

LAURENCE BALL

New York University

N. GREGORY MANKIW

Harvard University

DAVID ROMER

Princeton University

The New Keynesian Economics and the Output- Inflation Trade-off

IN THE EARLY 1980s, the Keynesian view of business cycles was in trouble. The problem was not new empirical evidence against Keynesian theories, but weakness in the theories themselves.¹ According to the Keynesian view, fluctuations in output arise largely from fluctuations in nominal aggregate demand. These changes in demand have real effects because nominal wages and prices are rigid. But in Keynesian models of the 1970s, the crucial nominal rigidities were assumed rather than

We thank Toshiaki Jinushi, David Johnson, and David Weil for research assistance; Ray Fair, Paul Wachtel, and members of the Brookings Panel for helpful discussions; and the National Science Foundation for financial support.

1. Keynesian models of wage and price adjustment based on Phillips curves provided poor fits to the data of the early-to-mid-1970s. But subsequent modifications of the models, such as the addition of supply shocks, have led to fairly good performances. See the discussions in Olivier J. Blanchard, "Why Does Money Affect Output? A Survey," Working Paper 2285 (National Bureau of Economic Research, June 1987); and Robert J. Gordon, "Postwar Developments in Business Cycle Theory: An Unabashedly New-Keynesian Perspective," Keynote Lecture, 18th CIRET Conference, Zurich, September 1987.

explained—assumed directly, as in disequilibrium models, or introduced through theoretically arbitrary assumptions about labor contracts.² Indeed, it was clearly in the interests of agents to eliminate the rigidities they were assumed to create. If wages, for example, were set above the market-clearing level, firms could increase profits by reducing wages. Microeconomics teaches us to reject models in which, as Robert Lucas puts it, “there are \$500 bills on the sidewalk.” Thus the 1970s and early 1980s saw many economists turn away from Keynesian theories and toward new classical models with flexible wages and prices.

But Keynesian economics has made much progress in the past few years. Recent research has produced models in which optimizing agents choose to create nominal rigidities. This accomplishment derives largely from a central insight: nominal rigidities, and hence the real effects of nominal demand shocks, can be large even if the frictions preventing full nominal flexibility are slight. Seemingly minor aspects of the economy, such as costs of price adjustment and the asynchronized timing of price changes by different firms, can explain large nonneutralities.

Theoretical demonstrations that Keynesian models can be reconciled with microeconomics do not constitute proof that Keynesian theories are correct. Indeed, a weakness of recent models of nominal rigidities is that they do not appear to have novel empirical implications. As Lawrence Summers argues:

While words like menu costs and overlapping contracts are often heard, little if any empirical work has demonstrated connection between the extent of these phenomena and the pattern of cyclical fluctuations. It is difficult to think of any anomalies that Keynesian research in the “nominal rigidities” tradition has resolved, or of any new phenomena that it has rendered comprehensible.³

The purpose of this paper is to provide evidence supporting new Keynesian theories. We point out a simple prediction of Keynesian

2. For disequilibrium models, see Robert J. Barro and Herschel I. Grossman, “A General Disequilibrium Model of Income and Employment,” *American Economic Review*, vol. 61 (March 1971), pp. 82–93; and E. Malinvaud, *The Theory of Unemployment Reconsidered* (Basil Blackwell, 1977). For contract models, see Stanley Fischer, “Long-Term Contracts, Rational Expectations and the Optimal Money Supply Rule,” *Journal of Political Economy*, vol. 85 (February 1977), pp. 191–205; and Jo Anna Gray, “On Indexation and Contract Length,” *Journal of Political Economy*, vol. 86 (February 1978), pp. 1–18.

3. Lawrence H. Summers, “Should Keynesian Economics Dispense with the Phillips Curve?” in Rod Cross, ed., *Unemployment, Hysteresis, and the Natural Rate Hypothesis* (Basil Blackwell, 1988), p. 12.

models that contradicts other leading macroeconomic theories and show that it holds in actual economies. In doing so, we point out a “new phenomenon” that Keynesian theories “render comprehensible.”

The prediction that we test concerns the effects of steady inflation. In Keynesian models, nominal shocks have real effects because nominal prices change infrequently. An increase in the average rate of inflation causes firms to adjust prices more frequently to keep up with the rising price level. In turn, more frequent price changes imply that prices adjust more quickly to nominal shocks, and thus that the shocks have smaller real effects. We test this prediction by examining the relation between average inflation and the size of the real effects of nominal shocks both across countries and over time. We measure the effects of nominal shocks by the slope of the short-run Phillips curve.

Other prominent macroeconomic theories do not predict that average inflation affects the slope of the Phillips curve. In particular, our empirical work provides a sharp test between the Keynesian explanation for the Phillips curve and the leading new classical alternative, the Lucas imperfect information model.⁴ Indeed, one goal of this paper is to redo Lucas’s famous analysis and dramatically reinterpret his results. Lucas and later authors show that countries with highly variable aggregate demand have steep Phillips curves. That is, nominal shocks in these countries have little effect on output. Lucas interprets this finding as evidence that highly variable demand reduces the perceived relative price changes resulting from nominal shocks. We provide a Keynesian interpretation of Lucas’s result: more variable demand, like high average inflation, leads to more frequent price adjustment. We then test the differing implications of the two theories for the effects of average inflation. Our results are consistent with the Keynesian explanation for the Phillips curve and inconsistent with the classical explanation.

In addition to providing evidence about macroeconomic theories, our finding that average inflation affects the short-run output-inflation trade-off is important for policy. For example, it is likely that the trade-off facing policymakers in the United States has changed as a consequence of disinflation in the 1980s. Our estimates imply that a reduction in

4. Robert E. Lucas, Jr., “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, vol. 4 (April 1972), pp. 103–24; Lucas, “Some International Evidence on Output-Inflation Tradeoffs,” *American Economic Review*, vol. 63 (June 1973), pp. 326–34.

average inflation from 10 percent to 5 percent substantially alters the short-run impact of aggregate demand.

The body of the paper consists of three major sections. The first discusses the new research that provides microeconomic foundations for Keynesian theories. The second presents a model of price adjustment. It demonstrates the connection between average inflation and the slope of the Phillips curve and contrasts this result with the predictions of other theories. The third section provides both cross-country and time series evidence that supports the predictions of the model.

New Keynesian Theories

According to Keynesian economics, fluctuations in employment and output arise largely from fluctuations in nominal aggregate demand. The reason that nominal shocks matter is that nominal wages and prices are not fully flexible. These views are the basis for conventional accounts of macroeconomic events. For example, the consensus explanation for the 1982 recession is slow growth in nominal demand resulting from tight monetary policy. The research program described here is modest in the sense that it seeks to strengthen the foundations of this conventional thinking, not to provide a new theory of fluctuations. In particular, its goal is to answer the theoretical question of how nominal rigidities arise from optimizing behavior, since the absence of an answer in the 1970s was largely responsible for the decline of Keynesian economics.

In the following discussion we first describe the central point of the recent literature: large nominal rigidities are possible even if the frictions preventing full nominal flexibility are small. We next describe some phenomena that greatly strengthen the basic argument, including rigidities in *real* wages and prices and asynchronized timing of price changes. We then discuss two innovations in recent models that are largely responsible for their success: the introduction of imperfect competition and an emphasis on price as well as wage rigidity. Finally, we argue that the ideas in recent work are indispensable for a plausible Keynesian account of fluctuations.⁵

5. Some of the ideas of this literature are discussed informally by earlier Keynesian authors. To cite just two examples, see the discussion of asynchronized timing of price changes in Robert J. Gordon, "Output Fluctuations and Gradual Price Adjustment,"

SMALL NOMINAL FRICTIONS AND LARGE NOMINAL RIGIDITIES

The recent literature on nominal rigidities enters an argument that Keynesians appeared to be losing. Members of the new classical school that developed in the 1970s challenged Keynesians to explain the rigidities in Keynesian models. In response, Keynesians sometimes cited costs of adjusting prices. But as the classicals pointed out, these costs, while surely present, appear small. Indeed, the frequently mentioned “menu costs”—the costs of printing new menus and catalogs, of replacing price tags, and so on—sound trivial. Thus the impediments to nominal flexibility in actual economies appear too small to provide a foundation for Keynesian models.

A common but mistaken response is that there are many obvious sources of large wage and price rigidities: implicit contracts, customer markets, efficiency wages, insider-outsider relationships, and so on. The problem is that these phenomena imply rigidities in *real* wages and prices, while the Keynesian theory depends on rigidities in *nominal* wages and prices. Real rigidities are no impediment to complete flexibility of nominal prices, because full adjustment to a nominal shock does not require any change in real prices. The absence of models of nominal rigidity reflects the microeconomic proposition that agents do not care about nominal magnitudes. The only apparent departures from this proposition in actual economies are the small costs of nominal adjustment.

Thus recent work begins with the premise that it is inexpensive to reduce nominal rigidity and asks how substantial rigidity nonetheless arises. The central answer of the literature is presented by Mankiw, Akerlof and Yellen, Blanchard and Kiyotaki, and Ball and Romer.⁶

Journal of Economic Literature, vol. 19 (June 1981), pp. 493–530, and the discussion of externalities from nominal rigidity in Charles L. Schultze, “Microeconomic Efficiency and Nominal Wage Stickiness,” *American Economic Review*, vol. 75 (March 1985), pp. 1–15.

6. N. Gregory Mankiw, “Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly,” *Quarterly Journal of Economics*, vol. 100 (May 1985), pp. 529–37; George A. Akerlof and Janet L. Yellen, “A Near-Rational Model of the Business Cycle, with Wage and Price Inertia,” *Quarterly Journal of Economics*, vol. 100 (1985, Supplement), pp. 823–38; Olivier Jean Blanchard and Nobuhiro Kiyotaki, “Monopolistic Competition and the Effects of Aggregate Demand,” *American Economic Review*, vol. 77 (September 1987), pp. 647–66; Laurence Ball and David Romer, “Are Prices Too Sticky?” Working Paper 2171 (NBER, February 1987).

Second-Order Private Costs and First-Order Business Cycles. Mankiw and Akerlof and Yellen make a simple but important point. They study imperfectly competitive economies and show that the cost of nominal rigidities to price setters can be much smaller than the macroeconomic effects. An example that illustrates the cost to price setters is a firm that initially sets its price at the profit-maximizing level but does not adjust after the money supply falls. We let $\pi(\bullet)$ denote the firm's profits as a function of its price and let P be the firm's predetermined price and P^* its profit-maximizing price, which it would set if it adjusted. Using a Taylor expansion, we can approximate the firm's profit loss from not adjusting as

$$(1) \quad \pi(P^*) - \pi(P) \approx \pi'(P^*)(P^* - P) - \frac{1}{2} \pi''(P^*)(P^* - P)^2.$$

But since P^* maximizes profits, $\pi'(P^*)$ is zero. Thus the profit loss from nonadjustment is second order—that is, proportional to the square of $(P^* - P)$. As long as the predetermined price is close to the profit-maximizing price, the cost of price rigidity to the firm is small.

But rigidity can have first-order macroeconomic effects. An increase in nominal money with nominal prices fixed leads to a first-order increase in real aggregate demand, and hence in real output. For example, if the aggregate demand curve is simply $Y = M/P$, rigid prices imply a change in output proportional to the change in money.

The effect on social welfare is also first order, as follows from the assumption of imperfect competition. Under imperfect competition, the profit-maximizing price is socially suboptimal. The price is too high and output is too low. Thus at P^* the first derivative of welfare with respect to the firm's price is negative: welfare would rise if the price fell below P^* . Nonadjustment to a fall in money implies P greater than P^* ; given the negative first derivative of welfare, the welfare loss is first order.

Because the cost of rigidity to a price setter is second order while the macroeconomic effects are first order, the latter can be much larger. This finding resolves the puzzle of why price setters refuse to incur the small costs of reducing the business cycle through more flexible prices. Despite the large macroeconomic effects, the private incentives are small.

Aggregate Demand Externalities. Blanchard and Kiyotaki provide an important interpretation of the result in Mankiw and Akerlof-Yellen:

the macroeconomic effects of nominal rigidity differ from the private costs because rigidity has an “aggregate demand externality.” A few equations make this clear. Suppose the demand for the product of firm i depends on aggregate spending and on the firm’s relative price:

$$(2) \quad Y_i^D = \left(\frac{P_i}{P} \right)^{-\epsilon} Y.$$

For simplicity, aggregate demand is given by a quantity equation⁷

$$(3) \quad Y = \frac{M}{P}.$$

Combining equations 2 and 3 yields

$$(4) \quad Y_i^D = \left(\frac{M}{P} \right) \left(\frac{P_i}{P} \right)^{-\epsilon}.$$

According to equation 4, firm i ’s demand depends on its relative price and on real money, which determines aggregate demand. Changes in real money shift the demand curve facing firm i , and the firm’s price determines its position on the demand curve.

If M falls and firm i does not adjust, the second-order cost to firm i is that P_i/P does not adjust to the new profit-maximizing level. The externality is that rigidity in firm i ’s price contributes to rigidity in the aggregate price level. Given the fall in nominal money, rigidity in P implies a first-order fall in real money, which reduces demand for all firms’ goods. In other words, there is an externality because adjustment of all prices would prevent a fall in real aggregate demand, but each firm is a small part of the economy and thus ignores this macroeconomic benefit.

The importance of the externality is illustrated by a firm in a recession caused by tight money. To the firm, the recession means an inward shift of its demand curve and a resulting first-order loss in profits. The firm would very much like to shift its demand curve back out, but of course it cannot do so by changing its price. Instead, price adjustment would yield only the second-order gain from optimally dividing the losses from

7. The only essential feature of equation 3 is the negative relation between Y and P . We can interpret M as simply a shift term in the aggregate demand equation. Thus, as we discuss below, the results in recent papers concern the effects of any shock to aggregate demand, not just changes in the money stock.

the recession between reduced sales and a lower price. The recession would end and everyone would be much better off if all firms adjusted. But each firm believes that it cannot end the recession and therefore may fail to adjust even if the costs of adjustment are much smaller than the costs of the recession.

This argument resembles standard microeconomic analyses of externalities. Consider the classic example of pollution. Pollution would be greatly reduced, and social welfare greatly improved, if each person incurred the small cost of walking to the trash can at the end of the block. But each individual ignores this when he throws his wrapper on the street because he is only one of many polluters. Because of externalities, economists do not find highly inefficient levels of pollution puzzling even though the costs of reducing pollution are small. For similar reasons, highly inefficient nominal rigidities are not a mystery even though menu costs are small.

Externalities from Fluctuations in Demand. Keynesians believe not only that shocks to nominal aggregate demand cause large fluctuations in output and welfare, but also that these fluctuations are inefficient, and thus that stabilization of demand is desirable. The models surveyed so far do not provide a foundation for this view. As explained above, nonadjustment of prices to a fall in demand leads to large reductions in output and welfare. But nonadjustment to a *rise* in demand leads to higher output and, because output is initially too low under imperfect competition, to higher welfare. Thus the implications of fluctuations for *average* welfare, and hence the desirability of reducing fluctuations, are unclear. Indeed, Ball and Romer show that the first-order welfare effects of fluctuations average to zero, which means that the first order–second order distinction is irrelevant to this issue.⁸

Nonetheless, Ball and Romer show, by comparing the average social and private costs of nominal rigidity, that small nominal frictions are sufficient for large reductions in average welfare. The private cost is fluctuations of a firm's relative price around the profit-maximizing level. The social cost is the private cost plus the cost of fluctuations in real aggregate demand. Greater flexibility would stabilize real demand, but each firm ignores its effect on the variance of demand, just as it ignores its effect on the level of demand after a given shock. Although both the

8. Ball and Romer, "Are Prices Too Sticky?"

average social and average private costs are second order, Ball and Romer show that the former may be much larger: fluctuations in aggregate demand can be much more costly than fluctuations in relative prices. As a result, small frictions can prevent firms from adopting greater flexibility even if business cycles are highly inefficient.

STILL LARGER RIGIDITIES

The papers discussed so far establish that nominal rigidities can be far larger than the frictions that cause them. But as we now describe, the simple models in these papers cannot fully explain nonneutralities of the size and persistence observed in actual economies. Therefore, we turn to more complicated models that incorporate realistic phenomena that magnify nominal rigidities. These phenomena include rigidities in *real* wages and prices and asynchronized timing of price changes by different firms.

Real Rigidities. As we argue above, real rigidities alone are no impediment to full nominal flexibility. But Ball and Romer show that a high degree of real rigidity, defined as small responses of real wages and real prices to changes in real demand, greatly increases the nonneutralities arising from small nominal frictions.⁹

This finding is important because, although models with nominal frictions but no real rigidities can in principle produce large nominal rigidities, they do so only for implausible parameter values. Most important, large rigidities arise only if labor supply is highly elastic, while labor supply elasticities in actual economies appear small. The role of labor supply is illustrated by a hypothetical economy with imperfect competition and menu costs in the goods market but a Walrasian labor market. If menu costs led to nominal price rigidity, then nominal shocks would cause large shifts in labor demand. But if labor supply were inelastic, these shifts in labor demand would cause large changes in the real wage and thereby create large incentives for price setters to adjust their prices. As a result, nominal rigidity would not be an equilibrium.

While for plausible parameter values nominal frictions alone produce little nominal rigidity, Ball and Romer show that considerable rigidity

9. Laurence Ball and David Romer, "Real Rigidities and the Non-Neutrality of Money," Working Paper 2476 (NBER, December 1987).

can arise if the frictions are combined with real rigidities arising from efficiency wages, customer markets, and the like. For example, substantial nominal rigidity can arise from a combination of real rigidity in the labor market and imperfect competition and menu costs in the goods market. If firms pay efficiency wages, for instance, then real wages may be set above the market-clearing level, so that workers are off their labor supply curves. In this situation a fall in labor demand can greatly reduce employment without a large fall in the real wage even if labor supply is inelastic.

The importance of real rigidities for explaining nominal rigidities is not settled, because there is no consensus about the sources and magnitudes of real rigidities in actual economies. In particular, phenomena like efficiency wages and customer markets increase nominal rigidity to the extent that they reduce desired responses of real wages and real prices to demand shifts, but economists are still unsure of the sizes of these effects. Further research on real rigidities will lead to a better understanding of nominal rigidities.

Staggered Price Setting. Even when real rigidities are added, the models surveyed so far cannot fully explain the size and persistence of the real effects of nominal shocks. In these models, the effects of shocks are eliminated when nominal prices adjust. In actual economies, recessions following severe demand contractions can last for several years, and while individual prices are fixed for substantial periods, these periods are generally shorter than several years. Thus models with sticky prices must explain why the effects of shocks persist after all prices are changed.

An explanation is provided by the literature on staggered price setting, which shows that if firms change prices at different times, the adjustment of the aggregate price level to shocks can take much longer than the time between adjustments of each individual price.¹⁰ The “price level inertia” caused by staggering implies that nominal shocks can have large and long-lasting real effects even if individual prices change frequently.

10. John B. Taylor, “Staggered Wage Setting in a Macro Model,” *American Economic Review*, vol. 69 (May 1979, *Papers and Proceedings*, 1978), pp. 108–13; Taylor, “Aggregate Dynamics and Staggered Contracts,” *Journal of Political Economy*, vol. 88 (February 1980), pp. 1–23; Olivier J. Blanchard, “Price Asynchronization and Price Level Inertia,” in Rudiger Dornbusch and Mario Henrique Simonsen, eds., *Inflation, Debt, and Indexation* (MIT Press, 1983), pp. 3–24; Olivier J. Blanchard, “The Wage Price Spiral,” *Quarterly Journal of Economics*, vol. 101 (August 1986), pp. 543–65.

A simple example makes clear the importance of the timing of price changes. Suppose first that every firm adjusts its price on the first of each month, so that price setting is synchronized. If the money supply falls on June 10, output is reduced from June 10 to July 1, because nominal prices are fixed during this period. But on July 1 all prices adjust in proportion to the fall in money, and the recession ends.

Now suppose that half of all firms set prices on the first of each month and half on the fifteenth. If the money supply falls on June 10, then on June 15 half the firms have an opportunity to adjust their prices. But in this case they may choose to make little adjustment. Because half of all nominal prices remain fixed, adjustment of the other prices implies changes in *relative* prices, which firms may not want. (In contrast, if all prices change simultaneously, full nominal adjustment does not affect relative prices.) If the June 15 price setters make little adjustment, then the other firms make little adjustment when their turn comes on July 1, because they do not desire relative price changes either. And so on. The price level declines slowly as the result of small decreases every first and fifteenth, and the real effects of the fall in money die out slowly. In short, price adjustment is slow because neither group of firms is willing to be the first to make large cuts.¹¹

As Blanchard emphasizes, if staggering occurs among firms at different points in a chain of production, its effects are strengthened.¹² A

11. A natural question is why firms change prices at different times if this exacerbates aggregate fluctuations. One obvious answer is that different firms receive shocks at different times and face different costs of price adjustment. Laurence Ball and David Romer, "The Equilibrium and Optimal Timing of Price Changes," Working Paper 2412 (NBER, October 1987), show that, because of externalities from staggering, idiosyncratic shocks can lead to staggering even if synchronized price setting is Pareto superior. But idiosyncratic shocks cannot explain all staggering. For example, some firms with two-year labor contracts set wages in even years and some set them in odd years, and this does not correspond to deterministic two-year cycles in the arrival of shocks. Another explanation for staggering is that it arises from firms' efforts to gain information. This source of staggering is discussed in Arthur M. Okun, *Prices and Quantities: A Macroeconomic Analysis* (Brookings, 1981), and formalized in Laurence Ball and Stephen G. Cecchetti, "Imperfect Information and Staggered Price Setting," Working Paper 2201 (NBER, April 1987). For example, a firm wants to set wages in line with the wages of other firms. If all wages are set simultaneously, each firm is unsure of what wage to set because it does not know what others will do. This gives each firm an incentive to set its wage shortly after the others. The desire of each firm to "bat last," as Okun puts it, can lead in equilibrium to a uniform distribution of signing dates.

12. Blanchard, "Price Asynchronization."

firm's profit-maximizing price is tied to both the prices of its inputs and the prices of goods for which its product is an input (the latter influence demand for the firm's product). Thus a firm does not want to adjust its price to a shock if these other prices do not adjust at the same time. This reluctance to make asynchronized adjustments causes price level inertia. Blanchard shows that the degree of inertia increases the longer the chain of production: it takes a long time for the gradual adjustment of prices to make its way through a complicated system.

The literature on staggered price setting complements that on nominal rigidities arising from menu costs. The degree of rigidity in the aggregate price level depends on both the frequency and the timing of individual price changes. Menu costs cause prices to adjust infrequently. For a given frequency of individual adjustment, staggering slows the adjustment of the price level. Large aggregate rigidities can thus be explained by a combination of staggering and nominal frictions: the former magnifies the rigidities arising from the latter.

Asymmetric Effects of Demand Shocks. We conclude this part of our discussion by mentioning a little-explored possibility for strengthening Keynesian models. The models surveyed imply symmetric responses of the economy to rises and falls in nominal aggregate demand. For example, in menu cost models the range of shocks to which prices do not adjust is symmetric around zero, and so is the range of possible changes in output. But traditional Keynesian models often imply asymmetric effects of demand shifts. In undergraduate texts, for example, the aggregate supply curve is often drawn so that decreases in demand lead to large output losses while the effects of increases are mostly dissipated through higher prices. Such asymmetries are intuitively appealing, and they greatly strengthen the Keynesian view that demand stabilization is desirable: stabilization raises the average levels of output and employment as well as reducing the variances. It is unclear whether plausible modifications of new Keynesian models can produce asymmetries. Asymmetric effects of shocks could arise from asymmetric price rigidity—prices that are sticky downward but not upward—but this is another appealing notion that is difficult to formalize.¹³

13. Timur Kuran shows that asymmetries in firms' profit functions, which many menu cost models ignore, can lead to asymmetric price rigidity. But the asymmetries appear small. See Timur Kuran, "Asymmetric Price Rigidity and Inflationary Bias," *American Economic Review*, vol. 73 (June 1983), pp. 373–82; Kuran, "Price Adjustment Costs,

THE NEW ASSUMPTIONS IN NEW KEYNESIAN MODELS

Aside from the specific arguments outlined above, recent research establishes the general point that nominal rigidities can result from optimizing choices of agents in well-specified models. This contrasts with the ad hoc imposition of rigidities in many of the Keynesian models of the 1970s. Recent progress is largely a result of two innovations in modeling: the introduction of imperfect competition and greater emphasis on price rather than wage rigidities.

Imperfect Competition. Microeconomists have long recognized that sticky prices and perfect competition are incompatible.¹⁴ In a competitive market, a firm does not set its price, but accepts the price quoted by the Walrasian auctioneer. Only under imperfect competition, when firms set prices, does it make sense to ask whether a firm adjusts its price to a shock. Nonetheless, Keynesian models of the 1970s, most clearly disequilibrium models, imposed nominal rigidities on otherwise Walrasian economies. The result was embarrassments in the form of unappealing results or the need for additional arbitrary assumptions. Many recent models simply generalize earlier models by allowing the firms' demand curves to slope down. This single modification sweeps away many of the problems with older models. Specifically, the new models with imperfect competition offer six advantages:

—*Private costs of rigidity are second order.* Under perfect competition, the gains from nominal adjustment are large. For example, if nominal demand rises and prices do not adjust, there is excess demand. In this situation, an individual firm can raise its price significantly and still sell as much output as before, which implies a large increase in profits. In contrast, under imperfect competition a higher price always implies lower sales. Starting from the profit-maximizing price-quantity combination, the gains from trading off price and sales after a shock are second order.

Anticipated Inflation, and Output," *Quarterly Journal of Economics*, vol. 101 (May 1986), pp. 407–18.

14. See, for example, Kenneth J. Arrow, "Toward a Theory of Price Adjustment," in Moses Abramovitz and others, eds., *The Allocation of Economic Resources: Essays in Honor of Bernard Francis Haley* (Stanford University Press, 1959), pp. 41–51.

—*Output is demand determined.* When price rigidity is imposed on a Walrasian market, so that the market does not clear, it is natural to assume that quantity equals the smaller of supply and demand, so that output falls below the Walrasian level when price is either above or below the Walrasian level. But Keynesians believe that when prices are rigid, increases in demand, which mean prices below Walrasian levels, raise output, just as decreases in demand reduce output. This result is built into many Keynesian models through the unappealing assumption that output is demand determined even if demand exceeds supply. For example, in the Gray-Fischer contract model, firms hire as much labor as they want, regardless of the preferences of workers.¹⁵ In contrast, under imperfect competition, demand determination arises naturally. Firms set prices and then meet demand. Crucially, if demand rises, firms are happy to sell more even if they do not adjust their prices, because under imperfect competition price initially exceeds marginal cost. Thus changes in demand always cause changes in output in the same direction.

—*Booms raise welfare.* Under perfect competition, the equilibrium level of output in the absence of shocks is efficient. Thus increases in output resulting from positive shocks, as well as decreases resulting from negative shocks, reduce welfare. In the Gray-Fischer model, for example, half the welfare loss from the business cycle occurs when workers are required to work more than they want. In actual economies, unusually high output and employment mean that the economy is doing well.¹⁶ And this is the case in models of imperfect competition. Since imperfect competition pushes the no-shock level of output below the social optimum, welfare rises when output rises above this level.

—*Wage rigidity causes unemployment through low aggregate demand.* In 1970s models with sticky nominal wages, unemployment occurs when prices fall short of the level expected when wages were set, so that real wages rise and firms move up their labor demand curves. In actual economies, however, firms often appear to reduce employment because demand for their output is low, not because real wages are high. This fact is not necessarily a problem for Keynesian theories if the goods market is imperfectly competitive. In this case, a firm's labor demand

15. Fischer, "Long-Term Contracts," and Gray, "On Indexation."

16. Of course economists worry that low unemployment may be inflationary. But sticky-price models with perfect competition imply that low unemployment is undesirable per se.

depends on real aggregate demand as well as the real wage, because changes in aggregate demand shift the firm's product demand (see equation 2).

—*Real wages need not be countercyclical.* Imperfect competition can remedy an embarrassing empirical failure of traditional models based on sticky nominal wages—the cyclical behavior of real wages. We can tautologically write $P = \mu W/MPL$, where P is the price level, W is the wage, MPL is the marginal product of labor, and μ is the markup of price over marginal cost. If the markup is constant and marginal product of labor is diminishing, as many 1970s models assumed, then the real wage, $W/P = MPL/\mu$, must be countercyclical. In actual economies, however, real wages appear acyclical or a bit procyclical. This fact can be explained if the marginal product of labor is constant, as suggested by Hall, or if the markup is countercyclical, as suggested by Rotemberg and Saloner and by Bils.¹⁷ Thus there need not be a link between changes in employment and changes in real wages.

—*Nominal rigidities have aggregate demand externalities.* As we have explained, since real aggregate demand affects the demand curves facing individual firms, nominal rigidities have externalities. Rigidity in one firm's price contributes to rigidity in the price level, which causes fluctuations in real aggregate demand and thus harms all firms. These externalities are crucial to the finding that small frictions can have large macroeconomic effects. The externalities depend on imperfect competition, for under perfect competition, aggregate demand is irrelevant to individual firms because they can sell all they want at the going price.

Product Market Rigidities. Keynes and most Keynesians emphasize rigidities in nominal wages. But recent work focuses largely on rigidities in product prices. The change offers two advantages.

—*Goods are sold in spot markets.* Although there is clearly much nominal wage rigidity in actual economies—in U.S. labor contracts, for example, wages are set up to three years in advance—the allocative effects of this rigidity are unclear. The implicit contracts literature shows that it may be efficient for contract signers to make employment

17. Robert E. Hall, "Market Structure and Macroeconomic Fluctuations," *BPEA*, 2: 1986, pp. 285–322; Julio J. Rotemberg and Garth Saloner, "A Supergame-Theoretic Model of Price Wars during Booms," *American Economic Review*, vol. 76 (June 1986), pp. 390–407; Mark Bils, "Cyclical Pricing of Durable Luxuries," Working Paper 83 (University of Rochester Center for Economic Research, May 1987).

independent of wages. That is, given long-term relationships with their workers, firms may choose the efficient amount of employment rather than moving along their labor demand curves when real wages change.¹⁸ In many product markets, on the other hand, buyers clearly operate on their demand curves. For example, the local shoe store has no agreement, explicit or implicit, from its customers to buy the efficient number of shoes regardless of the prices. Instead, rigidity in the store's prices affects its sales of shoes.

—*Real wages need not be countercyclical.* As we argue above, acyclical real wages are possible even if nominal rigidities occur only in wages. But it is easiest to explain acyclical or procyclical real wages if prices as well as wages are sticky. In this case, the effect of a shock on real wages depends on the relative sizes of the adjustments of prices and wages.

Despite the advantages of studying rigidities in goods markets, we are ambivalent about the deemphasis of labor markets, because the apparent rigidities in nominal wages may have important allocative effects. Further research on the relative importance of wage and price rigidities is needed.

DISCUSSION

We conclude this section by discussing several issues concerning the importance of recent theories and their plausibility.

The Importance of Nominal Rigidities. Nominal rigidities are essential for explaining important features of business cycles. As we have emphasized, real effects of nominal disturbances, such as changes in the money stock, depend on some nominal imperfection. The only prominent alternative to nominal rigidities is imperfect information about the aggregate price level, an explanation that many economists find implausible. It is possible, of course, to maintain that money is neutral in the short run—that Paul Volcker, for example, had nothing to do with the 1982 recession—but this also appears unrealistic to many economists.

18. Early expositions of this idea appear in Martin Neil Baily, "Wages and Employment under Uncertain Demand," *Review of Economic Studies*, vol. 41 (January 1974), pp. 37–50; Costas Azariadis, "Implicit Contracts and Underemployment Equilibria," *Journal of Political Economy*, vol. 83 (December 1975), pp. 1183–1202; and Robert E. Hall, "The Rigidity of Wages and the Persistence of Unemployment," *BPEA*, 2:1975, pp. 301–35.

Thus it is difficult to explain the relation of output to nominal variables without nominal rigidities.

Nominal rigidities are also important for explaining the effects of *real* shocks to aggregate demand, resulting, for example, from changes in government spending or in the expectations of investors. The point is clear if we interpret M in the aggregate demand equation, $Y = M/P$, as simply a shift term, in which case real disturbances that shift demand affect output through the same channels as changes in money.

Not all explanations for the output effects of real demand shocks depend on nominal imperfections. Robert Barro's model of government purchases, for one, does not.¹⁹ But such explanations invoke implausibly large labor supply elasticities. Thus nominal rigidities, while not the only explanation for the effects of real demand, are perhaps the most appealing.

In the models we have surveyed, slow adjustment of prices implies that shocks cause temporary deviations of output and employment from their "natural rates." Recently, however, models of hysteresis, in which shocks have permanent effects, have become popular. For example, Blanchard and Summers argue that the natural rate of unemployment in European countries changes when actual unemployment changes, so that there is no unique level to which unemployment returns.²⁰ If these theories are correct, then nominal rigidities cannot fully explain unemployment, because nominal prices eventually adjust to shocks; some additional explanation, such as the insider-outsider model in Blanchard and Summers, is needed for the persistence of unemployment. But nominal rigidities may be crucial for explaining the initial impulses in unemployment. For example, after rising during the late 1970s, unemployment in Britain has remained high, suggesting hysteresis. But the best explanation for the original increase is arguably a conventional one: slow adjustment of wages and prices to shocks like tight monetary policy and increases in import prices.

The Importance of Externalities from Rigidity. Externalities from nominal rigidity, the central element of menu cost models, are essential

19. Robert J. Barro, "Output Effects of Government Purchases," *Journal of Political Economy*, vol. 89 (December 1981), pp. 1086–1121.

20. Olivier J. Blanchard and Lawrence Summers, "Hysteresis and the European Unemployment Problem," in Stanley Fischer, ed., *NBER Macroeconomics Annual, 1986* (MIT Press, 1986), pp. 15–78.

for a plausible theory of rigidities. If rigidities exist, one of the following statements must be true: rigidities do not impose large costs on the economy; rigidities have large costs to the firms and workers who create them, but these are exceeded by the costs of reducing rigidities; or rigidities have small private costs, and so small frictions are sufficient to create them, but externalities from rigidity impose large costs on the economy. The problem with the first statement is the difficulty of explaining apparently costly events, such as rises in unemployment following monetary contractions, without nominal rigidities. The second seems implausible: it would not be costly for magazine publishers to print new prices every year rather than every four years, as they typically do.²¹ Thus the third statement is the best hope for explaining rigidities.

What Are Menu Costs? Models of nominal rigidity depend on some cost of full flexibility, albeit a small one. The term menu cost may be misleading because the physical costs of printing menus and catalogs may not be the most important barriers to flexibility. Perhaps more important is the lost convenience of fixing prices in nominal terms—the cost of learning to think in real terms and of computing the nominal price changes corresponding to desired real price changes. More generally, we can view infrequent revision of nominal prices as a rule of thumb that is more convenient than continuous revision. Thus, rather than referring to menus, we can state the central argument of recent papers as follows. Firms take the convenient shortcut of infrequently reviewing and changing prices. The resulting profit loss is small, so firms have little incentive to eliminate the shortcut, but externalities make the macroeconomic effects large.

At a somewhat deeper level, we can interpret the convenience of fixing nominal rather than real prices as that of using the medium of exchange, dollars, as a unit of account.²² Alternatively, following Akerlof and Yellen, we can view simple rules of thumb as arising from “near-

21. Stephen G. Cecchetti, “The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines,” *Journal of Econometrics*, vol. 31 (August 1986), pp. 255–74. The cost of reducing nominal wage rigidity may be significant if rigidity is reduced through shorter labor contracts, which require more frequent negotiations between unions and management. But wage rigidity can also be reduced through greater indexation or by having the nominal wage change more often over the life of a contract, neither of which appears to have large costs.

22. Bennett T. McCallum, “On ‘Real’ and ‘Sticky-Price’ Theories of the Business Cycle,” *Journal of Money, Credit, and Banking*, vol. 18 (November 1986), pp. 397–414.

rationality,” a small departure from full optimization.²³ In any case, the precise source of frictions is not important. The effects of nominal shocks are the same whether rigidity arises from printing costs, near-rationality, or something else.

Inflation, the Frequency of Adjustment, and the Phillips Curve

Recent research shows that nominal rigidity is possible in principle—that one can construct a model with firm microeconomic foundations in which rational agents choose substantial rigidity. But the validity of Keynesian theories is not thereby established. For these theories to be convincing, they must have empirical implications that contradict other macroeconomic theories, and these predictions must be confirmed by evidence. This section derives implications of recent Keynesian models, and the next section tests them. As explained in the introduction, the main prediction is that the real effects of nominal shocks are smaller when average inflation is higher. Higher average inflation erodes the frictions that cause nonneutralities, for example by causing more frequent wage and price adjustments.

This section studies a specific model of the class described in the previous section. In the model, a cost of price adjustment leads firms to change prices at intervals rather than continuously. In addition to providing a basis for the empirical tests of the next major section, the model is of theoretical interest. Previous models of nominal rigidity are highly stylized; for example, most menu cost models are static. Our model is dynamic and has the appealing feature that the price level adjusts slowly over time to a nominal shock. The speed of adjustment, which is treated as exogenous in older Keynesian models, is endogenous. It depends on the frequency of price adjustment by individual firms, which in turn is derived from profit-maximization.²⁴

We first present the model and show that high average inflation reduces the output effects of nominal shocks. We also show that highly variable aggregate demand reduces these effects. We then investigate

23. Akerlof and Yellen, “A Near-Rational Model.”

24. The speed of adjustment is also endogenous in Laurence Ball, “Externalities from Contract Length,” *American Economic Review*, vol. 77 (September 1987), pp. 615–29.

the model's quantitative implications by calculating the real effects of shocks for a range of plausible parameter values. The results suggest that the effects of average inflation and demand variability are large. Next we argue that the implications of our model are robust: they carry over to broad classes of other Keynesian models.

Finally, we compare the predictions of Keynesian theories with those of models in the new classical or equilibrium tradition, focusing on Lucas's model of imperfect information. Like our model, Lucas's predicts that the size of the real effects of shocks depends negatively on the variance of aggregate demand. Since this prediction is common to Keynesian and new classical theories, testing it empirically, as Lucas and others have done, is not useful for distinguishing between the two theories. Crucially, Lucas's model differs from ours by predicting that the effects of shocks do *not* depend on average inflation. This difference leads to the tests of the models in the next section.

THE MODEL AND QUALITATIVE RESULTS

Our model of price adjustment is similar in spirit to those of John Taylor and Olivier Blanchard.²⁵ The model is set in continuous time. The economy contains imperfectly competitive firms that change prices at discrete intervals rather than continuously, because adjustments are costly. Price setting is staggered, with an equal proportion of firms changing prices at every instant. The crucial departure from Taylor and Blanchard is that the length of time between price changes, and hence the rate at which the price level adjusts to shocks, is endogenous. Thus we can study the determinants of the speed of adjustment.

Consider the behavior of a representative firm, firm i . Rather than derive a profit function from specific cost and demand functions, we simply assume that firm i 's profits depend on three variables: aggregate spending in the economy, y ; firm i 's relative price, $p_i - p$; and a firm-specific shock, θ_i (all variables are in logs). The aggregate price level p is defined simply as the average of prices across firms. Aggregate spending y affects firm i 's profits by shifting the demand curve that it faces. When aggregate spending rises, the firm sells more at a given

25. Taylor, "Staggered Wage Setting" and "Aggregate Dynamics and Staggered Contracts"; and Blanchard, "Price Asynchronization" and "Wage Price Spiral."

relative price. The term $p_i - p$ affects the firm's profits by determining the position on the demand curve at which it operates. And θ_i is an idiosyncratic shock to either demand or costs (the presence of θ_i is not needed for our main qualitative results, but it strongly affects the quantitative results of the next section).

We assume that the elasticity of firm i 's profit-maximizing real price, $p_i^* - p$, with respect to y is a positive constant, v . Without loss of generality, we assume that the elasticity of $p_i^* - p$ with respect to θ_i is one, and that θ_i has zero mean. Thus we can write the profit-maximizing real price as

$$(5) \quad p_i^*(t) - p(t) = v[y(t) - \bar{y}(t)] + \theta_i(t), \quad v > 0,$$

where \bar{y} is the natural rate of output—the level at which, if θ_i equals its mean, the firm desires a relative price of one. (Relative prices equal to one is the condition for a symmetric equilibrium of the economy when prices are flexible.)²⁶

If price adjustment were costless, firm i would set $p_i = p_i^*$ at every instant. We assume, however, that an adjustment cost leads firms to change prices only at intervals of length λ , which for simplicity is constant over time (later in this section we discuss the implications of allowing λ to vary). Specifically, each price change has a fixed cost F , so adjustment costs per period are F/λ .

As noted above, an equal proportion of firms sets prices at every instant.²⁷ If firm i sets a price at t , it chooses the price and λ to maximize its expected profits, averaged over the life of the price (from t to $t + \lambda$). Maximizing profits is equivalent to minimizing profit losses from two sources: adjustment costs and deviations of price from the profit-

26. As an example of foundations for equation 5, suppose that firm i 's demand equation is $y_i = y - \epsilon(p_i - p)$ (demand depends on aggregate spending and the firm's relative price), and that its log costs are $\gamma y_i + (1 - \epsilon + \epsilon\gamma)\theta_i$. This implies equation 5 with $v = (\gamma - 1)/(1 - \epsilon + \epsilon\gamma)$ and $\bar{y} = [1/(\gamma - 1)] \ln [(\epsilon - 1)/\epsilon\gamma]$ (the coefficient on θ_i in the cost function is chosen to satisfy the normalization that the coefficient on θ_i in equation 5 is one). For deeper microfoundations, see Ball and Romer, "Equilibrium and Optimal Timing of Price Changes," where a price-setting rule like equation 5 is derived from utility and production functions.

27. We assume that price setting is staggered so that inflation is smooth. If all firms changed prices at the same times, the aggregate price level would remain constant between adjustments and then jump discretely. We do not model the sources of staggering explicitly; for explanations of staggering, see the references in note 11.

maximizing level. We approximate the latter by $\frac{1}{2} K (p_i - p_i^*)^2$, where K is the negative of the second derivative of profits with respect to $p_i^* - p_i$. Thus firm i 's loss per unit of time is

$$(6) \quad \frac{F}{\lambda} + \frac{1}{\lambda} \frac{1}{2} K \int_{s=0}^{\lambda} E_t [p_i^* (t + s) - p_i]^2 ds.$$

Minimization of equation 6 implies a simple rule for choosing p_i :

$$(7) \quad p_i = \frac{1}{\lambda} \int_{s=0}^{\lambda} E_t p_i^* (t + s) ds.$$

That is, a firm sets its price to the average of its expected profit-maximizing prices for the period when the price is in effect. We describe the more complicated determination of λ below.

To study the effects of nominal shocks, we must introduce a stochastic nominal variable. We assume that the log of nominal aggregate demand, $x \equiv y + p$, is exogenous and follows the continuous-time analogue of a random walk with drift:

$$(8) \quad x(t) = gt + \sigma_x W(t),$$

where $W(t)$ is a Wiener process. The first term in the expression for $x(t)$ captures trend growth of g per unit time; the second captures random walk innovations with variance σ_x^2 per unit time. Our analysis below focuses on the effects of the parameters g and σ_x on the economy. A monetarist interpretation of equation 8 is that $x(t) = m(t) + V$ —the velocity of money is constant, and aggregate demand is driven by random walk movements in the money stock. A more general interpretation is that a variety of exogenous variables—fiscal policy, the expectations of investors, and so on—drive $x(t)$.

We make two final assumptions. First, the natural rate of output grows smoothly at rate μ :

$$(9) \quad \bar{y}(t) = \mu t.$$

Along with the process for $x(t)$, this implies that average inflation is $g - \mu$. Second, the firm-specific disturbances, the θ_i 's, are uncorrelated across firms and follow continuous-time random walks whose innovations have mean zero and variance σ_θ^2 per unit time.²⁸

28. More precisely, we assume that θ_i follows a stationary process and consider the limit as this process approaches a random walk. (If θ_i is a random walk, its mean is undefined, which contradicts our earlier assumption that its mean is zero.)

The Behavior of the Economy for a Given Frequency of Price Changes. Below we show that average inflation, by influencing the interval between price changes, affects the output-inflation trade-off in our model. A preliminary step is to solve for the behavior of the economy for a given interval, λ . We do this by combining our assumptions about price setting by individual firms and then aggregating. The behavior of individual firms determines the behavior of the price level. As described above, the behavior of price level—which each firm, being small, takes as given—in turn determines the behavior of firms. The condition for equilibrium is that individual and aggregate behavior are consistent; that is, that profit-maximizing price-setting rules for individual firms given the behavior of the price level in fact yield that behavior of the price level. The details are complicated, so we leave them for the appendix. Here we simply present our main results.

The solution for the behavior of the price level takes the form

$$(10) \quad p(t) = (g - \mu)t + \int_{s=0}^{\infty} w(s;\lambda) dZ(t-s),$$

where $dZ(t-s) \equiv \sigma_x dW(t-s)$ is the innovation in aggregate demand at $t-s$. The first term in equation 10 captures average inflation of $g - \mu$, and the second captures the effects of shocks. The term $w(s;\lambda)$ gives the effect of a demand shock at $t-s$ on the price level at t .

The appendix derives the expression that defines $w(\bullet)$. We cannot find an analytic solution to the expression and therefore solve it numerically; the appendix describes how. We find when we solve for $w(\bullet)$ that, assuming $\nu < 1$, $w(s;\lambda)$ equals zero when $s=0$, increases with s , and approaches one as s approaches infinity. That is, the immediate effect of a shock on the price level is zero (because an infinitesimal proportion of firms changes prices at t); the effect of the shock grows over time; and asymptotically the shock is passed one-for-one into prices.

The crucial result about $w(\bullet)$ concerns the frequency of price changes: when $\nu < 1$, $w(s;\lambda)$ is decreasing in λ . A longer interval between changes in individual prices leads to slower adjustment of the aggregate price level—for any s , a smaller proportion of a shock at $t-s$ is passed into prices by t .²⁹

29. If $\nu > 1$, firms want to adjust their prices more than one-for-one with real output; as a result, in this case the approach to full adjustment is oscillatory. The response is again slower when the frequency of price adjustment is lower.

The behavior of real output follows directly from the behavior of the price level, the stochastic process for aggregate demand, and the identity $y = x - p$:

$$(11) \quad y(t) - \bar{y}(t) = \int_{s=0}^{\infty} [1 - w(s; \lambda)] dZ(t-s).$$

The sizes of the real effects of nominal shocks are given by $1 - w(\bullet)$; this is the theoretical counterpart of the parameter that we estimate in the following section.

Finally, equation 11 implies an expression for the variance of output:

$$(12) \quad E\{[y(t) - \bar{y}(t)]^2\} = \sigma_x^2 \int_{s=0}^{\infty} [1 - w(s; \lambda)]^2 ds.$$

The variance of output depends on the variance of demand shocks, σ_x^2 , and the size of the effects of shocks, $1 - w(\bullet)$. This result is also used in the empirical work of the next major section.

The Equilibrium Frequency of Price Changes. We now derive a condition defining the equilibrium interval between price changes. Consider firm i 's problem of choosing its interval, λ_i , given that all other firms in the economy choose an interval λ . The value of a firm's loss function, L (equation 6), is affected by both λ_i and λ ; the latter matters because it determines the behavior of the price level. Minimization of $L(\lambda_i, \lambda)$ with respect to λ_i yields the first-order condition $\partial L(\lambda_i, \lambda)/\partial \lambda_i = 0$. A symmetric Nash equilibrium for λ , λ^E is defined implicitly by setting $\lambda_i = \lambda$ in this condition:

$$(13) \quad \left. \frac{\partial L(\lambda_i, \lambda^E)}{\partial \lambda_i} \right|_{\lambda_i = \lambda^E} = 0.$$

In other words, an interval λ is an equilibrium if, when λ is chosen throughout the economy, it is in firm i 's interest to choose λ as well.³⁰

Because we can find $w(\bullet)$ only numerically, we must also find the

30. Solving for the equilibrium interval between price changes is different from the common exercise of solving for the socially optimal interval. See, for example, Gray, "On Indexation" (which focuses on the interval between wage changes—that is, the length of labor contracts). The equilibrium and optimal intervals differ because, as we stress in the first section, firms' choices of the frequency of price adjustment have externalities. See Ball, "Externalities from Contract Length," for a further discussion of this point.

equilibrium λ numerically, as described in the appendix. We find that λ is decreasing in $\bar{\pi}$, σ_x , and σ_θ , where $\bar{\pi} \equiv g - \mu$ is the average inflation rate. Thus the interval between price changes decreases the higher average inflation. High inflation causes a firm's profit-maximizing nominal price to change rapidly, which raises the benefits from frequent adjustment. The interval λ also decreases the greater the variances of aggregate and firm-specific shocks. When either variance is large, a firm's future profit-maximizing price is highly uncertain, so the firm does not wish to fix its price for long.

These results, along with the results about the effects of λ , imply that the Phillips curve is steeper when $\bar{\pi}$, σ_x , or σ_θ is larger. Higher average inflation reduces the interval between price changes, which in turn raises $w(\bullet)$, the proportion of a shock that is passed into prices. A larger variance of aggregate or firm-specific shocks also reduces λ and thus raises $w(\bullet)$. These results imply that increases in $\bar{\pi}$, σ_x , or σ_θ lead to *decreases* in $1 - w(\bullet)$, the real effects of shocks. These predictions lead to the empirical tests of the next section.

QUANTITATIVE RESULTS

We now ask whether the effects of inflation and demand variability identified above are quantitatively important. We do so by computing the interval between price changes and the real effects of shocks for a range of plausible parameter values.

Choice of Parameters. Since our focus is the effects of average inflation, $g - \mu$, and the standard deviation of demand, σ_x , we experiment with wide ranges of values of these parameters (g and μ affect the results only through their difference). This leaves three other parameters for which we need baseline values: F/K , the ratio of the cost of changing prices to the negative of the second derivative of the profit function (F and K enter only through their ratio); σ_θ , the standard deviation of firm-specific shocks; and ν , the elasticity of a firm's profit-maximizing real price with respect to aggregate output.

We choose baseline parameters by experimenting with values of F/K , σ_θ , and ν to find a combination that implies plausible sizes for the real effects of shocks. We then ask whether these parameter values are realistic. Finally, we investigate robustness by calculating the effects of changing each parameter from its baseline value.

It is difficult to measure F and K directly, so we take an indirect approach. In a model of steady inflation and no shocks ($\sigma_\theta = \sigma_x = 0$), F/K determines the frequency of price changes. With $F/K = 0.00015$, firms change prices every five quarters under steady 3 percent inflation and every two quarters under steady 12 percent inflation. Microeconomic evidence suggests that actual intervals between price changes typically average two years or more. Thus our baseline value of F/K is *conservative*. We certainly do not assume menu costs that are too large to be consistent with price setting in actual economies.³¹

To pick a value for σ_θ , we use the fact that σ_θ equals the standard deviation of movements in profit-maximizing prices, the p_i^* 's, across firms. This leads us to use data on relative price variability to gauge plausible values of σ_θ . Vining and Elwertowski report a 4–5 percent standard deviation of annual relative price movements across highly disaggregate (8-digit) components of the U.S. consumer price index; this is consistent with our assumption of $\sigma_\theta = 3$ percent.³²

Finally, there is little quantitative evidence concerning the size of v , the elasticity of profit-maximizing relative prices with respect to aggregate output. However, our choice of a small elasticity, 0.1, is consistent with the common view that relative prices vary little in response to aggregate fluctuations. Our baseline parameters are therefore $F/K = 0.00015$, $\sigma_\theta = 3$ percent, and $v = 0.1$.

Results. Table 1 shows the effects of average inflation, $g - \mu$, and the variability of demand, σ_x , when F/K , σ_θ , and v equal their baseline values. For wide ranges of $g - \mu$ and σ_x , the table shows two figures.

31. For microeconomic evidence on price behavior, see Cecchetti, "The Frequency of Price Adjustment"; Anil K. Kashyap, "Sticky Prices: New Evidence from Retail Catalogs" (MIT, November 1987); and W. A. H. Godley and C. Gillion, "Pricing Behavior in Manufacturing Industry," *National Institute Economic Review*, no. 33 (August 1965), pp. 43–47.

32. Daniel R. Vining, Jr., and Thomas C. Elwertowski, "The Relationship between Relative Prices and the General Price Level," *American Economic Review*, vol. 66 (September 1976), pp. 699–708. As a measure of σ_θ , Vining and Elwertowski's figure has both an upward and a downward bias. The upward bias occurs because staggered price adjustment causes actual prices, the p_i 's, to vary across firms even when profit-maximizing prices, the p_i^* 's, do not. As a result, the standard deviation of p_i , which Vining and Elwertowski measure, is greater than the standard deviation of p_i^* , which equals σ_θ . The negative bias in that variation across components of the CPI, even if these are highly disaggregated, is less than variation across individual prices. It is difficult to tell the relative magnitudes of these biases.

Table 1. Effect of a Nominal Shock on Real Output and the Equilibrium Interval between Price Changes^a

Average inflation rate, $g - \mu$ (percent)	Demand variability, σ_x (percent)					
	0	1	3	5	10	20
0	0.560 (31)	0.540 (30)	0.495 (28)	0.424 (25)	0.250 (17)	0.078 (10)
2	0.519 (29)	0.519 (29)	0.470 (27)	0.405 (24)	0.250 (17)	0.078 (10)
5	0.424 (25)	0.424 (25)	0.405 (24)	0.366 (22)	0.224 (16)	0.078 (10)
10	0.322 (20)	0.322 (20)	0.299 (19)	0.299 (19)	0.199 (15)	0.078 (10)
20	0.174 (14)	0.174 (14)	0.174 (14)	0.174 (14)	0.124 (12)	0.057 (9)
50	0.057 (9)	0.057 (9)	0.057 (9)	0.057 (9)	0.039 (8)	0.023 (7)
100	0.012 (6)	0.012 (6)	0.012 (6)	0.012 (6)	0.012 (6)	0.012 (6)
250	0.001 (4)	0.001 (4)	0.001 (4)	0.001 (4)	0.001 (4)	0.001 (4)

Source: Authors' calculations. See text description.

a. The table shows the effects of changing values of $g - \mu$ and σ_x when F/K , σ_θ , and v equal their baseline values. F/K is the ratio of the cost of changing prices to minus the second derivative of the profit function; σ_θ is the standard deviation of firm-specific shocks; and v is the elasticity of a firm's profit-maximizing real price with respect to aggregate output. Baseline values: $F/K = 0.00015$; $\sigma_\theta = 3$ percent; and $v = 0.1$. For each entry in the table, the first number is the percentage effect of a 1 percent nominal shock on real output after six months; the number in parentheses is the equilibrium interval between price changes, λ , in weeks.

The first is the percentage effect of a 1 percent change in demand on real output six months later. A value of zero would mean that prices adjust fully to the shock; a value of one would mean that prices do not adjust at all. We refer to this figure as simply the real effect of a shock. The figure in parentheses is the equilibrium interval between price changes, λ , in weeks. As we explain above, inflation and demand variability influence the real effects of shocks through their effects on λ .

Table 1 shows that realistic increases in average inflation have quantitatively important effects. With $\sigma_\theta = 3$ percent, roughly the standard deviation of nominal GNP growth for the postwar United States, the interval between price changes is 28 weeks if $g - \mu = 0$, but falls to 19 weeks if $g - \mu = 10$ percent and 6 weeks if $g - \mu = 100$ percent. As a result, the real effect of a shock is 0.50 for $g - \mu = 0$, 0.30 for $g - \mu = 10$ percent, and 0.01 for $g - \mu = 100$ percent.

Table 2. The Effects of Changes in F/K , ν , and σ_θ on the Slope of the Phillips Curve and the Equilibrium Interval between Price Changes

Parameter values			Average inflation rate, ^a $g - \mu$	
F/K	ν	σ_θ (percent)	5 percent	20 percent
0.00015	0.1	3	0.405 (24)	0.174 (14)
0.0003	0.1	3	0.519 (29)	0.274 (18)
0.00015	0.2	3	0.273 (24)	0.080 (14)
0.00015	0.1	6	0.174 (14)	0.100 (11)
0.00045	0.1	6	0.405 (24)	0.274 (18)
0.0003	0.285	3	0.406 (32)	0.104 (18)
0.00015	0.0295	6	0.403 (15)	0.367 (12)

Source: Authors' calculations. See text description.

a. For various combinations of F/K , ν , and σ_θ , the table shows the real effect of a nominal shock and the interval between price changes: for each entry in the last two columns of the table, the first number is the percentage effect of a 1 percent nominal shock on real output after six months; the number in parentheses is the equilibrium interval between price changes, λ , in weeks. Demand variability, σ_x , is set to 3 percent.

Table 1 shows that increases in σ_x also have important effects. With average inflation of 5 percent, raising σ_x from 3 percent to 10 percent reduces the interval between price changes from 24 to 16 weeks and the real effect of a shock from 0.41 to 0.22. These effects are similar to the effects of raising average inflation from 5 percent to 15 percent.

Table 2 shows the effects of varying F/K , σ_θ , and ν . For various combinations of these parameters, we show the real effect of a shock and the interval between price changes for $g - \mu = 5$ percent and $g - \mu = 20$ percent, assuming $\sigma_x = 3$ percent in both cases. The first line reproduces the results for the baseline parameters, and each of the following three lines shows the effects of doubling one parameter while holding the others constant. An increase in F/K raises the real effect of a shock, and an increase in σ_θ or ν reduces it. But for all combinations of F/K , σ_θ , and ν , our central result holds: higher average inflation reduces the real effect of a shock. The remaining three lines of the table show the effects of combinations of changes that leave the real effect of a shock

unchanged for $g - \mu = 5$ percent. These lines show how the parameters affect the strength of the link between average inflation and the real effect of a shock.

ROBUSTNESS

Traditional Keynesian models, such as textbook models of price adjustment or the staggered contracts models of Fischer and Taylor, do not share the key predictions of our model.³³ These older theories treat the degree of nominal rigidity (for example, the length of labor contracts or the adjustment speed of the price level) as fixed parameters; thus they rule out the channel through which average inflation affects the output-inflation trade-off. On the other hand, our central results appear to be robust implications of Keynesian theories in which the degree of rigidity is endogenous. The intuition for the effects of inflation on the frequency of price adjustment, and of this frequency on the size of nonneutralities, is not tied to the specific assumptions of our model.

One assumption of our model that requires attention is that the interval between price changes is constant over time. This assumption is *ad hoc*: given our other assumptions, firms could increase profits by varying the interval based on the realizations of shocks. In addition, the assumption is unrealistic, because firms in actual economies do not always change prices at fixed intervals.

We now consider the alternative assumption that firms can freely vary the timing of price changes. This assumption of complete flexibility is also far from realistic. Most wages are adjusted at constant intervals of a year. There appears to be greater flexibility in the timing of price changes, but the limited evidence suggests that it is not complete. Mail order companies change prices at fixed times during the year, even though they issue catalogs much more frequently than they change prices, and thus could vary the dates of adjustments without issuing extra catalogs. In addition, a broad range of industries appears to have a preferred time of the year, often January, for price changes.³⁴

It is not yet possible to solve a model like ours with flexible timing,

33. For a textbook model, see Rudiger Dornbusch and Stanley Fischer, *Macroeconomics*, 4th ed. (McGraw-Hill, 1987).

34. For evidence on mail order catalogs, see Kashyap, "Sticky Prices." For evidence on industries' preferred months for price changes, see Julio J. Rotemberg and Garth Saloner, "A 'January Effect' in the Pricing of Goods" (MIT, 1988).

but suggestive results are available for simpler models. In particular, a literature beginning with Sheshinski and Weiss presents partial-equilibrium models in which a firm chooses to follow an “*Ss*” rule for adjusting its price: whenever inflation pushes its real price outside some bounds, it adjusts its nominal price to return the real price to a target level. These models reproduce a crucial implication of our model: higher average inflation leads to more frequent price changes. High inflation causes a firm’s real price to change rapidly, so, for given *Ss* bounds, the price hits the bounds more often. High inflation also causes the firm to widen its bounds, which reduces the frequency of price changes, but does not fully offset the first effect.³⁵

For our main argument to hold, the more frequent changes in individual prices that result from higher inflation must lead in turn to faster adjustment of the aggregate price level. Intuition clearly suggests a link between the frequency of individual adjustment and the speed of aggregate adjustment, but the difficulty of studying general equilibrium with flexible timing precludes a definitive proof. Indeed, in one prominent special case, the link does not exist. Andrew Caplin and Daniel Spulber show that if we *assume* that firms follow *Ss* rules with constant bounds, and if aggregate demand is nondecreasing, then the aggregate price level adjusts immediately to nominal shocks—nominal shocks are neutral. Because aggregate adjustment is always instantaneous, its speed is obviously independent of the frequency of individual price changes.³⁶

Current research suggests that the Caplin-Spulber result does not hold under realistic conditions. There exist examples in which firms do *not* follow *Ss* rules with constant bounds, and so a shock to the money supply is not neutral, either if there is some persistence to inflation or if firms’ optimal nominal prices sometimes fall. And when nonneutralities exist, it appears plausible that their size depends on the frequency of individual price adjustment. Thus, overall, models of price adjustment with flexible timing appear consistent with the predictions of our model.³⁷

35. The link between inflation and the frequency of adjustment is established for the case of constant inflation in Eytan Sheshinski and Yoram Weiss, “Inflation and Costs of Price Adjustment,” *Review of Economic Studies*, vol. 44 (June 1977), pp. 287–303. An extension to the case of stochastic inflation is presented by Andrew S. Caplin and Eytan Sheshinski, “Optimality of (*s*, *S*) Pricing Policies” (Princeton University, 1987).

36. Andrew S. Caplin and Daniel F. Spulber, “Menu Costs and the Neutrality of Money,” *Quarterly Journal of Economics*, vol. 102 (November 1987), pp. 703–25.

37. For the implications of persistent inflation, see Daniel Tsiddon, “On the Stubbornness of Sticky Prices” (Columbia University, July 1987). For the implications of falling

Another robustness issue concerns the nature of the friction that prevents nominal flexibility. In our model, the friction is a fixed cost of price adjustment. An alternative view is that the technological costs of making prices highly flexible are negligible but that for some reason, such as convenience, the desire to avoid computation costs, or habit, price setters nonetheless follow rules that focus on nominal prices.³⁸ Without a theory that predicts the particular rules of thumb that price setters follow, theories of this type do not make precise predictions concerning the relationship between average inflation and the degree of price flexibility. But it appears that under reasonable interpretations these theories imply that higher inflation increases nominal flexibility. As average inflation rises, so does the cost of following a rule-of-thumb pricing policy stated in nominal terms, as does the evidence that keeping a fixed nominal price is not equivalent to keeping a fixed real price. Although price setters may continue to follow rules of thumb, they will increasingly think in terms of real rather than nominal magnitudes. Nominal price flexibility will thus increase.

THE PREDICTIONS OF NEW CLASSICAL THEORIES

The prediction of Keynesian models that average inflation affects the output-inflation trade-off is important because it is inconsistent with alternative macroeconomic models in the new classical tradition. We now review the predictions of new classical models, focusing on Lucas's imperfect information theory. Like Keynesian models of nominal rigidity, Lucas's model is designed to explain the effects of nominal shocks on output—that is, to generate a short-run Phillips curve. But Lucas's model has different implications about what determines the size of the effects.

In Lucas's model, agents wish to change their output in response to changes in their relative prices, but not in response to changes in the aggregate price level. When an agent observes a change in his price, however, he cannot tell whether it results from a relative or an aggregate

optimal prices, see Blanchard, "Why Does Money Affect Output?" and Tsiddon, "The (Mis)behavior of the Aggregate Price Level" (Columbia University, 1987). These authors establish results for the special case in which a firm's optimal price moves one-for-one with aggregate demand and is independent of the price level ($v = 1$ in our notation).

38. See Akerlof and Yellen, "A Near-Rational Model."

movement. He acts upon his best guess, which is that part of the change comes from each source. Since agents interpret any price change as partly relative, changes that in fact result from a nominal shock have effects on output.

In Lucas's model, the size of the effects of nominal shocks depends on the relative magnitudes of nominal and idiosyncratic real shocks. In particular, if nominal shocks are large, agents attribute most of the movements in their prices to nominal shocks, and respond little. Thus a large variance of nominal aggregate demand leads to a steep Phillips curve. Lucas presents cross-country evidence supporting this prediction in his famous 1973 paper. We show, however, that Keynesian models make the same prediction, although the reason—a large variance of aggregate demand causes more frequent price changes—is very different. Because both Keynesian and new classical theories explain Lucas's results, his test does not help to distinguish between them.

The effect of average inflation on the output-inflation trade-off *does* distinguish Keynesian and new classical models. Theories of nominal rigidities predict that high inflation makes the Phillips curve steeper. In Lucas's imperfect information model, average inflation is irrelevant to the output-inflation trade-off, because only the variances of random variables, not the means, affect the uncertainty that agents face. This difference between the theories is the basis for our empirical work. (A *simple* correlation between average inflation and the slope of the Phillips curve is consistent with Lucas's model, because average inflation is correlated with the variance of demand, which affects the slope. The issue is whether there is a relation between average inflation and the slope after we control for the variance of demand.)

Another difference between the predictions of Keynesian and new classical theories concerns the effects of idiosyncratic shocks. According to Lucas, a large variance of relative price shocks increases the real effects of nominal shocks, because it raises the proportion of these shocks that agents misperceive as real. Our model predicts that a large variance of idiosyncratic shocks, like a large variance of aggregate shocks, leads to more frequent price changes and thus *reduces* the effects of nominal shocks. If one could construct a measure of the variance of firm-specific shocks, which we do not attempt in this paper, then estimating the relation between this variable and the slope of the Phillips curve would be another test between the two competing theories.

The leading new classical alternative to Lucas's imperfect information

model is real business-cycle theory.³⁹ This theory attributes all fluctuations in output to real disturbances and assumes that nominal disturbances are simply passed into prices. Because nominal shocks have no causal role in output fluctuations, it is difficult for the theory to explain the observed positive correlations of real and nominal variables, much less the effect of average inflation on the strength of these correlations. King and Plosser have devised a real business-cycle model in which output moves with nominal money through reverse causality: the banking system creates inside money in anticipation of output movements. But as Mankiw points out, the model predicts that the aggregate price level *falls* when output rises.⁴⁰ Thus real business-cycle models do not appear to provide an alternative explanation of the results that we report below.

International Evidence

We examine here how the trade-off between output and inflation varies across countries. Our goal is to test the theoretical results discussed in the previous section. In particular, we wish to examine whether in countries with high rates of inflation, changes in aggregate demand have relatively small effects on output and instead are reflected quickly in prices.

Our analysis is divided into two parts. First we describe the data and present the basic results. We estimate the output-inflation trade-off for 43 industrialized countries and examine the relationship between the trade-off and average inflation and demand variability. Then we consider econometric issues raised by our procedure and examine variations on our basic test.

DATA AND BASIC RESULTS

The data we examine, originally from *International Financial Statistics* of the International Monetary Fund, are from the IMF databank of

39. For a recent real business-cycle model, see Edward C. Prescott, "Theory Ahead of Business Cycle Measurement," *Federal Reserve Bank of Minneapolis Quarterly Review*, vol. 10 (Fall 1986), pp. 9–22.

40. Robert G. King and Charles I. Plosser, "Money, Credit, and Prices in a Real Business Cycle," *American Economic Review*, vol. 74 (June 1984), pp. 363–80; N. Gregory Mankiw, "Real Business Cycles: A Neo-Keynesian Perspective," *Journal of Economic Perspectives*, forthcoming.

Data Resources, Inc. All the data are annual. Depending on the country, output is real GNP or real GDP, whichever is available. We denote the log of output as y and the log of the corresponding nominal quantity as x . The log of the price level is then $p = x - y$.

We wanted the most extensive possible sample of large, industrialized, free market economies. We used the following five criteria for choosing the sample of countries: the population had to be at least one million; at least 10 percent of output had to be in manufacturing; not more than 30 percent of output could be in agriculture; data had to be available at least back to 1963; the economy had to be largely unplanned. Information on the first three criteria was taken from the IMF's *Yearbook of National Account Statistics* and the *International Financial Statistics Yearbooks*; data for the year 1965 were used for these criteria. The fifth criterion is obviously open to interpretation. It led us to exclude such countries as Czechoslovakia, East Germany, and Yugoslavia.

The countries are listed in table 3, together with the period of time for which data are available. We present here some sample statistics for each country: the mean and standard deviation for real growth, inflation, and the growth in nominal demand. We see from this table that there is substantial variation in the macroeconomic experiences of these countries. For example, Panama had the lowest average inflation rate, less than 3 percent a year, while Argentina and Brazil each had average inflation exceeding 40 percent a year.

Estimating the Output-Inflation Trade-off. We express the short-run output-inflation trade-off by estimating the following equation:

$$(14) \quad y_t = \text{constant} + \tau \Delta x_t + \lambda y_{t-1} + \gamma \text{Time}.$$

The log of real GNP is regressed on its own lag, a time trend, and the change in nominal GNP. This sort of equation has been used widely, both by new classical macroeconomists such as Robert Lucas and by Keynesian macroeconomists such as Charles Schultze.⁴¹ Equation 14 is the empirical counterpart of equation 12 of our theoretical model. It differs from equation 12 by the use of discrete rather than continuous time and by summarizing the effects of past demand movements through

41. Lucas, "Some International Evidence"; Charles L. Schultze, "Cross-Country and Cross-Temporal Differences in Inflation Responsiveness," *American Economic Review*, vol. 74 (May 1984, *Papers and Proceedings*, 1983), pp. 160-65.

Table 3. Descriptive Statistics on Inflation and Output, Various Countries, Selected Periods, 1948–86

Country	Sample period	Real growth		Inflation		Nominal growth	
		Mean	Standard deviation	Mean	Standard deviation	Mean	Standard deviation
Argentina	1963–81	0.0262	0.04253	0.5439	0.42064	0.5702	0.40685
Australia	1949–85	0.0416	0.02446	0.0677	0.05043	0.1094	0.04880
Austria	1950–86	0.0396	0.02615	0.0526	0.04745	0.0923	0.04329
Belgium	1950–85	0.0329	0.02238	0.0424	0.03028	0.0754	0.03343
Bolivia	1958–83	0.0376	0.03839	0.2012	0.29272	0.2388	0.26608
Brazil	1963–84	0.0633	0.05500	0.4237	0.25825	0.4871	0.23597
Canada	1948–85	0.0436	0.02605	0.0480	0.03296	0.0917	0.03696
Colombia	1950–85	0.0465	0.01881	0.1371	0.06849	0.1836	0.06394
Costa Rica	1960–86	0.0459	0.03809	0.1241	0.13105	0.1701	0.10563
Denmark	1950–85	0.0318	0.02451	0.0634	0.02869	0.0953	0.02812
Dominican Republic	1950–86	0.0493	0.05084	0.0534	0.07553	0.1028	0.09106
Ecuador	1950–85	0.0569	0.04195	0.0877	0.10082	0.1447	0.09687
El Salvador	1951–86	0.0328	0.04140	0.0507	0.08044	0.0837	0.07675
Finland	1950–85	0.0429	0.03339	0.0749	0.05648	0.1178	0.06088
France	1950–85	0.0415	0.01987	0.0675	0.03798	0.1091	0.03118
Germany	1950–86	0.0455	0.03548	0.0374	0.02125	0.0830	0.03705
Greece	1948–86	0.0549	0.03872	0.0915	0.06568	0.1465	0.05949
Guatemala	1950–83	0.0425	0.02802	0.0374	0.05120	0.0799	0.05812
Iceland	1948–85	0.0353	0.05647	0.1977	0.14758	0.2331	0.13415
Iran	1959–85	0.0575	0.08270	0.0937	0.12213	0.1514	0.13114
Ireland	1948–85	0.0319	0.02375	0.0741	0.05432	0.1061	0.06096
Israel	1953–82	0.0745	0.04307	0.2119	0.24465	0.2865	0.22193
Italy	1950–85	0.0433	0.02775	0.0796	0.05715	0.1229	0.04683
Jamaica	1960–85	0.0153	0.04616	0.1113	0.08769	0.1267	0.06723
Japan	1952–85	0.0716	0.03642	0.0472	0.03772	0.1189	0.04482
Mexico	1948–85	0.0577	0.03069	0.1392	0.15016	0.1969	0.13479
Netherlands	1950–85	0.0388	0.02918	0.0493	0.03399	0.0882	0.03779
Nicaragua	1960–83	0.0377	0.08672	0.0806	0.09377	0.1184	0.09897
Norway	1950–86	0.0416	0.01618	0.0564	0.05166	0.0982	0.04576
Panama	1950–86	0.0537	0.03315	0.0299	0.03548	0.0837	0.04613
Peru	1960–84	0.0354	0.04342	0.2554	0.22662	0.2909	0.20236
Philippines	1948–86	0.0484	0.03216	0.0720	0.08221	0.1205	0.06948
Portugal	1953–82	0.0510	0.02771	0.0739	0.07211	0.1250	0.06844
Singapore	1960–84	0.0861	0.04680	0.0356	0.04768	0.1218	0.05245
South Africa	1948–86	0.0383	0.02279	0.0718	0.05457	0.1102	0.05012
Spain	1954–84	0.0455	0.03085	0.0992	0.04698	0.1448	0.04393
Sweden	1950–86	0.0301	0.01795	0.0635	0.03776	0.0937	0.03169
Switzerland	1948–86	0.0286	0.03415	0.0387	0.02672	0.0674	0.03217
Tunisia	1960–83	0.0621	0.04424	0.0596	0.05795	0.1217	0.05419
United Kingdom	1948–86	0.0243	0.01890	0.0664	0.04984	0.0909	0.04202
United States	1948–86	0.0315	0.02676	0.0415	0.02529	0.0731	0.03252
Venezuela	1950–85	0.0459	0.03757	0.0526	0.08237	0.0985	0.07665
Zaire	1950–84	0.0334	0.04248	0.2002	0.22374	0.2338	0.21167
Across-country values							
Mean		0.0441	0.03591	0.1048	0.09303	0.1489	0.08881
Standard deviation		0.0138	0.0147	0.0999	0.0843	0.1003	0.0776

Source: Authors' calculations with data from International Monetary Fund, *International Financial Statistics*. The data were obtained from the IMF data bank of Data Resources, Inc. All data are annual. Depending on the country, output is real GNP or real GDP, whichever is available. Growth rates are computed as differences in logarithms with the log of real output as y and the log of nominal output as x : the log of the price level is $p = x - y$. For information on the selection of countries in the sample, see the text description.

Table 4. Estimates of the Output-Inflation Trade-off, Various Countries, Selected Periods, 1948–86^a

Country	Sample period	Full sample		Data through 1972		Data after 1972	
		Trade-off parameter, τ	Standard error	Trade-off parameter, τ	Standard error	Trade-off parameter, τ	Standard error
Argentina	1963–81	-0.0047	0.0335	-0.1179	0.1140	0.0021	0.0322
Australia	1949–85	0.1383	0.0862	0.3029	0.0858	0.3196	0.1937
Austria	1950–86	-0.0196	0.1069	-0.0830	0.1219	0.6823	0.2058
Belgium	1950–85	0.4967	0.1035	0.3897	0.1036	0.2081	0.2950
Bolivia	1958–83	-0.0525	0.0424	0.1418	0.1567	-0.0621	0.0276
Brazil	1963–84	-0.0951	0.1037	-0.1999	0.2111	0.0770	0.1478
Canada	1948–85	0.4731	0.0899	0.5052	0.1151	0.4619	0.2333
Colombia	1950–85	0.0550	0.0879	-0.0233	0.0919	0.2089	0.2151
Costa Rica	1960–86	-0.2302	0.0911	0.4041	0.1937	-0.2912	0.0722
Denmark	1950–85	0.8486	0.1385	0.6762	0.1491	1.0091	0.5805
Dominican Republic	1950–86	0.3993	0.0750	0.5689	0.0591	-0.1173	0.0733
Ecuador	1950–85	0.1976	0.1148	0.4903	0.2042	-0.2062	0.1226
El Salvador	1951–86	0.3432	0.0822	0.4368	0.1162	0.3230	0.1127
Finland	1950–85	0.2417	0.0823	0.2242	0.1000	0.5835	0.1139
France	1950–85	-0.0648	0.0899	0.1046	0.0765	0.3858	0.3129
Germany	1950–86	0.6137	0.1005	0.6182	0.1369	1.0761	0.1333
Greece	1948–86	0.2577	0.0974	0.4528	0.0644	0.4583	0.2869
Guatemala	1950–83	0.3966	0.0772	0.3705	0.1036	0.5021	0.1377
Iceland	1948–85	0.0154	0.1173	0.3892	0.1978	-0.2487	0.2236
Iran	1959–85	0.3785	0.1097	0.1081	0.0834	0.5018	0.2084
Ireland	1948–85	0.2731	0.0710	0.3767	0.1074	0.1306	0.1733
Israel	1953–82	0.0015	0.0847	0.4082	0.0928	0.0901	0.0372
Italy	1950–85	0.2035	0.1007	0.5279	0.1363	0.5470	0.1276
Jamaica	1960–85	0.1399	0.1591	-0.0977	0.3055	0.2389	0.1169
Japan	1952–85	0.5065	0.1524	0.4812	0.1363	0.4119	0.2441
Mexico	1948–85	-0.1095	0.0530	0.3139	0.0491	-0.4304	0.0997
Netherlands	1950–85	0.4546	0.1244	0.3245	0.1802	0.5214	0.3035
Nicaragua	1960–83	0.5834	0.1551	0.8431	0.1833	0.6332	0.4505
Norway	1950–86	-0.0448	0.0625	-0.0875	0.0719	0.0402	0.1748
Panama	1950–86	0.5969	0.0858	0.5775	0.0900	0.6592	0.0811
Peru	1960–84	-0.0713	0.1171	0.1116	0.1696	0.0419	0.2337
Philippines	1948–86	0.0424	0.0762	0.2202	0.1069	-0.2266	0.0721
Portugal	1953–82	0.1769	0.1692	0.3533	0.2599	0.3291	0.2994
Singapore	1960–84	0.6022	0.1369	1.0316	0.3661	0.3166	0.0477
South Africa	1948–86	0.2017	0.0763	0.2615	0.0914	0.3203	0.1317
Spain	1954–84	0.3507	0.1255	0.5020	0.0945	0.3289	0.0699
Sweden	1950–86	0.0067	0.0971	0.1648	0.1015	0.4184	0.1732
Switzerland	1948–86	0.8264	0.1137	0.7475	0.1254	0.7940	0.1693
Tunisia	1960–83	0.5251	0.1703	0.7856	0.2896	0.1342	0.1536
United Kingdom	1948–86	-0.0199	0.0958	0.0793	0.1293	-0.0766	0.2197
United States	1948–86	0.6714	0.0771	0.7229	0.0598	0.8486	0.1915
Venezuela	1950–85	0.1146	0.0623	0.3252	0.1239	-0.0240	0.0784
Zaire	1950–84	0.0160	0.0414	-0.0188	0.0419	-0.0502	0.0984
Across-country values							
Mean		0.2419	0.0985	0.3422	0.1348	0.2761	0.1739
Standard deviation		0.2719	0.0326	0.2796	0.0695	0.3463	0.1089

Source: Authors' estimates using equation 14. The data used in the estimation are from IMF, *International Financial Statistics*.

a. The dependent variable is the log of real output, y_t . The output-inflation trade-off parameter, τ , is the coefficient of the change in nominal demand, expressed as differences in logarithms.

the term in lagged real output. We discuss this specification of the output-inflation trade-off further in the second part of this section.

The coefficient of the change in nominal demand, τ , is the parameter of central interest. It tells us how much of a shock to nominal GNP shows up in output in the first year. If $\tau = 1$, then all of the change in nominal GNP shows up in real GNP; if $\tau = 0$, then all the change in nominal GNP shows up in prices.

Table 4 presents the estimated value of τ for the 43 countries, together with the estimated standard errors. For each country, the entire available time series is used in the estimation. Table 4 also presents the estimated value of τ for two subsamples. We use 1972–73 as the cutoff between the two subsamples. The early 1970s are often considered a time of major structural change; certainly many empirical macroeconomic relationships broke down. We therefore wanted to see whether the trade-off parameter τ changed and, if so, whether the changes could be explained.


Table 4 shows substantial variation in the output-inflation trade-off across countries. The mean value of τ for our 43 countries is 0.242 and the standard deviation is 0.272. The trade-off parameter for the United States is 0.671, which is 1.6 standard deviations above the mean. Hence, relative to the typical country in our sample, the United States exhibits large effects of aggregate demand on output.

Table 4 shows that the trade-off parameter sometimes changes substantially from the period through 1972 to the period after 1972. For the United States, there is little change in the estimate. But in 63 percent of the countries, one can reject the hypothesis of no change in τ at the 5 percent level. Across countries, the correlation between τ estimated with the earlier data and the τ estimated with the later data is 0.36. It appears that there can be substantial change in the output-inflation trade-off over time.


The Determinants of the Trade-off: Cross-Section Results. We now wish to see whether the cross-country variation in the estimated trade-off τ can be explained. Our theoretical model suggests that τ should be low in countries where the variability of aggregate demand is high and in countries where the average level of inflation is high. Our primary attention is on these two hypotheses.

Figures 1 and 2 present scatterplots of the trade-off parameter τ against the mean level of inflation $\bar{\pi}$, the log of the mean level of inflation, and the standard deviation of the change in aggregate demand σ_{Δ} . Both pictures display the negative relation predicted by theory.

Figure 1. The Output-Inflation Trade-off and Mean Inflation




Source: The output-inflation trade-off parameter, τ , is from table 4. Mean inflation is from table 3.

Figure 2. The Output-Inflation Trade-off and the Variability of Demand

Source: Output-inflation trade-off parameter, τ , is from table 4. Standard deviation of nominal GNP growth is from table 3.

Figure 1 also suggests that the relation between the trade-off parameter and mean inflation is nonlinear. This result should not be surprising. As the rate of inflation grows larger, the trade-off parameter should decline. But we do not expect τ to decline below zero. The relation between $\bar{\pi}$ and τ should be convex. An increase in inflation from 5 percent to 10 percent should have a larger effect on τ than an increase from 10 percent to 15 percent. When we turn to formal estimation, therefore, a linear specification is likely to be inadequate.

Because our sample includes a few countries with extremely high inflation rates, it is difficult to gauge the relationship between the estimated trade-off and mean inflation among low- and moderate-inflation countries from the top portion of figure 1. The bottom portion therefore presents a scatterplot of τ against the log of mean inflation. That portion shows that the inverse relation between τ and $\bar{\pi}$ holds at both low and high inflation rates.

Figure 3. Mean Inflation and the Variability of Demand*Mean inflation*

Source: Table 3.

Figure 3, a scatterplot of mean inflation and the standard deviation of nominal GNP growth, shows a strong positive relation. This figure thus reestablishes the well-known fact that countries with high levels of inflation tend to have unstable aggregate demand. The correlation between these two variables is 0.92. We will see below, however, that multiple regression is able to identify separate effects of these two variables on the output-inflation trade-off.

Table 5 presents cross-sectional regressions of the estimated values of the trade-off parameter τ on the mean of inflation $\bar{\pi}$ and the standard deviation of aggregate demand growth σ_x . To account for the nonlinearity, the squares of these variables are also included in some regressions.⁴² The last column is the most general specification; it includes both variables and allows both to enter nonlinearly. As expected, the second-order term in mean inflation is statistically significant.⁴³

42. Including an interaction term does not affect the results.

43. Our measures of the short-run output-inflation trade-offs for the countries in our

Table 5. Determinants of the Output-Inflation Trade-off, Full Sample Period^a

Independent variable	Equation					
	5.1	5.2	5.3	5.4	5.5	5.6
Constant	0.384 (0.053)	0.388 (0.057)	0.389 (0.057)	0.600 (0.079)	0.516 (0.089)	0.589 (0.086)
Mean inflation	-1.347 (0.368)	...	-1.119 (0.919)	-4.835 (1.074)	...	-5.729 (1.973)
Square of mean inflation	7.118 (2.088)	...	8.406 (3.849)
Standard deviation of nominal GNP growth	...	-1.639 (0.482)	-0.322 (1.183)	...	-4.242 (1.512)	1.241 (2.467)
Square of standard deviation of nominal GNP growth	7.455 (4.118)	-2.380 (7.062)
<i>Summary statistic</i>						
R^2	0.228	0.201	0.210	0.388	0.243	0.359
Standard error	0.241	0.245	0.244	0.215	0.239	0.219

Source: Authors' calculations.

a. The dependent variable is the output-inflation trade-off parameter, τ (estimated in table 4). Numbers in parentheses are standard errors.

The estimates in table 5 suggest that mean inflation is a statistically significant determinant of the inflation-output trade-off, but that demand variability is not. The hypothesis that inflation and inflation squared do not enter regression 5.6 is rejected at the 5 percent level. The hypothesis that the standard deviation of aggregate demand and its square do not enter is not rejected even at the 20 percent level. An examination of the substantive implications of regression 5.6 also shows that only mean inflation is important. For example, an increase in mean inflation from 5 percent to 10 percent, as might be plausible for the United States, would reduce the trade-off by 0.22. An increase in σ_x from 5 percent to 10 percent increases the trade-off by 0.04. Hence, only the effects of inflation on the trade-off are substantively important.

sample are estimates. Because τ is the dependent variable in our cross-section regressions, this measurement error does not cause bias. But because the measurement errors are of different sizes, they cause heteroskedasticity in the cross-section regression. Using the estimated variances of the errors (from the standard errors of the estimated trade-offs), we can correct for the heteroskedasticity and therefore obtain more efficient estimates. Our estimates imply, however, that less than a quarter of the average variance of the residuals in the cross-section regression is due to the errors in estimating the τ 's. As a result, accounting for the heteroskedasticity has virtually no effect on the results.

Table 6. Determinants of the Output-Inflation Trade-off, through 1972^a

Independent variable	Equation					
	6.1	6.2	6.3	6.4	6.5	6.6
Constant	0.501 (0.053)	0.518 (0.075)	0.519 (0.069)	0.595 (0.085)	0.539 (0.165)	0.575 (0.175)
Mean inflation	-3.051 (0.730)	...	-2.783 (0.979)	-6.081 (2.287)	...	-5.803 (2.417)
Square of mean inflation	12.145 (8.696)	...	11.939 (8.946)
Standard deviation of nominal GNP growth	...	-3.589 (1.281)	-0.654 (1.570)	...	-4.295 (4.958)	0.583 (4.959)
Square of standard deviation of nominal GNP growth	3.845 (26.066)	-6.605 (24.440)
<i>Summary statistic</i>						
R^2	0.281	0.140	0.267	0.298	0.119	0.266
Standard error	0.239	0.262	0.242	0.237	0.266	0.242

Source: Authors' calculations.

a. The equations are specified exactly as in table 5, but are estimated with data only through 1972. Numbers in parentheses are standard errors.

Note that the estimated effect of aggregate demand variability on the trade-off is positive, not negative as predicted by theory. This result is not obtained when mean inflation and its square are left out of the regression. Yet when all the variables are included, the variability coefficients, although small and statistically insignificant, have the wrong sign. This result is puzzling, and we have no definite explanation (but see the discussion below).

Tables 6 and 7 present the same regressions for the data ending in 1972 and the data beginning in 1973. In both subsamples we find similar results. In high-inflation countries, aggregate demand has a smaller effect on output.

To make clear the implications of our regression results, table 8 presents the predicted values of τ for various inflation rates. We present results for each of our samples, in each case using the most general specification (regressions 5.6, 6.6, and 7.6) and assuming $\sigma_x = 3$ percent. The results show that the effects of average inflation are large. At a zero rate of inflation, fluctuations in aggregate demand are in the first year reflected two-thirds in output and one-third in prices. At a 5 percent rate of inflation, the first-year impact on output is between one-third and one-

Table 7. Determinants of the Output-Inflation Trade-off, after 1972^a

Independent variable	Equation					
	7.1	7.2	7.3	7.4	7.5	7.6
Constant	0.458 (0.071)	0.431 (0.069)	0.459 (0.072)	0.683 (0.099)	0.629 (0.103)	0.731 (0.107)
Mean inflation	-1.025 (0.296)	...	-0.888 (0.629)	-3.307 (0.809)	...	-2.571 (1.162)
Square of mean inflation	3.144 (1.051)	...	2.043 (1.390)
Standard deviation of nominal GNP growth	...	-1.852 (0.599)	-0.308 (1.244)	...	-6.564 (1.952)	-2.808 (2.550)
Square of standard deviation of nominal GNP growth	14.375 (5.700)	8.827 (7.192)
<i>Summary statistic</i>						
R^2	0.207	0.169	0.189	0.336	0.265	0.328
Standard error	0.312	0.319	0.316	0.286	0.300	0.287

Source: Authors' calculations.

a. The equations are specified exactly as in table 5, but are estimated using data after 1972. Numbers in parentheses are standard errors.

half. At a 20 percent rate of inflation, the estimated first-year impact on output is small and sometimes negative.

The Determinants of the Trade-off: Time Series Cross-Section Results. Table 9 presents cross-country regressions of the change in the trade-off from the first to the second subsample on the change in the mean level of inflation and the change in variability. These regressions test Keynesian and new classical theories by examining the differences across countries, not in the level of the output-inflation trade-off, but in the change in the trade-off over time. These regressions have the advantage of correcting for any fixed country effects. For example, the extent of wage and price rigidity and thus the output-inflation trade-off may depend on various country-specific institutions, such as the laws governing labor negotiations. If such institutions do not change substantially from our first to our second subsample, then they will not introduce a bias in these regressions, even if they do introduce a bias in the regressions in levels.⁴⁴

44. For these fixed country effects to bias the regressions in levels, the fixed effects must for some reason be correlated with the average level of inflation.

Table 8. Predicted Output-Inflation Trade-off at Various Inflation Rates^a

<i>Mean inflation (percent)</i>	<i>Regression</i>		
	<i>5.6</i>	<i>6.6</i>	<i>7.6</i>
0	0.62	0.59	0.65
5	0.36	0.33	0.53
10	0.14	0.13	0.42
15	-0.05	-0.02	0.32
20	-0.19	-0.10	0.22

Source: Authors' calculations based on regression equations 5.6, 6.6, and 7.6 from tables 5, 6, and 7.

a. These figures assume the standard deviation of nominal GNP growth, σ_{Δ} , is 3 percent.

The results in table 9 are qualitatively the same as those for the level regressions above, but the estimated effects are somewhat smaller. For example, regression 5.6 implies that an increase in inflation from 5 percent to 10 percent reduces the trade-off by 0.22. In contrast, regression 9.6 implies that such an increase in inflation reduces the trade-off by only 0.12. This finding may be due to the possibility of bias in the regression in levels discussed above.

There are two other reasons why the regression for the change in the trade-off might produce smaller effects of mean inflation, one statistical and one economic. The statistical reason is that the change in the sample mean inflation might be a very noisy estimator of the change in the true mean inflation. Such sampling error would tend to bias downward the coefficients. This downward bias is probably smaller in the levels regression, because the "signal-to-noise" ratio is greater. The noise is less because the sampling error for the level of inflation is less than it is for the change in inflation; the signal is greater because the variation in the level of mean inflation across countries is plausibly larger than the variation in the change in mean inflation across countries. Hence, measurement error in mean inflation due to sampling error is probably a more important problem for the regressions in table 9 than for those in tables 5, 6, and 7.

The economic reason is that the frequency with which prices are adjusted might not change immediately with changes in the mean level of inflation. For example, a company that issues a catalog once a year might not switch to issuing a catalog twice a year unless it were certain that the change in mean inflation were permanent. Hence, changes in mean inflation observed between our two subsamples might have been

Table 9. Explaining the Change in the Output-Inflation Trade-off^a

Independent variable	Equation					
	9.1	9.2	9.3	9.4	9.5	9.6
Constant	0.009 (0.078)	-0.037 (0.062)	0.009 (0.085)	0.172 (0.091)	-0.012 (0.059)	0.154 (0.104)
Change in mean inflation	-0.595 (0.442)	...	-0.603 (0.773)	-3.174 (0.968)	...	-2.619 (1.458)
Change in the square of mean inflation	3.094 (1.054)	...	1.895 (1.608)
Change in the standard deviation of nominal GNP growth	...	-0.852 (0.787)	0.016 (1.366)	...	-5.820 (2.095)	-2.384 (3.046)
Change in the square of standard deviation of nominal GNP growth	14.392 (5.678)	9.094 (8.589)
<i>Summary statistic</i>						
\bar{R}^2	0.019	0.004	-0.005	0.173	0.121	0.158
Standard error	0.359	0.362	0.364	0.329	0.340	0.333

Source: Authors' calculations.

a. The dependent variable is $\tau_{late} - \tau_{early}$, where τ_{late} is the estimate of τ using data after 1972 and τ_{early} is the estimate of τ using data through 1972. Numbers in parentheses are standard errors.

perceived as partly transitory and may have evoked smaller changes in price-setting behavior.

ECONOMETRIC ISSUES AND ROBUSTNESS

Having estimated the output-inflation trade-offs for different countries using standard specifications, we next discuss econometric issues raised by those specifications and examine a series of variations on our basic test. Our central finding, that the effect of nominal demand movements on real output falls as average inflation rises, is robust.

Supply Shocks. In both the theoretical model of equations 5–13 and the preceding empirical work, we assume that all aggregate shocks are demand shocks. In actual economies, of course, output movements result from supply as well as demand shocks. The residual of equation 14 reflects these supply shocks. We now investigate the effects of supply shocks on our results.

The presence of supply shocks can in principle cause several distinct problems. As we explain below, our estimates of the trade-off parameter τ are biased if changes in nominal GNP are correlated with supply shocks, as can occur either if aggregate demand is not unit elastic or if supply and demand shocks are correlated. Most important, supply shocks can bias our estimates of the key relationship between τ and average inflation. This is the case if average inflation is correlated with bias in the estimated τ 's.

Before considering these separate problems, we perform a simple check of their overall importance by examining whether the results change when we restrict our attention to the period through 1972, when supply shocks are generally thought to have been less significant. Comparison of tables 5 and 6 shows that ending the sample in 1972 has essentially no effect on the results. Focusing on the period after 1972, on the other hand, leads to weaker results, suggesting that if supply shocks have any effect on our results, it is to obscure the phenomenon for which we are testing.

We now turn to the specific problems caused by supply shocks, beginning with bias in the estimates of τ . This parameter gives the output effects of demand shocks, as measured by Δx . Supply shocks can be viewed as variables left out of the output equation, and so they cause bias if they are correlated with Δx . Supply shocks directly affect Δx as long as the aggregate demand curve is not unit elastic; only with unit elastic demand do the movements in p and y caused by supply shocks have exactly offsetting effects on x .⁴⁵ In addition, even if demand is unit elastic, so Δx reflects only demand shocks, movements in demand and supply may be correlated. This is the case, for example, if monetary policy accommodates supply shocks.⁴⁶

The importance of these problems is questionable. Available evidence suggests that an elasticity of aggregate demand of roughly one is realistic.⁴⁷ And the endogeneity of monetary policy can *reduce* bias:

45. See Marcelle Arak, "Some International Evidence on Output-Inflation Tradeoffs: Comment," *American Economic Review*, vol. 67 (September 1977), pp. 728–30.

46. This discussion suggests that instrumental variables estimation is unlikely to be useful here. Appropriate instruments are variables that affect nominal GNP growth and are uncorrelated with supply shocks. If monetary and fiscal policies respond to supply shocks, measures of the stance of these policies are not valid instruments. Both for this reason and because of data limitations, we do not pursue use of instrumental variables.

47. For example, the values that Mankiw and Summers suggest for the relevant parameters of the *IS* and *LM* curves imply an elasticity of the aggregate demand curve of

policymakers may respond to supply shocks in a way that eliminates the effects on nominal GNP. In any case, we empirically investigate the importance of biases in $\hat{\tau}$ in two ways. First, note that if aggregate demand has elasticity of $a \neq 1$, a supply shock affects $x = p + y$ but leaves $ap + y$ unchanged. The same is true if policymakers target $ap + y$ rather than $p + y$. If supply shocks leave $ap + y$ unaffected, the effect of aggregate demand movements can be estimated by regressing y_t on $\Delta(ap_t + y_t)$ rather than Δx_t .⁴⁸ The estimated coefficients can then be regressed on average inflation and the standard deviation of changes in $ap + y$. Thus a check for bias caused by supply shocks is to posit a range of values of a and examine whether the results are robust to the choice of a . We consider four values for a ranging from 0.5 to 2. For $a = 0.5$ the second-stage regression (with quadratic terms included) yields:

$$\hat{\tau}_{a=0.5} = 0.812 - 6.397\bar{\pi} + 10.935\bar{\pi}^2 + 8.630\sigma_{0.5p+y} - 58.140\sigma_{0.5p+y}^2,$$

$$(0.110) \quad (1.628) \quad (3.671) \quad (4.704) \quad (30.706)$$

$$\bar{R}^2 = 0.421; \text{ standard error} = 0.231,$$

where $\hat{\tau}_{a=0.5}$ denotes the coefficient on $\Delta(ap_t + y_t)$ with $a = 0.5$ from the first-stage regression (standard errors are in parentheses). For $a = 2$,

$$\hat{\tau}_{a=2} = 0.211 - 2.632\bar{\pi} + 3.569\bar{\pi}^2 + 0.167\sigma_{2p+y} + 0.116\sigma_{2p+y}^2,$$

$$(0.038) \quad (0.995) \quad (1.851) \quad (0.585) \quad (0.792)$$

$$\bar{R}^2 = 0.333; \text{ standard error} = 0.100.$$

The results are similar to those for our baseline case: the coefficients on the average inflation variables are of the predicted sign, quantitatively large, and statistically significant; the coefficients on the variability measures are small, wrong-signed, and insignificant. (Because the units of σ_{ap+y} and τ_a depend on a , the magnitudes of the coefficients from different regressions are not directly comparable.) The results for $a = 0.67$ and $a = 1.5$ are also similar.

A second approach to reducing bias in the estimates of τ is to include measures of supply shocks in our equation for the output-inflation trade-off—that is, to add the left-out variable. We focus on oil price changes, which are the most easily identifiable and perhaps the largest supply shocks during our sample period. We do this by including a dummy

slightly less than one. N. Gregory Mankiw and Lawrence H. Summers, "Money Demand and the Effects of Fiscal Policies," *Journal of Money, Credit and Banking*, vol. 18 (November 1986), pp. 415–29.

48. As in our main regressions, we also include a constant, a trend, and y_{t-1} .

variable in our estimation of the output-inflation trade-off (equation 14) that is equal to +1 in the years of major oil price increases (1974, 1979, and 1980) and -1 in the one year of a major price decrease (1986). (The natural alternative of entering a separate dummy for each of these years is equivalent to simply discarding these years from the sample, and would thus be similar to the previous strategy of stopping the sample in 1972.) Including the dummy has little effect on the results: the correlation of the τ 's estimated with and without the dummies is 0.98, and the regressions of the estimated τ 's on average inflation and the variability of demand growth are little changed; in fact, the magnitude and significance of the effect of average inflation are slightly larger.

So far we have addressed the problem of supply shocks by attempting to reduce the possible bias in the estimates of τ . We now turn to the implications of any remaining bias for our estimates of the cross-country relation between τ and average inflation. Bias in these estimates arises from biases in $\hat{\tau}$ only if the latter are correlated with average inflation. This correlation could occur if the variance of supply shocks or the degree to which they are accommodated, which affects the bias in $\hat{\tau}$, is correlated with average inflation.⁴⁹ But there is no strong reason to expect this. For example, suppose that one country expands aggregate demand in response to unfavorable supply shocks and contracts in response to favorable shocks (so that shocks fall mainly on prices), while another country does the reverse. The estimates of τ are biased in different directions for the two countries, but the bias is not correlated with average inflation because neither country is pursuing a systematically more expansionary policy. On the other hand, if one country always pursues more expansionary policies, stimulating demand after both favorable and unfavorable shocks, then the countries have different average inflation, but $\hat{\tau}$ is not biased because nominal growth and supply shocks are uncorrelated.

Although we do not think it likely that bias in τ caused by supply shocks is correlated with average inflation, we check for such a problem. The effects of supply shocks on the cross-country regression can be reduced by controlling for differences in the size of these shocks. We

49. Indeed, a correlation between the variance of supply shocks and average inflation causes bias even if $\hat{\tau}$ is unbiased: the size of supply shocks is a left-out variable in the equation for τ because a large variance of supply (like a large variance of demand) reduces the frequency of price adjustment.

experiment with two types of measures of the size of supply shocks. The first are country characteristics. Specifically, we use the degree of industrialization, measured by manufacturing output as a fraction of total output in 1965, and the degree of openness, measured by the ratio of imports to output in 1965.⁵⁰ Both variables probably affect a country's susceptibility to supply shocks. (Both could also affect the output-inflation trade-off in ways unrelated to supply shocks.) When these variables are included in regression 5.6, however, their coefficients are small and highly insignificant; more important, the coefficients on the remaining variables are virtually unchanged.

The second type of variable that we add to the cross-section regression is a crude measure of the magnitude of supply shocks. Since the residual of equation 14 reflects supply shocks, we use the variance of the residual, σ_ϵ^2 , as a measure of the variance of supply shocks. We measure the variance of demand shocks by the variance of nominal GNP growth, σ_x^2 ; thus $\sigma_\epsilon^2/\sigma_x^2$ is a crude proxy for the relative magnitudes of supply and demand shocks. Adding this ratio to our cross-section regression yields

$$\hat{\tau} = 0.163 - 5.421\hat{\pi} + 7.833\hat{\pi}^2 + 3.451\sigma_x - 5.317\sigma_x^2 + 1.339(\sigma_\epsilon^2/\sigma_x^2),$$

(0.143) (1.731) (3.377) (2.251) (6.245) (0.379)

$$\bar{R}^2 = 0.508; \text{ standard error} = 0.192.$$

The coefficient on $\sigma_\epsilon^2/\sigma_x^2$ is positive and significant. This could occur if supply shocks fall mainly on output rather than prices, thereby causing a positive bias in $\hat{\tau}$, with the size of the bias increasing in the relative size of supply shocks. In any case, the coefficients on the variables of central interest, $\hat{\pi}$ and $\hat{\pi}^2$, are essentially unchanged. The effect of demand variability remains wrong-signed and insignificant but is now somewhat larger than before.

In sum, a wide variety of tests fails to provide any evidence that supply shocks have an important effect on our results.

Specification of the Cross-Country Regression. Another issue concerning our specification, which is related to the possibility of supply shocks discussed above, is how to measure aggregate variability. Since the only aggregate shocks in our theoretical model are demand shocks, the model implies that the variance of nominal GNP growth is the

50. We are unable to obtain data on manufacturing output for Switzerland and Iceland; we therefore exclude these countries from the regression.

appropriate measure. But if the model were extended to include aggregate supply shocks, the variance of inflation might be a better measure. The frequency of price changes, and hence τ , would depend on the variances of both shocks. And while σ_x^2 might capture only the variance of demand shocks (for example, if demand is unit elastic), σ_π^2 would reflect both variances.

Replacing σ_x with σ_π in our cross-country regression leads to

$$\begin{aligned} \hat{\tau} &= 0.616 - 3.221\bar{\pi} + 4.072\bar{\pi}^2 - 2.235\sigma_\pi + 5.513\sigma_\pi^2, \\ &\quad (0.085) \quad (2.254) \quad (4.098) \quad (2.572) \quad (6.708) \\ \bar{R}^2 &= 0.367; \text{ standard error} = 0.218. \end{aligned}$$

Inflation and inflation squared are now not statistically significant. The point estimates, however, continue to suggest that average inflation has a large effect on the output-inflation trade-off; for example, an increase in $\bar{\pi}$ from 5 percent to 10 percent reduces τ by 0.13. In addition, the estimated effects of variability, although highly imprecise, are now in the direction predicted by our model and are quantitatively reasonable; they imply that an increase in σ_π from 5 percent to 10 percent reduces τ by 0.07. It may be that the puzzling results in our basic specification concerning the effects of variability result from an inappropriate variability measure.

It might appear that a natural extension of this consideration of alternative variability measures is to include both σ_x and σ_π in the cross-country regression. But this specification is unlikely to provide useful information: it is likely to produce small coefficients on average inflation regardless of whether average inflation truly affects τ . To see why, suppose that the correct model is

$$(15) \quad \Delta y = \tau \Delta x + \epsilon,$$

where ϵ is a supply shock that is uncorrelated with Δx . Then, since $\pi = \Delta x - \Delta y$, $\sigma_\pi^2 = (1 - \tau)^2 \sigma_x^2 + \sigma_\epsilon^2$, and so

$$(16) \quad (1 - \tau)^2 = \frac{\sigma_\pi^2}{\sigma_x^2} - \frac{\sigma_\epsilon^2}{\sigma_x^2}.$$

Expression 16 is an identity and thus holds regardless of how τ is determined. Regressing τ on σ_x and σ_π would fail to produce a perfect fit only because of cross-country variation in σ_ϵ and because the functional form of equation 16 is not linear. Adding average inflation to an equation

that is almost an identity does not yield a valid test of the importance of this variable. For a given variance of supply shocks, the variance of inflation is determined completely by the variance of demand shocks (σ_x^2) and the proportion of each shock that is reflected in inflation ($1 - \tau$). Thus regressing any one of τ , σ_x^2 , and σ_π^2 on the other two should provide an excellent fit and leave little room for other explanatory variables. Again, this is true regardless of whether our theory of the determinants of τ is correct.

We conclude that if there were a theory that suggested that τ was a function of both σ_x and σ_π , the type of test that we employ would be unable to discriminate between such a model and our model. We know of no such theory, however.⁵¹

Specification of the Output-Inflation Trade-off. In the first part of this section, following Lucas and others, we estimated the short-run output-inflation trade-off by regressing real output on the change in nominal GNP and other variables. Our model, however, predicts that real output in the current period will depend on the innovation in nominal GNP in the current period and on lagged innovations (see equation 12). Thus the theory suggests a specification of form

$$(17) \quad y_t = \tau(x_t - E_{t-1}x_t) + \sum_i \alpha_i (x_{t-i} - E_{t-(i+1)}x_{t-i}) + \beta'Z_t,$$

where E_t denotes an expectation at time t , x is nominal demand, and Z is a vector of other variables that affect output.

The equation that we use to estimate τ , equation 14, differs from equation 17 by omitting past innovations and by employing the change in nominal GNP rather than the current innovation. Because $x_t - E_{t-1}x_t$ reflects information learned in period t and is thus uncorrelated with

51. Our discussants consider the simple model that inflation and output growth are governed by independent processes and are therefore uncorrelated. This model is quite implausible. Simple real business cycle theories, for example, suggest that real output and nominal demand are determined independently, which implies that real growth and inflation are negatively correlated. More important, the discussants' model does *not* suggest a specification in which σ_π and σ_x are entered separately. Instead it predicts that the estimated τ should equal $1 - \sigma_\pi^2/\sigma_x^2$. When τ is regressed on this ratio, average inflation, and average inflation squared, the coefficients on the average inflation variables have the signs predicted by our theory and are significant at the 1 percent level. When the discussants, appealing to their simple model, regress τ on $\bar{\pi}$, σ_π , and σ_x and their squares (which, as we explain above, is not the specification implied by the model), they find, as we expect, a positive effect of σ_x , a negative effect of σ_π , and a small effect of $\bar{\pi}$.

all variables known at time $t - 1$, the omission of past innovations does not bias the estimate of τ . The use of Δx_t in place of the current innovation also poses no difficulties. A natural way to estimate equation 17 would be to employ a two-stage procedure, first regressing x_t on a set of variables known at time $t - 1$ and then using the fitted values from this regression as an estimate of $E_{t-1} x_t$ in equation 17.⁵² Estimating equation 14, however, is numerically identical to first regressing Δx_t on the other right-hand-side variables and then using the residuals from this regression rather than Δx_t in equation 14. Thus equation 14 can be thought of as a simple one-step way of implementing the two-stage procedure, with the right-hand-side variables for the first regression the same as the control variables in the second stage.

While our regression is in principle equivalent to a two-stage procedure, one could argue that our specification in equation 14 includes too few control variables to capture expected movements in demand. Specifically, in countries where expected inflation varies considerably over time, a large part of the variation in Δx will be predictable (on the basis of lagged Δx , for example), but cannot be predicted using only the other right-hand-side variables of equation 14. We have also estimated more elaborate versions of equation 14 in which output depends on two lags of output, current and two lagged values of nominal GNP, and a time trend. This specification appears to be rich enough for the residuals from regressing nominal GNP on the other right-hand-side variables to largely represent innovations; for example, the residuals do not exhibit serial correlation. For approximately half of the countries, one can reject the restrictions imposed by equation 14 in favor of the more general equation. Yet the estimate of the coefficient on nominal GNP, which is our main interest, is not substantially affected by these restrictions. For the United States, for example, one can reject the restrictions in favor of the more general equation at the 1 percent level; yet the estimate of τ changes only from 0.642 to 0.656. Across the 43 countries, the correlation between the τ estimated from equation 14 and the τ estimated from the more general equation is 0.88. Moreover, using the τ 's from the more elaborate equation for the cross-section regression has only minor effects on the results. We thus conclude that the simpler equation is sufficient for our purposes.

52. See Robert J. Barro, "Unanticipated Money, Output, and the Price Level in the United States," *Journal of Political Economy*, vol. 86 (August 1978), pp. 549–80.

The parsimony of our equation also has an important statistical advantage. When we divide the time series in half and estimate the inflation-output trade-off for data both through and after 1972, we are sometimes left with very short time series: less than a dozen years. The simpler equation, even if rejected by the data for the overall sample, may be preferred because it conserves on the scarce degrees of freedom.⁵³

The Sample. To examine the effects of restricting the sample of countries, we limit the sample in two ways. First, to check whether our results depend crucially on a few extreme observations, we examine the effects of excluding countries with extreme average inflation and extreme variability of demand growth. Specifically, we drop from the sample the six countries with average inflation rates and standard deviations of nominal GNP growth greater than 20 percent (Argentina, Bolivia, Brazil, Israel, Peru, and Zaire). Second, because there may be systematic differences between the major industrialized countries and the remaining countries in our sample that influence both the output-inflation trade-off and average inflation or nominal demand variability, we consider the effects of restricting the sample to OECD countries.⁵⁴ One of these countries, Iceland, has both average inflation and standard deviation of nominal GNP growth that are nearly double those of any of the other OECD countries; we therefore consider the results both with and without Iceland. Of course, restrictions on the sample have the disadvantage of discarding some of the variation in the right-hand-side variables, which could make the relationships for which we are testing more difficult to detect.

53. If the theory were extended to allow demand to follow a process other than a random walk, it appears that it would imply that information about future changes in demand would also affect current output. Indeed, for most of the countries in our sample, changes in nominal GNP are positively serially correlated rather than white noise; moreover, the degree of serial correlation is positively correlated with average inflation. Thus a conceivable alternative explanation of our finding of an inverse link between the estimated π 's and average inflation is that nominal GNP changes are more persistent in high-inflation countries and that they therefore lead to larger short-run price responses even though the frequency of price adjustment is constant across countries. We find this explanation implausible: a rough calculation using a discrete-time staggering model with persistent demand changes suggests that the magnitude of this effect is much too small, and testing this explanation directly by adding to the cross-section regression an estimate of the extent to which the current change in nominal GNP helps to predict future changes given information previously available leaves the results essentially unchanged.

54. There are 21 OECD countries in our sample: Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, Japan, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom, and United States.

When we estimate the cross-country regression with the six countries with high inflation and high demand variability excluded, we find:

$$\begin{aligned} \hat{\tau} = & 1.042 - 10.563\bar{\pi} + 28.998\bar{\pi}^2 - 5.933\sigma_x + 40.634\sigma_x^2, \\ & (0.253) \quad (4.485) \quad (21.968) \quad (7.032) \quad (45.214) \\ \bar{R}^2 = & 0.306; \text{ standard error} = 0.225. \end{aligned}$$

The results are similar to those for the full sample of countries: average inflation has a large and statistically significant effect on the output-inflation trade-off (inflation and inflation squared are jointly significant at the 1 percent level), while variability has a small and insignificant impact on the trade-off.

As before, we also regress the change in the trade-off between the period ending in 1972 and the period beginning in 1973 on the changes in average inflation and aggregate demand variability. This yields

$$\begin{aligned} \Delta\hat{\tau} = & 0.189 - 3.089\Delta\bar{\pi} - 0.514\Delta\bar{\pi}^2 - 12.163\Delta\sigma_x + 80.887\Delta\sigma_x^2, \\ & (0.126) \quad (2.554) \quad (5.522) \quad (5.113) \quad (32.505) \\ \bar{R}^2 = & 0.324; \text{ standard error} = 0.315. \end{aligned}$$

The change in average inflation has an important effect on the change in the trade-off; the null hypothesis that the coefficients on $\Delta\bar{\pi}$ and $\Delta\bar{\pi}^2$ are zero is rejected at the 5 percent level. The estimated relationship is essentially linear. Changes in variability also have an important effect on the trade-off. As σ_x rises, increases in σ_x first lower and then raise τ .

Table 10 presents the results for OECD countries. We estimate the cross-country equation for our entire sample of years and for the two subsamples. In all cases, the inflation coefficients have the predicted signs and are large. As expected, however, they are estimated much less precisely than for the larger sample of countries. Consequently, the results are often not statistically significant. For the entire sample of years, the inflation coefficients are jointly significant only when Iceland is excluded (column 4). This regression implies that an increase in inflation from 4.5 percent to 7.9 percent, which is from one standard deviation below the mean to one standard deviation above the mean, reduces the trade-off parameter τ by 0.32.

Finally, we estimate the determinants of the change in the trade-off for OECD countries. The results are little affected by the inclusion of quadratic terms or by whether Iceland is included in the sample. We therefore focus on the linear specification with Iceland included. We find:

Table 10. Determinants of the Output-Inflation Trade-off for OECD Countries^a

Independent variable	Equation					
	Iceland included			Iceland excluded		
	Trade-off, full sample	Trade-off through 1972	Trade-off after 1972	Trade-off, full sample	Trade-off through 1972	Trade-off after 1972
Constant	1.073 (0.413)	0.840 (0.673)	1.124 (0.532)	4.369 (1.082)	1.470 (0.955)	2.215 (0.832)
Mean inflation	-21.831 (17.608)	-19.665 (14.220)	-6.450 (4.941)	-68.672 (20.279)	-47.146 (32.716)	-17.881 (8.361)
Square of mean inflation	118.022 (126.624)	144.447 (101.199)	20.284 (21.167)	477.128 (150.937)	460.137 (353.076)	76.422 (39.482)
Standard deviation of nominal GNP growth	9.020 (30.008)	7.343 (28.282)	1.485 (40.461)	-77.961 (36.209)	4.061 (28.613)	-37.254 (45.024)
Square of standard deviation of nominal GNP growth	-140.708 (298.452)	-117.567 (320.871)	-258.212 (575.272)	804.893 (379.404)	-79.342 (324.757)	326.317 (651.012)
<i>Summary statistic</i>						
\bar{R}^2	0.147	-0.036	0.453	0.457	0.008	0.395
Standard error	0.264	0.243	0.246	0.210	0.244	0.235

Source: Authors' calculations. See text description.

a. The dependent variable is the output-inflation trade-off parameter, τ , estimated in table 4. Numbers in parentheses are standard errors.

$$\Delta \hat{\tau} = 0.192 - 2.580 \Delta \bar{\pi} - 7.428 \Delta \sigma_x,$$

$$(0.068) \quad (0.792) \quad (2.779)$$

$$\bar{R}^2 = 0.583; \text{ standard error} = 0.189.$$

Both coefficients have the expected sign and are large and statistically significant, and the fit is excellent.

Output Variability. If higher average inflation reduces the real effects of nominal disturbances, it also reduces the variability of output. In our theoretical model, equation 12 shows that the variance of y falls when $1 - w(\bullet)$ falls. Thus another test of Keynesian theory is to examine the link between average inflation and output variability. As before, new classical theories, because they attribute fluctuations to unanticipated nominal disturbances and to real shocks, predict no role for average inflation.

In estimating the relation between output variability and average inflation, it is of course necessary to control for the variance of nominal GNP growth. Both Keynesian theories and Lucas's imperfect informa-

tion theory predict that an increase in the size of nominal shocks will increase the variance of output (in both Lucas's theory and our model, the direct effect of increased shocks is only partially offset by a smaller responsiveness of output to shocks). Because the variance of nominal GNP growth is highly correlated with average inflation, omitting this variable could cause severe bias.

Table 11 reports the results.⁵⁵ Output variability is measured by the standard deviation of real GNP growth σ_y . The first two columns show that the variability of output growth is positively correlated with both mean inflation and demand variability. Regression 11.3 shows the effects of including both variables on the right-hand side. The estimated effect of $\bar{\pi}$ is negative and that of σ_x positive. As before, there are reasons to expect nonlinearities: the negative impact of mean inflation should diminish as mean inflation rises, and the positive impact of variability should also fall as variability rises. The final column of table 11 reports the results with quadratic terms included. All four variables have the signs predicted by Keynesian theory, and all of the t -statistics exceed 3. The implied effects are large: an increase in mean inflation from 5 percent to 10 percent reduces the standard deviation of output growth, σ_y , by 1.2 percentage points, while an increase in the standard deviation of nominal growth from 5 percent to 10 percent increases σ_y by 2.1 percentage points. This finding of a strong inverse link between mean inflation and output variability confirms the predictions of Keynesian theories and contradicts those of new classical theories.⁵⁶

Reexamining Previous Evidence. Numerous studies have examined Lucas's proposition that the variability of aggregate demand affects the output-inflation trade-off. Here we reexamine several such studies and show that the evidence provided by other authors largely supports our claim that average inflation is an important determinant of the trade-off.

55. The links between output variability, average inflation, and demand variability are also examined in Lawrence H. Summers and Sushil B. Wadhvani, "Some International Evidence on Labour Cost Flexibility and Output Variability," Working Paper 981 (Centre for Labour Economics, June 1987).

56. The results for the two subperiods and for the change between the two subperiods point in the same direction as the results for the full sample but are less clear-cut. Both for the period through 1972 and for the change between the two subperiods the point estimates imply large negative effects of inflation on output variability. The estimates are highly imprecise, however; the null hypothesis of no effect cannot be rejected in either case. For the post-1972 period the point estimates imply that the effect of average inflation on output variability is not monotonic: it is negative at low inflation rates but becomes positive near the sample mean.

Table 11. Determinants of Output Variability^a

Independent variable	Equation					
	11.1	11.2	11.3	11.4	11.5	11.6
Constant	0.0315 (0.0032)	0.0290 (0.0032)	0.0290 (0.0031)	0.0270 (0.0053)	0.0174 (0.0047)	0.0202 (0.0044)
Mean inflation	0.0418 (0.0220)	...	-0.0817 (0.0507)	0.1150 (0.0719)	...	-0.3254 (0.1009)
Square of mean inflation	-0.1494 (0.1398)	...	0.6294 (0.1967)
Standard deviation of nominal GNP growth	...	0.0777 (0.0270)	0.1740 (0.0653)	...	0.3143 (0.0786)	0.6581 (0.1261)
Square of standard deviation of nominal GNP growth	-0.6779 (0.2140)	-1.572 (0.361)
<i>Summary statistic</i>						
\bar{R}^2	0.058	0.148	0.180	0.062	0.302	0.429
Standard error	0.0144	0.0137	0.0135	0.0144	0.0124	0.0112

Source: Authors' calculations. See text description.

a. Dependent variable is output variability, measured by the standard deviation of real GNP growth, σ_y . Numbers in parentheses are standard errors.

Lucas originally examined 18 countries with data from 1952 to 1967.⁵⁷ Our cross-country regression using the trade-offs estimated by Lucas and his sample statistics yields:

$$\hat{\tau} = 0.929 - 6.034\bar{\pi} + 19.641\bar{\pi}^2 - 8.397\sigma_x + 32.369\sigma_x^2,$$

(0.304) (5.698) (24.967) (9.435) (43.212)

$$\bar{R}^2 = 0.446; \text{ standard error} = 0.185.$$

Alberro redid Lucas's study with a sample of 49 countries.⁵⁸ The cross-country regression estimated with Alberro's figures is:

$$\hat{\tau} = 0.752 - 6.620\bar{\pi} + 14.763\bar{\pi}^2 + 0.158\sigma_x - 5.230\sigma_x^2,$$

(0.089) (2.021) (7.150) (1.875) (4.325)

$$\bar{R}^2 = 0.418; \text{ standard error} = 0.242.$$

Both Lucas's and Alberro's estimates provide strong support for our hypothesis that high rates of inflation reduce the real impact of nominal demand. The estimated effects of average inflation are large. An increase in average inflation from 5 percent to 10 percent implies a fall in τ of 0.15

57. Lucas, "Some International Evidence."

58. Jose Alberro, "The Lucas Hypothesis and the Phillips Curve: Further International Evidence," *Journal of Monetary Economics*, vol. 7 (March 1981), pp. 239–50.

when Lucas's figures are used and of 0.22 when Alberro's figures are used. These results are quite similar to those we obtained in table 8, even though the periods of time and samples of countries are substantially different. Hence, the results of Lucas and Alberro provide support for new Keynesian rather than new classical theories of the business cycle.

We have also examined the evidence of Kormendi and Meguire on the output effects of unanticipated money.⁵⁹ A cross-country regression of their estimate of the impact of unanticipated changes in the money supply, which they call χ , on their estimate of the variance of unanticipated money and on the average rate of inflation yields:

$$\chi = 0.906 - 3.508\bar{\pi} + 19.422\bar{\pi}^2 - 15.908\sigma_{um} + 94.093\sigma_{um}^2,$$

$$(0.230) \quad (5.390) \quad (33.844) \quad (6.055) \quad (47.834)$$

$$\bar{R}^2 = 0.303; \text{ standard error} = 0.140.$$

(Since Kormendi and Meguire do not provide the average rate of inflation, we obtain it from our sample. We restrict the regression to the 26 countries for which we can match their sample period exactly.)

In contrast to the results of Lucas and Alberro, the estimates of Kormendi and Meguire do not provide support for our hypothesis. Inflation has little effect on χ , while the variance of unanticipated money appears an important determinant. There are a variety of ways to reconcile this finding with the previous results.

One possibility is that money is measured with error and that the extent of measurement error varies across countries. Greater measurement error in the money supply would tend to reduce the estimate of χ while increasing the variability of unanticipated money. Measurement error would also make it more difficult to isolate the effect of inflation on the trade-off.

A second possibility is that money may be a bad measure of aggregate demand. If different countries follow different policies regarding the extent to which they offset exogenous demand shocks, such as changes in monetary velocity, then χ will again be a noisy measure.

An examination of the estimated χ suggests it is not a good measure of the real impact of aggregate demand. In particular, the correlation between the estimates of χ provided by Kormendi and Meguire and the

59. Roger C. Kormendi and Philip G. Meguire, "Cross-Regime Evidence of Macroeconomic Rationality," *Journal of Political Economy*, vol. 92 (October 1984), pp. 875-908.

estimates of τ we obtain for the same 26 countries over the same time samples is only 0.38. Moreover, the United States appears a very unusual country. The estimated χ for the United States is 0.96, the largest in the sample. The second biggest χ is 0.38 for Belgium. Since it seems implausible that the United States is such an extreme outlier, we conclude that χ is probably not a good measure of the output-inflation trade-off.

A third possibility is that high average inflation, by increasing price flexibility, reduces the effects of aggregate demand on output but *increases* the effect of money on aggregate demand. In this case, the influence of average inflation on the net effect of money on output is ambiguous. DeLong and Summers present a model in which price flexibility increases the effect of monetary shocks on demand by increasing the variability of expected inflation.⁶⁰ Future research might use cross-country data to untangle the effects of price flexibility on the money-demand and demand-output links.

Conclusion

We have examined the short-run trade-off between output and inflation using international data. A robust finding is that this trade-off is affected by the average rate of inflation. In countries with low inflation, the short-run Phillips curve is relatively flat—fluctuations in nominal aggregate demand have large effects on output. In countries with high inflation, the Phillips curve is steep—fluctuations in demand are reflected quickly in the price level. The same finding emerges when we examine the change in the trade-off over time. Countries that experience an increase in average inflation also typically experience an increased responsiveness of prices to aggregate demand.

Our finding has three important implications. First, it provides evidence against new classical theories of the output-inflation trade-off. In his classic study, Lucas found that international differences in the trade-off were related to differences in the variability of aggregate demand and interpreted his finding as evidence for the imperfect information theory of the business cycle. This theory, however, predicts that the trade-off

60. J. Bradford De Long and Lawrence H. Summers, "Is Increased Price Flexibility Stabilizing?" *American Economic Review*, vol. 76 (December 1986), pp. 1031–44.

should not be related to average inflation. It is therefore inconsistent with our empirical results.

Second, our finding supports new Keynesian theories of the business cycle that derive nominal rigidities from optimizing behavior. Our theoretical model of price adjustment shows that macroeconomic effects of the sort we observe can result from empirically plausible microeconomic parameters. In particular, average inflation can strongly influence the output-inflation trade-off through its effects on the frequency of price changes.

Third, our finding implies that the trade-off faced by macroeconomic policymakers depends on the average rate of inflation and that it changes when the average rate of inflation changes. This effect is substantial even for moderate rates of inflation. Our estimates using the entire sample imply that the real impact of aggregate demand is twice as great at 5 percent inflation as at 10 percent inflation. Perhaps the short-run Phillips curve Alan Greenspan is facing today is not the same one that Paul Volcker faced a decade ago.

APPENDIX

THIS APPENDIX describes how we solve our model of nominal rigidity.

The Behavior of $p(t)$ for a Given λ

The first step is to derive the behavior of the aggregate price level for a given interval between price changes, λ . Substituting the formula for a firm's profit-maximizing price, equation 5, into the price-setting rule, equation 7, yields

$$(A.1) \quad p_i = \frac{1}{\lambda} \int_{s=0}^{\lambda} E_t \{p(t+s) + v[y(t+s) - \bar{y}(t+s)]\} ds + \theta_i(t).$$

Let $Q(t)$ be the average of all prices that are *set* at t . Equation A.1 implies

$$(A.2) \quad Q(t) = \frac{1}{\lambda} \int_{s=0}^{\lambda} E_t \{p(t+s) + v[y(t+s) - \bar{y}(t+s)]\} ds,$$

where the θ_i 's average to zero because they are uncorrelated across firms. Using the facts that $y(t) = x(t) - p(t)$, $E_t x(t + s) = x(t) + gs$ (because x follows a random walk with drift), and $\bar{y}(t) = \mu t$, we can rewrite equation A.2 as

$$(A.3) \quad \begin{aligned} Q(t) = (1 - v) \frac{1}{\lambda} \int_{s=0}^{\lambda} E_t p(t + s) ds \\ + vx(t) + v \frac{gs}{2} - v\mu \left(t + \frac{\lambda}{2} \right). \end{aligned}$$

The aggregate price level, $p(t)$, is the average of prices in effect at t . Given our assumption of staggered price setting, this means that $p(t)$ is the average of prices set from $t - \lambda$ to λ :

$$(A.4) \quad p(t) = \frac{1}{\lambda} \int_{s=0}^{\lambda} Q(t - s) ds.$$

Substituting equation A.4 into equation A.3 yields

$$(A.5) \quad \begin{aligned} Q(t) = (1 - v) \frac{1}{\lambda} \int_{s=0}^{\lambda} E_t \frac{1}{\lambda} \int_{r=0}^{\lambda} Q(t + s - r) dr ds \\ + vx(t) + v \frac{gs}{2} - v\mu \left(t + \frac{\lambda}{2} \right) \\ = (1 - v) \frac{1}{\lambda^2} \left[\int_{s=0}^{\lambda} (\lambda - s) Q(t - s) ds \right. \\ \left. + \int_{s=0}^{\lambda} (\lambda - s) E_t Q(t + s) ds \right] \\ + vx(t) + v \frac{gs}{2} - v\mu \left(t + \frac{\lambda}{2} \right). \end{aligned}$$

According to equation A.4, prices set at t depend on prices set between $t - \lambda$ and t , which are still in effect, and on expectations of prices set between t and $t + \lambda$, the period when prices set at t are in effect.

Equation A.4 gives an expression for Q in terms of its own past and expected future values. Our goal is to solve for Q as a function of the underlying demand disturbances. To do this we use the method of undetermined coefficients. That is, we posit a solution of form

$$(A.6) \quad Q(t) = A + Bt + \int_{s=0}^{\infty} q(s; \lambda) dZ(t - s),$$

and then solve for A , B , and $q(\bullet)$. The term $q(s;\lambda)$ is the fraction of a nominal shock at $t - s$ that is passed into newly set prices at t . Substituting equation A.6 into equation A.5 and simplifying, we obtain

$$\begin{aligned}
 (A.7) \quad A + Bt + \int_{s=0}^{\infty} q(s;\lambda) dZ(t-s) \\
 &= (1-\nu)A + (1-\nu)Bt \\
 &+ (1-\nu) \frac{1}{\lambda^2} \int_{s=0}^{\infty} \left[\int_{r=0}^{\min(s,\lambda)} (\lambda-r)q(s-r;\lambda) dr \right. \\
 &\quad \left. + \int_{r=0}^{\lambda} (\lambda-r)q(s+r;\lambda) dr \right] dZ(t-s) + \nu gt \\
 &+ \nu \int_{s=0}^{\infty} dZ(t-s) + \frac{\nu g \lambda}{2} - \nu \mu \left(t + \frac{\lambda}{2} \right).
 \end{aligned}$$

For equation A.7 to hold, the constant terms, the coefficients on t , and the coefficients on $dZ(t-s)$ on the two sides of the equation must be equal. Setting them equal leads to

$$\begin{aligned}
 (A.8) \quad A &= (g - \mu) \frac{\lambda}{2}; \\
 B &= g - \mu; \\
 q(s;\lambda) &= \nu + (1-\nu) \frac{1}{\lambda^2} \left[\int_{r=0}^{\min(s,\lambda)} (\lambda-r)q(s-r;\lambda) dr \right. \\
 &\quad \left. + \int_{r=0}^{\lambda} (\lambda-r)q(s+r;\lambda) dr \right].
 \end{aligned}$$

The last equation in equation A.8 defines $q(\bullet;\lambda)$ implicitly. We cannot find a closed-form solution for $q(\bullet;\lambda)$.⁶¹ Finding $q(\bullet)$ numerically, however, is straightforward. We do this by making an initial guess of $q(\bullet;\lambda)$, substituting this guess of the function into the right-hand side of equation A.7 (approximating the integrals numerically), thereby obtaining a new $q(\bullet;\lambda)$, and then iterating this procedure until $q(\bullet;\lambda)$ converges.⁶²

61. Because $q(s;\lambda)$ equals a constant plus a weighted sum of values of $q(\bullet;\lambda)$ with sum of weights less than one, a solution to equation A.7 exists and is unique.

62. The specifics of the algorithm are as follows. Since we know that $q(\bullet;\lambda)$ converges to one as s approaches infinity, we assume that $q(\bullet;\lambda) = 1$ for $s \geq T$, $T = 9$ (time is measured

All that remains to describe the behavior of the economy for a given λ is to translate the results concerning the behavior of newly set prices into a description of the behavior of the price level. Substituting equation A.6 and the solutions for A and B into equation A.4 leads to

$$(A.9) \quad p(t) = (g - \mu)t + \int_{s=0}^{\infty} w(s; \lambda) dZ(t-s),$$

$$w(s; \lambda) = \int_{r=0}^{\min(s, \lambda)} q(s-r; \lambda) dr.$$

The first line in equation A.9 is the formula for the aggregate price level in the text (equation 10). We compute $w(\bullet)$ using the second line.⁶³

As described in the text, the solution for $p(t)$ leads directly to a solution for the real effects of nominal shocks.

The Equilibrium λ

We now describe how we solve for the equilibrium interval between price changes, λ^E . Recall that λ^E is the solution to equation 13,

$$\left. \frac{\partial L(\lambda_i, \lambda^E)}{\partial \lambda_i} \right|_{\lambda_i = \lambda^E} = 0,$$

where $L(\lambda_i, \lambda)$ gives firm i 's loss as a function of its interval between price changes and the common interval of other firms. Using the formula for the loss function, equation 6, and assuming without loss of generality that p_i is set at time zero, we can write $L(\lambda_i, \lambda)$ as

in years). We then consider $q(s; \lambda)$ for $s = 0, h, 2h, \dots, T - 2h, T - h$, where $h = 1/52$; thus we divide time into periods of a week. We begin by setting all of these $q(s; \lambda)$'s to one. New values are found by substituting this initial guess into the right-hand side of equation A.7. An integral from $r = 0$ to $r = nh$ is approximated by averaging the values of the integral for $r = h, 2h, \dots, (n-1)h$ and then multiplying by nh . (As described below, we consider only values of λ that are integral multiples of h .) The new $q(s; \lambda)$'s are then substituted back into equation A.7, and another set of values is computed. We continue this process until the mean of the absolute changes in the $q(s; \lambda)$'s from one iteration to the next is less than δ , $\delta = 0.00002$. In all cases the algorithm appeared to converge to the solution without difficulty.

63. We approximate the integral using the procedure described in note 62.

$$(A.10) \quad L(\lambda_i, \lambda) = \frac{F}{\lambda_i} + \frac{1}{\lambda_i} \frac{1}{2} K \int_{s=0}^{\lambda_i} E [p_i^*(s) - p_i]^2 ds.$$

Differentiating equation A.10 with respect to λ_i and evaluating at $\lambda_i = \lambda$, we find

$$(A.11) \quad \left. \frac{\partial L(\lambda_i, \lambda)}{\partial \lambda_i} \right|_{\lambda_i = \lambda} = -\frac{F}{\lambda^2} - \frac{1}{\lambda^2} \frac{1}{2} K \int_{s=0}^{\lambda} E [p_i^*(s) - p_i]^2 ds \\ + \frac{1}{\lambda} \frac{1}{2} K E [p_i^*(\lambda) - p_i]^2.$$

(The Envelope Theorem allows us to eliminate the term

$$\frac{1}{\lambda_i} \frac{1}{2} K \frac{\partial}{\partial p_i} \left\{ \int_{s=0}^{\lambda_i} E [p_i^*(s) - p_i]^2 ds \right\} \frac{\partial p_i}{\partial \lambda_i} \text{ from } \frac{\partial L(\lambda_i, \lambda)}{\partial \lambda_i}:$$

because the firm chooses p_i optimally given λ_i , the effect of the marginal change in p_i when λ_i changes on the objective function is zero.)

Evaluating equation A.11 requires that we find $E [p_i^*(s) - p_i]^2$ for $0 \leq s \leq \lambda$. Equations A.1 and A.2 imply that the deviation of p_i from $Q(0)$, the average of prices set at the same time as p_i , is $\theta_i(0)$. Thus, using the solution for $Q(t)$ (equations A.6 and A.8), we obtain

$$(A.12) \quad p_i = (g - \mu) \frac{\lambda}{2} + \int_{r=0}^{\infty} q(r; \lambda) dZ(-r) + \theta_i(0).$$

Substituting the solution for $p(t)$, equation A.9, the process for $x(t)$, equation 8, and the identity $y = x - p$ into the formula for p_i^* , equation 5, we obtain

$$(A.13) \quad p_i^*(s) = (g - \mu)s + \int_{r=0}^{\infty} [v + (1 - v)w(r; \lambda)] dZ(s - r) + \theta_i(s).$$

Finally, equations A.12 and A.13 lead to

$$(A.14) \quad E [p_i^*(s) - p_i]^2 = (g - \mu) \left(s - \frac{\lambda}{2} \right)^2 + s \sigma_{\theta}^2 \\ + \sigma_x^2 \left\{ \int_{r=0}^s [v + (1 - v)w(r; \lambda)]^2 dr \right. \\ \left. + \int_{r=s}^{\infty} [v + (1 - v)w(r; \lambda) - q(r - s; \lambda)]^2 dr \right\},$$

where we have used the fact that changes in Z and in θ_i are uncorrelated both over time and with each other.

We can now solve numerically for the equilibrium value of λ . We begin by guessing a value of λ . After finding $q(\bullet; \lambda)$ and $w(\bullet; \lambda)$ as described above, we evaluate equation A.14 numerically for s between 0 and λ . We then use these results to evaluate the derivative in equation A.11. If the derivative is negative, so the representative firm can reduce its loss by increasing its interval between price changes, we raise our guess of λ ; if the derivative is positive, we reduce λ . We continue until the derivative converges to zero.⁶⁴

64. We find $E[p_i^*(s) - p_i]^2$ for $s = 0, h, 2h, \dots, \lambda - h, \lambda$, where, as before, $h = 1/52$. Since we assume $q(r; \lambda) = 1$ for $r \geq T$, and thus $w(r; \lambda) = 1$ for $r \geq T + \lambda$, we truncate the second integral in equation A.14 at $r = T + \lambda$. Otherwise the integrals in equation A.14 and equation A.11 are approximated as described above. We stop when we find an n such that equation A.11 is negative for $\lambda = nh$ and positive for $\lambda = (n + 1)h$; we then set λ at whichever value yielded the smaller absolute value of equation A.11. All of the functions involved appeared to be well-behaved and convergence occurred without difficulty.

Comments and Discussion

George A. Akerlof, Andrew Rose, and Janet Yellen: The Ball, Mankiw, Romer paper makes several contributions. First, it provides an excellent survey of the new Keynesian macroeconomics. Second, it develops further the implications of menu cost theory and thereby devises an ingenious test of the Keynesian model against the new classical alternative. The new classical economics and also real business-cycle theory are based on the dual assumptions of rational expectations and continual market clearing. As a result, according to these theories, unemployment is voluntary, and anticipated money shocks have no real effects. The new Keynesian economics is particularly welcome because it makes common sense and, unlike the new classical theory, it fits the most important stylized facts of the business cycle: that cyclical unemployment is largely *involuntary* and that anticipated demand shocks *do* affect real output.

The purported advantage of the new classical theory over Keynesian theory was its greater theoretical coherence. Economists found it difficult to build “optimizing models” in which money shocks, if anticipated, are not neutral; and they have found it equally difficult to explain theoretically how there can be equilibria in which labor markets do not clear. But these theoretical difficulties, which revealed more about the lack of imagination of economists than about the functioning of the real world, have now been solved. Efficiency wage models explain why labor markets need not clear. The menu cost–near-rationality models surveyed by the authors show that small transactions costs (or rule of thumb behavior that imposes second-order losses on its practitioners) result in significant monetary nonneutrality. As a result, changes in monetary policy, even if expected, may result in changes in equilibrium output and employment. This development is important because it means that

the Keynesian assumption that has been pejoratively labeled as “money illusion” involves no gross departure from rationality.

The authors not only survey this literature; they take it one step further. Lucas’s classic paper “Some International Evidence on Output-Inflation Tradeoffs” derives one implication of the Lucas-Phelps equilibrium model of the economy: the response of real output to an innovation in the growth rate of nominal GNP— τ —should decline with the *variance* of aggregate demand shocks.¹ Alternatively, the Phillips curve relating unemployment to unanticipated inflation should be steeper for countries where the variance of aggregate demand shocks is larger. In a dramatic example, Lucas showed that τ for the United States, where inflation is low, is 10 times its value for Argentina, where inflation is high.²

Ball, Mankiw, and Romer have shown that a menu cost theory of the frequency of price changes yields the identical prediction of a steeper Phillips curve (a smaller value of τ) as the variance of nominal income changes increases. In addition, however, menu cost theory predicts that increases in the level of inflation, or the mean rate of nominal income growth, should steepen the Phillips curve. Why? With staggered contracts, the higher the rate of inflation, the greater is the cost of infrequent price changes (since relative prices get out of line more rapidly). As a result, with menu costs, the optimal frequency of price changes increases as the inflation rate rises. And the greater the frequency of price changes, the steeper the Phillips curve. The authors thus show that Lucas’s observations are consistent with new Keynesian theory, but that new Keynesian theory generates an additional implication as well: the slope of the Phillips curve should become steeper (τ should fall) as inflation rises, for a given variance in nominal demand. The theory is well worked out and coincides with common sense. For example, indexation became more prevalent in the United States as inflation rose; countries with higher average inflation appear to have greater indexation of contracts.

It may be useful to mention, at least parenthetically, that equilibrium models of the business cycle have already failed many tests. First, quits

1. Robert E. Lucas, Jr., “Some International Evidence on Output-Inflation Tradeoffs,” *American Economic Review*, vol. 63 (June 1973), pp. 326–34.

2. See, for example, Richard T. Froyen and Roger N. Wurd, “Further International Evidence on Output-Inflation Tradeoffs,” *American Economic Review*, vol. 70 (June 1980), pp. 409–21.

rise and vacancies decline as the unemployment rate falls. On average, people appear to value their jobs more highly in bad times than in good, as is exactly consistent with the Keynesian view that sticky nominal wages cause jobs to pay higher rents in cyclical downturns than in expansions. Second, the demand shocks of World War II, the Vietnam War, and Mrs. Thatcher's successful policy of disinflation, as well as Mr. Volcker's similar policy, all produced changes in equilibrium output long after the policies were announced (and seen to be credible). The changes persisted sufficiently long after their announcement that it is extraordinarily difficult to believe that the changes in employment and output they caused were due to the slow propagation of unanticipated shocks, as would be consistent with Lucas's explanation.

The authors' contribution is to give an additional test of the Lucas model—a test that is in Lucas's exact framework. It depends on being able to differentiate the effect on the steepness of the Phillips curve that is due to the variance in nominal GNP growth from the effect that results from the level of inflation. Unfortunately, these two variables are highly correlated (the correlation equals 0.92). In addition, the slopes of the Phillips curves for different countries are difficult to estimate. As a consequence, the authors' test is inevitably a weak test of their theory against Lucas's theory, as well as against other Keynesian alternatives.

Ball, Mankiw, and Romer face three types of problems in their empirical analysis. First, there are difficulties in the way that they estimate the output-inflation trade-off for a given country. Second, the way in which they model the dependency of τ on inflation and the volatility of nominal income growth across countries is also problematic. Finally, their interpretation of the results may not be the only admissible explanation; there may be other reasons why the output-inflation trade-off varies across countries.

A number of problems potentially affect equation 14, the equation used to generate estimates of the output-inflation trade-off. Most important, τ is treated as a constant, when in fact it probably varies over time. Also, there is reason to believe that different Phillips curves specifications would generate different estimates of τ .

The authors' regressions, at least in one respect, fail to correspond to their theory. They assume that τ_i is *constant* for each country i in the regressions used to estimate the Phillips curve slope. A more accurate representation of their theory implies a Phillips curve in which τ_i varies

over time. In each period, τ_i depends on the expected inflation rate and the variance of inflation in the country at that time. Realistically, expected inflation and the volatility of nominal GNP growth have probably not been constant during the sample period in most countries; indeed, this is the original motivation for Engle's work on ARCH (autoregressive conditional heteroscedasticity).³

If the conditional mean and variance of nominal demand growth were not constant over time, then the optimal interval between price changes would also have changed in each country over time; consequently, τ_i could not be treated as a constant for a given country. Quarterly postwar U.S. inflation is an extremely autoregressive variable (the Q -test for the hypothesis of a flat twelfth-order correlogram is over 350). Cecchetti provides evidence for the United States that suggests that as inflation and the mean expected rate of inflation rose during the seventies, the interval between price changes fell.⁴

If one assumes that expected inflation in each country was not constant over the sample period, a more natural way of estimating σ_x might be to use the standard deviation of the residual in an autoregression predicting Δx . Grubb, Jackman, and Layard have calculated such a measure of σ_x for a subsample of Ball, Mankiw, and Romer's countries by taking the residual variance in an AR(2) predicting Δx . Their measure of σ_x has a correlation of only 0.27 with the authors' estimate of σ_x .

The Phillips curve can be estimated in many ways. One would like the estimates of τ_i to be insensitive to precise specification. Grubb, Jackman, and Layard have estimated an alternative measure of the slope of the Phillips curve using a model that distinguishes between real and nominal wage rigidity.⁵ The correlation of their measure of nominal wage rigidity and Ball, Mankiw, and Romer's measure of τ_i across countries is 0.30. This suggests that the theoretical constructs needed for the authors' test are both difficult to pinpoint and sensitive to a wide range

3. Robert F. Engle, "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, vol. 50 (July 1982), pp. 987–1007.

4. Stephen G. Cecchetti, "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines," *Journal of Econometrics*, vol. 31 (August 1986), pp. 255–74.

5. Dennis Grubb, Richard Jackman, and Richard Layard, "Wage Rigidity and Unemployment in OECD Countries," *European Economic Review*, vol. 21 (March–April 1983), pp. 11–39.

of specification choices (we should note, however, that when equation 5.6 is estimated using the measures of Grubb, Jackman, and Layard, the authors' conclusions still obtain). It also suggests that a more efficient way to estimate τ would be to estimate equation 14 pooling the data across both countries and time, using, for example, a fixed-effects panel data estimator; this would also alleviate problems with low degrees of freedom. Further, the dependency of τ_i on π_i and σ_{xi} could be nested directly within equation 14.

Nonlinearity plays a significant role in the authors' second-stage regressions between τ and π (among other variables). It is therefore worth noting that the functional form of equation 14 implies that the short-run Phillips curve relating the inflation rate to the *level* of real output is concave rather than convex, as is commonly assumed. That is, according to equation 14 the short-run Phillips curve is assumed to become flatter rather than steeper as Y rises.

Many macroeconomists have worried recently about unit roots in macroeconomic time series. While it is hard to believe that Δx has a unit root (standard Dickey-Fuller tests easily reject the hypothesis for post-war quarterly American data), many economists (Mankiw among them) have shown that the behavior of real output is consistent with the existence of a unit root. Unless the lagged dependent variable has a coefficient of unity in equation 14, the unit root in real output implies a random walk residual term in equation 14. If this residual, which presumably represents "unimportant" omitted variables, is not orthogonal to Δx , then estimates of τ will be biased and inconsistent.

The authors' equation 5.6 models the variation of τ across countries as a function of the average values over time of π and σ_x and their squares. The authors include the squares so that τ_i can depend on π_i and σ_{xi} in a nonlinear fashion; however, this choice of functional form has several unfortunate consequences. In particular, for moderate values of π_i , the relationship between τ_i and π_i is ambiguous and generates negative fitted values for τ_i , so that the quadratic functional form is not very successful, either in generating a negative monotonic relationship between τ and π or in preventing the prediction of a negative τ_i . Further, the effects of innovations in σ_{xi} are insignificant and incorrectly signed. However, the quadratic functional form seems to be critical in establishing that the level of inflation affects the output-inflation trade-off.

The functional form that the authors use to relate π and τ has the disturbing property that the relationship between π and τ potentially changes sign, while theory suggests an unambiguous negative relation between these two variables. When the authors estimate equation 5.6 for only the OECD countries without Iceland, it is worrisome to note that the sign of $d\tau/d\pi$ turns positive at a 7.2 percent inflation rate, which is very close to the midpoint of inflation for the OECD sample (with Iceland, the sign becomes positive at 9.2 percent inflation).

A related problem concerning the specification of equation 5.6 is that the fitted value of τ reported in table 8 is negative for an inflation rate between 10 percent and 15 percent. The variable τ is presumably positive, despite the fact that in the sample, τ is estimated to be negative for one-quarter of the countries; the posited nonnegativity of τ also suggests that equation 5.6 should have been estimated using some sort of Tobit-like procedure, incorporating the fact that τ is measured with error (Tobit cannot be used without modification, given the presence of negative τ_i terms).

The authors argue for the inclusion of the square terms in equation 5.6 because τ is always positive and therefore the τ equation should have a nonlinear form. A comparison of equations 5.6 and 5.5 and of equations 5.3 and 5.2 shows how very important including the quadratic inflation term is in establishing the statistical relevance of inflation for τ . For the results to be considered robust, alternative nonlinear specifications should also yield the result that π has a significant effect on τ . Perhaps a simpler nonlinear form consistent with the nonnegativity of τ would be:

$$1/\tau = a + b\pi + c\sigma_x.$$

We estimated this equation with their data and obtained the following results (with standard errors in parentheses):

$$1/\tau = 8.95 - 621\pi + 806\sigma_x$$

(2.5) (416) (536)

$$\bar{R}^2 = 0.06; \text{ standard error} = 111; \text{ number of observations} = 43.$$

Note that the coefficient of π is statistically insignificant and also has the wrong sign. The F -test of the hypothesis that inflation and its square do not enter the regression when both σ_x and its square are included also indicates that inflation does not determine τ [$F(2, 37) = 1.71$].

The precision with which τ is estimated in equation 14 varies country by country; however, all τ_i observations are weighted equally in equation 5.6. The resulting heteroskedasticity in equation 5.6 leads to inefficient estimates, and also could lead one to conclude, for example, that π is significant in equation 5.6, when in fact it is not.

The authors interpret their finding of a negative relation between τ and π as reflecting more frequent price changes due to menu costs at higher levels of inflation. Alternative interpretations are possible. For example, Keynesian theory has traditionally assumed that Phillips curves become steeper as the economy approaches full employment because nominal wages are sticky downward but not upward. Countries with high values of inflation may be those with more of a taste for high employment, and therefore may have been more frequently on the steeper part of the short-run Phillips curve. Implicitly, the authors assume that their countries were at the natural rate of unemployment on average over the sample period, so that differences in τ across countries reflect changes in the Phillips curve slope at the same level of output. This assumption may be particularly difficult to maintain in the shorter periods. Movements along a given Phillips curve, rather than shifts of Phillips curves, may explain why the authors find a correlation between changes in inflation and changes in steepness of Phillips curves in the period through 1972 and the post-1972 period.

The country-by-country changes in the steepness of the Phillips curves (countries with greater increases in inflation had greater increases in the Phillips curve slope) are perfectly consistent with the authors' hypothesis that higher inflation induces more frequent price changes. But the period through 1972 and the post-1972 period were quite different structurally, and other explanations abound. Not only were there increases in inflation between the two periods, but also in the later period there were more supply shocks, less dependence on fixed exchange rates, a rise and then leveling off of developing country debt, a rise in developed country unemployment, and slower productivity growth. Any of these factors could impart a bias to the correlation between τ_i and π_i if it has a differential impact on countries with different rates of inflation. For example, the switch from fixed to floating exchange rates can be expected to shift and perhaps steepen the Phillips curve. With fixed rates, wage earners may be restrained in their bargaining by the thought that an increase in wages will result in balance of payments deficits and contrac-

tionary demand policies. It is extremely unlikely that the data will allow us to discriminate between the hypothesis of menu costs and alternative explanations for the changes in τ .

An important assumption of the authors' test is that Δx is a process that can be considered to be exogenous in equation 14. Clearly this assumption is special. If Δx is correlated with the error term in equation 14—that is, if changes in nominal demand are correlated with supply shocks—biases will be generated in the estimation of the τ 's (the slopes of the Phillips curves). The authors have argued that there is no special reason for such biases to occur since there is no particular reason for any bias in the τ 's to be correlated with mean inflation.

The only way to see whether such a bias is likely, or likely to be serious, is to propose an alternative model, to see whether significant simultaneity bias will occur, and to test whether the model fits the data. We shall therefore consider an alternative model in which income is purely random, so that there is *no* Phillips curve by construction. Not only does such a model generate the correlations found by the authors; in fact, it offers a superior explanation of the data.

Suppose that real income in country i is a white noise process (so there is no Phillips curve whatsoever):

$$y_{it} = \epsilon_{it} \quad \epsilon_{it} \sim \text{iid } (0, \mu_i^2).$$

Assume that inflation is also a white noise process, independent of ϵ :

$$\pi_{it} = \eta_{it} \quad \eta_{it} \sim \text{iid } (0, \sigma_{\pi i}^2).$$

Then

$$\Delta x_{it} = \pi_{it} + \Delta y_{it} = \eta_{it} + \epsilon_{it} - \epsilon_{it-1}.$$

If one runs the regression:

$$y_{it} = \tau_i \Delta x_{it},$$

the plim of τ_i is:

$$\tau_i = \mu_i^2 / (2\mu_i^2 + \sigma_{\pi i}^2) = \frac{\sigma_{xi}^2 - \sigma_{\pi i}^2}{2\sigma_{xi}^2}.$$

Clearly, τ_i gives, even asymptotically, a biased estimate of the responsiveness of real output to an innovation in nominal GNP growth, which is zero in this example. The value of τ_i depends (positively) on the

author's variable σ_{xi} (the standard deviation of Δx in country i) since σ_{xi}^2 is just $2\mu_i^2 + \sigma_{\pi i}^2$, and also (negatively) on $\sigma_{\pi i}$, the standard deviation of the inflation rate in country i . However, τ_i does not depend on $\bar{\pi}_i$, the mean inflation rate in country i .

Consider, however, the slightly modified process:

$$\pi_{it} = \bar{\pi}_i + \eta_{it} \quad \eta_{it} \sim \text{iid}(0, \sigma_{\pi i}^2).$$

In this case, the estimates of τ_i will be unchanged. If $\sigma_{\pi i}^2$ and $\bar{\pi}_i$ are positively correlated, regressions of τ_i on $\bar{\pi}_i$ and σ_{xi}^2 will likely yield a negative coefficient on $\bar{\pi}_i$, because of its positive correlation with the omitted variable $\sigma_{\pi i}^2$. Such a positive correlation between $\sigma_{\pi i}^2$ and $\bar{\pi}_i$ is in fact consistent with the authors' data, for which the sample correlation of $\bar{\pi}_i$ and $\sigma_{\pi i}^2$ is 0.92.

A test of the authors' process against the null hypothesis postulated here (in which there is no Phillips curve) can be conducted by including both $\bar{\pi}_i$ and $\sigma_{\pi i}$ in equation 5.6. The example just presented implies that the coefficient on $\sigma_{\pi i}$ should be negative and the coefficient on $\bar{\pi}_i$ should be zero; the opposite is true if the authors' hypothesis is relevant.

In fact, adding $\sigma_{\pi i}$ and $\sigma_{\pi i}^2$ to the authors' key regression (equation 5.6) shows the two new variables to be very significant while the coefficients of $\bar{\pi}_i$ and its square are insignificant:

$$\tau = 0.441 - 0.332\bar{\pi} + 0.181\bar{\pi}^2 + 25.58\sigma_x - 48.48\sigma_x^2 - 27.42\sigma_{\pi} + 49.93\sigma_{\pi}^2$$

(0.09) (2.08) (3.79) (5.67) (24.62) (5.93) (22.81)

$$\bar{R}^2 = 0.592; \text{ standard error} = 0.175; \text{ number of observations} = 43.$$

The F -test of the hypothesis that the coefficients of σ_{π} and σ_{π}^2 are both zero is resoundingly rejected [$F(2, 35) = 11.5$]. On the other hand the F -test of the hypothesis that both the $\bar{\pi}$ term and its square are zero is quite consistent with the data [$F(2, 35) = 0.05$]. Note the dramatic improvement in the fit of the equation in comparison with the authors' equation 5.6. Parenthetically, both the authors' results in table 11 and Lucas's results are similarly affected by inclusion of σ_{π} (and its square, where relevant).

We do not wish to say that we prefer our null hypothesis to that of the authors. We find their theory plausible and their tests consistent with the theory. We only wish to point out that their test is weak. Our example shows that one could believe that τ is actually independent of $\bar{\pi}$, yet still be able to explain the authors' results.

The authors have provided an extremely clear and valuable survey of the new Keynesian microeconomic foundations. They have also developed an interesting hypothesis regarding the frequency of price changes in new Keynesian models. The faster prices rise, the more frequently prices are changed, and hence the smaller the real effects of nominal demand shocks; the output-inflation trade-off should be negatively affected by the rate of inflation.

The prediction is intuitively believable and testable with microeconomic data. We are sympathetic with both the spirit and the conclusions of this paper. However, aggregate data just do not seem to discriminate between the hypothesis of interest and very different alternatives. This does not indicate that the authors are incorrect; it merely shows how ambiguous the results of aggregate regressions tests of subtle hypotheses almost always are.

Christopher A. Sims: Laurence Ball, Gregory Mankiw, and David Romer discuss a developing literature that aims at making price rigidity and money illusion respectable again in macroeconomics. Anyone who has bought from a catalog or served on a university departmental salary-setting committee knows the world is full of price rigidity and money illusion. It is therefore important to find a way to let macroeconomists discuss it without embarrassment.

Nonetheless I disagree with some of the assessments in the paper as to what has been achieved in this line of work and as to what the most promising lines of advance are. Part of my disagreement stems from a different view of how Keynesian economics and its market-clearing rational expectations antithesis succeeded. The paper asserts that Keynesian economics fell into disrepute because it could not provide a justification based on optimizing behavior for its assumption of price rigidity. This has the matter almost backwards. It is one of the accomplishments of Keynesian economics and a main reason for its original appeal that it proceeded boldly to legitimize money illusion and nonclearing markets as primitive assumptions in economic models. At the time it was obvious not only that these things existed, but that they were quantitatively important to the functioning of the economy. The fact that microeconomics then had no way to discuss them was a reason for rejecting any insistence that macroeconomics be built up from microeconomic first principles. It was and still is a strength of Keynesian

economics that it does not make a fetish of modeling based on optimizing behavior.

Rational expectations macroeconomics put at center stage the reaction of people to the uncertainty in their economic environment. While Keynes wrote at length about the importance and volatility of expectations, he paid little attention to the endogeneity of price and wage expectations. In the high-inflation, high-unemployment environment of the 1970s, it became clear that endogeneity of expectations, always as evident to the naked eye as money illusion, was also quantitatively important to macroeconomics. Precisely the nominal rigidities that Keynesianism legitimized as primitive assumptions seemed to be shifting in response to policy as rational expectations theory suggested.

But like Keynesianism's strategy of freely postulating price rigidities, rational expectations' coherent tracing out of the implications of dynamic optimization under uncertainty is both its central strength and its main weakness. It is still as obvious that people do not dynamically optimize as it was when Keynes wrote. And while rational expectations could explain why unemployment could rise even in the presence of inflation, it could not explain why it should have stayed so high for so long, either recently or during the Great Depression.

Keynesianism declined because it failed to offer a *true* explanation about how nominal rigidities respond to the economic environment, not because it failed to offer an explanation based on optimizing behavior.

It may not be the time yet for a new wave in macroeconomic theory, but if it is, I expect that it will come from rediscovering the Keynesian willingness to confront irrationality and market failure, while restricting the license Keynesianism seems to give to ad hockery. This will require going beyond rational expectations, to consider the technology of calculation and communication. Useful theory that really does this is not yet available. However, the idea of models in which individual behavior deviates from rationality in amounts and directions that cost the individual little is in this spirit and is promising.

From this point of view, a new Keynesian model should be confronted with three questions. Where is the irrationality in the model? Why there and not somewhere else? Would this suboptimal behavior plausibly hold still under the shifts in environment that are being considered? Unlike the authors, I view it as no particular virtue that irrationality be confined to a narrow aspect of behavior, with agents in the model perfect

calculating machines in every other respect. It is a virtue, though, for a model to make explicit how its irrationality shifts or disappears when it becomes costly.

I don't think the authors' model does very well by these tests. It takes as an unexamined assumption the idea that it is costly to change the level of prices. My view is that this is true in low-inflation economies because of the convenience in relying on aggregate price stability in making contracts. But if there were such a thing as an economy with a rock-solid inflation rate of 40 percent, plus or minus 2 percent, per year, institutions would surely adapt, so that prices would be announced in catalogs and wage contracts with smooth growth paths paralleling the smooth aggregate price path. Nominal rigidity would set in about this price path in much the same form as we see around the zero inflation rate in low-inflation economies.

The data presented by the authors in fact provide no support at all for their claim that the level of inflation affects the degree of price rigidity. The central problem is that the leading alternatives to their theory all agree that the *variability* of inflation should be related to the strength of the relation between nominal and real output. In the authors' sample, the standard deviation and the level of inflation have a correlation of 0.92. The question is then whether they have found anything more than the tendency—expected from neoclassical equilibrium theories—for economies with highly variable inflation to have a lower regression coefficient for output on aggregate demand.

The authors put aside any direct test of this proposition by pointing to the identity, equation 16, in the paper. The equation shows that in any given country, τ can be found as a function of σ_π^2 , σ_x^2 , and a third number σ_ϵ^2 . However, this identity does not imply any identity in the relation between τ , σ_π^2 , and σ_x^2 , because there is no more reason to suppose that σ_ϵ^2 is constant across countries than that any of the other terms in equation 16 is constant. The authors are right in pointing out that if $\bar{\pi}$ has no additional explanatory power once both σ_x^2 and σ_π^2 are entered in the regression, their theory can accommodate that result by postulating that “supply shocks” (ϵ in the paper's equation 15) have the same variance in all countries. But the proposition does not seem particularly attractive a priori.

I have verified that in both equations 11.6 and 5.6 the addition of the standard error of inflation as a regressor makes the coefficients on mean

inflation statistically insignificant at any reasonable significance level. Thus there is no statistical evidence for a relation of τ to $\bar{\pi}$, once the effects of σ_{π}^2 and σ_x^2 are both accounted for. The data are consistent with fairly large effects of $\bar{\pi}$ in such a regression, because of the high collinearity of $\bar{\pi}$ with σ_{π}^2 , but they give more significant coefficients to the other variables than to $\bar{\pi}$. The authors are then consistent in claiming that there is no strong evidence against their theory in the data, but there is also no evidence in favor of it for anyone who believes that the variability of inflation affects the response of output to demand.

The authors argue that there is no plausible theory that suggests τ should be a function of both σ_x and σ_{π} . There are really two classes of modern market-clearing theories with which to compare the authors' model. One, "Lucas I," is represented by Robert Lucas's 1973 paper mentioned by the authors. That model suggests that high variance in the aggregate price level, by making it more difficult to distinguish relative price movements, reduces the response of output to a given change in the aggregate price level. Because in that theory the basic relationship is between the aggregate price level and aggregate output, the coefficient in a regression of output growth on nominal output growth would be expected to vary with the variance of inflation.

The other class of alternative theories, represented by a 1987 paper by Lucas and Nancy Stokey, suggests that only anticipated changes in the price level (which influence nominal interest rates) have real effects, and that unanticipated variation in money stock is reflected one-for-one in unanticipated variation in prices that has no causal influence on output.¹ This latter type of theory probably best represents the generic implication of rational expectations. Rational expectations models can deliver the "Lucas I" type of causal effect of nominal surprises on real variables only by sustaining differences in information across agents in the economy in equilibrium. This requires special assumptions difficult to justify from the rational expectations perspective.

Correlations between prices and output could take varying forms in a Lucas-Stokey type of model, and certainly in this type of model τ could depend on both σ_x and σ_{π} . Unpredictable monetary policy fluctuations would generate price variation with no corresponding real fluctuations.

1. Robert E. Lucas, Jr., and Nancy L. Stokey, "Money and Interest in a Cash-in-advance Economy," *Econometrica*, vol. 55 (May 1987), pp. 491-513.

A simple version of such a model, then, might have real output approximately a random walk and price level an independent random walk. (The random walk behavior of prices would emerge if the monetary authority aimed at stabilizing the inflation rate, but did so erratically.) Since nominal aggregate demand in the Ball-Mankiw-Romer data is price change plus output change, this simple model would imply in equation 14 that $\gamma = 0$, $\lambda = 1$, and $\tau = \sigma_y^2/\sigma_x^2$. Regression of the authors' estimated τ 's on a constant and the "theoretical" σ_y^2/σ_x^2 alone produces an \bar{R}^2 of 0.74, compared with a maximum \bar{R}^2 in table 5 of 0.39. This highly simplified classical theory can be rejected as an exact theory by further tests, but it is a reasonable rough approximation. There appears to be correlation of price variability with output variability, in other words, that is consistent with more complicated market-clearing theories, but there is no evidence of deviation from the simple theory in the direction favored by the authors—dependence of τ or σ_y on $\bar{\pi}$.

General Discussion

Several panelists questioned the menu cost explanation of price stickiness. Martin Baily suggested that menu costs are fixed costs, much like advertising. Therefore large firms, for example firms in concentrated industries, have the most to gain by paying the menu cost and adjusting their price. Yet a reading of the menu cost literature would lead one to the opposite conclusion—that firms in a concentrated industry do not have flexible prices because their profit function is flat for a wide range around their optimum price. Baily also felt that a menu cost model could not explain why prices adjust quickly to changes in costs, as they did when oil prices rose in 1973, but not to changes in demand. He suggested that a model was needed to explain sticky markups of price over cost rather than sticky prices. Robert Gordon suggested that menu costs could not easily explain the stickiness of prices in mail order catalogues where the menu cost is clearly the cost of printing a catalog. That cost is the same however many prices are changed. Yet, typically, only some fraction of prices are changed in successive issues of a catalog, suggesting that something other than menu costs is keeping prices sticky.

Robert Hall acknowledged that the evidence on the stickiness of magazine and mail order prices is impressive, but argued that most prices

are much more flexible. No one buys a car without haggling over the price, and department store prices can be changed every day. Thus he was skeptical that the U.S. economy can be characterized as a collection of firms that set price and meet demand at that price. Gordon felt that Arthur Okun's distinction between auction markets and markets with a preset price had been overlooked in the authors' paper. He thought that the examples Hall had in mind would lie somewhere between the extremes identified by Okun.

There was general discussion of the distinction between two types of models of sticky wage and price dynamics: the periodic-review model and the continuous-review model in which the time intervals between price or wage changes are random. John Taylor's wage model is of the periodic type in which, because of administrative costs, wages are changed at fixed intervals. Taylor found that the average fixed interval between wage setting is 12 months, asynchronized across firms. The authors' menu cost model is of this type, but with the length of the periodic review determined endogenously. The Caplin and Spulber model is an example of the continuous-review type; a firm raises its price whenever it gets a certain distance away from the profit-maximizing level. Robert Hall argued that the Caplin-Spulber model is more natural than the authors' model, and, in sharp contrast, generates neutrality of money. The average price level increases in proportion to the increase in money, as larger price increases by some firms are just balanced by unchanged prices in other firms. Hall felt that the authors' model dismisses this neutrality result by an ad hoc assumption that firms change prices at fixed intervals. Ball replied that a model with fixed timing was more realistic, since wages are almost always adjusted at fixed intervals and, according to Rotemberg's findings, prices are most often adjusted in January. He also noted that the Caplin and Spulber neutrality results required rather strong assumptions, such as uncorrelated money shocks. Gordon cited Arthur Okun's evidence that FIFO pricing is more popular than LIFO pricing as evidence that there are large costs to management's continuously overseeing price changes.

Edmund Phelps noted that the empirical evidence in the authors' paper showing that prices are changed more frequently in high-inflation countries is consistent with either type of model. He saw no advantage to the more complicated, endogenous-interval menu cost models of wage and price setting for the United States, where inflation has always been

comparatively low. According to Phelps both models have the same failing: they do not explain why wage and price indexing is not more common.

Hall remarked that the authors' model was an "old Keynesian" model as opposed to the new Keynesian model of Greenwald and Stiglitz also in this volume. The Ball-Mankiw-Romer model emphasizes rigid nominal prices that make monetary disturbances the driving force of the business cycle. The Greenwald and Stiglitz model assumes many disturbances to the system and emphasizes how these disturbances get amplified by the structure of the economy. Monetary nonneutrality enters into the new Keynesian model only through possible distributional effects. Hall preferred this type of model because he found it unrealistic to assume that prices are rigid.

Ball disagreed with Christopher Sims's comment that the success of Keynesian economics comes from acknowledging irrational behavior. Ball argued that, as emphasized by Akerlof and Yellen, new Keynesian models give the insight that a very little irrationality can produce Keynesian results. The amplification of a little irrationality is due to an externality in the pricing decision of a monopolistic competitor. It is often the case that a firm gains little by lowering its price unilaterally in response to weak demand, because the increased profits from greater sales are almost completely offset by the loss from the lower price. However, through the real balance effect, demand at a given price, and hence the firm's profits, would be increased substantially if all other firms would lower their price. Baily noted that this argument that firms are more likely to lower their price if other firms lower theirs is contrary to the standard model of the firm in an oligopolistic industry. The representative firm in such an industry will make more profits from lowering its price if other firms in the industry do not change theirs. If the firms in the authors' model are reinterpreted as industries, with the externalities being between industries, firms in an industry have to be regarded as colluding perfectly.

James Tobin felt that many of the participants in the current debate about menu costs misrepresent Keynes and Keynesian economics. Both Keynes and the classical economists with whom he debated acknowledged that the economy would often be driven out of equilibrium by shocks to aggregate demand. That is, neither side believed in a Walrasian auctioneer who would continuously find wages and prices to clear all

markets. The issue that divided Keynes and the classical economists was whether the market mechanisms of prices, wages, and interest rates could return the economy to a full-employment equilibrium in a reasonable time. Keynes believed the market mechanism would be slow at best, and only one of his reasons involved nominal wages that are sticky downward. The more important reason is that real aggregate demand is neutral to the level of prices and wages. The real balance effect countered the second explanation, but the effect of falling prices on real interest rates could easily cancel out the real balance effect. Tobin doubted that the Depression would have ended sooner if prices and wages had fallen even faster. Explaining rigid wages and prices is therefore not the most important task of Keynesian macroeconomics.

There were several comments on the labor market in the authors' model. Hall emphasized Ball and Romer's earlier finding that real rigidities, such as an elastic labor supply locus, are crucial for the effect of small menu costs to be amplified. He felt that it was misleading for the authors to criticize Barro's assumption that labor supply is elastic, when they required an elastic supply locus themselves. They should have criticized Barro for assuming a competitive labor market, which makes elastic labor supply implausible, whereas the efficiency wage model generates elastic supply quite naturally.

Baily questioned how it was possible that the costs to firms of price or wage rigidities could be small if they resulted in layoffs and unemployment of their workers. Assuming that there is a large gap between the wage and the value of time to an unemployed worker, there is a first-order loss if firms do not cut their wage and price with a shift in demand. Baily reasoned that wage rigidity is not evidence against this being a large cost; if wages and employment are part of an implicit contract, firms that do not adjust their prices so as to maintain employment will have to compensate workers with a higher expected wage.