if there were no shocks).

B. A One-Time Shock

Here we demonstrate the asymmetry in our model in the simple case of a one-time shock. We assume that, in period zero, the firm acts as if the probability of future shocks is zero. Nonetheless, in period one, a shock occurs. Examining how the firm responds to the shock in this simple environment provides intuition for the more realistic cases considered below.

We begin by considering price setting in period zero. We assume that the desired relative price θ is zero in that period and is expected to remain constant. Since the firm does not anticipate a shock, it expects its nominal price for period zero to last through period one. The firm's optimal nominal price, $p+\theta$, equals zero in period zero and is expected to rise to π in period one. Given its quadratic loss, the firm chooses a price of $\pi/2$, the average of the two desired prices.

In period one, there is a surprise: the firm's desired relative price changes to $\theta \neq 0$. To determine the firm's response,

³This assumption is equivalent to taking a second-order approximation to a general profit function. See, for example, Ball and Romer (1989).

we first consider its desired price adjustment (ignoring, for the moment, the menu cost). The firm's ex post optimal price is $\pi + \theta$, and its actual price entering the period is $\pi/2$. The desired adjustment is the difference between the two:

$$\text{Desired adjustment} = \frac{\pi}{2} + \theta. \quad (1)$$

Note that the desired adjustment is asymmetric: since $\pi/2$ is positive, the absolute value of the desired adjustment is larger for a positive θ than for a negative θ . The intuition is that, in the absence of shocks, inflation in period one pushes the firm's desired price q^* past its actual price q . A positive θ pushes q^* even farther above q , creating a large desired adjustment. By contrast, a negative θ reduces q^* ; in this case, the shock offsets the need to catch up with inflation, and the firm desires a relatively small price change.

This asymmetry in desired price adjustments leads to an asymmetry in actual adjustments. With a menu cost, the firm either makes its desired adjustment or keeps its price fixed. If the firm does not adjust, the value of its loss function is $(q - q^*)^2 = (\pi/2 + \theta)^2$. If the firm does adjust, its loss is simply the menu cost C ; thus the firm fails to adjust if $C > (\pi/2 + \theta)^2$. This condition is equivalent to

$$\theta \in [-\sqrt{C} - \frac{\pi}{2}, \sqrt{C} - \frac{\pi}{2}], \quad (2)$$

where $\sqrt{C} - \frac{\pi}{2}$ is positive by the assumption that $C > \pi^2/2$. Not surprisingly, the firm fails to adjust for a range of shocks. Most important, this range is asymmetric: the lower bound is larger in

absolute value than the upper bound. Relatively small positive shocks trigger adjustment, whereas prices are sticky for a larger range of negative shocks. Finally, note that even if θ lies outside the range in (2), so the firm adjusts, the price change is larger for a positive θ than for the corresponding negative θ . Both the asymmetry in the non-adjustment range and the asymmetry in the size of adjustments are important for our macroeconomic results below.⁴

C. A Distribution of Shocks

To prepare for the applications below, we now assume that the firm faces a known distribution of shocks. Specifically, in period zero the desired relative price θ is zero; in period one it is drawn from a zero-mean, symmetric, and single-peaked distribution with cumulative distribution function $F(\theta)$. The knowledge that a shock may trigger adjustment in period one influences the price the firm sets in period zero. We solve jointly for the initial price, denoted x , and for the firm's price-adjustment rule.

We begin with behavior in period one, when the firm takes x as given. The price-adjustment rule is derived as before. Since the firm's desired price is $\pi + \theta$ and its initial price is x , its desired

⁴The asymmetry in price adjustment does not depend on the assumption that prices are set for only two periods. In particular, our results generalize to a more realistic continuous-time setting. In this case, in the absence of shocks a firm's relative price varies smoothly from a point above the desired price to a point equally far below. A shock can arrive when the firm's current price is either too high or too low. The decision whether to make a special adjustment, however, depends not on the current price, but on the average relative price until the next regular adjustment. This average price is below the desired level at all points in time, except at adjustment dates. Therefore, the incentive to make a special adjustment is greater for positive shocks than for negative shocks.

adjustment is $(\pi + \theta - x)$. The firm fails to make this adjustment if $C > (\pi + \theta - x)^2$; that is, if

$$\theta \in [\underline{\theta}, \bar{\theta}],$$

$$\text{where } \underline{\theta} = x - \pi - \sqrt{C}; \quad (3)$$

$$\bar{\theta} = x - \pi + \sqrt{C}.$$

This equation holds for any given x . Notice that positive shocks induce adjustment more quickly than negative shocks as long as $x < \pi$.

Consider now behavior in period zero, when the firm chooses the initial price x . The firm minimizes:

$$\text{Loss} = x^2 + \int_{\underline{\theta}}^{\bar{\theta}} (\pi + \theta - x)^2 dF(\theta) + [1 - (F(\bar{\theta}) - F(\underline{\theta}))] C. \quad (4)$$

The first term in this expression is the firm's loss in period zero, when its optimal nominal price is zero and its actual price is x . The other terms are the expected loss in period one: the loss is $(\pi + \theta - x)^2$ if $\theta \in [\underline{\theta}, \bar{\theta}]$, and C if $\theta \notin [\underline{\theta}, \bar{\theta}]$. The first-order condition for minimizing (4) with respect to x leads to

$$x = \frac{1}{1 + F(\bar{\theta}) - F(\underline{\theta})} [\pi [F(\bar{\theta}) - F(\underline{\theta})] + \int_{\underline{\theta}}^{\bar{\theta}} \theta dF(\theta)]. \quad (5)$$

One can show that the initial price x is a weighted average of the firm's optimal price in period zero and its expected optimal price in period one, conditional on the initial price remaining in effect. The weights are the probabilities that x is in effect in the two periods, which are one for period zero and $F(\bar{\theta}) - F(\underline{\theta})$ for period one. Equations (3) and (5) define x and the bounds $\bar{\theta}$ and $\underline{\theta}$.

Using equations (3) and (5), one can show that

$$0 < x < \frac{\pi}{2}. \quad (6)$$

Not surprisingly, x is bounded below by zero, the desired price in period zero. More interesting is the result that the firm's initial price is less than $\pi/2$, its price when it does not expect a shock. There are two reasons for this. First, in choosing x the firm gives greater weight to period zero, when q^* is zero, than to period one, when the expected q^* is one, because x may not be in effect in period one. Second, the firm knows that, with asymmetric adjustment, x is more likely to stay in effect when it receives a negative shock.

The fact that the initial price x is less than $\pi/2$ means that ex post price adjustment is even more asymmetric here than in the case of a one-time shock. This result can be seen formally by comparing equations (3) and (2): the range of non-adjustment is shifted farther to the left. The intuition is that the lower initial price exacerbates the tendency of inflation to push the desired price above the initial price and, therefore, exacerbates the asymmetry.

III. FLUCTUATIONS IN AGGREGATE DEMAND

Having examined the pricing decisions of individual firms, we can now consider the aggregate economy. In this section we develop a general-equilibrium model and show that output responds asymmetrically to fluctuations in aggregate demand. We also show

that, despite this asymmetric response, average output is invariant to the distribution of demand.

A. Assumptions

The economy contains a continuum of firms. Each firm's desired nominal price is

$$q^* = m, \quad (7)$$

where m is the log of the money supply. That is, firms simply wish to keep their prices in line with nominal money.⁵ The money supply follows a random walk with drift π . Finally, output is given by a quantity equation:

$$y = m - p, \quad (8)$$

where y is the log of aggregate output and p the log of the price level, defined as the mean of firms' prices. Equation (8) is the economy's aggregate demand curve.

As in the previous section, each firm sets a price for two periods but can adjust in extra periods by paying a menu cost C . The cost varies across firms with distribution function $G(C)$. (This heterogeneity ensures that output is a smooth function of the monetary shock.) Price setting is staggered across firms: half have regular adjustments in even periods, and half in odd periods.

⁵Equation (7) is a special case of the usual expression for profit-maximizing prices in menu-cost models; in the general case, profit-maximizing prices depend on the aggregate price level as well as the money supply (see Ball, Mankiw, and Romer). The general case is more complicated, because price adjustment by different firms is interdependent. We believe, however, that our results carry over to the more general setting.

B. Prices and Output for a Given Shock

We determine the behavior of an individual firm and then aggregate to find output in terms of the monetary shock. Consider a firm setting an initial price at time zero, and let $m_0=0$. As in the partial-equilibrium model, the change in the firm's optimal price in period one has mean π and a known distribution. That is, the firm faces the same problem as before, but now π is interpreted as trend money growth, and the shock θ is interpreted as a monetary shock (the deviation of money growth from its mean). Thus the expressions for x , $\underline{\theta}$, and $\bar{\theta}$ in equations (3) and (5) carry over from the previous section, with $F(\cdot)$ interpreted as the distribution of the monetary shock. Since the menu cost C varies across firms, individual behavior varies; equations (3) and (5) define $x(C)$, $\underline{\theta}(C)$, and $\bar{\theta}(C)$. A larger C implies a larger x and a larger difference between $\underline{\theta}$ and $\bar{\theta}$.

We now derive the aggregate price level in period one for a given shock θ . To do so, we must consider both the period-zero price setters and the period-one price setters. The first group of firms sets initial prices $x(C)$ in period zero. Within this group, a firm's initial price remains in effect in period one if the shock lies within the range $[\underline{\theta}(C), \bar{\theta}(C)]$. For a given θ , this condition holds for C sufficiently large; specifically, C must exceed a bound $C^*(\theta)$, defined implicitly by $C^*=(\pi+\theta-x(C^*))^2$.⁶ Thus a firm's

⁶To ensure that a well-defined C^* exists, we must assume that $x(C)$ does not increase too rapidly in C . In particular, we assume that $dx(C)/d\sqrt{C} < 1$. This condition does not hold in general, but for the normal distribution assumed below it holds for a sufficiently large variance of θ .

price equals its initial price $x(C)$ if $C > C^*(\theta)$, and adjusts to the current optimum, $\pi + \theta$, if $C < C^*(\theta)$. Finally, the other group of firms sets new initial prices in period one. These firms face the same problem as period-zero price setters, except that the current money supply is $\pi + \theta$ (reflecting trend growth and the shock) rather than zero. Within this group, a firm with cost C sets a price of $x(C) + \pi + \theta$. Combining the prices of all the firms, we find that the overall price level in period one is

$$\begin{aligned}
 p = & \frac{1}{2} \{ G(C^*(\theta)) (\pi + \theta) + \int_{C^*(\theta)}^{\infty} x(C) dG(C) \\
 & + \int_{C=0}^{\infty} [x(C) + \pi + \theta] dG(C) \}.
 \end{aligned}
 \tag{9}$$

In equation (9), the first term reflects period-zero price setters with $C < C^*$; the second, period-zero price setters with $C > C^*$; and the third, period-one price setters. Finally, aggregate output in period one is the difference between the money supply, $\pi + \theta$, and the price level in equation (9).

To examine how output varies with the monetary shock, we calibrate the model and then compute y numerically. We assume that the shock θ is distributed $N(0, \sigma^2)$ and that \sqrt{C} is distributed uniformly over $[0, \Gamma]$. For any given σ^2 , Γ , and π , one can use the equations above to find y as a function of θ .

We attempt to choose parameter values that are empirically plausible. We assume that prices are regularly adjusted once a year, so that a period in the model corresponds to six months. We let $\Gamma = .25$, which means that the mean of \sqrt{C} across firms is .125. This parameter implies that the average firm tolerates a 12.5

percent deviation between its actual and optimal prices before making a special adjustment. Firms with the largest menu costs tolerate a 25 percent deviation. These results appear consistent with microeconomic evidence on the frequency of price adjustment. (See, for example, Blinder, 1991, and Cecchetti, 1986). We let $\sigma = .025$, which implies that the standard deviation of the growth in nominal spending is 3.5 percent per year; this is close to the standard deviation for the postwar United States. Finally, we let π equal 0, 2.5, 5.0, and 7.5 percent, which correspond to annual inflation rates of 0, 5, 10, and 15 percent.

For these parameter values, Table 1 gives output as a function of θ . The table shows that monetary shocks have asymmetric real effects if inflation is positive. For π equal to 5.0 percent, for example, a 5 percent positive shock to the money supply raises output by 1.52 percent above its level when $\theta = 0$. By contrast, a 5 percent negative shock reduces output by 2.12 percent. The asymmetry is stronger for larger shocks: a 10 percent positive shock raises output by 2.03 percent, whereas a 10 percent negative shock lowers output by 3.73 percent. Table 1 also shows that this asymmetric response does not exist under zero inflation and is more pronounced at higher rates of inflation.

These asymmetries follow from the asymmetries in price adjustment. One can show that the cutoff $C^*(\theta)$ is smaller for a positive θ than for the corresponding negative θ . Thus, for a shock of any size, a larger proportion of firms are pushed beyond their range of inaction when the shock is positive than when it is negative. In addition, as described in Section II, the firms that

do adjust change their prices by $(\pi-x)+\theta$, which is larger in absolute value for a positive θ than for a negative θ . Because both the number and the size of individual price adjustments are larger, a positive shock has a larger effect on the aggregate price level and a smaller effect on output.

C. The Distribution of Output

Skewness: Another way to see the asymmetry in the model is to determine the distribution of aggregate output, given that output is $y(\theta)$ and θ is distributed normally. We perform this transformation numerically. Figure 1 presents distributions of y for the same parameters as in Table 1. As discussed below, average output is zero in each case. Although the monetary shock is symmetric, the distribution of output is skewed to the left if inflation is positive; that is, the negative tail is larger than the positive tail. Since the effects of monetary shocks are asymmetric, a deep recession is more likely than an equally large boom.

Average Output: Although the distribution of output varies with parameter values, its mean is always zero. Indeed, this property of the model follows almost immediately from our assumption that firms minimize a quadratic loss function. As discussed above, the first-order condition for minimizing the loss implies that a firm sets its price equal to the average optimal price, where the average is taken over the times and states in which the price is in effect. Therefore, the deviation between a firm's optimal and actual prices, q^*-q , has an unconditional expectation of zero. This, in turn, implies that the expectation

of $m-p$ equals zero, because m equals q^* and p is the average of individual q 's. Finally, since $y=m-p$, expected output is zero.

Because average output is invariant to the distribution of demand, our model provides a counterexample to a conjecture offered by DeLong and Summers (1988) and Ball-Mankiw-Romer (1988). These papers argue that reducing fluctuations in aggregate demand would raise average output, because, with asymmetric price adjustment, the output losses in recessions exceed the gains in booms. This argument implicitly assumes that the level of output when there is no shock is independent of the distribution of demand. In our model, however, demand variability reduces initial prices x : price setters realize that x may last only through the initial period, when they want a relatively low price, and also that x is more likely to stay in effect if θ is negative. The lower initial prices raise $y(0)$, the level of output in the absence of a shock. (In particular, $y(0)$ is positive when $s^2 > 0$, whereas it is zero when $s^2 = 0$). The higher level of $y(0)$ offsets the lower average output from asymmetric responses to shocks. Again, the exact cancellation of these effects follows directly from first-order conditions for optimal price-setting. The invariance of average output to demand fluctuations appears to be a robust result of models in which price setters minimize quadratic loss functions.⁷

⁷If firms' loss functions are not quadratic, then the variance of aggregate demand generally does influence average output. The reason is that certainty equivalence fails: firms do not set prices equal to expected optimal prices. See Kimball (1989) for details.

IV. SECTORAL SHOCKS

In this section we consider a variation on our model in which firms face sectoral shocks. We find that in periods of large sectoral dispersion, inflation rises above trend, and output falls. That is, sectoral shocks influence the economy much the same as adverse shocks to aggregate supply.

A. Assumptions

We now assume that a firm's desired price depends on a sectoral shock (arising from demand or costs) as well as the money supply:

$$q_i^* = m + \theta_i, \quad (10)$$

where q_i^* is firm i 's desired price and θ is now interpreted as a sectoral shock. For simplicity, we assume here that money growth is steady: $m_t = \pi t$. In addition, the menu cost C is the same for all firms. Our interest here is in analyzing how the aggregate economy responds to sectoral shocks.

We assume that for each firm i the sectoral shock θ_i follows a random walk without drift. As a normalization, we let $\theta_i = 0$ for all i in period zero. Innovations in θ_i are distributed $N(0, s^2)$ across firms. The cross-sectional variance s^2 is time-varying: each period, s^2 is drawn from a distribution $H(\cdot)$. That is, some periods have a large amount of sectoral dispersion, and others have greater stability. A large- s^2 period could be interpreted as a year with an oil shock.

B. Prices and Output

Once again, each period-zero price setter faces essentially the same problem as in the partial-equilibrium model of Section II. Initial prices and the bounds for non-adjustment are defined by equations (3) and (5); these rules are the same for all firms because the menu cost C is homogeneous. In equation (5), $F(\cdot)$ is the unconditional distribution of an innovation in θ . It is defined by the normal distribution conditional on s^2 and by $H(\cdot)$, the distribution of s^2 .

We again derive the price level in a representative period, period one. Among firms that set initial prices at zero, those with period-one realizations of θ_i within $[\underline{\theta}, \bar{\theta}]$ maintain the initial price of x . The others--those with large sectoral shocks--adjust to $\pi + \theta_i$, the current desired price. Firms that set initial prices in period one set them equal to $x + \pi + \theta_i$. Since θ_i has mean zero across firms, the average of new initial prices is $x + \pi$. Combining these results, we find that the price level in period one is

$$p = \frac{1}{2} [x(\Phi(\bar{\theta}) - \Phi(\underline{\theta})) + \int_{\underline{\theta}}^{\infty} (\pi + \theta) d\Phi(\theta)] + \int_{\underline{\theta}}^{\bar{\theta}} (\pi + \theta) d\Phi(\theta) + (x + \pi), \quad (11)$$

where Φ is the normal cumulative distribution function, given the realization of s^2 in period one. In equation (11), the first term reflects period-zero price setters who do not adjust in period one; the two integrals reflect those who do adjust; and the last term

reflects period-one price setters. Finally, we can simplify equation (11) to

$$p = \pi + \frac{x}{2} + \frac{x-\pi}{2} [\phi(\bar{\theta}) - \phi(\underline{\theta})] + \frac{s^2}{2} [\phi(\bar{\theta}) - \phi(\underline{\theta})], \quad (12)$$

where ϕ is the normal density, given s^2 . This equation shows how sectoral dispersion s^2 influences the overall price level p . Given the quantity equation and the fact that $m_1 = \pi$, output in period one is $\pi - p$.

In equation (12), the price level p depends on the initial price x . Recall that x depends on the unconditional distribution of θ_1 , which in turn depends on $H(\cdot)$, the distribution of s^2 . One could specify $H(\cdot)$ and the parameters π and C , and then derive x from (3) and (5). Instead, we take the simpler approach of choosing x directly. In doing this, we are implicitly assuming a distribution $H(\cdot)$ that yields the given value of x . We assume that, as in the previous sections, x lies in the range $(0, \pi/2)$. (One can generate any x in this range by varying $H(\cdot)$.) Our approach allows us to derive the price level and output as functions of s with parameters π , C , and x .

C. The Influence of s

Once again, we calibrate the model with empirically plausible parameter values and then solve it numerically. We set \sqrt{C} equal to .125, which implies that firms tolerate a 12.5 percent deviation of prices from optimal prices. We set the initial price x equal to $\pi/4$, which is the midpoint of the permissible range. We then examine $p(s)$ for alternative values of π .

Figure 2 plots the price level in period one against s for π equal to 0.0, 0.025, 0.05, and 0.075. In each case in which inflation is positive, p rises monotonically with s : prices are higher in periods of greater sectoral dispersion. When $s=0$, there are no special adjustments, and so p is the average of period-zero initial prices, x , and period-one initial prices, $x+\pi$; thus $p(0)=x+\pi/2$. As s approaches infinity, all period-zero price setters adjust to an average of π ; thus p approaches $x/2+\pi$.⁸

The price level rises with the amount of sectoral dispersion for two reasons. First, because the range of non-adjustment is asymmetric, firms that receive positive θ 's are more likely to adjust than those that receive negative θ 's. Therefore, the price increases triggered by sectoral dispersion outnumber the decreases. Note, however, that this effect alone is not monotonic: as s approaches infinity, all firms adjust, and this asymmetry disappears. Second, the size of price changes, $(\pi-x+\theta)$, is asymmetric. Even if the range of non-adjustment were symmetric in θ , the adjustment triggered by sectoral shocks would cause prices to rise by an average of $\pi-x$. The intuition is that when sectoral dispersion triggers adjustment, firms catch up with inflation as well as adjusting to their individual shocks, and so prices rise on average.

These results for a representative period determine how the economy behaves over time. Trend inflation equals trend money growth π . As in Section III, average output is zero: given the

⁸We have not been able to show analytically that $p(s)$ is monotonic, but extensive numerical calculations have not produced a counterexample.

distribution of s , the initial price x adjusts to ensure this result. The price level rises above trend in periods of high sectoral dispersion and falls below trend with low dispersion; given steady money growth, output moves in the opposite direction. Hence, a high realization of s affects the economy as an adverse supply shock does in textbook models; similarly, a low realization of s is like a beneficial supply shock.

V. OPTIMAL INFLATION

In his 1972 Presidential Address, Tobin argues that a positive rate of inflation is desirable. His conclusion is based on the assumption that nominal prices can rise more easily than they can fall. Tobin reasons that inflation allows relative prices to fall without reductions in nominal prices and, therefore, permits the economy to reallocate resources among sectors more readily. Our model is suited for evaluating Tobin's argument because it exhibits the asymmetric price adjustment that Tobin assumes. In this section, therefore, we determine the socially optimal inflation rate in our model.

We measure welfare by the average loss of a representative firm. Again, the loss in a period is $(q-q^*)^2+DC$, where D equals one if the firm pays the menu cost. The first term in the loss captures the costs of deviations of prices from their frictionless levels, and the second term captures the costs of price adjustment.⁹ We consider a simple version of our model with

⁹In analyzing social welfare, the "firms" in the model should be interpreted as yeoman farmers who both produce and consume. These farmers are the only agents in the model, and so it is

constant money growth, a constant variance of sectoral shocks, and homogeneous menu costs.

Equation (4) in Section II gives a firm's expected loss for an arbitrary inflation rate π . The socially optimal inflation rate is simply the value of π that minimizes this loss. Optimal inflation is therefore defined by the first-order condition¹⁰

$$\frac{d(\text{Loss})}{d\pi} = \int_{\theta-\bar{\theta}}^{\bar{\theta}} 2(\pi+\theta-x)dF(\theta) = 0. \quad (13)$$

The solution to this condition is $\pi=0$. To see this, note that equations (3) and (5) imply $x=0$ and $\underline{\theta}=-\bar{\theta}$ when $\pi=0$. Thus, for $\pi=0$, the integral in (13) reduces to $\int_{\theta-\bar{\theta}}^{\bar{\theta}} 2\theta dF(\theta)$, which is zero by the symmetry of the distribution of θ . Contrary to Tobin, we find that the optimal inflation rate is zero.

The key to this result is the fact that price adjustment is symmetric at zero inflation. When $\pi=0$, a firm's desired price adjustment, $\pi-x+\theta$, is simply θ ; its range of non-adjustment, $[\underline{\theta}, \bar{\theta}]$, reduces to $[-\sqrt{C}, \sqrt{C}]$. These results explain why our conclusion differs from Tobin's. Tobin implicitly assumes that the asymmetry in price adjustment is exogenous, and thus impedes downward relative adjustments at zero inflation. In our model, asymmetric

appropriate to equate their losses with social welfare. Ball and Romer (1989) derive a social loss function similar to the one in this paper as an approximation to a farmer's expected utility. Note that this approximation is valid only if comparing welfare across regimes with the same average level of output. In our model, average output is in fact constant across regimes.

¹⁰The variables x , $\underline{\theta}$, and $\bar{\theta}$ are functions of π . Since firms choose these variables optimally, however, the envelope theorem allows us to hold them constant when differentiating the loss with respect to π .

adjustment arises from the tendency of inflation to push desired prices beyond actual prices. With no inflation, this tendency disappears, and adjustment becomes symmetric; thus, decreases in relative prices do not become more difficult.

Zero inflation is optimal here because positive inflation not only lacks the benefit suggested by Tobin but also imposes a cost. Inflation introduces a new reason (in addition to shocks) that desired nominal prices change between periods. As a result, actual prices deviate more from desired prices; specifically, the initial price x is too high in period zero and, on average, is too low in period one. By contrast, with $\pi=0$, the initial price of zero equals the optimal price in period zero and the expected optimum in period one. Since initial prices are closer to optimal prices under zero inflation, each firm faces a more favorable tradeoff between non-optimal prices and costs of price adjustment.

VI. CONCLUSION

Economists frequently argue that price adjustment is asymmetric. In this paper we use a menu-cost model to explore a possible explanation for such an asymmetry. In most modern economies, trend inflation is positive, and so firms' relative prices decline automatically over time. Negative shocks to desired prices offset the tendency for actual prices to fall below the desired level, whereas positive shocks exacerbate this tendency. In this environment, it is optimal for firms to respond asymmetrically to shocks.

Previous authors who discuss the implications of asymmetric price adjustment often simply assume the existence of the asymmetry. One theme of this paper is that the implications often change when the asymmetry is derived from optimization. When fixed asymmetries are assumed, demand stabilization raises average output (DeLong-Summers); relative-price dispersion is inflationary (Fischer); and optimal inflation is positive (Tobin). Of these three conclusions, only the second survives in our model of endogenous asymmetry.

A variety of evidence suggests that the asymmetries predicted by our model in fact hold in actual economies. Cover (1988) and DeLong-Summers (1988) find that shifts in aggregate demand have asymmetric effects on output in U.S. data. Sichel (1989) reports that detrended output has a distribution that is skewed to the left. Sichel's finding accords with the skewed distributions of output shown in our Figure 1.

Our model also explains the empirical finding in Fischer (1981) and Ball and Mankiw (1992) that inflation rises in periods of high relative-price variability. Most important, our model can explain this finding more easily than can the standard view that causality runs from trend inflation to relative-price variability. According to the standard view, the mechanism is staggered price setting--the fact that different firms adjust their prices at different times. Yet Fischer and Ball-Mankiw find a relationship between inflation and relative-price variability when variability is measured across industries, each of which includes many firms. It is likely that price setting is staggered within industries, and

that different industries do not systematically adjust at different times. If so, then staggered adjustment cannot explain why inflation is associated with relative-price variability at the industry level. By contrast, our model can explain this finding, because it predicts that a high dispersion of industry-specific shocks raises the price level.

An aspect of our model that might be examined in future empirical work is the relation between price adjustment and trend inflation. With positive inflation, relative prices are adjusted upward more quickly than they are adjusted downward. This asymmetry, however, disappears under price stability. Moreover, during periods of sustained deflation, the asymmetry is reversed: relative prices are adjusted downward more quickly than they are adjusted upward. Thus, our model may prove useful for understanding differences in price adjustment under different monetary regimes.

REFERENCES

- Ball, Laurence, and N. Gregory Mankiw, "Relative-Price Changes as Aggregate Supply Shocks," 1992.
- Ball, Laurence, N. Gregory Mankiw, and David Romer, "The New Keynesian Economics and the Output-Inflation Trade-off," Brookings Papers on Economic Activity 1988, no.1: 1-65.
- Ball, Laurence, and David Romer, "The Equilibrium and Optimal Timing of Price Changes," Review of Economic Studies 56 (April, 1989): 179-198.
- Blinder, Alan S., "Why are Prices Sticky? Preliminary Evidence from an Interview Survey," American Economic Review 81 (May 1991): 89-96.
- Bertola, Giuseppe, and Ricardo J. Caballero, "Kinked Adjustment Costs and Aggregate Dynamics," NBER Macroeconomics Annual, 1990: 237-288.
- Caballero, Ricardo J., and Eduardo Engel, "The Output-Inflation Tradeoff Revisited," MIT, October 1991.
- Caplin, Andrew S., and John Leahy, "State-Contingent Pricing and the Dynamics of Money and Output," Quarterly Journal of Economics 106, August 1991, 683-708.
- Caplin, Andrew S., and Daniel F. Spulber, "Menu Costs and the Neutrality of Money," Quarterly Journal of Economics 102 (November 1987): 703-725.
- Cecchetti, Stephen G., "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines, 1953-1979," Journal of Econometrics 20 (April 1986): 255-274.

- Cover, James P., "Asymmetric Effects of Positive and Negative Money-Supply Shocks, University of Alabama, 1988.
- DeLong, J. Bradford, and Lawrence H. Summers, "How Does Macroeconomic Policy Affect Output?" Brookings Papers on Economic Activity 1988:2: 433-480.
- Fischer, Stanley, "Relative Shocks, Relative Price Variability, and Inflation," Brookings Papers on Economic Activity 1981:2: 381-431.
- Friedman, Milton, The Optimum Quantity of Money and Other Essays, Chicago: Aldine, 1969.
- Kimball, Miles S., "The Effect of Demand Uncertainty on a Precommitted Monopoly Price," Economics Letters 30, no. 1, (1989): 1-5.
- Lucas, Robert E., Jr., Models of Business Cycles, New York: Basil Blackwell, 1987.
- Sheshinski, Eytan, and Yoram Weiss, "Inflation and Costs of Price Adjustment," Review of Economic Studies 44 (June 1977): 287-303.
- Sichel, Daniel, "Business Cycle Asymmetry: A Deeper Look," Federal Reserve Economic Activity Working Paper No. 93, 1989.
- Tobin, James, "Inflation and Unemployment," American Economic Review 62 (March 1972): 1-18.
- Tsiddon, Daniel, "On the Stubbornness of Sticky Prices," International Economic Review 32 (February 1991a): 69-75.
- Tsiddon, Daniel, "The (Mis)behavior of the Aggregate Price Level," Hebrew University, October 1991b.

Table 1

Output as a Function of θ

$\sigma=2.5\%$, $\Gamma=25\%$

All values expressed in percent.

θ	$\pi=0$		$\pi=2.5$	
	$Y(\theta)$	$Y(\theta)-Y(0)$	$Y(\theta)$	$Y(\theta)-Y(0)$
-10.00	-3.00	-3.00	-3.38	-3.45
-7.50	-2.63	-2.63	-2.89	-2.96
-5.00	-2.00	-2.00	-2.16	-2.23
-2.50	-1.13	-1.13	-1.07	-1.14
0.00	0.00	0.00	0.07	0.00
2.50	1.13	1.13	1.04	0.97
5.00	2.00	2.00	1.80	1.73
7.50	2.63	2.63	2.31	2.24
10.00	3.00	3.00	2.56	2.49

θ	$\pi=5.0$		$\pi=7.5$	
	$Y(\theta)$	$Y(\theta)-Y(0)$	$Y(\theta)$	$Y(\theta)-Y(0)$
-10.00	-3.67	-3.73	-4.00	-4.06
-7.50	-3.16	-3.22	-2.88	-2.93
-5.00	-2.05	-2.12	-1.75	-1.81
-2.50	-0.92	-0.98	-0.82	-0.87
0.00	0.07	0.00	0.05	0.00
2.50	0.93	0.87	0.84	0.78
5.00	1.59	1.52	1.39	1.33
7.50	1.97	1.91	1.65	1.60
10.00	2.10	2.03	1.65	1.60

Figure 1
Distribution of Output ($\sigma=0.025, \Gamma=0.25$)

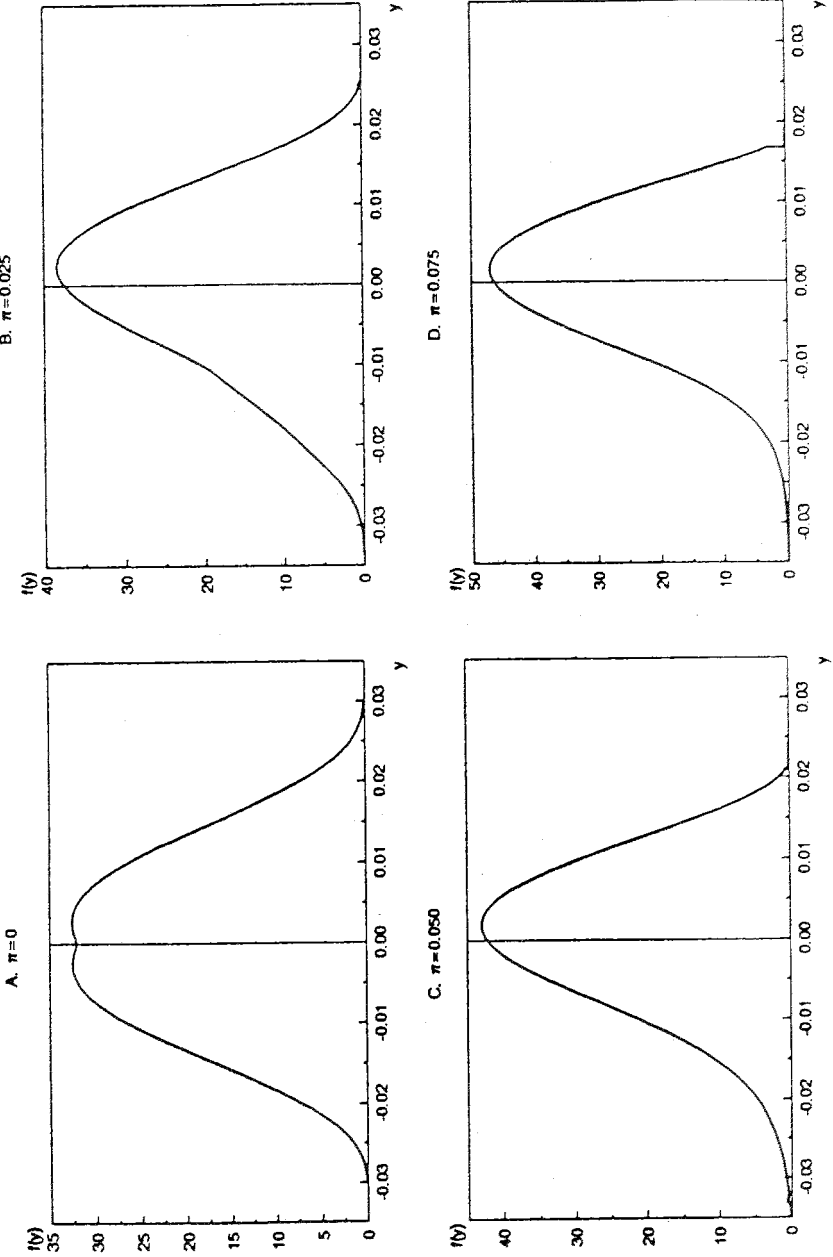


Figure 2
Price Level as a Function of Sectoral Dispersion
($\sqrt{C}=0.125, \chi=\pi/4$)

