STATE-DEPENDENT PRICING AND THE GENERAL EQUILIBRIUM DYNAMICS OF MONEY AND OUTPUT*

MICHAEL DOTSEY
ROBERT G. KING
ALEXANDER L. WOLMAN

Economists have long suggested that nominal product prices are changed infrequently because of fixed costs. In such a setting, optimal price adjustment should depend on the state of the economy. Yet, while widely discussed, state-dependent pricing has proved difficult to incorporate into macroeconomic models. This paper develops a new, tractable theoretical state-dependent pricing framework. We use it to study how optimal pricing depends on the persistence of monetary shocks, the elasticities of labor supply and goods demand, and the interest sensitivity of money demand.

I. INTRODUCTION

If there are fixed costs to changing prices, the timing and magnitude of an individual firm's price adjustment depends on the state of the economy. The effects of nominal disturbances on aggregate real activity will, therefore, also be state-dependent, since the price level depends on the fraction of firms that adjust and the prices that these firms set.

While state-dependent pricing is intuitively appealing, its modeling has proved technically difficult, forcing prior analyses to focus on special cases. Some researchers work in static models (e.g., Blanchard and Kiyotaki [1987] and Ball and Romer [1990]). Others analyze dynamic price setting, but make restrictive assumptions about forcing processes and macroeconomic equilibrium (e.g., Benabou [1992] and Caplin and Leahy [1991]). Although this previous work made important theoretical advances, the resulting models are ill-suited for empirical applications and policy analysis.

In this paper we make state-dependent pricing broadly

* This research has been improved by contributions from Marianne Baxter, Ricardo Caballero, Dorsey Farr, Marvin Goodfriend, Andrew John, John Leahy, David Romer, Julio Rotemberg, Stephanie Schmitt-Grohe, Brian Scholl, Alan Stockman, Julia Thomas, and Michael Woodford, and from the comments of seminar participants at Yale University, the University of Pennsylvania's Wharton School, the Federal Reserve Banks of Richmond and San Francisco, the University of California at Berkeley, the University of Virginia, and the 1996 NBER Summer Institute Monetary Economics Workshop. The views expressed here are solely those of the authors and do not necessarily reflect the views of the Federal Reserve System or the Federal Reserve Bank of Richmond.

The Quarterly Journal of Economics, May 1999

operational, providing a theoretical approach that can be used in a wide variety of macroeconomic models and for a wide range of applied purposes. At the core of our model is a real monopolistic competition setup, based on that in Blanchard and Kiyotaki [1987] and Rotemberg [1987], in which firms have the power to set prices. We posit that the fixed costs of changing prices are random across firms and drawn from a continuous distribution so that the resulting macroeconomic equilibrium involves some but not all firms opting to adjust. Since the fixed adjustment costs of individual firms are assumed independent over time, adjusting firms all choose the same price, as in many time-dependent models of price adjustment. As in such models, a vector of predetermined prices becomes part of the relevant history of the economy. In addition, our framework makes a distribution of firms—indexed by when they last adjusted price—part of the endogenously evolving state of the economy. Nevertheless, our framework leads to a manageably sized state space. When the driving shocks are sufficiently small so that the dynamic equilibrium can be approximated linearly around an inflationary steady state, the nature of state-dependent price dynamics can easily be analyzed and readily compared with the extensive prior work on time-dependent pricing.[1]

Time-dependent pricing models, such as those constructed by Taylor [1980] and Calvo [1983], generally imply that an unanticipated permanent change in money will temporarily affect aggregate real activity but will ultimately alter only the price level. As Caplin and Leahy [1991] point out, however, time-dependent models also imply that "between [exogenously specified] price adjustments firms are not allowed to respond even to extreme changes of circumstance. This makes it difficult to know whether the qualitative effects of money in these models are the result of nominal rigidities per se or of the exogenously imposed pattern of price changes." This criticism is forceful because Caplin and Leahy [1991] find that monetary shocks affect either output or the price level, but not both, in a general equilibrium setting with

1. One cost of our current focus on linear approximation dynamics is that there is no history dependence: a monetary shock has the same effect on output and prices if the economy is in an expansion or contraction. To study the accuracy of our linear approximation and to evaluate the quantitative importance of state dependence, it will be necessary to use nonlinear solution methods. In such a setting, the fact that our theoretical approach leads to a relatively small number of state variables will be essential since nonlinear solution methods rapidly encounter the "curse of dimensionality."

state-dependent pricing. Further, this real or nominal effect is expected to be permanent, in the sense of a limiting impulse response. However, to generate a rational expectations equilibrium, Caplin and Leahy [1991] make many other strong assumptions about the nature of the economy and the money supply process.[2] While their unusual results on the nonneutrality of money are provocative, it is therefore an open question whether these results are inherently related to state-dependent pricing or are artifacts of other special features of their framework.

We find that the time-dependent approach captures the main mechanisms that lead to monetary nonneutrality under state-dependent pricing, although the magnitude of nonneutralities is often overstated when adjustment timing is assumed invariant to shocks.[3] The conventionally defined long-run effect of money on output—in the sense of a limiting impulse response—is zero here, as in many other macroeconomic models. In these responses, however, the evolving distribution of price setters plays a central role in dictating how monetary shocks affect the dynamics of prices and real activity. Finally, elements outside the price block—such as factor market equilibrium, the nature of the money supply rule, and the response of interest rates—that are suppressed in other work on state-dependent pricing play an important quantitative role in the analysis.

The paper proceeds as follows. Section II introduces the dynamic macro model and then describes its steady-state characteristics. Section III describes the model's responses when it is subjected to money supply shocks. Section IV concludes with a summary and suggestions for future research.

II. THE MODEL

The model has three main elements. First, following Blanchard and Kiyotaki [1987] and Rotemberg [1987], we specify that households value a range of consumption goods and that firms are monopolistically competitive suppliers of differentiated products.

2. In particular, Caplin and Leahy [1991] require that the money supply is a continuous time random walk with zero drift (so that the average inflation rate is zero), that real marginal cost is independent of the level of output, and that money demand is interest insensitive. In Section III we contrast our results with those of Caplin and Leahy [1991] and Caplin and Spulber [1987].

3. Generally, the steady-state (average) pattern of price adjustment—which is sometimes used to calibrate time-dependent models—is a misleading guide to the marginal pattern of adjustment that is relevant for considering shocks to exogenous variables.

Households and firms otherwise behave competitively in markets for labor and for credit. Second, we make the conventional assumption that firms face fixed costs of adjusting their nominal prices and that they satisfy all demand at posted prices. In contrast to Blanchard and Kiyotaki and others, however, we assume that the fixed cost is random and varies across firms according to a continuous distribution. Third, we assume that households have a demand for money that takes a conventional semi-logarithmic form, to facilitate comparisons with the literature.[4] The model abstracts from the process of physical capital accumulation, although there is a fixed capital stock held by households and allocated each period among firms. This simplification makes the model's dynamics relatively transparent and keeps the basic structure close to Blanchard and Kiyotaki and Ball and Romer [1990].

II.1. The Underlying Real Monopolistic Competition Setup

The households in the model are identical infinitely lived agents who value the many different consumption goods produced by firms. These households demand consumption goods and supply factors of production on a competitive basis. The firms in our economy are monopolistically competitive in goods markets and competitive demanders in factor markets.

Households. The households' preferences for goods and leisure are represented by the standard time separable objective,

$$E_t \sum_{j=0}^{\infty} \beta^j u(c_{t+j}, l_{t+j}), \tag{1}$$

with the momentary utility function given by $u(c,l) = (1/(1-\sigma)) \cdot c^{1-\sigma} - \chi(1-l)^{\zeta}$. As in Blanchard and Kiyotaki [1987], the consumption good that enters agents' utility is a Dixit-Stiglitz aggregate of many individual goods. There is a continuum of goods, so the consumption index (c_t) is related to the component goods $(c_t(z), z \in$

4. Our results would not change in any important way if money demand were derived from household optimization given an explicit cash-in-advance constraint, or a "shopping time" technology that involved a plausible resource cost level. The reduced-form approach allows us to isolate the role of interest sensitivity, without introducing any "shoe leather" considerations.

[0,1]) according to

$$c_t = \left[\int_0^1 c_t(z)^{(\epsilon-1)/\epsilon}\, dz\right]^{\epsilon/(\epsilon-1)}. \tag{2}$$

Cost minimization on the part of consumers thus implies that demand for the zth good is $c_t(z) = (P_t(z)/P_t)^{-\epsilon}c_t$, where P_t is an index of the cost of buying a unit of c_t:

$$P_t \equiv \left[\int_0^1 P_t(z)^{1-\epsilon}\, dz\right]^{1/(1-\epsilon)}. \tag{3}$$

It is simplest to assume that households directly own the two factors of production, labor (n) and capital (k), and rent these to firms. Households also own a diversified portfolio of claims to the profits earned by the monopolistically competitive firms. We are not explicit about the pattern of assets available to households, since our macroeconomic equilibrium will imply that there is no accumulation. There is nonetheless a real shadow price of a unit of the consumption aggregate, which we denote λ_t, that can be used to establish the values of assets as necessary in the discussion below.

Firms. Taking for the moment the price that firm z charges as given, we can determine the composition of inputs that it will choose in order to minimize the cost of meeting demand. Each firm produces output according to a Cobb-Douglas production function,

$$c_t(z) = a_t(k_t(z))^{1-\gamma}(n_t(z))^{\gamma}, \tag{4}$$

where $k_t(z)$ is capital, $n_t(z)$ is labor, and a_t is a productivity factor common to all firms. Capital is a fixed factor in the aggregate, but is allocated among firms through an economywide market. Cost minimization implies that

$$\psi_t a_t \gamma \left[\frac{k_t(z)}{n_t(z)}\right]^{1-\gamma} = w_t, \tag{5}$$

and that

$$\psi_t a_t (1-\gamma) \left[\frac{k_t(z)}{n_t(z)}\right]^{-\gamma} = q_t, \tag{6}$$

where ψ_t is real marginal cost, w_t is the real wage, and q_t is the real rental price of capital. Since the real wage and rental rate on capital are not firm-specific, marginal cost and the capital-labor ratio will be identical for all firms. Any one firm can produce at constant marginal cost, although marginal cost will increase with aggregate output. If all prices are flexible, then the preferences specified above imply that all firms will choose to hold their markup of price over marginal cost fixed at $\epsilon/(\epsilon - 1)$.

II.2. Price Adjustment Costs and Price Dynamics

Each period, each firm faces a fixed cost of adjusting its nominal price. Firms adjust their price only if the gains from doing so outweigh the costs. This assumption of fixed costs of price adjustment has a lengthy tradition in macroeconomics, beginning with the work of Barro [1972] and Sheshinski and Weiss [1983]. It gives rise to adjustment rules of the (S,s) form and therefore lies in the background of Caplin and Leahy [1991], where these strategies are imposed on firms.[5]

As discussed above, the novel aspect of our state-dependent pricing framework is that each firm faces a different fixed cost, which is drawn independently over time from a continuous distribution. This implies that there is a marginal firm indifferent to changing its price. The common nature of firms' demand and cost conditions means that all adjusting firms choose the same price, because the fixed costs are time independent. Our model is thus related to other recent work taking a "generalized (S,s) approach" (e.g., Caballero and Engel [1994] and Cooper, Haltiwanger, and Power [1995]).[6]

The basic mechanics of our model. The key features of price dynamics in our economy are highlighted in Figure I. Within each period, some firms will adjust their price, and all adjusting firms will choose the identical value which we call P^*_t. At the start of each period there is a discrete distribution of firms, with fractions $\theta_{jt}(j = 1,2, \ldots, J)$ which last adjusted its price j periods ago, to P^*_{t-j}; these firms will charge P^*_{t-j} at date t if they do not adjust.[7] The number J of firm types, which we will henceforth refer to as vintages, is determined endogenously and will vary with factors

5. Caplin and Leahy [1997] derive optimal (S,s) pricing strategies in a general equilibrium setting.

6. Caballero and Engel [1994] use continuously distributed i.i.d shocks to produce models in which some, but not all, otherwise identical firms adjust. Their study of investment focuses on how these generalized (S,s) rules interact with rich underlying heterogeneity in the discrepancy between target and actual capital stocks. In another study of investment, Cooper, Haltiwanger, and Power [1995] use discrete firm-specific productivity shocks and restrict the heterogeneity so that all adjusting firms choose the same action.

7. Since all firms are in one of these situations, $\Sigma_{i=1}^{J} \theta_{it} = 1$.

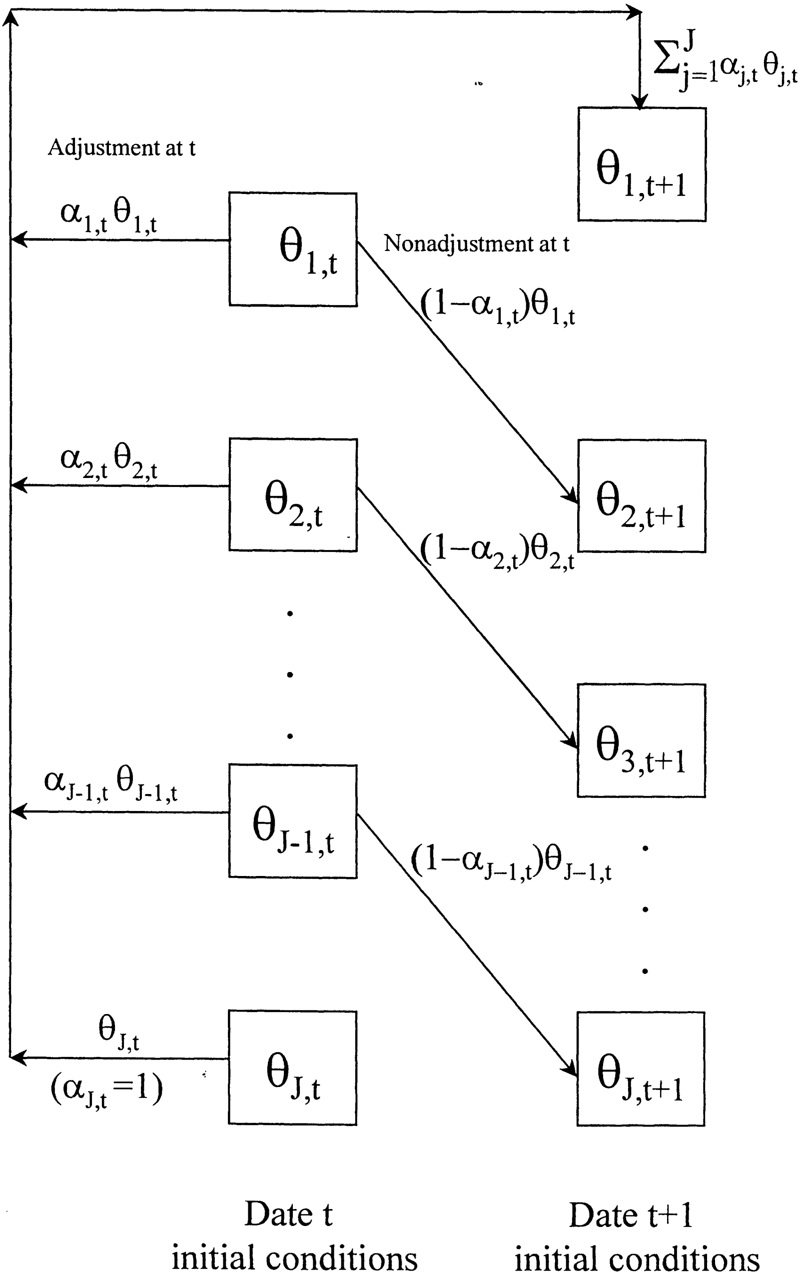

FIGURE I
Evolution of "Vintages" of Price Setters

such as the average inflation rate and the elasticity of product demand.

Figure I illustrates that the distribution of prices, which is part of the relevant state of the economy, evolves through time according to some simple mechanics. In period t a fraction α_{jt} of vintage j firms decides to adjust its price, and a fraction $1 - \alpha_{jt}$ decides not to adjust its price (all vintage J firms choose to adjust). The total fraction of adjusting firms (ω_{0t}) satisfies

$$\omega_{0t} = \sum_{j=1}^{J} \alpha_{jt}\theta_{jt}. \tag{7}$$

There are corresponding fractions of firms,

$$\omega_{jt} = (1 - \alpha_{jt}) \cdot \theta_{jt}, \tag{8}$$

that remain with a price set at period $t - j$. These "end-of-period" fractions are useful because they serve as weights in various contexts later in the paper. The "beginning-of-period" fractions are mechanically related to the "end-of-period" fractions:

$$\theta_{j+1,t+1} = \omega_{jt} \qquad \text{for } j = 0,1, \ldots, J - 1. \tag{9}$$

If the adjustment fractions α_j are treated as fixed through time, then Figure I summarizes the mechanics of models of randomized price-setting opportunities like those time-dependent models developed by Calvo [1983] and Levin [1991].[8] In this interpretation, α_j plays two roles: it is the fraction of firms given the opportunity to adjust within a period, and it is also the probability of an individual firm being allowed to adjust after j periods, conditional on not having adjusted for $j - 1$ periods.

The adjustment decision and the adjustment rate. We employ randomized fixed costs of adjustment to induce discrete adjustment by individual firms, while allowing for an adjustment rate that responds smoothly to the aggregate state of the economy.

In particular, in each period each firm faces a random fixed labor cost of changing its price. The fixed cost ξ is i.i.d. across firms and over time, with c.d.f. $G(\cdot)$ and p.d.f. $g(\cdot)$. We assume that $G(0) = 0 < G(x) < 1 = G(B)$, for $x \in (0,B)$, $B < \infty$. Thus, adjustment by a higher fraction of vintage j firms (α_j) corresponds to adjustment by firms with higher costs, and the cost of adjusting

8. Calvo assumes that $\alpha_j = \alpha$; whereas Levin allows α_j to depend on j.

price is bounded. With positive average inflation the benefit to adjusting price becomes arbitrarily large over time, insuring that the number of vintages is finite.

The adjustment decision for an individual firm is based on three considerations: its value if it adjusts (gross of adjustment costs); its value if it does not adjust; and the current realization of its fixed cost of adjustment. Let $v_{0,t}$ be the real value of a firm if it adjusts, gross of the adjustment cost, and let $v_{j,t}$ be the real value of a firm that last set its price j periods ago. Let $\pi_{j,t}$ be the firm's current period real profits if it has nominal price P^*_{t-j}. Then the value of a price-adjusting firm is given by

$$(10)\quad v_{0,t} = \max_{P_t^*} \left\{ \pi_{0,t} + \beta \cdot E_t \frac{\lambda_{t+1}}{\lambda_t} \cdot (1 - \alpha_{1,t+1}) \cdot v_{1,t+1} + \beta \cdot E_t \frac{\lambda_{t+1}}{\lambda_t} \cdot \alpha_{1,t+1} \cdot v_{0,t+1} - \beta \cdot E_t \frac{\lambda_{t+1}}{\lambda_t} \cdot \Xi_{1,t+1} \right\},$$

where λ_{t+1}/λ_t is the ratio of future to current marginal utility, which is the appropriate discount factor for future real profits, and the term $\beta E_t(\lambda_{t+1}/\lambda_t)\Xi_{j+1,t+1}$ represents the present value of next period's expected adjustment costs and is described further below. The value of a firm that maintains its price at P^*_{t-j} is given by

$$(11)\quad v_{j,t} = \left\{ \pi_{j,t} + \beta \cdot E_t \frac{\lambda_{t+1}}{\lambda_t} \cdot (1 - \alpha_{j+1,t+1}) \cdot v_{j+1,t+1} + \beta \cdot E_t \frac{\lambda_{t+1}}{\lambda_t} \cdot \alpha_{j+1,t+1} \cdot v_{0,t+1} - \beta \cdot E_t \frac{\lambda_{t+1}}{\lambda_t} \cdot \Xi_{j+1,t+1} \right\}$$

for $j = 1,2, \ldots, J - 1$. There is no max operator in (11) because the only decision made by nonadjusting firms is their input mix and the profit term already incorporates this cost minimization.

The assumption that there is a continuous distribution of fixed costs implies that there will be a marginal firm of each vintage $j = 1,2, \ldots, J - 1$ which is just indifferent given its fixed cost realization, i.e., a firm for which $v_{0,t} - v_{j,t} = w_t\xi$ since the costs are assumed to be in units of labor. Combined with the distribution of fixed costs, this condition determines the fraction $(\alpha_{j,t})$ of vintage j firms that adjust to P^*_t,

$$(12)\qquad \alpha_{j,t} = G\,((v_{0,t} - v_{j,t})/w_t)),$$

as shown in Figure II. To induce more firms to adjust, the value

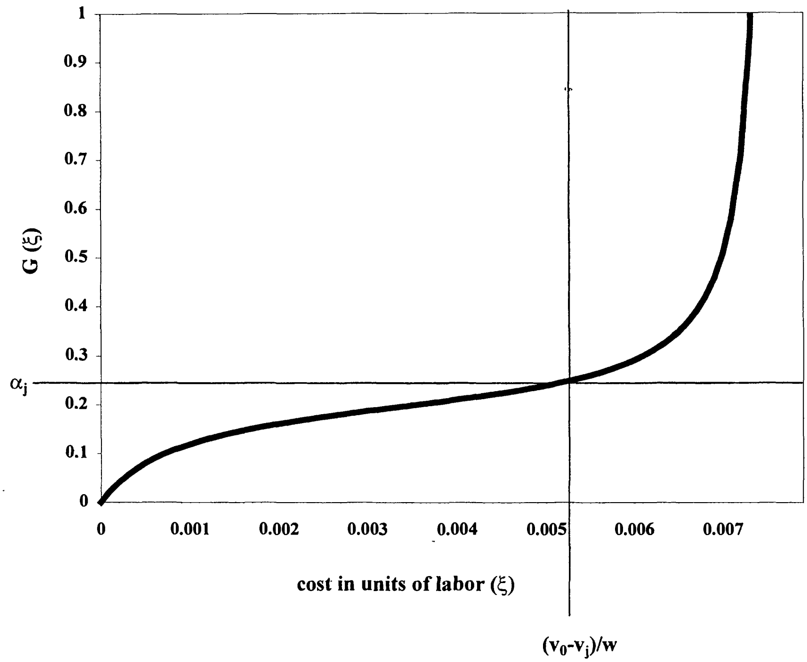

FIGURE II
The Distribution of Fixed Costs, and Determination of the Marginal Firm

difference must be higher, or the wage rate must be lower.[9] Given the fraction of firms that adjust and the wage rate, total resources associated with adjustment by the jth vintage, $\Xi_{j,t}$, are given by

$$\Xi_{j,t} = w_t \int_0^{G^{-1}(\alpha_{jt})} xg(x)\,dx. \tag{13}$$

Looking one period ahead, the expected discounted cost is $\beta E_t(\lambda_{t+1}/\lambda_t)\Xi_{j+1,t+1}$.[10]

Optimal pricing in a state-dependent adjustment model. The dynamic program (10) implies that the optimal price satisfies an

9. Alternatively, the costs could be specified in units of output. In that case (12) would be modified to $\alpha_{jt} = G(v_{0,t} - v_{j,t})$.

10. In the value functions above, the probability of adjustment is $\alpha_{j+1,t+1}$ and the average adjustment cost to be paid conditional on adjustment is $\Xi_{j+1,t+1}/\alpha_{j+1,t+1}$. Thus, expected discounted cost is $E(\lambda_{t+1}/\lambda_t)\alpha_{j+1,t+1}(\Xi_{j+1,t+1}/\alpha_{j+1,t+1})$, which simplifies to the expression in the text.

Euler equation that involves balancing pricing effects on current and expected future profits. That is, as part of an optimal plan, the current price decision requires that

$$0 = \frac{\partial \pi_{0,t}}{\partial P_t^*} + \beta E_t \frac{\lambda_{t+1}}{\lambda_t} (1 - \alpha_{1,t+1}) \frac{\partial v_{1,t+1}}{\partial P_t^*},$$

By iterating this Euler equation and using the implications of (11) for terms like $(\partial v_{1,t+1})/\partial P_t^*$, the optimal price P_t^* can be written as an explicit function of current and expected future variables:

$$(14)\quad P_t^* = \frac{\epsilon}{\epsilon - 1} \frac{\sum_{j=0}^{J-1} \beta^j E_t\{(\omega_{j,t+j}/\omega_{0,t}) \cdot (\lambda_{t+j}/\lambda_t) \cdot \psi_{t+j} \cdot P_{t+j}^{\epsilon} \cdot c_{t+j}\}}{\sum_{j=0}^{J-1} \beta^j E_t\{(\omega_{j,t+j}/\omega_{0,t}) \cdot (\lambda_{t+j}/\lambda_t) \cdot P_{t+j}^{\epsilon-1} \cdot c_{t+j}\}},$$

where $(\omega_{j,t+,j}/\omega_{0,t}) = ((1 - \alpha_{j,t+j}) \cdot (1 - \alpha_{j-1,t+j-1}) \ldots (1 - \alpha_{1,t+1}))$ is the probability of nonadjustment from t through $t + j$.

The pricing rule (14) is a natural generalization of that which obtains in time-dependent settings with exogenous adjustment probabilities (see, for example, King and Wolman [1996] and Yun [1996]). According to (14), the optimal relative price is a fixed markup over real marginal cost ($P^*/P = \epsilon/(\epsilon - 1)\psi$) if real marginal cost and the price level are expected to be constant over time. More generally, (14) illustrates that the optimal price varies with current and expected future demands, aggregate price levels, real marginal costs, discount factors, and adjustment probabilities. All except the last are also present in time-dependent models. Intuitively, firms know that the price they set today may also apply in future periods, so the expected state of the economy in those future periods affects the price that they choose today. If marginal cost is expected to be high next period, for example, a firm will set a high price in the current period, so as not to sell at a loss next period. Similarly, if demand is expected to be high next period, the firm will set a higher price today so that one period of inflation leaves it closer to maximizing static profits next period. The conditional probability terms $(\omega_{j,t+j}/\omega_{0,t})$ are present in time-dependent models, but they are not time-varying. In our setup these conditional probability terms effectively modify the discount factor in a time-varying manner: a very low expected probability of nonadjustment in some future period leads the firm to set a price that heavily discounts the effects on profits beyond that period.

II.3. General Equilibrium

The general equilibrium of our economy involves optimization by firms and households as well as aggregate consistency conditions, such as market clearing and rational expectations. We have previously discussed how consumers allocate income obtained from factor markets across goods and how firms determine factor demands, output, and prices. To complete the description of general equilibrium, we discuss the way money enters in our economy and aggregate market clearing.

Money demand and monetary equilibrium. In the current analysis we have introduced the demand for money simply as an assumption, rather than by deriving it from deeper assumptions about the monetary structure of the economy such as a cash-in-advance requirement or a shopping time technology in which money enters as a productive input. We have done so not out of religious conviction or to obtain a computational advantage, but simply to allow us to parametrically vary the interest elasticity of money demand without having any direct effect on resource utilization or on monetary distortions. In particular, we assume that money demand has a unit elasticity with respect to consumption and a constant semi-elasticity with respect to the nominal interest rate:

$$\ln (M_t/P_t) = \ln (c_t) - \eta \cdot R_t.$$

We vary the interest semi-elasticity (η) in our analysis below. The money supply is an exogenous driving process whose parameters we also vary in our subsequent analysis.

Asset market equilibrium conditions. The nominal interest rate in our economy is given by a Fisher equation,

$$1 = E_t\{\beta(1 + R_t) \cdot (\lambda_{t+1}/\lambda_t) \cdot (P_t/P_{t+1})\}, \tag{15}$$

and the real interest rate is given by

$$1 = E_t\{\beta(1 + r_t) \cdot (\lambda_{t+1}/\lambda_t)\}, \tag{16}$$

since it is defined in units of the composite consumption good.

Aggregate output and consumption. The market-clearing conditions in goods, capital, and labor markets are standard, although the demand for labor derives from both final production and price adjustment sources. However, it is important to note that since c_t is a nonlinear aggregate, it is not equal to the simple sum of all firms' outputs. For some purposes it is useful to keep track of the linear aggregate $\int_0^1 c_t(z)\,dz$, which we call y_t. While the linear output aggregate and the utility-based consumption aggregate do not move identically in the model, in most contexts the difference is quantitatively unimportant. Yun [1996] shows that by defining an auxiliary price index,

$$\overline{P}_t = \left[\int_0^1 P_t(z)^{-\epsilon}\,dz\right]^{-1/\epsilon},$$

one can express the relationship between the two aggregates as

$$y_t = (P_t/\overline{P}_t)^{\epsilon} c_t. \tag{17}$$

Alternatively, one can define the model's implicit deflator from $D_t y_t = P_t c_t$ so that $D_t = P_t^{(1-\epsilon)}\overline{P}_t^{\epsilon}$. Our general finding that y_t and c_t move closely together indicates that the perfect price index in the model (P_t) is well approximated by the implicit deflator.

The price level. Individual firms choose their prices according to (10), taking as given the prices charged by other firms, but all adjusting firms choose the same price. Accordingly, (3) simplifies to

$$P_t = \left[\sum_{j=0}^{J-1} \omega_{jt} \cdot (P^*_{t-j})^{1-\epsilon}\right]^{1/(1-\epsilon)}, \tag{18}$$

where ω_{jt} is the fraction of firms at time t charging price P^*_{t-j}.

The aggregate state vector. In the current economy there are only two types of state variables that are sufficient to describe the past history of the economy at date t. First, there is the price distribution ($\{P^*_{t-j}\}_{j=1}^{J}$ and $\{\theta_{jt}\}_{j=1}^{J}$).[11] Second, there are the exogenous state variables that describe the money supply process.

II.4. Computation and Calibration

Although the preceding description makes clear that the state space of our model is finite, it is too large for the implementation of conventional nonlinear solution techniques. We therefore compute the rational expectations equilibrium by (i) solving a

11. Since $\Sigma_{j=1}^{J}\theta_{jt} = 1$, we can drop θ_{Jt} from the state vector. Since no firm will ever charge P^*_{t-J}, we can drop it as well.

system of nonlinear equations for the model's steady state,[12] and (ii) linearizing the model's behavioral equations around the steady-state equilibrium.[13] The approximate behavioral equations are detailed in an appendix that also serves as a guide to replication materials that can be downloaded by interested readers.[14]

The calibration of the economy is generally standard. Labor's share is two-thirds, the discount factor implies a real return of 6.5 percent annually ($\beta = .984$), momentary utility is logarithmic in consumption ($\sigma = 1$), the elasticity of demand is such that the markup in a flexible price economy is 1.3 ($\epsilon = 4.33$), and agents work 20 percent of their time endowment.

We also assume that there is infinitely elastic labor supply, in the sense of a Frisch labor supply elasticity, by setting $\zeta = 1$.[15] This unrealistic assumption is made for two reasons. First, many studies of price adjustment assume that marginal cost is insensitive to changes in output (for example, this is the implicit assumption in Caplin and Leahy [1991]), and $\zeta = 1$ minimizes the cyclical variation in marginal cost. Second, as in Hansen [1985] and Rogerson [1988], this restriction can be interpreted as arising from optimal labor contracts in the presence of indivisible labor.[16] However, there are implications for the model's dynamics, which we discuss in subsection III.4 below.

Our introduction of random adjustment costs, which convexifies the economy so as to permit its ready solution, introduces a new set of calibration issues. Ultimately, we think that models along the lines developed in this paper will need to choose the parameters of the distribution function $G(\ )$ so as to match microeconomic data on the frequency of price adjustment. In particular, one would want to match the model's hazard rate for price adjustment as displayed in Figure III with estimates of price adjustment hazards for particular commodities in the U. S. economy. However, in the present paper we have built a very

12. In the steady state there is no aggregate uncertainty, although individual firms face uncertainty about their fixed costs of price adjustment.

13. We use the solution methodology of King and Watson [forthcoming, 1995] to compute the resulting linear rational expectations equilibrium. The King-Watson procedures are generalizations of the Blanchard-Kahn [1980] method suitable for large systems such as ours that have a mixture of dynamic and nondynamic behavioral equations.

14. http:/wsrv.clas.virginia.edu/~rgk4m/

15. The labor supply elasticity and the requirement that individuals work 20 percent of the time endowment jointly imply that $\chi = 2.57$.

16. Given the assumption that 20 percent of time is devoted to work, the labor supply elasticity has no effect on the model's steady state.

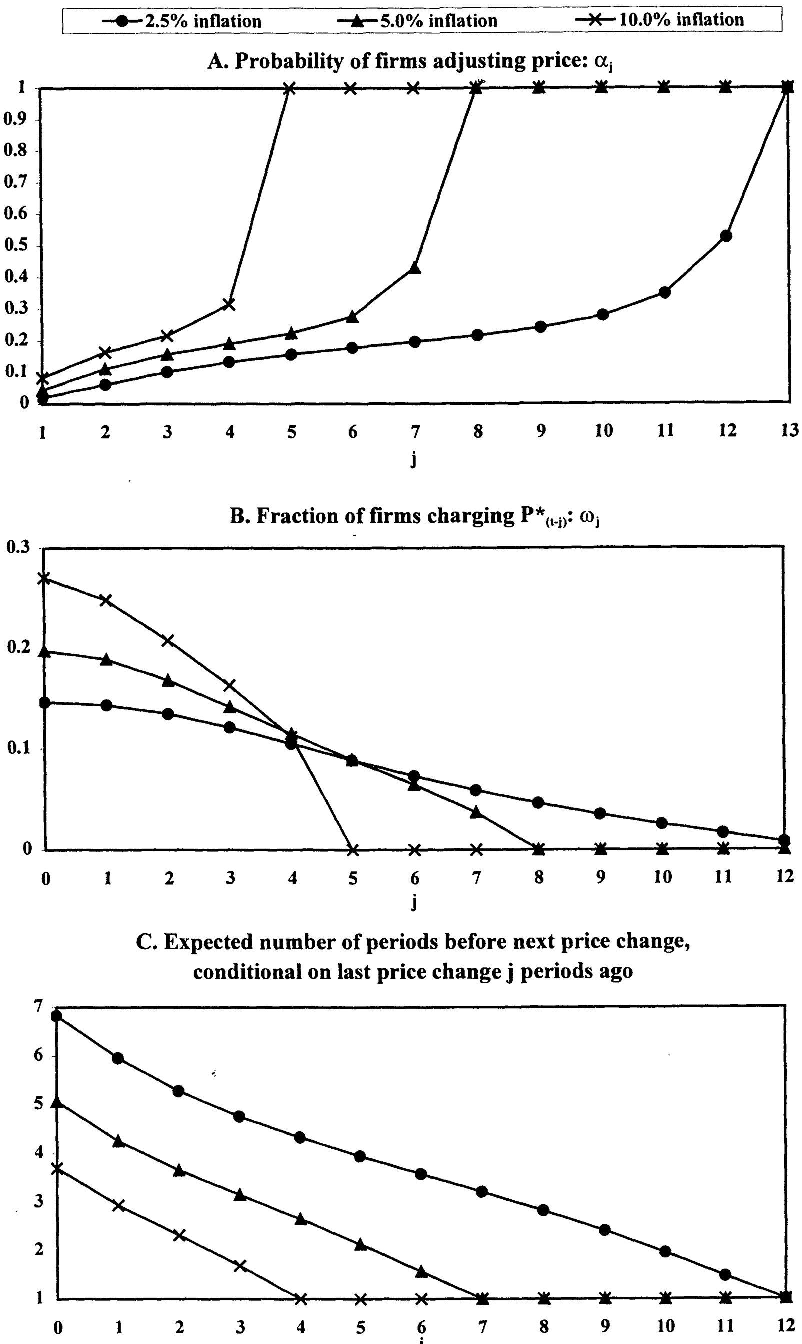

FIGURE III
The Steady State

simple model that highlights the central features of our approach to price adjustment. We have omitted many features of actual economies, such as capital accumulation and a flexible price sector. For this reason, and to facilitate comparison with time-dependent models below, we adopt a simpler approach that allows us to illustrate the quantitative implications of our theory. We assume that the distribution of fixed costs takes the form,

$$G(\xi) = c_1 + c_2 \tan\,(c_3 \cdot \xi - c_4). \tag{19}$$

This four-parameter function is quite flexible. It can be concave, convex, almost linear, or S-shaped, depending on the chosen values of the parameters.

The specific adjustment cost parameterization that we choose is displayed in Figure II.[17] In general, the shape of the distribution is an important factor in determining both the steady-state adjustment probabilities and the dynamic responses discussed below.[18] Before turning to analysis of various inflation rates, it is useful to discuss some implications of the choices that we have made. From the distribution $G(\xi)$, the mean price adjustment cost as a fraction of labor input is $E(\xi)/n = 0.005795/0.20 = 0.029$, and the most expensive price adjustment for any firm is $\max(\xi)/n = .0075/.20 = .0375$, where $n = .20$ is the total fraction of time allocated to market work. The most expensive price adjustment thus costs the adjusting firm 3.8 percent of its steady-state labor time. However, the social cost of price adjustment will be much lower because only a fraction of firms are adjusting prices, and adjusting firms pay less than the maximum. This social cost depends on the endogenously determined time pattern of price adjustment, which we study next.

II.5. The Effect of Steady-State Inflation on Price Adjustment

Our analysis of the model's quantitative properties begins by investigating the relationship between the steady-state rate of inflation and steady-state patterns of price adjustment; we compare steady states for annual inflation rates of 2.5, 5, and 10 percent. The time period is assumed to be one-quarter of one year throughout the steady-state and business cycle analysis.

17. The parameter values for the C.D.F. are $c_1 = 0.1964$, $c_2 = 0.0625$, $c_3 = 367.44$, and $c_4 = 1.2626$.

18. An earlier version of this paper [Dotsey, King, and Wolman 1997a] used a uniform distribution of fixed costs. Holding constant the maximum level of fixed costs, a uniform distribution generates more variation in the adjustment probabilities in response to shocks than the distribution used here.

Effects of inflation on the steady-state distribution of firms. Because inflation erodes a firm's relative price, firms choose to maintain a given price for fewer periods when inflation is high. Thus, higher inflation endogenously generates a smaller value for J, the number of vintages. For 10 percent inflation all firms change their price at least once every five quarters, while for 2.5 percent inflation firms may not change their price for thirteen quarters. Figure III, panel A, shows how the endogenous conditional adjustment probabilities (α_j)—which are treated as primitives in time-dependent models—vary with the inflation rate. *Positive* inflation means that the benefits of adjusting price are higher for firms whose price was set further in the past, and this translates into higher adjustment probabilities for such firms ($\alpha_{j+1} > \alpha_j$). The *higher* is the inflation rate, the greater are the benefits to adjusting for any j, and hence the higher is α_j. With fewer vintages panel B (Figure III) shows that the fraction of firms adjusting their price (charging P^*) is 0.27 with 10 percent inflation, 0.197 with 5 percent inflation, and 0.146 with 2.5 percent inflation. Also, the expected time to the next price change falls dramatically with an increase in inflation, as shown in panel C. A firm that changed its price last period would expect to wait 6.8 quarters before changing its price under 2.5 percent inflation, but would only expect to wait 3.7 quarters under 10 percent inflation.[19] Note that the adjustment pattern under 5 percent inflation is consistent with Blinder's [1994] finding that firms typically set prices for about a year.

Effects of inflation on the average markup of price over marginal cost. Nonadjusting firms see their markups erode more quickly with higher inflation, leading them to desire a higher markup when they do adjust. This effect is mitigated by the endogenous decline in expected time until next adjustment, but the optimal markup for an adjusting firm still arises with inflation. The average markup can be written in terms of the marginal markup (the markup chosen by adjusting firms) and the ratio of the price level to the price chosen by adjusting firms: $P/\Psi = (P^*/\Psi) \cdot (P/P^*)$, where Ψ is nominal marginal cost. Holding the number of vintages constant, higher inflation would reduce the ratio of P to P^*, as in King and Wolman [1996]; using the

19. In panel C we follow the convention that a firm adjusting its price every period is waiting one period before its next price change. Thus, the smallest number on the vertical axis is unity.

adjustment probabilities associated with 5 percent inflation, for example, the ratio P/P^* is 0.9844 at 2.5 percent inflation, 0.9690 at 5 percent inflation, and 0.9386 at 10 percent inflation.[20] Intuitively, with positive inflation, P is less than P^*, and at higher inflation rates the *range* of prices charged increases if the longest period of price stickiness is held constant. This effect too is mitigated by the endogenous decrease in the number of vintages and the expected time until next adjustment; with adjustment probabilities sensitive to inflation, P/P^* is much *less* sensitive, just falling from 0.9769 at 2.5 percent inflation to 0.9690 at 5 percent inflation, and 0.9610 at 10 percent inflation. The changes in (P/P^*) and the marginal markup combine to generate an extremely small increase in the average markup—from 1.300 to 1.301 to 1.302—as inflation rises from 2.5 to 5 to 10 percent.[21] These values for the average markup are not much different from what would be found holding the adjustment probabilities fixed at the values corresponding to 2.5, 5, or 10 percent inflation. We saw that with adjustment probabilities fixed, P/P^* varies more with inflation than when those probabilities are endogenous, but fixing the adjustment probabilities also generates more variation in P^*/Ψ, and these two effects are roughly offsetting.

Effects of inflation on relative price variability. The standard deviation of relative prices is also positively related to inflation, increasing from 1.83 percent at 2.5 percent inflation to 3.17 percent at 10 percent inflation.[22] This illustrates that the decrease in the number of vintages at higher inflation—which would lead to a lower standard deviation—is more than offset by the higher relative price chosen by adjusting firms.

Effects of inflation on the resources allocated to price adjustment. There is a variety of ways to evaluate the magnitude of the costs of price adjustment that are present in our economy. Previously, we calculated that the highest price adjustment cost a firm could face was 3.75 percent of its labor input. However, the distributions shown in Figure III imply that the average cost of adjustment is only 0.74 percent at 2.5 percent inflation.

20. This can be shown by writing (14) in steady state and taking the derivative of P/P^* with respect to the inflation rate.

21. In a model with customer search and fixed costs of price adjustment, Bénabou [1992] finds that the markup falls with inflation, because higher inflation raises each search intensity and decreases firm's market power.

22. The standard deviations are calculated by weighting relative prices according to the fraction of firms charging those prices. Alternatively, one could weight the prices by quantities of goods sold.

Our specification of $G(\xi)$ implies that these costs increase significantly with inflation. The amount of resources devoted to price adjustment increases from 0.74 percent of an average adjusting firm's labor input with 2.5 percent inflation to 1.61 percent of its labor input with 10 percent inflation. However, given the preceding information about the fractions of firms that adjust each period in alternative steady states, the amount of society's labor input used in price adjustment is $1.61 \times 0.27 = 0.44$ percent with 10 percent inflation and $0.74 \times 0.146 = 0.11$ percent with 2.5 percent inflation.

III. The Effect of Money on Output and Prices

In this section we describe the model's dynamic response to money supply shocks, starting from the 5 percent inflation steady state described above. We first analyze a permanent, unanticipated increase in the quantity of money, assuming that velocity is constant. Extensions involve varying the persistence of the shock process and the interest elasticity of money demand, and varying the elasticities of labor supply and product demand.

III.1. An Unanticipated, Permanent Monetary Expansion

The baseline experiment is a positive innovation to a random walk money stock driving process, under the assumption that the interest elasticity of money demand is zero and the consumption elasticity is one. This first example is close to the central experiment in Caplin and Leahy [1991], except for differences in the assumed trend rates of money growth.

The impulse response functions for the state-dependent model are shown as the solid lines in Figure IV. Output responds strongly to the money shock, rising by 0.45 percent above its steady state on impact. Prices also rise, by about 0.52 percent on impact, and there is inflation on the transition path.[23] Thus, our state-dependent pricing model delivers a mixture of price and output effects.[24]

23. The reader may wonder why the price and output effects do not sum to unity, given that velocity is constant. The explanation is that while consumption velocity is constant, we are referring here to the behavior of output. The relationship between consumption and output is given by (17).

24. As was discussed in subsection II.4, the distribution of fixed costs is an important determinant of the response to money shocks. A uniform distribution generates significantly smaller output and larger price level effects than those shown in Figure IV.

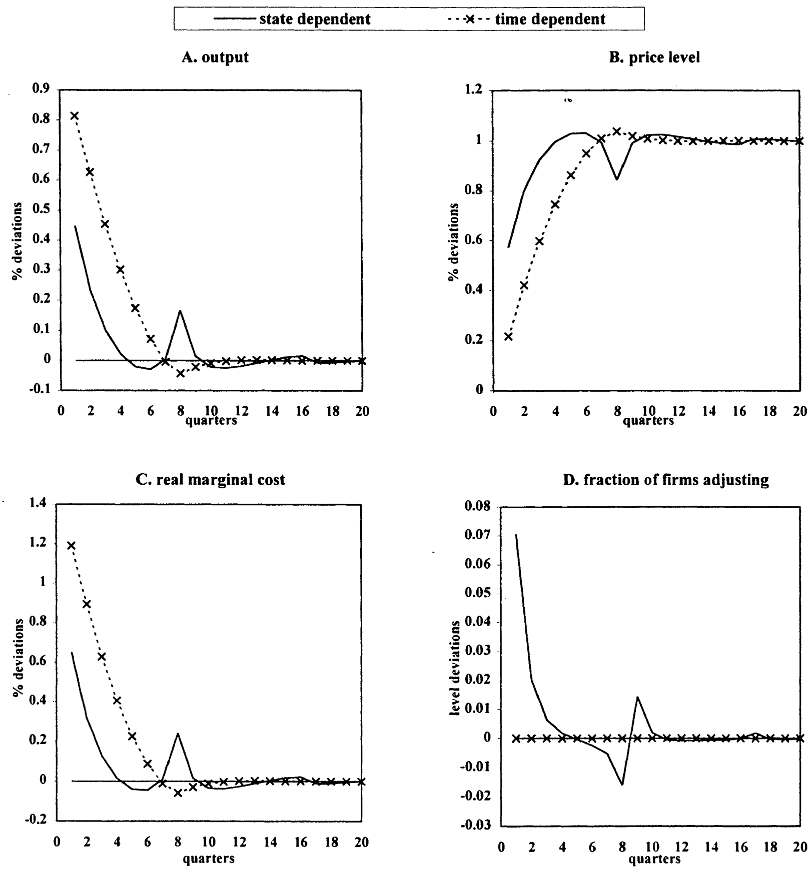

FIGURE IV
Baseline Experiment: Response to a Permanent Money Supply Shock

What explains the evolution of the price level in our model? From (18) variations in the price level can be attributed to variations in the end-of-period fractions of firms of each vintage, and the prices they set previously. Similarly, from (14) variations in P^* can be attributed to variations in current and expected future interest rates, adjustment probabilities, real marginal costs, real aggregate demands, and price levels. Figure V indicates which of these influences are strongest, using log-linearized versions of (14) and (18).

The equation for the optimal price P_t^* that we discussed above

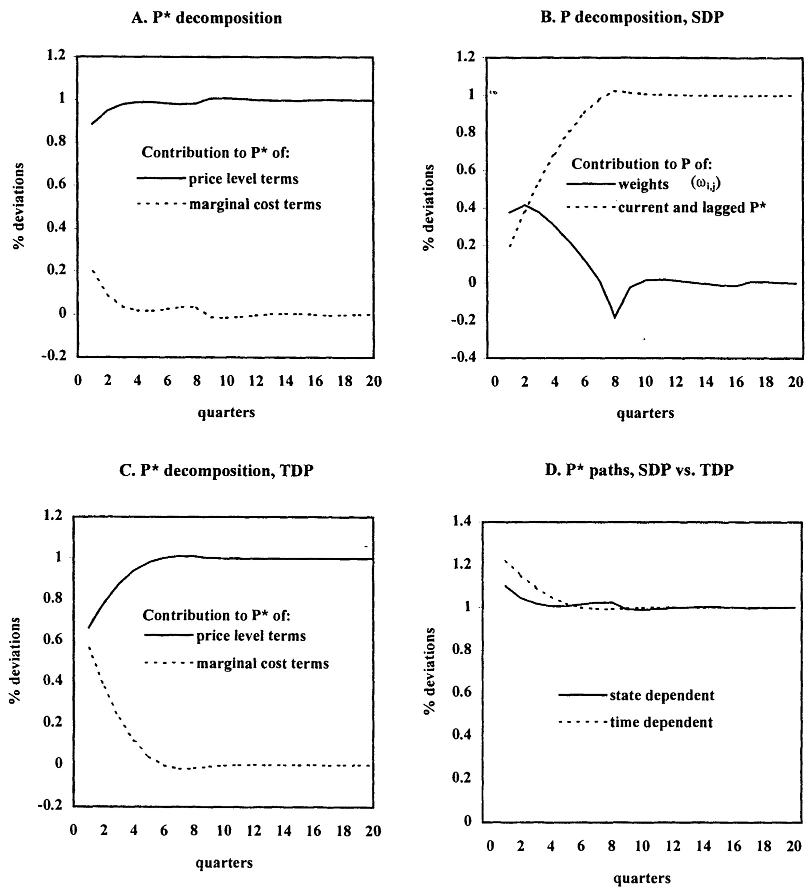

FIGURE V
Pricing Decompositions

can be written as

$$P_t^* = \frac{\epsilon}{\epsilon - 1} \frac{\sum_{j=0}^{J-1} E_t\{(\omega_{j,t+j}/\omega_{0,t}) \cdot (\lambda_{t+j}/\lambda_t) \cdot [\psi_{t+j} P_{t+j}] \cdot [P_{t+j}^{\epsilon-1} c_{t+j}]\}}{\sum_{j=0}^{J-1} E_t\{(\omega_{j,t+j}/\omega_{0,t}) \cdot (\lambda_{t+j}/\lambda_t) \cdot [P_{t+j}^{\epsilon-1} c_{t+j}]\}}.$$

The first bracketed term is nominal marginal cost $[\Psi_{t+j} = \psi_{t+j} \cdot P_{t+j}]$, and the second bracketed term is the demand shift arising due to aggregate economic factors $[P_{t+j}^{\epsilon-1} c_{t+j}]$, i.e., the direct effect of real aggregate demand and the demand-switching effect of movements in the price level. To study the determinants of P_t^* in more detail, we totally differentiate this expression, yielding

$$
\begin{aligned}
(20) \qquad d \ln P_t^* = & \sum_{j=0}^{J-1} \rho_j E_t\{d \ln \psi_{t+j} + d \ln P_{t+j}\} \\
& + \sum_{j=0}^{J-1} \delta_j E_t\{[d \ln (\omega_{j,t+j} - d \ln \omega_{0,t}]\} \\
& + \sum_{j=0}^{J-1} \delta_j E_t\{[d \ln \lambda_{t+j} - d \ln \lambda_t] \\
& + (\epsilon - 1) \cdot d \ln P_{t+j} + d \ln c_{t+j}\}.
\end{aligned}
$$

That is, the current change in the optimal price can be decomposed into (i) effects of the time path of marginal cost, arising from its real and price level components, and (ii) effects of the time paths of the probabilities of nonadjustment, discount factors, and demand factors.[25] For inflation rates close to zero, these weights are approximately $\rho_j = [\beta^j \omega_j]/[\Sigma_{h=0}^{J-1} (\beta^h \omega_h)$ and $\delta_j = 0$.

The solid line in Figure V, panel A, shows the price level terms' contribution to the behavior of P_t^*, while the dashed line shows the contribution of the real marginal cost terms.[26] It is interesting to note that the other factors (interest rates, output, and future adjustment probabilities) do not have an important effect on the optimal price, because the δ_j are small.[27] At the impact date (quarter 1) the optimal price increases for two reasons, as shown in panel A. First, current and future real marginal costs are higher, meaning that a firm's markup would fall if it raised its price one-for-one with inflation. Second, the price level is expected to increase faster than in steady state, so that any markup a firm sets will be eroded unusually quickly. A rational price-setting firm, recognizing that its price will likely be fixed for several periods, reacts to these two factors by setting a higher price, in much the same manner that firms respond to higher steady-state inflation. Focusing on the P^* decomposition in

25. The weights attached to these various factors are $\rho_j = [\beta^j \omega_j \Pi^{j\epsilon}]/\Sigma_{h=0}^{J-1} \cdot \beta^h \omega_h \Pi^{h(1+\epsilon)}$ and $\delta_j = \rho_j - [\beta^j \omega_j \Pi^{(\epsilon-1)J}]/\Sigma_{h=0}^{J-1} \beta^h \omega_h \Pi^{(\epsilon-1)h}$, where Π is the steady-state gross inflation rate.

26. The decomposition we use in Figure V is as follows. Suppose that $dy_t = a \cdot dx_{1,t} + b \cdot dx_{2,t}$. Then the contribution of $x_{1,t}$ to the behavior of y_t is $a \cdot dx_{1,t}$.

27. That is, one could approximate the response with a simpler distributed lead:

$$P_t^* = \frac{\epsilon}{\epsilon - 1} \cdot \sum_{j=0}^{J-1} E_t\left\{\left(\frac{1}{1+R}\right)^j \cdot \left(\frac{\omega_j}{\omega_0}\right) \cdot \psi_{t+j} \cdot P_{t+j}\right\}.$$

isolation, the dynamics are very similar to those generated by time-dependent models.[28] In this state-dependent model, however, the behavior of the aggregate price level is also heavily influenced by changes in the distribution.

Figure V, panel B is based on the log-linearized version of the price level (18):

$$(21)\quad d\ln P_t = \left(\frac{P^*}{P}\right)^{1-\epsilon} \cdot \sum_{j=0}^{J-1} \Pi^{j(\epsilon-1)} \left[\frac{1}{1-\epsilon} \cdot d\omega_{jt} + \omega_j \cdot d\ln P^*_{t-j}\right].$$

The solid line shows the P^*_{t-j} terms' contribution to the behavior of P_t, while the dashed line shows the contribution of the weights. The monetary expansion causes the fraction of adjusting firms to rise from 20 percent in the steady state to 27 percent in the quarter during which the monetary shock occurs; this accounts for almost two-thirds of the initial rise in the price level.

As is perhaps not surprising in a model of state-dependent pricing, changes in the distribution of price setters play a role in shaping the dynamic relationship between money and output in the early periods, but then are increasingly less important as the economy approaches the long run. However, it is notable that the changing distribution exerts an influence on price dynamics with a periodicity that is related to the longest time that it takes a firm to adjust a price. Above, we noted that the fraction of firms adjusting on the impact date jumps by 7 percent. In Figure IV we see that as these firms cycle through the pricing process, there ends up being an unusually high fraction of firms in the low end of the price distribution in period 8. This means that after eight quarters the price level is relatively low and output is correspondingly high. After cycling through the adjustment process for a second time (i.e., after sixteen quarters), the economy essentially reaches its new steady state. As can be seen from Figure V, the variation in the distribution is also important for the decomposition of the price level. Generally, as the firms cycle through the mechanism illustrated in Figure I and change their price, more and more of the change in P is due to changes in P^*. However, the decline in P relative to trend in period 8 is entirely accounted for by the behavior of the adjustment process, i.e., by the fact that an unusually large fraction of firms was induced to change price by the monetary shock.

28. See, for example, the dynamics in King and Wolman [1996], which uses the infinite lead Calvo specification.

III.2. State- versus Time-Dependent Pricing

It is natural to contrast the dynamic behavior of money and output under state-dependent pricing with that in a similar model of time-dependent pricing, in which the fraction of firms changing price is held fixed at the steady-state values. The results of these experiments are shown in Figure IV as the lines marked with "x"s, as well as in panels C and D of Figure V. On impact, prices are much stickier in the time-dependent model than in the state-dependent model, so that the response of output is much stronger. Because there is no departure from steady-state adjustment behavior, there is no longer the upward spike in output in period 8. The P and P^* decompositions are also revealing. In the time-dependent model, P^* rises by more than it does in the state-dependent model (panel D); firms raise their optimal price by more when their adjustment pattern is inflexible. Also, marginal cost contributes more to the increase in P^* in the time-dependent model, as a direct consequence of the greater response in output. However, the qualitative nature of optimal price behavior is quite similar across the two models: the optimal price is essentially a distributed lead of current and expected future price levels and real marginal costs. The fundamental difference between the two models is illustrated implicitly by the B panels in Figures IV and V. Under time-dependent pricing the price level is a fixed distributed lag of optimal prices. In contrast, under state-dependent pricing the lag weights change as the distribution changes. Figure V, panel B, shows that variation in the weights accounts for much of the initial price increase under state-dependent pricing.

III.3. Varying Interest Elasticities and Driving Processes

We now explore the sensitivity of our results to changes in the interest elasticity of money demand and the persistence of the money supply process.

Figure VI displays the effect of altering the persistence of money, maintaining the zero interest elasticity assumption. The autocorrelation parameter is on the horizontal axis, and the impact effects on y and P are plotted as the solid and dashed lines, respectively. In panel A we look at shocks to the level of the money supply, starting with a white noise process increasing the autocorrelation parameter up to 1. In panel B we turn to studying growth

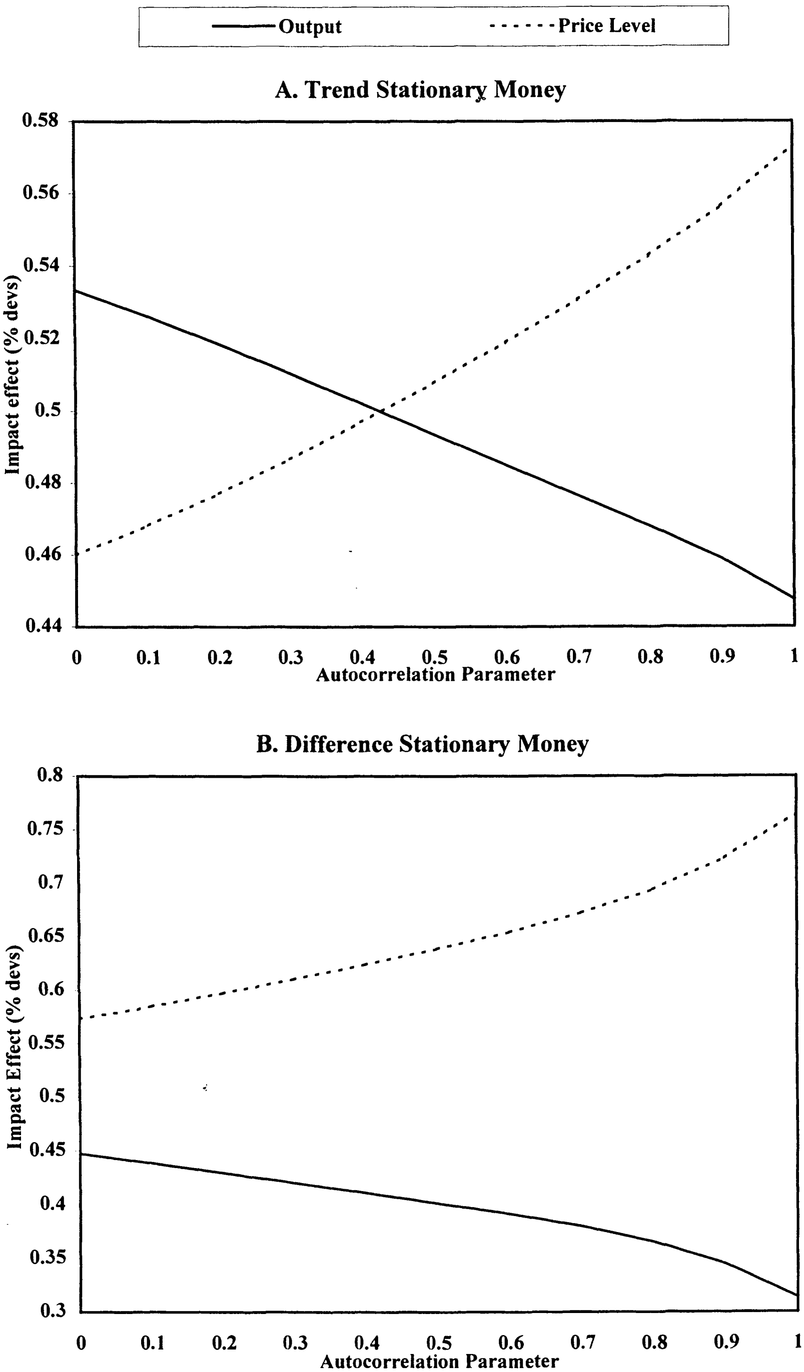

FIGURE VI
Varying the Shock Persistence with Interest Inelastic Money Demand

rate shocks. We begin with a white noise growth rate shock and again increase the autocorrelation up to 1.[29]

For a purely transitory change in the money stock, there is a negligible response of the price level because few additional firms are willing to pay a fixed cost of adjusting and those that do adjust correctly understand that the effect of the shock is temporary. There is correspondingly a large response of output, given that the quantity equation must be satisfied. As the persistence of the process increases, more firms find it worthwhile to change their price, and they change their price by more. Therefore, the impact effect on prices increases with persistence, and the output effect declines. For money growth rates, the result is much the same. Increasing persistence magnifies the impact effect on prices and reduces the response of output.

Figure VII involves the same experiment, except that money demand is highly interest elastic.[30] The results are strikingly different. In panel A of Figure VII for a purely transitory change in the money stock, there is little effect on output and a small effect on prices. Prices rise a little on impact, and agents expect that there will be deflation, so the nominal interest rate falls, raising the demand for money by enough so that almost no output response is required to clear the money market. As the persistence of the processes increases, implying expectations of inflation rather than deflation at short horizons, these results are altered sharply. With a higher nominal interest rate, money demand falls, and a larger increase in output is required to clear the money market. In panel B of Figure VII, where the persistence is moved into money growth rates, the output and price responses become larger than the unit impact effect on money.

The interaction between the money supply process and the interest sensitivity of money demand is explored further in Figure VIII, which depicts the impulse responses of output, the price level, the nominal interest rate, and real marginal cost for a money growth process that has an autoregressive coefficient of 0.5. The dynamic responses for an economy with a zero interest

29. In all of these examples the steady-state inflation rate is kept at 5 percent, which implies positive trend growth in money. Note in Figure VI that a white noise shock-to-the growth rate is identical to a random walk shock to the level, the benchmark of subsections III.1 and III.2.

30. Figure VII was generated with the interest semi-elasticity η set to 17.65, which corresponds to an interest elasticity of money demand equal to $-\frac{1}{2}$ at the steady state of 5 percent annual inflation.

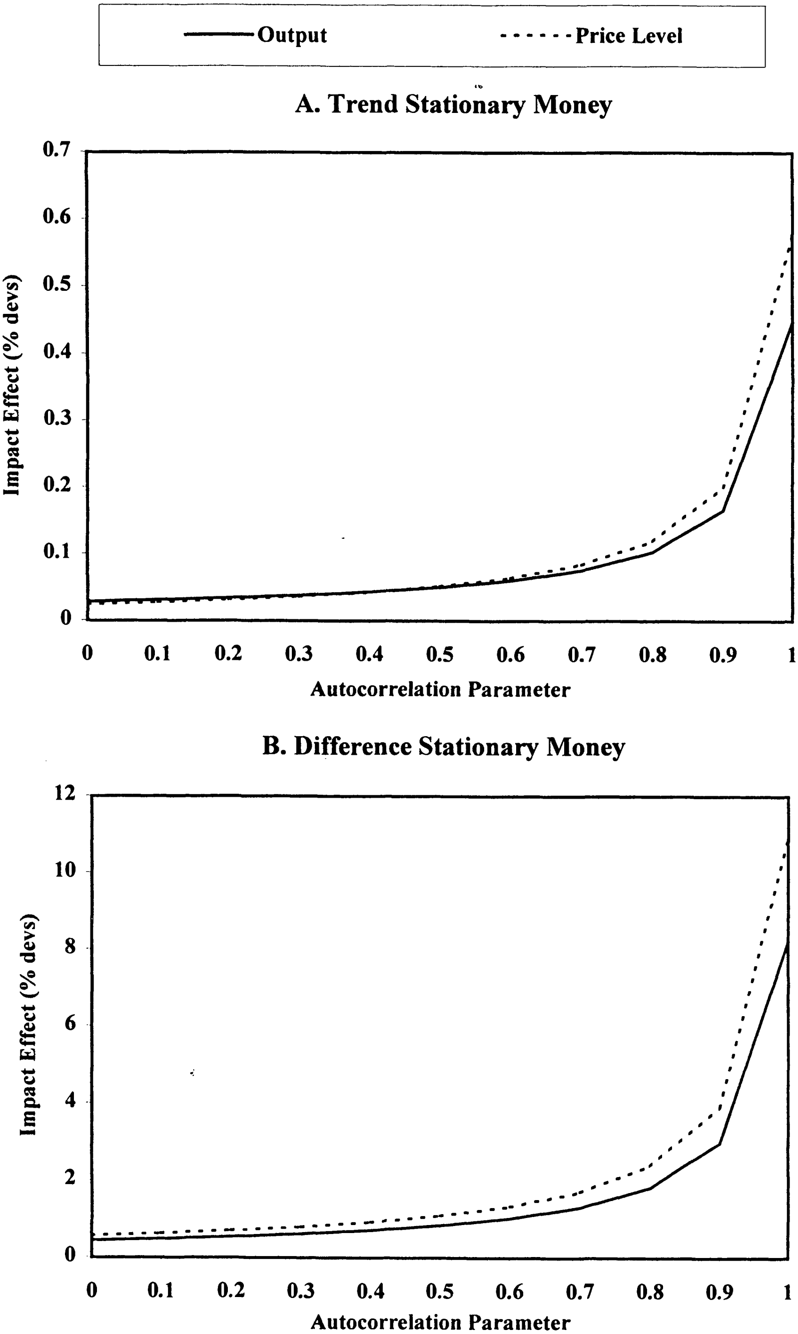

FIGURE VII
Varying the Shock Persistence
Interest Elasticity of Money Demand = $-\frac{1}{2}$

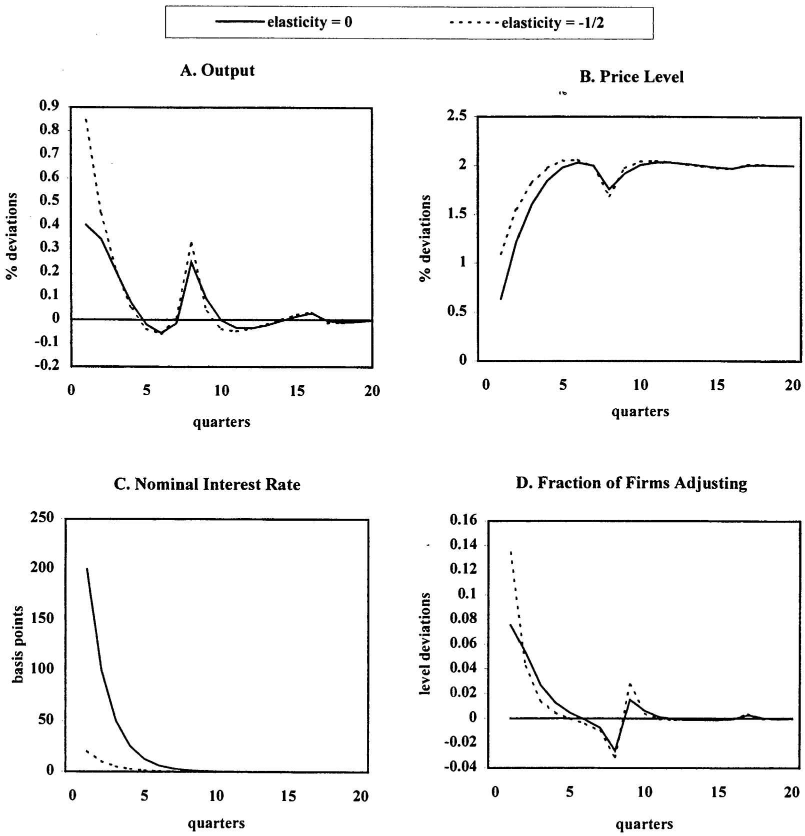

FIGURE VIII
Varying the Interest Elasticity of Money Demand
Response to a Persistent Money Growth Shock ($\rho = 0.5$)

elasticity are the solid lines, and those for an economy with an interest elasticity of 0.5 are the dashed lines.

First, notice the different impact effects on output. An interest elasticity of 0.5 doubles the impact effect on output *and* almost doubles the impact effect on the price level. The presence of expected inflation lowers the demand for money in the interest-sensitive case, implying that the effect on nominal income ($P \cdot y$) must be greater to restore money market equilibrium. Consequently, both prices and output respond more strongly to the shock. The price level by itself cannot completely absorb the higher level of nominal demand, because that would imply low real balances, no change in the nominal interest rate, and a fall in output.

These results illustrate that there can be significant differences in the behavior of prices depending on the structure of the economy (in this case the interest elasticity) and the persistence of the driving process (the money growth process) even when firms are behaving in the same state-dependent way.

III.4. Varying the Elasticities of Labor Supply and Goods Demand

The labor supply elasticity used above is counterfactually large if it is interpreted as describing the effect of a change in the wage rate on the supply of hours for an individual. In light of Ball and Romer's [1990] work, we examine the model's sensitivity to this parameter. Table I displays the impact effects on output and prices for the baseline experiment of subsection III.1, varying ζ between 1 and 7. Table I's basic message is that holding other parameters constant, the extent of short-run nonneutrality is quite sensitive to the labor supply elasticity. When ζ rises above 6, which is equivalent to the elasticity falling below 0.2, the impact effect on output actually becomes negative. The economics of the effect of altering the labor supply elasticity is straightforward: in order for firms to be willing to hold their price fixed in the face of an increase in nominal demand, their marginal cost must not rise too much. When the elasticity of labor supply is large, real wages and hence marginal cost are relatively insensitive to output. But when the elasticity of labor supply is low, any increase in output carries with it much higher marginal cost; instead of producing at the implied lower markup, firms choose to raise their price, and output does not in fact increase.[31] Thus, our dynamic pricing model replicates a key result derived in static menu cost models.

TABLE I
SENSITIVITY TO THE LABOR SUPPLY ELASTICITY

ζ	1	2	3	4	5	6	7
Implied elasticity	∞	1	0.5	0.33	0.25	0.2	0.1667
y impact effect	0.45	0.26	0.16	0.09	0.04	0.01	−0.02
P impact effect	0.57	0.76	0.86	0.92	0.97	1.01	1.03

31. See Romer [1996, pp. 280–287] for a good discussion of these issues.

TABLE II
SENSITIVITY TO THE GOODS DEMAND ELASTICITY

ϵ	2	3	4	5	6	7
Implied static markup	2.00	1.50	1.33	1.25	1.20	1.17
y impact effect	0.22	0.52	0.54	0.61	0.50	0.13
P impact effect	0.77	0.50	0.49	0.43	0.53	0.87

There are two feasible directions of research suggested by this finding. First, modification of production technologies and preferences may enhance intertemporal substitution, leading marginal cost to increase less strongly with output.[32] Second, modifications of the labor market structure, perhaps along efficiency wage lines, could make prices respond less strongly to demand.[33]

Another important parameter in the model is the elasticity of demand for the differentiated goods, which determines how close substitutes the goods are and, in turn, influences the size of the equilibrium markup. Table II displays the impact effects on output and prices for the baseline experiment of subsection III.1, varying ϵ between 2 and 7 (the implied variation in the static markup is between 1.17 and 2). These experiments reveal a nonmonotonicity in the relationship between ϵ and the extent of nonneutrality. When ϵ is very low, adjusting firms raise their price by a lot to exploit their market power. Thus, for low ϵ the fact that the price level soaks up most of the monetary shock is accounted for by the behavior of P^*. When ϵ is very high, adjusting firms do not choose as high a price, but all firms have a stronger incentive to adjust; the penalty for charging a price below marginal cost is especially high in this case, because consumers are more sensitive to price differentials. For high ϵ then, the small degree of nonneutrality is mainly explained by an increase in the number of adjusting firms. For intermediate levels of ϵ, neither the degree of market power nor the incentive to adjust is high enough to prevent a substantial impact effect on output.

III.5. Inflation and Monetary Shocks

With state-dependent pricing, higher inflation increases the frequency of price adjustment, but what are the consequences for

32. For example, Dotsey, King, and Wolman [1997a, 1997b] include capital formation, which allows for intertemporal substitution of work effort. In these settings, prices respond less strongly to money.

33. See Ball and Romer [1990], Romer [1996], and Blanchard and Kiyotaki [1987].

the effect of money on output? To answer this question, it helps first to consider what would happen in a time-dependent model with the identical steady-state pattern of adjustment. That is, imagine varying the steady-state inflation rate, and for each inflation rate analyzing the time-dependent model with α_j's and ω_j's found by solving the state-dependent model. This comparison across inflation rates is unambiguous, and it is implicitly the one made in Ball, Mankiw, and Romer [1988]. Since a higher fraction of firms adjust their prices at higher inflation, and the fraction adjusting does not respond to shocks, the impact effect of money on output is lower, the higher is the inflation rate. However, with state-dependent pricing the story is not so simple, because the adjustment structure (the α_j's and ω_j's) is not fixed. In fact, the deviation from steady state of the adjustment probabilities is decreasing in the inflation rate. Given that with high inflation firms will be adjusting in the near future, it does not pay for them to change their adjustment probabilities much in response to a shock. This leads higher inflation to raise the impact effect of money on output, making the overall effect of higher inflation ambiguous. In our calibrated model, the relationship between inflation and the impact effect of money on output turns out to be nonmonotonic. At a 2.5 percent inflation rate the impact effect corresponding to the experiment in Figure IV is 0.591 percent, and the cumulative effect over four quarters is 0.851 percent. At a 5 percent inflation rate these effects falls to 0.448 percent and 0.813 percent, respectively. Finally, with 10 percent inflation the impact effect rises to 0.481 percent, while the cumulative effect continues to fall, to 0.764 percent.

III.6. Relationship to Other Work on State-Dependent Pricing

In this subsection we review the earlier work of Caplin and Spulber [1987] and Caplin and Leahy [1991, 1997] on state-dependent pricing. We describe the key findings of these studies, highlight differences between our results and theirs, and suggest perturbations of the various model economies that might reconcile these findings. While our conceptual framework and computational method are unlike those employed in these earlier studies, we think that the difference in results mainly arises from alternative assumptions about costs of price adjustment and the driving process for money.

Constructing a continuous time theoretical model in which each of a continuum of firms faces the same real cost of changing nominal prices, so that (S,s) adjustment rules are appropriate, Caplin and Spulber derive the striking implication that individual price stickiness does not lead to either aggregate price level stickiness or the nonneutrality of money. The money supply is governed by a continuous time stochastic process that is nondecreasing, so that money growth is nonnegative, and the money supply is not allowed to "jump." A beautiful attribute of their model is that an initial uniform distribution of nominal prices is preserved in the face of random monetary expansions. When such monetary expansions occur, there is individual price stickiness in the sense that most firms do not change prices, but the aggregate price level varies proportionately with the quantity of money. This occurs because firms that do adjust make a large change, moving from the lowest price in the distribution to set a new highest price.

An extension of our framework to accommodate discrete switches of the money growth rate and a limiting distribution of adjustment costs would likely make our model yield results similar to those of Caplin and Spulber [1987]. While the steady state of our model is one in which there is sustained inflation, in general, our price distribution is not uniform: as in Figure III, a higher fraction of firms currently adjust their price ($j = 0$) or have recently adjusted ($j = 1$) than have a price from x quarters ago ($j = x > 1$). To move our model's steady state toward Caplin and Spulber, we would need to assume a distribution of fixed costs that places most of the mass close to the highest value (called B above). Then, there would be a discrete distribution of prices that would be close to uniform in steady state; i.e., $\theta_j \approx (1/J)$. Next, suppose that we also assumed that the money supply growth rate could take on only two values, $g > 0$ with probability μ and 0 with probability $(1 - \mu)$.[34] With this pair of modifications, a Caplin-Spulber type of equilibrium would plausibly arise: if there was positive money growth, the price level would increase by g percent and $1/J$th of the population would adjust; while if there was zero money growth, no adjustment would occur. However, the value of g would have to be restricted so as to produce just the right set of incentives for price adjustment (just $1/J$ of the firms). If money growth was larger than some critical level, then neutrality would not occur because too many firms would choose to adjust; and if it was smaller, too few would choose to adjust. In fact, Caplin and

34. This modification would be nontrivial, as it would require solving the complete dynamic model nonlinearly, conditional on the two-state driving process.

Spulber take care to point out that "large" shocks would deform the distribution and invalidate their neutrality result; for that reason they rule out jumps in the money stock.[35] In our setting, even with smooth cost distributions that are aimed at generating a near-uniform price distribution, changes in the level of money are not neutral because money growth is not drawn from the two-state process described above. We therefore explore territory near that of the Caplin-Spulber model: we find that the neutrality result is fragile with respect to both the distribution of adjustment costs and the form of the driving process for money.

Caplin and Leahy [1991, 1997] develop an alternative (S,s) model that preserves the self-replicating uniform distribution of prices. In their setup the money stock may either increase or decrease, but it is required to be a continuous time stochastic process that does not jump and has a zero mean growth rate. Their basic model [1991] demonstrates that Caplin and Spulber's neutrality result is fragile with respect to the form of the money supply process, while maintaining other features of the earlier model. They also highlight a feature of nominal and real interactions that is surely a general implication of state-dependent pricing models: the effect of a monetary disturbance depends on the existing distribution of prices. Depending on the distribution of prices relative to the money stock, changes in the quantity of money affect *either* output *or* prices in Caplin and Leahy's model.

With its generalized (S,s) structure, our model economy would surely exhibit similar behavior: if the rate of money growth was high enough, for example, all firms would choose to adjust every period in the steady state and in response to transitory shocks.[36] Our present computational approach—linearization—abstracts from those features in studying the effect of monetary shocks: we assume that the disturbances being studied are small enough so that the economy never encounters "corners" in which adjustment rates (the α_j above) are driven to unity. It would be useful, but difficult, to extend our model to allow for such nonlinearity and history dependence. However, with our assumption of a generalized (S,s) structure and our use of linear system methods, we have begun to analyze how richer state-dependent models may work, going beyond the initial explorations of Caplin

35. Shocks that are "too small" cannot arise in the continuous time setting.

36. In addition, if we assumed a positive lower bound on the fixed cost of price adjustment, low enough positive rates of money growth would generate no price adjustment by some vintages of firms.

and Spulber and Caplin and Leahy. In particular, we construct a state-dependent pricing model in which monetary shocks affect both output and the price level, in contrast to Caplin and Leahy. Our distribution of fixed costs means that firms have time-varying patterns of price adjustment, which serve to smooth out Caplin and Leahy's discrete regions where money affects only the price level or output. In addition, our state-dependent pricing framework can be used with a wide range of assumptions about the sensitivity of product demand to relative price and output, the sensitivity of marginal cost to firm and aggregate conditions, the sensitivity of money demand to income and interest rates, as well as the form of the monetary policy rule.

IV. Conclusions and Extensions

We have developed a framework for incorporating state-dependent pricing into general equilibrium macroeconomic models. Individual firms adjust prices discretely at intervals of random length, and the frequency of price adjustment varies with the average inflation rate and the business cycle. Nevertheless, the resulting specification of price dynamics is roughly as tractable as the time-dependent pricing rules that are standard in small rational expectations business cycle models. This tractability stems from our assumption that the fixed costs of adjustment are randomly distributed across firms in a continuous manner, but are independent across time for a given firm.

In a basic macroeconomic model, this form of state-dependent pricing yields a short-run real effect of money on output and prices, but also conventional long-run neutrality. This type of dynamic response pattern is one that is shared with existing time-dependent pricing models, but not with previous work on state-dependent pricing. Yet, we also find that the extent of price stickiness and the extent of nonneutrality is very responsive to the nature of agents' beliefs about the permanence of monetary disturbances, because these beliefs affect the incentives that agents have to adjust the timing pattern of their price adjustments. From the perspective of our model then, time-dependent models have been appropriately criticized for treating the pattern of price adjustment as exogenous.

Our framework could be applied and extended in various ways. First, we can analyze the consequences of alternative shocks and policy rules. In the current paper we have focused on a single shock, viewed as a change in monetary policy with a money stock instrument. However, it is easy to study the response of prices and output to other shocks (including productivity shocks) and under different policy regimes (including interest rate rules). Second, the nature of changes in trend inflation with state-dependent pricing is of considerable interest. Third, it is feasible to explore the implications of a range of modifications of the macroeconomic equilibrium; for example, adding capital accumulation or mechanisms that generate real rigidities.

FEDERAL RESERVE BANK OF RICHMOND
UNIVERSITY OF VIRGINIA, FEDERAL RESERVE BANK OF RICHMOND, AND NATIONAL BUREAU OF ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF RICHMOND

REFERENCES

Ball, Laurence, N. Gregory Mankiw, and David Romer, "The New Keynesian Economics and the Output-Inflation Trade-off," *Brookings Papers on Economic Activity* (1988), 1–65.
Ball, Laurence, and David Romer, "Real Rigidities and the Non-Neutrality of Money," *Review of Economic Studies,* LVII (1990), 183–203.
Barro, Robert J. "A Theory of Monopolistic Price Adjustment," *Review of Economic Studies,* XXXIV (1972), 17–26.
Bénabou, Roland, "Inflation and Efficiency in Search Markets," *Review of Economic Studies,* LIX (1992), 299–329.
Blanchard, Olivier J., and Charles M. Kahn, "The Solution of Linear Difference Models under Rational Expectations," *Econometrica,* XLVIII (1980), 1305–1311.
Blanchard, Olivier J., and Nobuhiro Kiyotaki, "Monopolistic Competition and the Effects of Aggregate Demand," *American Economic Review,* LXXVII (1987), 647–666.
Blinder, Alan S., "On Sticky Prices: Academic Theories Meet the Real World," in N. Gregory Mankiw, ed., *Monetary Policy* (Chicago: University of Chicago Press, 1994), pp. 117–150.
Caballero, Ricardo J., and Eduardo M.R.A. Engel, "Explaining Investment Dynamics in U. S. Manufacturing: A Generalized (S,s) Approach," NBER Working Paper No. 4887, 1994.
Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics,* XII (1983), 383–398.
Caplin, Andrew, and John Leahy, "State-Dependent Pricing and the Dynamics of Money and Output," *Quarterly Journal of Economics,* CVI (1991) 683–708.
Caplin, Andrew, and John Leahy, "Aggregation and Optimization with State-Dependent Pricing," *Econometrica,* LXV (1997), 601–625.
Caplin, Andrew, and Daniel Spulber, "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics,* CII (1987), 703–725.
Cooper, Russell, John Haltiwanger, and Laura Power, "Machine Replacement and the Business Cycle: Lumps and Bumps," NBER Working Paper No. 5260, 1995.
Dotsey, Michael, Robert G, King, and Alexander L. Wolman, "State-Dependent Pricing and the Dynamics of Business Cycles," Federal Reserve Bank of Richmond Working Paper, February 1997a.
Dotsey, Michael, Robert G. King, and Alexander L. Wolman, "Menu Costs, Staggered Price-Setting and Elastic Factor Supply," manuscript, June 1997b.
Hansen, Gary D., "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics,* LVI (1985), 309–327.

King, Robert G., and Mark W. Watson, "The Solution of Singular Linear Difference Models under Rational Expectations," *International Economic Review,* forthcoming.

King, Robert G., and Mark W. Watson, "System Reduction and Solution Algorithms for Singular Linear Difference Systems under Rational Expectations," manuscript, University of Virginia, 1995.

King, Robert G., and Alexander L. Wolman, "Inflation Targeting in a St. Louis Model of the 21st Century," *Federal Reserve Bank of St. Louis Review,* LXXVIII (1996), 83–107.

Levin, Andrew, "The Macroeconomic Significance of Nominal Wage Contract Duration," University of California, San Diego Discussion Paper 91-08, 1991.

Rogerson, Richard, "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics,* XXI (1988), 3–16.

Romer, David, *Advanced Macroeconomics* (New York: McGraw-Hill, 1996).

Rotemberg, Julio J., "The New Keynesian Microfoundations," *NBER Macroeconomics Annual* (Cambridge, MA: MIT Press, 1987).

Sheshinski, Eytan, and Yoram Weiss, "Optimum Pricing Policy under Stochastic Inflation," *Review of Economic Studies,* L (1983), 513–529.

Taylor, John, "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy,* LXXXVIII (1980), 1–24.

Yun, Tack, "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles," *Journal of Monetary Economics,* XXXVII (1996), 345–370.