A Phillips Curve with an *Ss* Foundation

Mark Gertler and John Leahy

New York University and National Bureau of Economic Research

We develop an analytically tractable Phillips curve based on state-dependent pricing. We consider a local approximation around a zero inflation steady state and introduce infrequent idiosyncratic shocks. The resulting Phillips curve is a simple variant of the conventional time-dependent Calvo formulation with important differences. First, the model is able to match the micro evidence on the magnitude and timing of price adjustments. Second, our state-dependent model exhibits greater flexibility in the aggregate price level than the time-dependent model. With real rigidities present, however, our model can exhibit nominal stickiness similar to a conventional time-dependent model.

I. Introduction

In recent years there has been considerable progress in developing structural models of inflation and output dynamics. A common aspect of this approach is to begin with the individual firm's price-setting problem, obtain optimal decision rules, and then aggregate behavior. The net result is a simple relation for inflation that is much in the spirit of a traditional Phillips curve: Inflation depends on a measure of real activity, as well as expectations of future inflation. This relationship differs from the traditional Phillips curve in its forward-looking nature and in that all the coefficients are explicit functions of the primitives of the model.

We would like to thank Andrew Caplin, Jordi Gali, Mikhail Golosov, Per Krusell, Robert Lucas, Julio Rotemberg, Lars Svensson, Andrea Tambalotti, Michael Woodford, the editor, and three anonymous referees. Thanks to Andrea Ferrero for excellent research assistance. We thank the National Science Foundation for financial support.

[*Journal of Political Economy,* 2008, vol. 116, no. 3]

To date, these new Phillips curves (often grouped under the heading of "New Keynesian") reflect a pragmatic compromise between theoretical rigor and the need for empirical tractability.[1] While they evolve from optimization at the firm level, they typically restrict pricing behavior to time-dependent strategies in which the frequency of adjustment is given exogenously. A leading alternative, of course, is state-dependent pricing, where the firm is free to adjust whenever it would like, subject to a fixed adjustment cost. This latter approach, however, leads to "*Ss*" pricing policies that are, in general, difficult to aggregate.[2] For this reason, the time-dependent approach has proved to be the most popular, despite the unattractiveness of arbitrarily fixing the degree of price rigidity.

Besides tractability considerations, however, there have been two additional justifications for the time-dependent approach. First, Klenow and Kryvtsov (2005) have shown that during the recent low-inflation period in the United States, the fraction of firms that adjust their prices in any given quarter has been reasonably stable, which is certainly consistent with time-dependent pricing. Second, in this spirit, it is often conjectured that time-dependent models are the natural reduced forms of a state-dependent framework for economies with relatively stable inflation. Indeed Klenow and Kryvtsov provide support for these notions by showing that a conventional state-dependent pricing model (Dotsey et al. 1999) and a conventional time-dependent model (Calvo 1983) yield very similar dynamics when calibrated to recent U.S. data.

An interesting recent paper by Golosov and Lucas (2007) challenges this rationalization. The authors first note that to reconcile the evidence on the large size of individual firm price adjustments in the Klenow-Kryvtsov data with the low U.S. inflation rate, it is necessary to introduce idiosyncratic shocks that create sufficient dispersion in price adjustments. They then observe that in this environment, even if price adjustment frequencies are stable (because of moderate inflation variability), there remains an important difference between state dependence and time dependence: Under state-dependent pricing, the firms that find themselves farthest away from their target price adjust, whereas under time dependence there is no such relation. The authors then go on to show numerically that an exogenous shock to the money supply has a much stronger effect on the price level and a much weaker effect on real output in a state-dependent model with idiosyncratic productivity

[1] Examples include Gali and Gertler (1999), Gali, Gertler, and Lopez-Salido (2001), Sbordone (2002), and Eichenbaum and Fisher (2004).

[2] See Caplin and Spulber (1987), Benabou (1988), Caballero and Engel (1991), and Caplin and Leahy (1991, 1997) for early analyses of dynamic *Ss* economies. Dotsey, King, and Wolman (1999) place *Ss* policies within a standard dynamic stochastic general equilibrium model.

shocks than it does in a standard time-dependent model calibrated to have a similar degree of price stickiness at the firm level. In particular, they find that the "selection effect" associated with state-dependent pricing may lead to quantitatively important differences with time-dependent pricing models. Overall, their numerical exercise is reminiscent of the theoretical example in Caplin and Spulber (1987), where state dependence can turn the nonneutrality of money resulting from time dependence on its head.

Because pricing behavior in their model is very complex, Golosov and Lucas restrict attention to numerical solutions, as is typical in the *Ss* literature. In addition, they keep the other model features as simple as possible. Perhaps most significant, they abstract from interactions among firms that can lead to strategic complementarities in price setting. These complementarities—known in the literature as "real rigidities"—work to enhance the overall nominal inertia that a model of infrequent nominal price adjustment can deliver.[3] It is now well known that to obtain an empirically reasonable degree of nominal stickiness within a time-dependent pricing framework, it is critical to introduce real rigidities. Accordingly, abstracting from real rigidities makes it difficult to judge in general whether state dependence undoes the results of the conventional literature.

Our paper addresses this controversy by developing a simple state-dependent pricing model that allows for both idiosyncractic shocks and real rigidities. We differ from the existing *Ss* literature by making assumptions that deliver a model that is as tractable as the typical time-dependent framework. As with the standard time-dependent frameworks and the Dotsey et al. state-dependent framework, we focus on a local approximation around the steady state. We differ from Dotsey et al. by introducing idiosyncratic shocks, as in Golosov and Lucas's paper. We differ from Golosov and Lucas, in turn, by introducing several restrictions and technical assumptions that permit an approximate analytical solution. The end result is a Phillips curve built up explicitly from state-dependent pricing at the micro level that is comparable in simplicity and tractability to the standard New Keynesian Phillips curve that arises from the time-dependent pricing.

Because we restrict attention to a local approximation around a zero inflation steady state, our analysis is limited to economies with low and stable inflation. We thus cannot use our *Ss* framework to analyze the effect of large regime changes (which of course is also a limitation of the time-dependent approach). However, our framework does capture the "selection effect" of state-dependent pricing: those farthest away

[3] Ball and Romer (1990) first noted that for sticky price models to generate sufficient nominal inertia, real rigidities are critical. See Woodford (2003) for a recent discussion.

from the target tend to adjust more frequently, a feature that does not arise in time-dependent pricing. We can thus use our model to assess quantitatively how much extra price flexibility state dependence adds relative to time dependence, after allowing for the kinds of real rigidities thought to be important in the time-dependent literature.

In Section II we lay out the basic features of the model: a simple New Keynesian framework, but with state-dependent as opposed to time-dependent pricing. Firms face idiosyncratic productivity shocks. Borrowing some insights from Danziger (1999), we restrict the distribution of these shocks in a way that facilitates aggregation.

In Section III we discuss our approximation strategy, and in Section IV we characterize the firm's optimal pricing policy. We make assumptions on the size of the adjustment costs that make it reasonable to restrict attention to a second-order approximation of the firm's objective function. Given a quadratic firm objective and the restrictions on the idiosyncratic shock distribution, we are then able to derive the key theoretical result of the paper: a "simplification theorem" that makes the state-dependent pricing problem as easy to solve as the conventional time-dependent pricing problem, up to a second order.

In Section V we characterize the complete model and present a log-linear approximation about the steady state. In Section VI we discuss some of the properties of the model. Among other things, we derive a Phillips curve relation that is very similar in form to the New Keynesian Phillips curve, except of course that it is based on state-dependent pricing. The slope coefficient on the real activity measure in the Phillips curve reflects this distinction. The slope coefficient in our state-dependent Phillips curve is larger than under time dependence. In this respect, our state-dependent framework exhibits greater price flexibility than the corresponding time-dependent framework. As in Golosov and Lucas's paper, the selection effect is at work. How much extra price flexibility state dependence induces depends on the model calibration and whether real rigidities are present.

In Section VII we calibrate the model to match the Klenow-Kryvtsov evidence on the frequency and absolute magnitude of price adjustments and also evidence on the costs of price adjustment. We then show that the model can deliver the kind of aggregate price-level stickiness emphasized in the time-dependent literature and yet remain consistent with the microeconomic evidence on price adjustment. Key to this result, as we show, is allowing for real rigidities. Concluding remarks are in Section VIII. Finally, we present some numerical simulations in the appendices to demonstrate that our approximate analytical solution provides a reasonably accurate description of local model dynamics.

II. The Model

We begin with a conventional New Keynesian model (see, e.g., Woodford 2003). The basic features of the standard model include monopolistic competition, money, and nominal price stickiness. Also, for convenience, there are only consumption goods. We add three features to this familiar baseline framework: real rigidities, idiosyncratic productivity shocks, and state-dependent pricing. It is of course incorporating this latter feature that poses the biggest challenge.

In particular, state-dependent pricing raises two difficult modeling issues. The first we alluded to in the introduction: In the most general state-dependent pricing model, the entire distribution of prices is a state variable. This gives rise to an intractable fixed-point problem. Inevitably, there is a need for some kind of simplifying assumptions or shortcuts. As we noted, our strategy is to make restrictions on the distribution of idiosyncratic shocks to simplify the distributional dynamics.[4] Specifically, we use a formulation in the spirit of Danziger (1999), who by assuming a uniform distribution of shocks was able to solve a carefully parameterized *Ss* economy in closed form. As we show, however, the effects of money on output are small (i.e., second-order and above) for the case he was able to solve. We thus alter Danziger's formulation by allowing for a more flexible parameterization of the idiosyncratic shock process, one that makes possible significant nominal inertia and hence a significant first-order effect of money on output. Specifically, each period a firm receives a uniformly distributed idiosyncratic productivity shock with probability $1 - \alpha$ and no shock with probability α. Though we cannot solve for an exact solution with this more flexible process, we can obtain an approximate analytical solution by considering a local expansion of the model around a zero inflation steady state (as is done in the time-dependent literature; see, e.g., Woodford 2003). Aside from having these very convenient features, the resulting unconditional distribution of the productivity shock has a simple fat-tailed form of the type that Midrigan (2006) argues is broadly consistent with the evidence.[5]

The second modeling issue arises from the discontinuities and nondifferentiabilities associated with *Ss* adjustment. This issue potentially

[4] Caplin and Spulber (1987), Benabou (1988), and Caplin and Leahy (1991, 1997) also make distributional assumptions that reduce the state space. Dotsey et al. (1999) make assumptions that limit the number of prices observed in the economy to a finite number. Willis (2002) and Midrigan (2006) follow Krusell and Smith (1998) and approximate the distribution by a finite number of moments. Golosov and Lucas (2007) avoid the fixed-point problem by setting variables (except for the wage, which they take to be exogenous) at their steady-state values when computing firm decision rules.

[5] Given that the idiosyncratic shock does not arrive with probability α, the unconditional distribution of productivity shocks the firm faces has a single peak at zero and then is uniformly distributed about zero.

complicates finding a log-linear approximation of the model since Taylor's theorem does not apply to functions that are not differentiable. Fortunately, this technical problem is applicable to only a small percentage of firms that happen to lie near the *Ss* bands. In particular, for a firm near either of the *Ss* boundaries that does not receive an idiosyncratic shock in the current period, an aggregate shock in one direction may cause it to adjust, whereas a shock in the other direction will take it deeper into the inaction region. In this instance there is a kink in the firm's response to the aggregate state. We address the issue by assuming that in addition to the fixed cost of adjusting the price, there is a small "decision cost" to contemplating a price adjustment prior to the decision whether to adjust. If aggregate shocks are sufficiently small relative to idiosyncratic shocks, there will be a range of decision costs for which the firm considers adjusting only when an idiosyncratic shock hits. We assume that the decision costs are in this range. This assumption leads to smooth behavior of firms as they approach the *Ss* bands, eliminating any complication to log-linearizing our model. In Appendix C we confirm that the decision cost is not of quantitative importance for the local dynamics of the model.

In the remainder of this section we lay out the basic ingredients of the model. There are three types of agents: households, final goods firms, and intermediates goods firms. We describe each in turn.

A. *Households*

Households consume, supply labor, hold money, and hold bonds. The latter are zero in net supply. We assume a segmented labor market in order to generate strategic complementarities in price setting as in Woodford (2003). In particular, we assume a continuum of "islands" of mass unity. On each island, there is a continuum of households of mass unity. Households can supply labor only on the island they live on. There is perfect consumption insurance across islands, and any firm profits are redistributed lump-sum to households.

Time is discrete and is indexed by t. Let C_t be consumption; M_t end-of-period nominal money balances; P_t the nominal price index; $N_{z,t}$ labor supply on island z; $W_{z,t}$ the nominal wage on island z; $\Gamma_{z,t}$ lump-sum transfers (including insurance, dividends, and net taxes); B_t one-period nominal discount bonds; and R_t^n the nominal interest rate from t to $t+1$. Then the objective for a representative household on island z is given by

$$\max E_t \sum_{i=0}^{\infty} \beta^i \left\{ \log\left[C_{t+i} \cdot \left(\frac{M_{t+i}}{P_{t+i}} \right)^{\nu} \right] - \frac{1}{1+\varphi} N_{z,t+i}^{1+\varphi} \right\} \tag{1}$$

subject to the budget constraint

$$C_t = \frac{W_{z,t}}{P_t} N_{z,t} + \Gamma_{z,t} - \frac{M_t - M_{t-1}}{P_t} - \frac{(1/R_t^n)B_t - B_{t-1}}{P_t}. \tag{2}$$

We index labor supply and the nominal wage by z because the island z labor market is segmented. Since there is perfect consumption insurance, there is no need to similarly index the other variables, except for lump-sum transfers, which may be island specific.

The first-order necessary conditions for labor supply, consumption/saving, and money demand are given by

$$\frac{W_{z,t}}{P_t} = \frac{N_{z,t}^{\varphi}}{1/C_t}, \tag{3}$$

$$E_t\left\{\beta \frac{C_t}{C_{t+1}} R_t^n \frac{P_t}{P_{t+1}}\right\} = 1, \tag{4}$$

and

$$\frac{M_t}{P_t} = \nu C_t \frac{R_t^n}{R_t^n - 1}. \tag{5}$$

B. *Final Goods Firms*

Production occurs in two stages. Monopolistically competitive intermediate goods firms employ labor to produce input for final goods. There is a continuum of mass unity of these intermediate goods firms on each island. Final goods firms package together all the differentiated intermediate inputs to produce output. These firms are competitive and operate across all islands.

Let Y_t be output of the representative final good firm, $Y_{z,t}^j$ be input from intermediate goods producer j on island z, and $Q_{z,t}^j$ be the associated nominal price. The production function for final goods is the following constant elasticity of substitution (CES) aggregate of intermediate goods:

$$Y_t = \left[\int_0^1 \int_0^1 (Y_{z,t}^j)^{(\varepsilon-1)/\varepsilon} dj dz\right]^{\varepsilon/(\varepsilon-1)}, \tag{6}$$

where $\varepsilon > 1$ is the price elasticity of demand for each intermediate good.

From cost minimization, the demand for each intermediate good is given by

$$Y^j_{z,t} = \left(\frac{Q^j_{z,t}}{P_t}\right)^{-\varepsilon} Y_t, \tag{7}$$

and the price index is the following CES aggregate of intermediate goods prices:

$$P_t = \left[\int_0^1 \int_0^1 (Q^j_{z,t})^{1-\varepsilon} dj dz\right]^{1/(1-\varepsilon)}. \tag{8}$$

C. *Intermediate Goods Firms*

Each intermediate goods firm produces output that is a linear function of labor input:

$$Y^j_{z,t} = X^j_{z,t} \cdot N^j_{z,t}. \tag{9}$$

Here $X^j_{z,t}$ is an idiosyncratic productivity factor for producer j on island z. (For simplicity we abstract from aggregate productivity shocks, though we can easily add them.)

Islands are occasionally subject to turbulence in the form of productivity shocks. These shocks follow a compound Poisson process. Each period a shock hits an island with probability $1 - \alpha$. These shocks are independent across islands.

When a shock hits an island, two things happen. First, a random fraction $1 - \tau$ of firms die and are replaced with an equal number of new entrants. New entrants are born with $X^j_{z,t} = 1$ and set price optimally within the period in which they enter. Second, the remaining τ firms are hit with multiplicative independent and identically distributed productivity shocks. Let $\xi^j_{z,t}$ denote the shock to firm j on island z at date t. We assume that $\xi^j_{z,t}$ is distributed uniformly with mean zero and a density $1/\phi$. The evolution of productivity is therefore

$$X^j_{z,t} = \begin{cases} X^j_{z,t-1} & \text{if no productivity shock occurs} \\ X^j_{z,t-1} e^{\xi^j_{z,t}} & \text{if a productivity shock occurs} \\ & \text{and the firm survives.} \end{cases} \tag{10}$$

The only purpose of the death probability is to make the distribution of productivity in the economy stationary.[6] It will drop out of the aggregate equations.

Each producer faces a fixed cost of adjusting its price. We assume

[6] Random walk shocks are convenient since it is possible to make the *Ss* bands homogeneous in the shock. The cost is that they are nonstationary.

that the firm incurs a cost equal to $b(X^j_{z,t})^{\varepsilon-1}$ if and only if it chooses to alter its nominal price. The adjustment cost $b>0$ will lead to *Ss*-style price adjustment policies. There will be a range of inaction in which firms keep their price fixed. Firms with prices outside of this range will adjust to a new optimum. We scale the adjustment cost by the factor $(X^j_{z,t})^{\varepsilon-1}$ to keep the firm's decision problem homogeneous as it size varies. This adjustment cost is in units of the final consumption good.

As we discussed in the introduction of this section, we introduce a cost to gathering and processing information to address a potential problem of nondifferentiability, which we refer to as a "decision cost." We assume the following: Firms know when idiosyncratic turbulence hits their island, but to gather information about the precise value of the shock ξ they receive, and also to organize this information and information regarding the aggregate economy, they must pay a decision cost, $d \cdot (X^j_{z,t})^{\varepsilon-1}$. If a firm elects to pay the decision cost, it can then decide whether to adjust its price. If it chooses to adjust, then it also incurs the fixed cost $b(X^j_{z,t})^{\varepsilon-1}$.

D. *Money Supply*

We close the model by characterizing monetary policy. In particular, we assume that the nominal money stock M_t obeys a simple exogenous first-order Markov process. It is straightforward to extend the model to allow for a richer characterization of monetary policy (see, e.g., Clarida, Gali, and Gertler 2000). However, we stick with this simple process because it provides the best way to illustrate the implications of our state-dependent framework for nominal stickiness.

We consider "small" monetary shocks in the approximation below. We introduce a set of relationships that will allow us to be explicit about what we mean by small. As we have just discussed, our approximation method will rely on the aggregate shock not having "too large" an effect on the firm's desired price in between periods in which an idiosyncratic shock hits. What this requires is that the money growth rate be bounded over a horizon long enough to have the fraction of firms that survive the interval without receiving an idiosyncratic shock be small. A simple condition that will ensure this is

$$|\ln M_{t+k} - \ln M_t| \leq m \qquad \text{for all } k < k(m), \tag{11}$$

where $k(m)$ is the least integer that satisfies $\alpha^{k(m)} \leq cm$ and c is a fixed positive constant. The parameter m scales the money shock. As m becomes small, two things happen. First, movements in the money supply will become small over $k(m)$ periods. Second, $k(m)$ rises and the probability that a firm survives $k(m)$ periods without an idiosyncratic shock will also become small. This ensures that both money growth and the survival probability over the relevant horizon will be first-order in m. For reasonable model parameter values, the implied bound on money growth is not overly restrictive.[7]

III. The Approximation

We approximate the model around a zero inflation steady state, with no shocks and no price adjustment costs. We assume that the ranges of the money shock m and the idiosyncratic shock ϕ are first-order. At the same time, we capture the notion that the adjustment cost b and the decision cost d are small by assuming that they are second-order. A virtue of having second-order adjustment costs is that they lead to *Ss* bands that are first-order (e.g., Akerlof and Yellen 1985; Mankiw 1985).

Formally, let $\varsigma = \{m, \phi, \sqrt{b}, \sqrt{d}\}$. When $\varsigma = \mathbf{0}$, the economy experiences no shocks and no price adjustment frictions, as well as zero inflation (since M_t is fixed over an infinite horizon when $m = 0$). We will consider a first-order perturbation in ς about a value $\varsigma = \mathbf{0}$. Note that since $\sqrt{b}$ and $\sqrt{d}$ are first-order, b and d will be second-order. Since m is first-order, $\alpha^{k(m)}$ will be first-order (see eq. [11]).

In addition, our aggregation strategy relies on the fact that in the data idiosyncratic shocks appear to be large relative to aggregate shocks. We use this to motivate certain relationships among the four elements of ς. First, in order for the uniformity assumption to induce simple distributional dynamics, we need there always to be a productivity shock that causes a firm to raise its price and a productivity shock that causes a firm to lower its price. For this we need ϕ to be large relative to m and $\sqrt{b}$. A sufficient condition is

$$\phi > 4\omega + 2m, \tag{12}$$

where ω is the distance between either price adjustment trigger and the target.[8] To interpret this condition, note that the first term on the right-hand side is equal to twice the steady-state log distance between the

[7] What is critical is that the variation in aggregate nominal demand is bounded, which within this model translates into exogenous variation in the money supply. A better measure of nominal aggregate demand in the real world is nominal GDP, since there is considerable movement in the money supply that reflects factors such as financial innovation opposed to movements in nominal demand. In App. C we show that for empirically reasonable variation in nominal GDP, our restriction on the magnitude of the aggregate shock is satisfied as are our other parameter restrictions.

Under our baseline calibration, α is roughly 0.5. With $cm = 0.05$, $k = 4$, which implies that the bound applies to the annual growth rate of the money supply.

[8] We will see below that

$$\omega = \sqrt{2\frac{1-\alpha\beta}{1-\varepsilon}\left(\frac{\varepsilon}{\varepsilon-1}\right)^{1/(\varphi+1)} b}.$$

upper and lower price adjustment triggers. Since the shock is mean zero in logs, it must have a support greater than twice this amount in order to take a firm at the upper trigger below the lower trigger. The second term is the maximal cumulative nominal disturbance over the past $k(m)$ periods. This bounds how far the bands may move for the fraction $1 - cm$ of firms that set their prices during the past $k(m)$ periods. This term may be made tighter depending on the exact relationship between nominal marginal cost and the money supply.

Second, in order to solve the nondifferentiability problem involving firms near the *Ss* bands, we require that the decision cost be large enough to prevent these firms from adjusting to aggregate conditions in the absence of an idiosyncratic shock. This will eliminate any asymmetric response of these firms to aggregate shocks. Note that this requires that the decision cost be in a specific range: large enough that firms are not willing to pay it in the absence of idiosyncratic turbulence but small enough to make them willing to pay it when they receive news that an idiosyncratic shock has hit. These considerations place further restrictions on the parameters of the model: ϕ must be large relative to d and movements in the money supply must be small. Since the implied bounds on d depend on the firm's value function, we delay presenting them until after solving the model. At this point we just assume that d is in the required range, so that firms adjust only after an idiosyncratic shock. The firm delays adjusting to aggregate conditions until there is a sufficiently large idiosyncratic shock.

In Appendix C, we derive the explicit bounds on d and show that for our calibrated model there is indeed an admissible range d that is empirically reasonable. In particular, d needs to be between about one and three times the size of the menu cost b.[9]

In addition, Appendix C also presents a numerical solution to a version of the model that omits the decision cost and shows that it yields dynamics nearly identical to those of our baseline model that includes the decision costs. Intuitively, since aggregate shocks tend to be small relative to the *Ss* bands, the decision cost affects only a small number of firms near the price adjustment triggers and affects the aggregate dynamics only through the effect that these firms have on the price level.[10]

[9] While we introduce d for technical reasons, Zbaracki et al. (2004) find that the managerial costs of information gathering and decision making are larger than the physical costs of adjusting prices b. Moreover, Fabiani et al. (2004) find that firms in the Euro area review their prices more often than they change them.

[10] Appendix C shows that under our baseline calibration, the fraction of firms close to the *Ss* bands is very small. Indeed, at the end of any given period most firms are concentrated in the center of the *Ss* bands. Intuitively, all firms that adjusted price in the period move to the center of the bands. Further, a fraction of firms that were around the center at the beginning of the period do not receive a new idiosyncratic shock.

IV. The Firm's Optimal Pricing Decision

In this section we first characterize the firm's objective function. We argue that given our approximation, it is reasonable to consider a second-order approximation of the objective function. We then show that our restriction on adjustment costs in conjunction with the uniform distribution of the shock leads to considerable simplification of the objective, up to a second order. We use this quadratic objective to derive a log-linear approximation of the decision rules about the steady state. Throughout, we make use of our assumptions about the decision cost to restrict firms to adjusting price only during periods of idiosyncratic turbulence.

A. The Firm's Objective

At this point we drop the (j, z) subscripts except where clarity is a concern. Real profits net of the menu cost, $\tilde{\Pi}^j_{z,t}$, are given by

$$\tilde{\Pi}_t = \left(\frac{Q_t}{P_t} - \frac{W_t}{P_t X_t}\right) Y_t - b_t X_t^{\varepsilon-1}, \tag{13}$$

where b_t is equal to b if $Q_t \neq Q_{t-1}$ and zero otherwise. Note that the firm's real marginal cost is $W_t/(P_t X_t)$.

It is convenient to recast the firm's problem in terms of the markup $\mu^j_{z,t}$:

$$\mu_t = \frac{Q_t X_t}{W_t}. \tag{14}$$

There are two advantages of using the markup as the decision variable as opposed to the posted price, Q_t. The first is that all firms that reset their price in period t will wind up choosing the same value of μ_t. By contrast, given that the labor market is competitive within the island (which implies that all firms on the island face the same wage, W_t), the optimal reset value of Q_t will vary inversely with X_t. The second advantage is that the markup is stationary.

Note that because μ_t depends on X_t and W_t, it may change even if the firm keeps its nominal price constant. That is, if we let μ_t^* be the optimal reset value of μ_t,

$$\mu_t = \begin{cases} \mu_t^* & \text{if } Q_t \neq Q_{t-1} \\ \mu_{t-1}\dfrac{X_t}{X_{t-1}}\dfrac{W_{t-1}}{W_t} & \text{if } Q_t = Q_{t-1}. \end{cases} \tag{15}$$

Effectively, the firm's decision problem is to judge whether its markup in the absence of a price adjustment has drifted sufficiently away from the optimal reset price to justify the fixed adjustment cost. Restating period profits in terms of the markup and making use of the firm's demand function (7) yields

$$\tilde{\Pi}_t = X_t^{\varepsilon-1}\left[\left(\frac{P_t}{W_t}\right)^{\varepsilon-1} Y_t\mu_t^{-\varepsilon}(\mu_t - 1) - b_t\right].$$

We let Π_t denote the term in brackets, so that $\tilde{\Pi}_t \equiv X_t^{\varepsilon-1}\Pi_t$.

At this point it is useful to define several variables:

$$A_t = P_t^{\varepsilon-1}W_t^{1-\varepsilon}Y_t,$$

$$\Lambda_{t,t+i} = \frac{U'(C_{t+i})}{U'(C_t)},$$

$$\tilde{\mu}_t = \frac{Q_{t-1}}{W_t/X_t} = \mu_{t-1}\frac{X_t}{X_{t-1}}\frac{W_{t-1}}{W_t},$$

$$\tilde{\alpha} = \alpha + (1-\alpha)\tau.$$

The term $A_{z,t}$ summarizes the effect of the economy on current profits, and $\Lambda_{t,t+i}$ is the stochastic discount factor. The term $\tilde{\mu}_t$ is the markup that is inherited from the prior period. It depends on last period's nominal price Q_{t-1} and is the markup in the case of nonadjustment. The term $\tilde{\alpha}$ is the probability that the firm survives to the next period.

The firm's value function is then

$$\begin{aligned} V(\tilde{\mu}_t, X_t, \Omega_t) &= \max E_t\sum_{i=0}^{\infty}(\tilde{\alpha}\beta)^i\Lambda_{t,t+i}\tilde{\Pi}_{t+i} \\ &= \max X_t^{\varepsilon-1}E_t\sum_{i=0}^{\infty}(\tilde{\alpha}\beta)^i\Lambda_{t,t+i}\left(\frac{X_{t+i}}{X_t}\right)^{\varepsilon-1}\Pi_{t+i}. \end{aligned} \tag{16}$$

The term V depends on the inherited markup $\tilde{\mu}_t$, the level of productivity X_t, and Ω_t, which summarizes variables exogenous to the firm and depends on the current values of C_t, Y_t, $W_{z,t}$, and P_t, as well as their future evolution.[11]

Given that gross profits and adjustment costs are homogeneous in $X_t^{\varepsilon-1}$, it is convenient to define the normalized value function $v(\tilde{\mu}_t, \Omega_t)$:

$$V(\tilde{\mu}_t, X_t, \Omega_t) = X_t^{\varepsilon-1}\cdot v(\tilde{\mu}_t, \Omega_t), \tag{17}$$

[11] Note that we suppress the decision cost. Since it is paid whenever the idiosyncratic turbulence hits, it is effectively exogenous.

with

$$v(\tilde{\mu}_t, \Omega_t) = \max E_t \sum_{i=0}^{\infty} \left(\frac{X_{t+i}}{X_t}\right)^{\varepsilon-1} (\tilde{\alpha}\beta)^i \Lambda_{t,t+i}[A_{t+i}\mu_{t+i}^{-\varepsilon}(\mu_{t+i} - 1) - b_{t+i}].$$

We now express the normalized value function in a recursive form and incorporate our assumption that because of the decision cost the firm adjusts only in periods of idiosyncratic turbulence. Let $\bar{v}(\mu_t, \Omega_t)$ denote the value after price adjustment. In deciding whether or not to adjust in period t, the firm compares $\bar{v}$ at the inherited markup to $\bar{v}$ at the optimal markup:

$$v(\tilde{\mu}_t, \Omega_t) = \max\{\bar{v}(\tilde{\mu}_t, \Omega_t), \ \max_{\mu_t} \bar{v}(\mu_t, \Omega_t) - b\}. \quad (18)$$

Then $\bar{v}$ is equal to current profits plus the present value of an optimal policy tomorrow:

$$\bar{v}(\mu_t, \Omega_t) = A_t\mu_t^{-\varepsilon}(\mu_t - 1) + \alpha\beta E_t\Lambda_{t,t+1}\bar{v}\left(\mu_t\frac{W_t}{W_{t+1}}, \Omega_{t+1}\right)$$
$$+ (1-\alpha)\tau\beta E_t\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}}v\left(\mu_t\frac{W_t}{W_{t+1}}e^{\xi_{t+1}}, \Omega_{t+1}\right). \quad (19)$$

With probability α, no idiosyncratic shock hits tomorrow and the firm maintains its price. The continuation value is $\bar{v}$ where the markup changes with the change in the wage rate. With probability $(1-\alpha)\tau$, the firm survives idiosyncratic turbulence and chooses whether or not to adjust its price. The continuation value is v, and the inherited markup reflects both the change in the wage rate and the innovation in productivity.

B. *Approximate Value Function*

Our goal in this subsection is to derive an approximate value function that leads to a tractable (approximate) solution to the decision problem. We approximate the firm's problem for small ς. We do so in two steps. First, we show that second-order b in conjunction with the uniform distribution of the productivity shock implies that the continuation value contingent on an idiosyncratic shock at date $t+1$ is independent of μ_t up to a second order. Second, we use the fact that second-order adjustment costs lead to a first-order range of inaction to justify taking a second-order approximation of profits.

It is convenient to define the target and trigger in logarithmic terms. Let $\ln \mu_t^*$ denote the natural log of the target markup and let $\ln \mu_t^L$ and $\ln \mu_t^H$ be the natural logs of the upper and lower triggers. Under the *Ss* policy, the firm adjusts to $\ln \mu_t^*$ if $\ln \mu_t \notin [\ln \mu_t^L, \ln \mu_t^H]$.

1. A Simplification Theorem

What complicates the firm's problem is that it must take account of the continuation value conditional on an idiosyncratic shock, $E_t\{\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}}v(\cdot)\}$, in equation (19). Without this consideration, the choice of the target price at time t would just involve taking into account discounted profits in states in which the firm's price and productivity remain fixed. In this respect, the choice of the target is no more difficult than in the conventional time-dependent framework. The choice of the triggers is also simplified.

We now show that under our assumptions, $E_t\{\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}}v(\cdot)\}$ is independent of the firm's period t choice of the target, up to a second-order approximation. The decision problem will then be simplified, along the lines we have just suggested.

Proposition 1. Suppose that (a) b is second-order in ς (implying that $\ln\mu_t^L - \ln\mu_t^*$ and $\ln\mu_t^H - \ln\mu_t^*$ are first-order), and (b) ϕ satisfies (12). Then the expected value at date t of an optimal policy after an idiosyncratic shock at date $t+1$, $E\{e^{(\varepsilon-1)\xi_{t+1}}v(\tilde{\mu}_{t+1}, \Omega_{t+1})\}$, is independent of the current value of μ_t to a second order. In particular, the firm can treat its objective as

$$\bar{v}\left(\frac{Q_t}{W_t/X_t}, \Omega_t\right) = \Pi_t + \alpha\beta E_t\left\{\Lambda_{t,t+1}v\left(\frac{Q_t}{W_{t+1}/X_{t+1}}, \Omega_{t+1}\right)\right\} + [\text{terms independent of } \mu_t] + \mathcal{O}^3(\|\varsigma\|). \quad (20)$$

The main insight of the proposition is that in future states in which the idiosyncratic shock will hit, history will be erased. The subsequent continuation value $E_t\{\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}}v(\cdot)\}$ is irrelevant to the current pricing decision to a second order.

In Appendix A we provide a formal proof of the proposition. Here we present the intuition, which follows from figure 1. Consider a firm with log markup equal to $\ln\mu_t$ that receives an idiosyncratic shock in period $t+1$. The shock leaves the log markup uniformly distributed over the interval

$$\left[\ln\mu_t + \ln\frac{W_t}{W_{t+1}} - \frac{\phi}{2},\ \ln\mu_t + \ln\frac{W_t}{W_{t+1}} + \frac{\phi}{2}\right] \equiv [A,\ B].$$

Now in period $t+1$ the firm follows a pricing strategy characterized by the triplet $\{\ln\mu_{t+1}^L,\ \ln\mu_{t+1}^*,\ \ln\mu_{t+1}^H\}$. Since ϕ satisfies (12), $[\ln\mu_{t+1}^L, \ln\mu_{t+1}^H] \subset [A, B]$. Given the policy, $\ln\mu_{t+1}$ will be uniformly distributed over $(\ln\mu_{t+1}^L,\ \ln\mu_{t+1}^H)$ if the firm does not adjust (the dark gray region in fig. 1). If the firm does adjust (the light gray regions of fig. 1), then

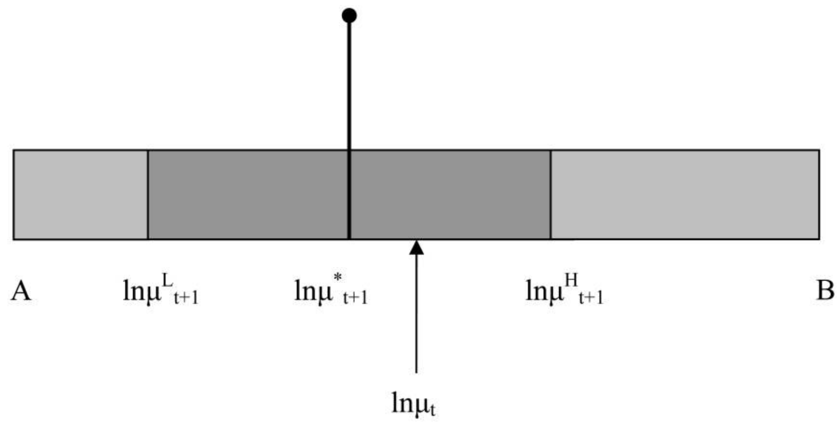

Fig. 1.—A firm's response to a productivity shock

$\ln \mu_{t+1} = \ln \mu^*_{t+1}$. Since the triplet $\{\ln \mu^L_{t+1},\ \ln \mu^*_{t+1},\ \ln \mu^H_{t+1}\}$ is independent of μ_t (it depends only on the state at $t+1$), it follows that the distribution of μ_{t+1} and hence $\bar{v}(\mu_{t+1}, \Omega_{t+1})$ is independent of μ_t. This can be seen from figure 1: a shift in $\ln \mu_t$ shifts the entire interval $[A, B]$. This does not affect the distribution after adjustment (the mass at $\ln \mu^*_{t+1}$ and the dark gray region); it affects only the states in which the firm adjusts up and the states in which the firm adjusts down (the light gray region).

Accordingly, the only way that μ_t could possibly affect $E_t\{\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}}v(\cdot)\}$ is by affecting the correlation between $\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}}$ and $v(\cdot)$. Given our assumption on b, this correlation is second-order and its dependence on μ_t is third-order.

The proposition rests on two critical assumptions. The first is that the idiosyncratic shock is uniform and has a wide enough support that both price increases and price decreases are possible. This assumption implies that the distribution of prices within the *Ss* bands is independent of μ_t. The second is that b is second-order, which makes the correlation between the decision to change price and μ_t third-order.

2. Approximate Optimal Pricing Policy

Let bars above variables denote values in the nonstochastic steady state with $\varsigma = \mathbf{0}$. Let hats above variables denote log deviations from these steady-state values.

Armed with the preceding proposition, we now take a second-order approximation of the profit function about the frictionless optimal markup $\bar{\mu} = \varepsilon/(\varepsilon - 1)$:

$$\Pi_t = \chi_1 \bar{Y}\left(\frac{\overline{W}}{\overline{P}}\right)^{1-\varepsilon} - \chi_2 \bar{Y}\left(\frac{\overline{W}}{\overline{P}}\right)^{1-\varepsilon}(\ln \mu_t - \ln \bar{\mu})^2 - b_t$$
$$+ \text{[terms independent of } \mu_t] + \mathcal{O}^3(\|\varsigma\|),$$

where χ_1 and χ_2 are constants, with $\chi_2 = \frac{1}{2}[\varepsilon^{1-\varepsilon}/(\varepsilon - 1)^{-\varepsilon}]$. Because we are approximating the profit function about the frictionless optimal markup, the first-order term is zero.[12] This proposition has the flavor of an envelope theorem. The "terms independent of μ_t" include effects of aggregate variables on firm profits.

Proposition 1 implies that we can ignore the continuation values in all states in which the idiosyncratic shock arrives. Since the decision cost d implies that firms consider adjustment only following the idiosyncratic shock, it follows that Q and X remain fixed in all states in which the idiosyncratic shock does not hit. We can therefore write $\bar{v}$ as

$$\bar{v}\left(\frac{Q_t X_t}{W_t}, \Omega_t\right) = E_t \sum_{i=0}^{\infty} (\alpha\beta)^i \left[\chi_1 \bar{Y}\left(\frac{\overline{W}}{\overline{P}}\right)^{1-\varepsilon} - \chi_2 \bar{Y}\left(\frac{\overline{W}}{\overline{P}}\right)^{1-\varepsilon}\left(\frac{Q_t X_t}{W_{t+i}} - \ln \bar{\mu}\right)^2\right]$$
$$+ \text{[terms independent of } \mu_t] + \mathcal{O}^3(\|\varsigma\|). \quad (21)$$

The first term on the right-hand side gives the quadratic approximation to profits in the states in which the idiosyncratic shock does not hit. These are weighted by α^i, the probability that there is no shock for i periods in succession. The Q and X terms in this expression are dated t since in these states no idiosyncratic shock hits and the price remains constant.

It is now straightforward to derive the optimality conditions for the target and the two triggers. The first-order necessary condition for the target is that the expected markup over the life of the price is zero:

$$E_t \sum_{i=0}^{\infty} (\alpha\beta)^i (\ln \mu_t - \ln \bar{\mu}) = 0. \quad (22)$$

Given $\mu = QX/W$, it follows immediately that the optimal nominal price $\ln Q_t^*$ satisfies

$$\ln Q_t^* = \ln \bar{\mu} - \ln X_t + (1 - \alpha\beta) E_t \sum_{i=0}^{\infty} (\alpha\beta)^i \ln W_{t+i}. \quad (23)$$

As in the pure time-dependent model, the target depends on a discounted stream of future values of nominal marginal cost. In the time-

[12] Linear-quadratic approximations are often inappropriate when the first-order terms in the objective are nonzero. See the discussion in Woodford (2002).

dependent framework, however, future marginal cost in each period is weighted by the probability that the price remains fixed. In our state-dependent framework, the relevant weight is the probability α^i that a new idiosyncratic shock has not arisen, which in general is a number smaller than the probability that the price has stayed fixed.

The price adjustment triggers are given by a value-matching condition that equates the gain from not adjusting to the gain from adjustment, net of the adjustment cost. For $J = H, L$,

$$\bar{v}(\mu_t^J, \Omega_t) = \bar{v}(\mu_t^*, \Omega_t) - b. \tag{24}$$

Given our quadratic approximation, we restate this condition on the optimal markup in terms of the corresponding pricing policy $\{Q_t^L, Q_t^*, Q_t^H\}$, since, unlike the markup, price remains fixed over the relevant horizon:

$$E_t\sum_{i=0}^{\infty}(\alpha\beta)^i\left[\chi_2\bar{Y}\left(\frac{\overline{W}}{\bar{P}}\right)^{1-\varepsilon}\left(\ln\frac{Q_t^J X_t}{W_{t+i}} - \ln\bar{\mu}\right)^2\right] =$$

$$E_t\sum_{i=0}^{\infty}(\alpha\beta)^i\left[\chi_2\bar{Y}\left(\frac{\overline{W}}{\bar{P}}\right)^{1-\varepsilon}\left(\ln\frac{Q_t^* X_t}{W_{t+i}} - \ln\bar{\mu}\right)^2\right] + b.$$

Rearranging, we get

$$\ln Q_t^H = \ln Q_t^* + \sqrt{(1-\alpha\beta)\frac{b}{\chi_2\bar{Y}(\overline{W}/\bar{P})^{1-\varepsilon}}} = \ln Q_t^* + \sqrt{2\frac{1-\alpha\beta}{\varepsilon-1}\frac{b}{\bar{Y}}} \tag{25}$$

and

$$\ln Q_t^L = \ln Q_t^* - \sqrt{(1-\alpha\beta)\frac{b}{\chi_2\bar{Y}(\overline{W}/\bar{P})^{1-\varepsilon}}} = \ln Q_t^* - \sqrt{2\frac{1-\alpha\beta}{\varepsilon-1}\frac{b}{\bar{Y}}}. \tag{26}$$

The second equality follows from noting that $\bar{P} = \bar{\mu}\overline{W}$ and substituting for χ_2 and $\bar{\mu}$ in terms of ε. Note that the bands $\ln\mu^H - \ln\mu_t^*$ and $\ln\mu^L - \ln\mu^*$ are first-order in $\sqrt{b}$, as we maintained earlier. Note also that they are symmetric about the optimum. This will prove useful in calculating the price index below.

The comparative statics of the *Ss* bands are straightforward. Increases in the menu cost b lead to wider bands for the obvious reason. Increases in ε increase the concavity of the profit function. This increases the cost of deviations from the optimum and leads to narrower bands. Increases in $\bar{Y}$ allow the menu cost to be spread over more units of output. This leads to narrower bands. Increases in α allow the menu cost to be spread over a longer time horizon and thus to narrower bands.[13]

V. The Complete Model

In this section we put together the complete model. We restrict attention to a log-linear approximation about the steady state. We begin with the price index and the "state-dependent" Phillips curve and then turn our attention to the rest of the model.

A. *The Price Index*

At this point we reintroduce the j and z indexes. We divide both sides of the definition of the price index (8) by P_{t-1} and raise both sides by the power $1-\varepsilon$:

$$\left(\frac{P_t}{P_{t-1}}\right)^{1-\varepsilon} = \int \int \left(\frac{Q_{z,t}}{P_{t-1}}\right)^{1-\varepsilon} djdz.$$

Firms on a fraction α of islands maintain their prices unchanged. Since the idiosyncratic shocks are independent across islands, these firms' prices are representative of P_{t-1}. Let J_t denote the set of islands that receive shocks. Then we have

$$\left(\frac{P_t}{P_{t-1}}\right)^{1-\varepsilon} = \alpha + (1-\alpha)\int \int_{J_t} \left(\frac{Q^{j*}_{z,t}}{P_{t-1}}\right)^{1-\varepsilon} djdz. \tag{27}$$

Let $\pi_t = \ln P_t - \ln P_{t-1}$. Log-linearizing (27) about the zero-inflation steady state, we get

$$\pi_t = \alpha + (1-\alpha)\int \int_{J_t} (\ln Q^{j*}_{z,t} - \ln P_{t-1}) djdz.$$

Now let $\ln Q^*_t$ without a j or z index denote the sum $\ln Q^{j*}_{z,t} + \ln X^j_{z,t}$. Since every island that receives a shock is in essentially the same position, it follows from (23) that $\ln Q^*_t$ is the same for all (j, z). Since the $\ln X^j_{z,t-i}$ are independent and mean zero to a first order, $\ln Q^*_{t-i}$ is also the average price of firms that adjust in period $t-i$. Now given that the *Ss* bands in (25) and (26) are symmetric about the optimum and

[13] One interesting feature of the model is that the *Ss* bands are independent of the elasticity of labor supply. Dotsey and King (2005) find that local labor markets lead to tighter bands, more frequent price adjustment, and less persistence in the effects of monetary shocks. The reason is that in Dotsey and King's study the labor market is local to the firm, so that the firm sees an upward-sloping labor supply curve. This makes the firm's profit function more concave in its own price. Here there are a large number of firms on each island, so that each firm faces an elastic labor supply curve. As output on the island rises, the wage will rise, but each firm treats this increase as exogenous.

that the distribution of shocks is uniform, $\ln Q^*_{t-i}$ is also the average price of those firms that did not adjust on islands that receive the idiosyncratic shock.[14] It follows that

$$\pi_t = (1-\alpha)(\ln Q^*_t + \ln P_{t-1}). \tag{28}$$

B. *The State-Dependent Phillips Curve*

Let $\widehat{mc}_t$ be the deviation of average economywide real marginal cost from its steady state and let $\hat{Y}_t$ be the log deviation of output from its steady-state value. After (i) making use of the relations for the optimal reset price and the price index (eqq. [23] and [28], respectively) and (ii) log-linearizing the rest of the model, it is straightforward to derive the following Phillips curve relation for inflation (details are provided in App. B):

$$\pi_t = \lambda\widehat{mc}_t + \beta E_t \pi_{t+1}, \tag{29}$$

with

$$\lambda = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\Psi, \tag{30}$$

where Ψ is a measure of the degree of real rigidities,

$$\Psi = \frac{1}{1+\varphi\varepsilon},$$

and $\widehat{mc}_t = (1+\varphi)\hat{Y}_t$.[15]

It should be clear that this state-dependent Phillips curve (29) has the same form as the canonical time-dependent Phillips curve as originally formulated by Calvo (1983). The key (and only) difference lies with the parameterization of the slope coefficient λ, as we discuss shortly.

[14] The quadratic approximation of the profit function leads to the symmetric bands. If we had log-linearized about the exact solution to the firm's problem in the steady state, the bands would have been asymmetric, and we would have had to introduce another state variable associated with the distribution of firms that did not adjust. Of course, this alternative linearization differs from ours by terms that are second-order in ς.

[15] If we had aggregate productivity shocks, then we would replace $\hat{Y}$ with the output gap.

C. The Rest of the Model

Given that there are only consumption goods and utility is logarithmic, we can log-linearize the household's intertemporal condition to obtain the following "IS" curve:

$$\hat{Y}_t = -(\hat{R}^n_t - E_t\pi_{t+1}) + E_t\hat{Y}_{t+1}. \tag{31}$$

Next, log-linearizing the first-order condition for money demand and taking into account that consumption equals output yields

$$\ln M_t - \ln P_t - \ln\left(\overline{\frac{M}{P}}\right) = \hat{Y}_t - \zeta\hat{R}^n_t. \tag{32}$$

Equations (29), (31), and (32) determine the equilibrium aggregate dynamics, conditional on a monetary policy rule.

VI. Properties of the Model

Before proceeding to some numerical exercises, we first characterize some general properties of the model.

A. Relationship to the Calvo Model

As we have noted, our state-dependent Phillips curve differs from the Calvo formulation only in the parameterization of the slope coefficient on output. In the Calvo formulation, the exogenously given probability of no price adjustment enters the slope coefficient on the output gap in place of α, the probability of no idiosyncratic shock, the relevant primitive for the state-dependent case. In particular, for the time-dependent Phillips curve, the slope coefficient, λ_{td}, is given by

$$\lambda_{td} = \frac{(1-\theta)(1-\beta\theta)}{\theta}\Psi, \tag{33}$$

where θ is the probability of survival with no price adjustment.

In our setting, θ satisfies

$$1-\theta = (1-\alpha)\left[\tau\left(1 - \frac{\ln\bar{\mu}^H + \ln\bar{\mu}^L}{\phi}\right) + (1-\tau)\right].$$

The probability of price adjustment, $1-\theta$, is equal to the fraction of firms that receive the Poisson shock, survive, and receive an idiosyncratic shock large enough to trigger adjustment, $(1-\alpha)\tau[1-(\ln\bar{\mu}^H + \ln\bar{\mu}^L)/\phi]$ and the fraction of new firms that must set an initial price

$(1 - \alpha)(1 - \tau)$.[16] Since a fraction of firms that receive an idiosyncratic shock may not adjust, θ is, in general, greater than α. This implies that $\lambda \geq \lambda_{td}$. The implication is that inflation is more sensitive to movements in the output gap in the state-dependent framework. Intuitively, within the state-dependent model there is a selection effect at work. In contrast to the time-dependent case, firms not adjusting are those that are already close to the target. In fact, the firms that do not adjust on an island that receives an idiosyncratic shock have an average price equal to the target price. Thus, in general, the state-dependent formulation will yield greater nominal flexibility than the time-dependent formulation.

How much nominal flexibility state dependence delivers overall, however, depends on the other key primitive parameters of the model. Note in particular that λ also depends on the degree of real rigidities as measured by Ψ. With strong pricing complementaries (i.e., real rigidities) present, λ could be small, indicating considerable nominal stickiness, even within the state-dependent framework.[17]

B. Two Polar Cases

Our state-dependent formulation of inflation is quite flexible. It nests two polar extremes, along with a continuous range of intermediate outcomes. In particular, at one extreme the model can generate the kind of complete flexibility suggested by Caplin and Spulber (1987). At the other, it can perfectly mimic the degree of nominal stickiness in the pure time-dependent Calvo model.

When the idiosyncratic productivity shock hits each firm each period, the model behaves exactly like a flexible price model. In this instance, $\alpha = 0$, implying that the slope coefficient, λ, on the output gap in the *Ss* Phillips curve goes to infinity. According to (23), $\hat{\mu}_t^* = 0$. The economy is always at its frictionless optimum and money is neutral. Neutrality holds in spite of the fact that a fraction θ of firms do not adjust their prices in each period.[18]

What is the source of this neutrality? It is instructive to analyze it from the perspective of both a firm and the economy as a whole. Consider

[16] This fraction is equal to the fraction of firms that receive an idiosyncratic shock and die (see eq. [10]). Since turnover of firms leads to price adjustment, in calibrating the model we treat price changes that stem from new entrants as price adjustments in the data that come from new product substitutions. See Nakamura and Steinsson (2007) and Klenow and Kryvtsov (2008).

[17] Introducing other complementarities such as firm-specific capital (Woodford 2003) or a chain of production (Basu 1995) may reduce λ further.

[18] Note that Danziger (1999) does not find neutrality in his model even though he assumes that $\alpha = 0$. The reason is that he presents an exact analytic solution, whereas we log-linearize. The effects of money on output that Danziger finds are second-order in our framework.

first a firm that is contemplating price adjustment. It faces an expected path for the nominal wage. In a time-dependent model, the firm would set its price equal to a markup over a weighted average of future wages, where the weights represent the discounted probability that the firm has not yet had an opportunity to alter its price. The weights would have the form $(\beta\theta)^i$. How can the state-dependent firm ignore the future path of wages and set its price as a markup only of the current wage? The answer is that the state-dependent firm can use its future price adjustment decision to bring its costs in line with whatever price it sets today. Suppose that the wage is expected to rise in the next period. A time-dependent firm would find that its price is too low. The state-dependent firm shifts the set of future productivities for which it maintains its price so that its average markup is unchanged. The resulting distribution of markups is unaffected by the increase in the wage. It is important to note that this stark neutrality result depends crucially on the assumption of a uniform distribution with wide support and that the shock hits the firm each period. This assumption allows the firm each period to alter its adjustment triggers without altering the resulting distribution of the markups.

From the perspective of the economy as a whole, this neutrality result is similar to the neutrality result of Caplin and Spulber. In their paper, an increase in the nominal wage causes a few firms to raise their prices by a discrete amount, so that the aggregate real wage remains constant. Here what changes is the mix of firms that raise and lower their prices. When a shock causes the nominal wage to rise, the set of firms that maintain their prices fixed changes. Some that had marginally low productivities decide to raise their prices, and some that have marginally high productivities decide not to lower theirs. The result is an unchanging distribution of markups: uniform between two fixed triggers and a fixed mass at the target.

In the general case in which a subset of firms each period do not get hit with an aggregate shock $(0 < \alpha < 1)$, the slope coefficient λ is less than infinity, implying nominal stickiness at the aggregate level. In this instance, the unconditional distribution of idiosyncratic shocks is not uniform, and hence monetary policy will affect the distribution of markups.[19] How important these effects are depends on the model calibration. We turn to this issue in the next section.

At the other extreme, as long as $0 < \alpha < 1$, the model converges to the pure time-dependent case as the menu cost b goes to zero. In this case the *Ss* bands go to zero, implying that firms adjust whenever an idiosyncratic shock hits. In this instance the probability of price adjust-

[19] As we noted in Sec. II, with $\alpha > 0$, the unconditional distribution of idiosyncratic shocks has extra mass at zero.

ment $1 - \theta$ converges to $1 - \alpha$. In this case our *Ss* model behaves exactly like the Calvo model: the slope coefficients on marginal cost in the respective Phillips curves are identical in each case.

One interesting feature of the polar case of no costs of price adjustment is that the small decision cost we introduced earlier for technical considerations emerges to play an important role. Without menu costs, the decision cost provides the only friction to preclude price adjustment in response to aggregate shocks. Recall that the decision cost (i.e., the cost of contemplating a price change and gathering information) precludes price adjustments in the absence of idiosyncratic shocks. As we show in Appendix C, under our baseline calibration of menu costs, the decision cost serves only to simplify the analytics and does not materially affect the quantitative properties. As menu costs disappear, however, the decision cost plays the key role in delivering nominal stickiness. While the focus of this paper is on the role of menu costs, in future work we plan to develop more thoroughly the role of decision costs.

C. *Relationship to the Klenow-Kryvtsov Evidence*

Klenow and Kryvtsov (2008) show that for the recent low-inflation decade in the United States, (i) the proportion of firms that adjust their prices is fairly constant, and (ii) the variation in the inflation has been driven almost entirely by variation in the size of price adjustment, not by variation in the frequency of price adjustment.

In our framework, the probability of price adjustment is fixed up to a second order. It can be seen from equation (28) that all of the variation in inflation is the result of variation in the average size of price adjustment. Hence our model is consistent with the Klenow-Kryvtsov facts even though it is based on *Ss* pricing.

VII. Calibration and Some Simulations

In this section we explore the response of the model economy to a monetary shock as a way to evaluate the effects of *Ss* pricing. We begin by calibrating the model. Where possible we choose standard parameters. The time period is a quarter. We set the discount rate β at .99. We set the elasticity of substitution between goods, ε, equal to 11, which implies a steady-state markup of 10 percent. We set the Frisch elasticity of labor supply (the inverse of φ) at 1.0, which is a reasonable intermediate range value in the literature.

Next we turn to the key parameters of price adjustment: the probability of no idiosyncratic shock, α, the density of the idiosyncratic shock, ϕ, the adjustment cost relative to average steady-state firm output, $b/\bar{Y}$, and the probability that the firm remains in the market τ conditional on an idiosyncratic shock. Note first that the *Ss* band, $\omega = \ln \mu^H - \ln \mu^*$, is a function of these parameters. From equation (25),

$$\omega = \sqrt{2 \frac{1 - \alpha\beta}{\varepsilon - 1} \frac{b}{\bar{Y}}}.$$

The average size of price adjustment will depend on ω and the range of the idiosyncratic productivity shock (which depends on ϕ). The frequency of price adjustment conditional an idiosyncratic shock and survival, $1 - (2\omega/\phi)$, will also depend on ω and ϕ. We proceed to derive a system of relations that pin down the vector $(\alpha, \phi, b/\bar{Y}, \tau)$ using evidence on (i) the frequency of price adjustment with and without substitutions, (ii) the absolute size of price adjustments, and (iii) the costs of price adjustment.

Nakamura and Steinsson (2007) report the median time a price is fixed in consumer price index data. They report numbers for two periods, 1995–97 and 1998–2005, and they report numbers including and excluding substitutions. We interpret the substitution rate as the firm turnover rate in our model, $(1 - \alpha)(1 - \tau)$. When their numbers are averaged, the median duration is 1.51 quarters excluding substitutions and 1.49 quarters including substitutions. We set the probability of price adjustment conditional on firm survival,

$$\psi \equiv \frac{(1 - \alpha)\tau[1 - (2\omega/\phi)]}{1 - (1 - \alpha)(1 - \tau)},$$

equal to .369, the implied probability of price adjustment per quarter excluding substitutions.[20] We set the overall frequency of price adjustment $1 - \theta$ to .380, the implied probability of price adjustment including substitutions. Next we require that the model match Klenow and Kryvtsov's evidence that the average absolute size of price adjustments is about 8.5 percent: we therefore set $(\phi/4) + (\omega/2) = .085$. Finally, we set the steady-state resources devoted to price adjustment equal to 0.55 percent of revenue, on the basis of the evidence in Zbaracki et al. (2004) and Levy et al. (1997).[21] This implies $\psi b/\bar{Y} = .0055$.

[20] As Cogley and Sbordone (forthcoming) note, given that adjustment is a binomial random variable, the time until the next adjustment can be approximated as a continuous-time exponential random variable, implying a median waiting time equal to $-\ln(2)/\ln(\theta)$, where θ is the Poisson arrival rate. Note that the median waiting time is less than the mean duration of prices ($1/0.4 = 2.5$ quarters, roughly 7.5 months) since the exponential distribution implies that some prices may not change for a very long time.

[21] Zbaracki et al. (2004) quantify the physical costs of price adjustment for a large manufacturing firm. They find b/Y to be about .004. Levy et al. (1997) in a study of four grocery stores find that resources devoted to the price adjustment are slightly higher, approximately 0.7 percent of revenue. However, Golosov and Lucas (2007) find that a value of 0.24 allows their model to best match certain properties of the data. Our quantitative results are robust to alternative choices of this parameter.

TABLE 1
Imputed Values of the Parameters

θ	α	ω	ϕ	$b/\bar{Y}$	τ
.620	.4550	.0405	.2590	.0149	.9667

TABLE 2
The Effect of Complementarities on the Coefficient on Marginal Cost

φ	λ	λ_{td}
0	.658	.237
1	.055	.020

Table 1 shows the values of α, θ, ϕ, $b/\bar{Y}$, ω, and τ implied by this parameterization. The parameter $\alpha > \theta$, reflecting the fact that not all firms adjust after an idiosyncratic shock. The probability of price adjustment conditional on the idiosyncratic shock and survival, $1 - (2\omega/\phi)$, is equal to .69. This reflects the fact that $\phi > 4\omega$, implying that the support of the idiosyncratic shock is large enough that the idiosyncratic shock leads to both price increases and price decreases for small aggregate shocks. The death probability $1 - \tau$ implies that each period about 2 percent of firms leave the market. This is in the range reported by Klenow and Kryvtsov.

Table 2 shows the implied value of λ, the slope coefficient on marginal cost in the Phillips curve given by equation (29). For comparison, we also report the implied slope coefficient for a conventional time-dependent Calvo formulation λ_{td}, with a similar frequency of price adjustment. In addition, we also report slope coefficients for the case in which real rigidities are absent (global labor markets). In this case $\varphi = 0$.

For the case with local labor markets ($\varphi = 1$), the slope coefficient is .055 for our *Ss* model, whereas for the Calvo model the parameter shrinks to about .020. The value of λ exceeds λ_{td} by a factor of almost 3, indicating greater nominal flexibility with state dependence. However, the absolute difference is small. Further, λ lies in the upper range of estimates reported by Gali and Gertler (1999). Eliminating real rigidities raises both λ and λ_{td} by a factor of 12. In this case the absolute difference between the two cases is large. However, both slope coefficients lie well above estimates in the literature.

Figures 2 and 3 illustrate the response of the model economy to an unanticipated monetary shock and clearly demonstrate the effect of adding real rigidities. Figure 2 illustrates the response of the model

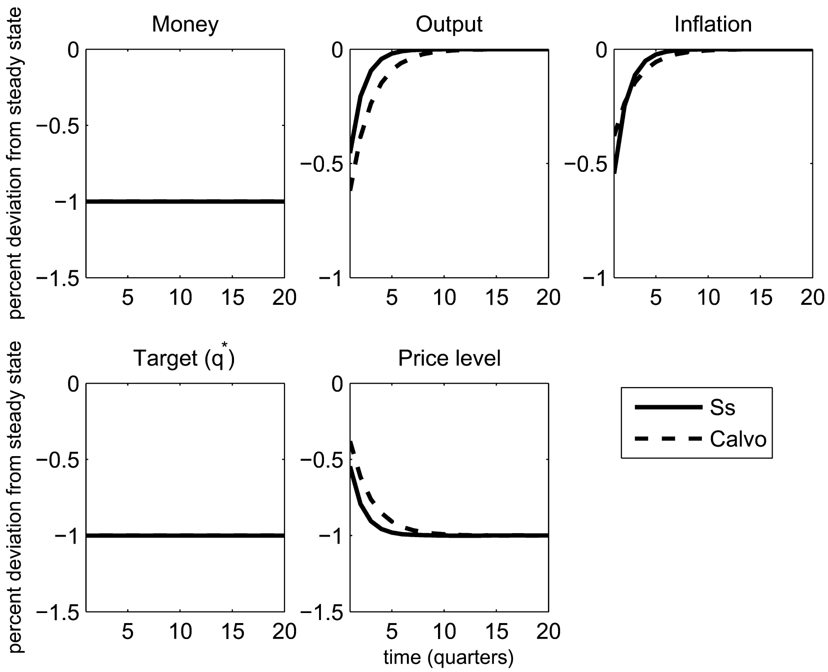

Fig. 2.—Response to a 1 percent shock to the money supply (no complementarities)

economy without real rigidities ($\varphi = 0$) to a permanent 1 percent decrease in the money stock. This scenario corresponds closest to the policy experiment considered by Golosov and Lucas (2007). The solid line is the response of our state-dependent pricing model, and the dotted line shows the response of the time-dependent Calvo model. For the state-dependent model there is only a transitory decrease in real output that lasts about 3 quarters. The initial response of the price level, further, is slightly greater in percentage terms than the response of real output, suggesting considerable nominal flexibility. Indeed, consistent with the findings of Golosov and Lucas, the state-dependent model also exhibits greater nominal flexibility than the Calvo model. For the Calvo model, the initial output response is roughly 20 percent larger, and the overall response lasts several quarters longer. Conversely, the overall movement in the price level is smaller.

The absence of complementarities is evident in figure 2. The optimal target price $\ln Q^*$ immediately falls by 1 percent, mimicking the path followed by the money supply. There is no effect of the firms that do not change their prices on the firms that do. By the time that all firms have adjusted, the transition to the new steady state is complete.

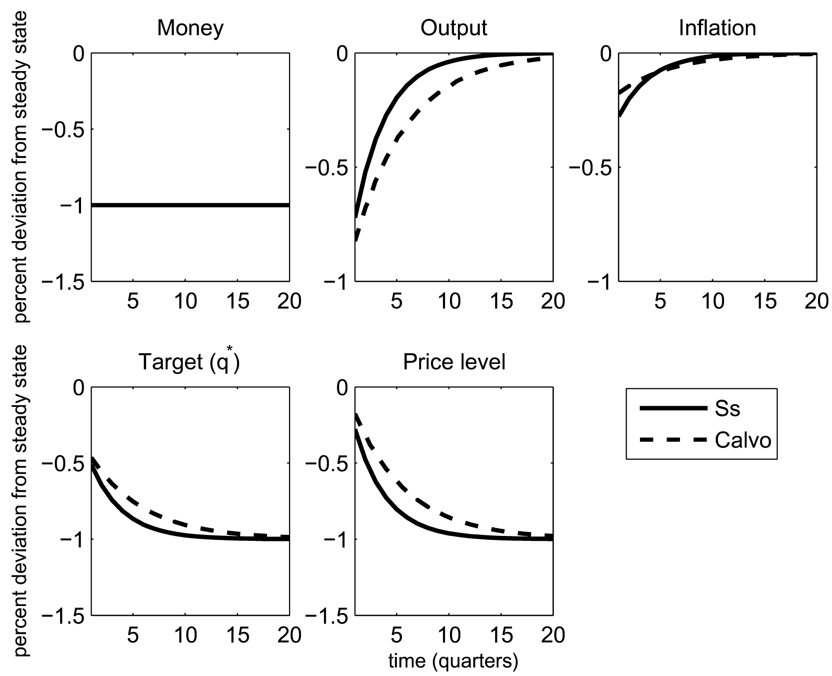

Fig. 3.—Response to a 1 percent shock to the money supply (with strategic complementarities).

When we add complementarities in the form of local labor markets, it is still the case that the state-dependent model exhibits the most flexibility, but the percentage difference from the Calvo model becomes smaller. Figure 3 illustrates the response of the model economy for this case. As we would expect, there is a stronger response of output and a weaker response of the price level for both the state- and time-dependent models. For the state-dependent model, the percentage output response is now roughly triple the response of the price level. Further, output does not return to trend for over 10 quarters. Importantly, the addition of real rigidities reduces the percentage difference in the output response across the state- and time-dependent models. Now, for example, the initial output response for the state-dependent model is only about 10 percent less than for the Calvo model.

We can see the effect of the complementarities in the response of $\ln Q^*$. The initial response of the target price is only half the size of the money shock, and it takes about 10 quarters for $\ln Q^*$ to adjust to the steady state. In this case even after the majority of firms have adjusted their prices, the economy will not have returned to the steady state. Some of these firms will have adjusted to a non-steady-state price. This is the source of sluggishness in the price level that generates greater and more persistent real effects of money.

To be sure, while our model is useful for exploring the implications of state-dependent pricing and is capable of capturing qualitatively the relative strong response of output and weak response of inflation to a monetary policy shock, it is clearly too simple to closely match the evidence (e.g., Christiano, Eichenbaum, and Evans 2005). For example, it cannot capture the delayed and hump-shaped response of real output. However, it is straightforward to add a number of features (e.g., habit formation, capital, investment with delays and adjustment costs, and so on) that have proved useful in improving the empirical performance of such models.

VIII. Conclusion

We have developed a simple macroeconomic framework that features an analytically tractable Phillips curve relation based on state-dependent pricing. At the micro level, firms face idiosyncratic shocks and fixed costs of adjusting price. We cut through the usual difficulties in solving and aggregating *Ss* models with restrictions on the distribution of idiosyncratic shocks and also by focusing on a local approximation around a zero inflation steady state, as is done in the time-dependent pricing literature. In the end, our model is able to match the micro evidence on the frequency and size of price adjustment. At the same time, the resulting Phillips curve is every bit as tractable as the Calvo relation based on time-dependent pricing.

Consistent with the numerical exercises in Golosov and Lucas (2007), we find that for a given frequency of price adjustment, the *Ss* model exhibits greater nominal flexibility than a corresponding time-dependent framework because of a selection effect: firms farthest away from the target adjust in the *Ss* model, whereas this is not the case within the time-dependent framework. However, with the introduction of real rigidities, our *Ss* model is capable of generating considerable nominal stickiness, as we demonstrate with a simple calibration model. Also key to the result is that, under our baseline calibration, the idiosyncratic shock distribution has most firms concentrated well within the *Ss* bands, so that individual firms adjust prices typically only in the wake of large idiosyncratic shocks. This is what gives our *Ss* model the strong flavor of a time-dependent model, even though the adjustment decision is purely endogenous.

An interesting recent paper by Midrigan (2006) achieves considerable nominal stickiness by modifying the Golosov-Lucas framework by introducing a fat-tailed distribution of idiosyncratic shocks along with scale economies in making multiple price changes. As in Golosov and Lucas's paper, he solves the model numerically. We similarly use a fat-tailed unconditional distribution of idiosyncratic shocks (though a simpler one), but instead introduce real rigidities and find an approximate analytical solution. Nonetheless, it would be interesting to consider blending features of the two approaches.

While our model is capable of capturing the basic features of the micro data, it is too simple at this stage to capture the cyclical dynamics of output and inflation. It is straightforward to add some features that have proved useful in explaining performance, such as habit formation, investment, and adjustment costs. Accounting for the persistence of inflation may prove trickier, given that the simple Calvo model also has difficulty on this account. Specifically, the evidence suggests that a hybrid Phillips curve that allows for lagged inflation as well as expected future inflation to affect inflation dynamics is preferred over the pure forward-looking model.[22] However, at this point we suspect that some of the strategies employed in the time-dependent literature to address this problem, such as dynamic indexing, information lags, and/or learning, may prove useful in this context as well.

Appendix A

Proof of Proposition 1

Suppose that the firm has a current markup of μ_t such that $\ln \mu_t \in [\ln \mu_t^L, \ln \mu_t^H]$. We are interested in the expected value of an optimal policy conditional on an idiosyncratic productivity shock in period $t+1$. Also let μ_{t+1}^* denote the optimal choice of μ_{t+1} in the event of adjustment.

Consider $E\{\Lambda_{t,t+1} \exp[(\varepsilon - 1)\xi_{t+1}] v(\mu_{t+1}, \Omega_{t+1})\}$ over the states of the world in which the idiosyncratic shock hits. Given the assumption on ϕ, $\xi^H > \mu_{t+1}^H - \ln(W_t/W_{t+1}) - \ln \mu_t$, and $\xi^L < q_{t+1}^L - \ln(W_t/W_{t+1}) - \ln \mu_t$,

$$
\begin{aligned}
&E\{\Lambda_{t,t+1} \exp[(\varepsilon - 1)\xi_{t+1}] v(\mu_{t+1}, \Omega_{t+1})\} = \\
&E\Bigg\{\Lambda_{t,t+1} \frac{1}{\phi} \int_{\ln \mu_{t+1}^H - \ln(W_t/W_{t+1}) - \ln \mu_t}^{\xi^H} \exp[(\varepsilon - 1)\xi_{t+1}]\, \bar{v}(\mu_{t+1}^*, \Omega_{t+1}) d\xi_{t+1} \\
&+ \Lambda_{t,t+1} \frac{1}{\phi} \int_{\ln Q_{t+1}^L - \ln(W_t/W_{t+1}) - \ln \mu_t}^{\ln Q_{t+1}^H - \ln(W_t/W_{t+1}) - \ln \mu_t} \exp[(\varepsilon - 1)\xi_{t+1}]\, \bar{v}\left(\frac{\mu_t W_t}{W_{t+1}} \exp(\xi_{t+1}), \Omega_{t+1}\right) d\xi_{t+1} \\
&+ \Lambda_{t,t+1} \frac{1}{\phi} \int_{\xi^L}^{\ln Q_{t+1}^L - \ln(W_t/W_{t+1}) - \ln \mu_t} \exp[(\varepsilon - 1)\xi_{t+1}]\, \bar{v}(\mu_{t+1}^*, \Omega_{t+1}) d\xi_{t+1}\Bigg\}.
\end{aligned}
$$

[22] For example, Gali and Gertler (1999) find that a hybrid model with a coefficient of roughly 0.65 on expected future inflation and 0.35 on lagged inflation is preferred over the pure forward-looking model.

Rearranging yields

$$E\{\Lambda_{t,t+1}\exp[(\varepsilon-1)\xi_{t+1}]v(\mu_{t+1},\ \Omega_{t+1})\} =$$

$$E\Bigg\{\Lambda_{t,t+1}\bar{v}(\mu^{*}_{t+1},\ \Omega_{t+1})$$

$$+\ \Lambda_{t,t+1}\frac{1}{\phi}\int_{\ln Q^{L}_{t+1}-\ln(W_t/W_{t+1})-\ln\mu_t}^{\ln Q^{H}_{t+1}-\ln(W_t/W_{t+1})-\ln\mu_t}\exp[(\varepsilon-1)\xi_{t+1}]\Bigg[\bar{v}\left(\frac{\mu_t W_t}{W_{t+1}}\exp(\xi_{t+1}),\ \Omega_{t+1}\right)$$

$$-\ \bar{v}(\mu^{*}_{t+1},\ \Omega_{t+1})\Bigg]d\xi_{t+1}\Bigg\}.$$

A change of variable, $\Phi_{t+1} = \ln\mu_t + \ln(W_t/W_{t+1}) + \xi_{t+1}$, gives

$$E\{\Lambda_{t,t+1}\exp[(\varepsilon-1)\xi_{t+1}]v(\mu_{t+1},\ \Omega_{t+1})\}$$

$$= E\Bigg\{\Lambda_{t,t+1}\bar{v}(\mu^{*}_{t+1},\ \Omega_{t+1})$$

$$+\ \Lambda_{t,t+1}\frac{1}{\phi}\int_{\ln\mu^{L}_{t+1}}^{\ln\mu^{H}_{t+1}}\exp[(\varepsilon-1)(\Phi_{t+1}-\ln\mu_t)][\bar{v}(\exp(\Phi_{t+1}),\ \Omega_{t+1})$$

$$-\ \bar{v}(\mu^{*}_{t+1},\ \Omega_{t+1})]d\Phi_{t+1}\Bigg\}. \tag{A1}$$

Note that $\ln\mu^{H}_{t+1}$ and $\ln\mu^{L}_{t+1}$ are chosen optimally in period $t+1$. They depend on the period $t+1$ state. The term $\exp(\Phi_{t+1})$ and the aggregate variables are independent of μ_t. The only place that μ_t enters is the exponential term inside the integral. Now, by the assumption on b, $\ln\mu_t$ is equal to $\ln\mu^{*}_t$ plus a first-order term, and given the limits of integration, Φ_{t+1} is equal to $\ln\mu^{*}_{t+1}$ plus a first-order term. The exponential term is therefore equal to $\exp[(\varepsilon-1)(\ln Q^{*}_{t+1}-\ln Q^{*}_t)]$ plus a first-order term. The term in brackets inside the integral is bounded by b. By the assumption on b, this term is second-order. Hence,

$$\Lambda_{t,t+1}\exp[(\varepsilon-1)(\Phi_{t+1}-\ln Q_t)][\bar{v}(\exp(\Phi_{t+1}),\ \Omega_{t+1})-\bar{v}(Q^{*}_{t+1},\ \Omega_{t+1})] =$$

$$\Lambda_{t,t+1}\exp[(\varepsilon-1)(\ln Q^{*}_{t+1}-\ln Q^{*}_t)][\bar{v}(\exp(\Phi_{t+1}),\ \Omega_{t+1})-\bar{v}(Q^{*}_{t+1},\ \Omega_{t+1})] + \mathcal{O}^3.$$

Further, the assumptions that $\phi > 2(\ln Q^{H}_t - \ln Q^{L}_t)$ and that the range of integration is first-order imply

$$E\{\Lambda_{t,t+1}\exp[(\varepsilon-1)\xi_{t+1}]v(Q_t\exp(\xi_{t+1}),\ \Omega_{t+1})\}$$

$$= E\Bigg\{\Lambda_{t,t+1}\bar{v}(Q^{*}_{t+1},\ \Omega_{t+1})$$

$$+\ \Lambda_{t,t+1}\frac{1}{\phi}\int_{\ln\mu^{L}_{t+1}}^{\ln\mu^{H}_{t+1}}\exp[(\varepsilon-1)(\ln\mu^{*}_{t+1}-\ln\mu^{*}_t)][\bar{v}(\exp(\Phi_{t+1}),\ \Omega_{t+1})$$

$$-\ \bar{v}(Q^{*}_{t+1},\ \Omega_{t+1})]d\Phi_{t+1}\Bigg\} + \mathcal{O}^3$$

It follows that $E\{\exp[(\varepsilon - 1)\xi_{t+1}]v(\mu_{t+1}, \Omega_{t+1})\}$ is independent of μ_t to a second order. QED

Appendix B

Derivation of the Log-Linear Phillips Curve

Equation (23) in the text implies that the average optimal reset price $\ln Q_t^*$ (net of the $\ln X_{z,t}^j$, which are mean zero across firms) is the following discounted stream of future nominal wages:

$$\ln Q_t^* = \ln \bar{\mu} + (1 - \beta\alpha)E_t \sum_{i=0}^{\infty} (\beta\alpha)^i \ln W_{z,t+i}.$$

Note that $\ln Q_t^*$ depends on the island-specific wage $W_{z,t}$. As a step toward aggregation, we would like to derive this relation in terms of the economywide average wage, W_t.

Log-linearizing the household's first-order condition for labor supply yields

$$\ln W_{z,t} - \ln P_t + \ln \bar{\mu} = \varphi \hat{N}_{z,t} + \hat{C}_t. \tag{B1}$$

Averaging over this condition yields $\ln W_t - \ln P_t + \ln \bar{\mu} = \varphi \hat{N}_t + \hat{C}_t$, implying the following relation between the island z relative wage and the relative employment levels:

$$\ln W_{z,t} - \ln W_t = \varphi(\hat{N}_{z,t} - \hat{N}_t). \tag{B2}$$

Making use of the demand function and the production function leads to a relationship between the relative wage and the relative price of firms that adjust at time t:

$$\ln W_{z,t+i} = \ln W_{t+i} - \varphi\varepsilon(\ln Q_t^* - \ln P_{t+i}). \tag{B3}$$

Notice that $\ln W_{z,t+i}$ depends inversely on $\ln Q_t^*$. Raising prices on an island reduces output and labor demand. Since the labor market is segmented, it also reduces wages on the island, thus moderating the need to raise prices in the first place. As emphasized in Woodford (2003), this factor segmentation thus introduces a strategic complementarity or "real rigidity" that gives adjusting firms a motive to keep their relative prices in line with the relative prices of nonadjusting firms.[23] This strategic complementarity, in turn, contributes to the overall stickiness in the movement of prices. Combining (B3) with the expression for Q^* yields

$$\ln Q_t^* = (1 - \alpha\beta)[\Psi(\ln W_t - \ln P_t) + \ln P_t] + \beta\alpha E_t \ln Q_{t+1}^*, \tag{B4}$$

with $\Psi = 1/(1 + \varphi\varepsilon)$.

In equilibrium, the real wages of adjusting firms, $\ln W_{z,t} - \ln P_t$, move less than one for one with the aggregate real wage, implying similarly sluggish movement in the target price $\ln Q_t^*$. In this respect, the strategic complementarity measured inversely by the coefficient Ψ dampens the adjustment of prices. With economywide labor markets, Ψ equals unity, implying that $\ln W_{z,t}$ simply is equal to $\ln W_t$.

[23] For a menu of alternative ways to introduce real rigidities, see Kimball (1995) and Woodford (2003).

We are now in a position to present the Phillips curve. Let $\pi_t = \ln P_t - \ln P_{t-1}$ denote inflation. Combining the equation for the target price (B4) with the price index (28) yields

$$\pi_t = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\Psi(\ln W_t - \ln P_t) + \beta E_t \pi_{t+1}. \tag{B5}$$

It should be clear that this state-dependent Phillips curve has the same form as the canonical time-dependent Phillips curve as originally formulated by Calvo (1983). The key difference is that in our formulation the primitive parameter entering the slope coefficient on marginal cost is the probability α of no idiosyncratic shock, whereas in the time-dependent framework it is the exogenously given probability of no price adjustment.

The rest of the model is standard. Log-linearizing the first-order condition for labor supply, averaging across households, and taking into account that consumption equals output yields a linear relation between the aggregate real wage and output:

$$\ln W_t - \ln P_t + \ln \bar{\mu} = \kappa \hat{Y}_t, \tag{B6}$$

where $\kappa = 1 + \varphi$ is the elasticity of marginal cost (with log utility). Combining equations (B5) and (B6) then yields the Phillips curve in terms of the output gap:

$$\pi_t = \lambda\kappa \hat{Y}_t + \beta E_t \pi_{t+1}, \tag{B7}$$

with

$$\lambda = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\Psi.$$

Appendix C

On the Accuracy of the Approximation

We have made a number of approximations and assumptions in order to arrive at an analytically tractable model. Theoretically our approximation holds if (1) there exists a decision cost such that firms adjust if and only if the idiosyncratic shock arrives, (2) the other parameter restrictions hold, (3) the second-order approximation of the value profit function is accurate, (4) the continuation value following an idiosyncratic shock is third-order, and (5) the aggregate shocks are sufficiently small that the log-linearization is accurate. In this appendix, we evaluate the first four assumptions. The last is standard in the literature on Calvo pricing. We also evaluate whether removing the decision cost has a large effect on the model's dynamics. In summary, the model holds up well for our calibration.

The Bounds on the Decision Cost

In this subsection, we calculate the range of the decision costs $[\underline{d}, \bar{d}]$ that ensure that a firm adjusts only after receiving an idiosyncratic shock. We show that this range is nonempty for our parameterization.

We first amend the firm's problem to include the decision cost. We assume that the firm knows the aggregate state and whether or not an idiosyncratic shock has occurred (and whether or not it has survived). The firm must pay the decision cost d in order to learn the value of the idiosyncratic shock and to process information regarding both the aggregate and idiosyncratic variables. The firm must then pay the menu cost b in order to adjust its price.

Suppose that the firm last paid the decision cost at date s. Let $\mu_s = Q_s X_s / W_s$ be the markup at that time. Let $z_{s,t} \in \mathcal{N}$ denote the number of times that an idiosyncratic shock has hit since date s (including date t). We define three value functions, each corresponding to a different stage of the period t decision. Let $\hat{v}(\mu_s, \Omega_t, z_{s,t})$ denote the optimal policy at the beginning of date t given that the markup was set to μ_s at date s, the current aggregate information set is Ω_t, and idiosyncratic turbulence has hit the firm $z_{s,t}$ times since s. Let $v(\tilde{\mu}_t, \Omega_t)$ denote the optimal policy at date t conditional on having paid the decision cost but before paying the menu cost b. The function $\hat{v}$ depends on the inherited markup $\tilde{\mu}_t$ and the current aggregate information Ω_t. Finally, let $\bar{v}(\mu_t, \Omega_t)$ denote the value of an optimal policy after price adjustment. The function $\bar{v}$ depends on the postadjustment markup μ_t and the current aggregate state Ω_t. Now

$$\begin{aligned}\hat{v}(\mu_s, \Omega_t, z_{s,t}) = \max\{&E[v(\tilde{\mu}_t, \Omega_t)|z_{s,t}, \Omega_t] - d,\\ &\bar{\Pi} + E[-\tfrac{1}{2}(\varepsilon - 1)\bar{Y}(\ln \tilde{\mu}_t - \ln \bar{\mu})^2\\ &+ \beta \hat{v}(\mu_s, \Omega_{t+1}, z_{s,t+1})|z_{s,t}, \Omega_t]\}. \qquad \text{(C1)}\end{aligned}$$

If the firm pays the decision cost d, then it receives the expected value of $v(\tilde{\mu}_t, \Omega_t)$. This expectation is taken with respect to the current aggregate information Ω_t and $z_{s,t}$, which indexes information regarding the distribution of the productivity shock. If the firm does not pay the decision cost, then it cannot alter its price and it does not update its information regarding X_t. It thus receives the expected profit Π_t defined in equation (20) and the discounted value of $v(\mu_s, \Omega_{t+1}, z_{s,t+1})$, where the information is now tomorrow's information. The discount rate is β since the arrival of the idiosyncratic shock is encoded in z. Again the expectation is taken with respect to the current information.

In the second step, the firm that pays the decision cost has the option of altering its price

$$v(\tilde{\mu}_t, \Omega_t) = \max\{\bar{v}(\tilde{\mu}_t, \Omega_t), \ \max_\mu \bar{v}(\mu, \Omega_t) - b\}. \qquad \text{(C2)}$$

As the final step, we define $\bar{v}$:

$$\bar{v}(\mu_t, \Omega_t) = \bar{\Pi} + E[-\tfrac{1}{2}(\varepsilon - 1)\bar{Y}(\ln \mu_t - \ln \bar{\mu})^2 + \tilde{\alpha}\beta\hat{v}(\mu_t, \Omega_{t+1}, z_{t,t+1})|\Omega_t], \qquad \text{(C3)}$$

where $\tilde{\alpha}$ is the survival probability. Note that if the firm pays the decision cost if and only if $z = 1$, then the decision problem in equations (C1), (C2), and (C3) is equivalent to the decision problem in equations (18) and (19) of the text.

We begin with the upper bound $\bar{d}$. We need d low enough that a firm experiencing a productivity shock always chooses to pay the decision cost. Rearranging (C1) and using (C2) and (C3) yields the following condition:

$$d \leq E\{\max\{0, \ \max_\mu \bar{v}(\mu, \Omega_t) - b - \bar{v}(\tilde{\mu}_t, \Omega_t)\}|1, \Omega_t\}. \qquad \text{(C4)}$$

Here $z_{s,t}$ is equal to one since the idiosyncratic shock has hit once. The equation states that d must be less than the expected gain to adjustment.

Given the concavity of the profit function, a firm in the center of the bands is the least willing to pay the decision cost. Therefore, consider a firm that (1) knows that $\hat{\mu} = 0$ when it experiences an idiosyncratic productivity shock and (2) believes that the aggregate state will remain unchanged over the foreseeable future so that $\hat{\mu} = 0$ is the optimal markup. It is sufficient that such a firm wishes to adjust.

We now calculate the gain to paying the decision cost for this firm after an idiosyncratic shock. The benefit of paying the decision cost is the option to adjust price. The firm will adjust if and only if $|\xi| > \omega$. In Section IV, we showed that the normalized per-period loss to nonadjustment is approximately $\frac{1}{2}(\varepsilon - 1)\bar{Y}\hat{\mu}^2$. If $|\xi| > \omega$, the firm will therefore gain $\frac{1}{2}(\varepsilon - 1)\bar{Y}(\xi^2 - \omega^2)$, where we have used the fact that when $\xi = \omega$ the firm is indifferent between adjusting and nonadjusting. We integrate this potential benefit over the region of price adjustment $[-\phi/2, -\omega] \cup [\omega, \phi/2]$ using the density $1/\phi$. We may assume that $d \in [\underline{d}, \bar{d}]$ in all subsequent periods, so we take the present value until the next idiosyncratic shock. The simplification theorem implies that we can ignore payoffs beyond that point. Condition (C4) becomes

$$d \leq \frac{(\varepsilon - 1)\bar{Y}}{(1 - \alpha\beta)\phi} \int_{\omega}^{\phi/2} (\xi^2 - \omega^2) d\xi = \frac{1}{2} \frac{(\varepsilon - 1)\bar{Y}}{1 - \alpha\beta} \left(\frac{\phi}{12} - \omega^2 + \frac{4\omega^3}{3\phi} \right) \equiv \bar{d}. \quad \text{(C5)}$$

Given our parameterization, $\bar{d} = 0.039\bar{Y}$, or about 2.5 times the menu cost.

We now turn to the lower bound $\underline{d}$. We need d large enough that a firm will not pay the decision cost in the absence of an idiosyncratic disturbance. Substituting (C2) and (C3) into (C1) and using the fact that with $z = 0$ the firm knows its current markup, we arrive at the following condition:

$$\max_{\mu} \bar{v}(\mu, \Omega_t) - (b + d) < \bar{v}(\tilde{\mu}_t, \Omega_t).$$

When there is no idiosyncratic shock, it is as though the firm faces an adjustment cost of $b + d$. This leads to wider *Ss* bands and potentially prevents a firm that set its markup at $\hat{\mu}^H$ or $\hat{\mu}^L$ from adjusting in subsequent periods. We require that this condition hold for all states of the model.

To calculate $\underline{d}$, we need to calculate $\bar{v}$ and construct a worst-case scenario. To do so we need to be explicit about the process for the aggregate shock. We obtain a reasonable lower bound by feeding into the model shocks based on an estimated process for nominal demand and calculating for each period the decision cost required to prevent firms that have not received the idiosyncratic shock from adjusting. Specifically, we fit an AR(1) to quarterly nominal GDP growth over the sample 1984:1–2007:1, a period of moderate inflation to which our model is most relevant. (As we discussed Sec. II.D, n. 12, our disturbance is best thought of as a variation in nominal demand, which can be captured in the data by movements in nominal GDP.) This yields an autoregressive parameter $\rho = .37$ and a sequence of shocks. We use these to simulate our calibrated model. We then compute at each date t the decision cost required to discourage adjustment by firms whose prices were at the upper and lower adjustment triggers at dates $t - 1$, $t - 2$, and $t - 3$. These are the firms with the greatest desire to adjust. We go back three periods because in our calibration with $\alpha = .455$, only 4 percent of firms operating at date t last received an idiosyncratic shock prior to date $t - 3$. We take $\underline{d}$ to be the maximum of these costs. This procedure yields $\underline{d} = .018\bar{Y}$. If we drop the largest outlier shock (the third quarter of 2001, i.e., 9/11), $\underline{d}$ falls to $.0147\bar{Y}$, roughly the size of the menu cost.[24] We conclude that the range of admissible decision costs is nonempty in our parameterization.

Other Parameter Restrictions

Condition (12) requires that ϕ be greater than $4\omega + 2m$. Given our parameterization, $\omega = .0405$. We need a reasonable value for m. Given $\alpha = .455$, only 4 percent of firms survive 1 year. We therefore choose $k(m) = 4$. We associate m with nominal GDP. The standard deviation of nominal GDP growth over the period 1984:1–2007:1 is 1.28 percent. We take m to be two standard deviations, or about 2.56 percent. This implies that ϕ must be greater than 0.2132, which it is.

The Third-Order Terms in the Approximation

We take a quadratic approximation of the period profit function. This requires that third-order terms are negligible relative to the second-order terms. To evaluate this assumption, we compute the ratio of the third-order term in the profit function to the second-order term. We use our calibrated parameters and evaluate $\ln Q - \ln Q^*$ at the bands. The ratio is .28. Hence it is not obvious that the third-order terms are small.[25] Below, we solve a nonlinear version of the firm's problem and find that these terms do affect the position of the bands but have a negligible effect on the dynamics of the price level relative to the steady state.

The Continuation Value Following Idiosyncratic Shocks

When we took the second-order approximation of the value function, we ignored all terms involving the arrival of the idiosyncratic shock. According to proposition 1, these terms were third-order. We now show that these terms are indeed small in our calibration.

Note that we can rewrite equation (A1) as follows:

$$E\{\Lambda_{t,t+1} \exp[(\varepsilon - 1)\xi_{t+1}] v(\mu_{t+1}, \Omega_{t+1})\} =$$

$$C + E_t\left\{\Lambda_{t,t+1} \frac{1}{\phi} \int_{\ln \mu^L_{t+1}}^{\ln \mu^H_{t+1}} \exp[(\varepsilon - 1)\Phi_{t+1}][\bar{v}(\exp(\Phi_{t+1}), \Omega_{t+1})\right.$$

$$\left. - \bar{v}(\mu^*_{t+1}, \Omega_{t+1})] d\Phi_{t+1}\right\} \exp[-(\varepsilon - 1)\ln \mu_t],$$

where C and the coefficient on $\exp[-(\varepsilon - 1)\ln \mu_t]$ are independent of μ_t. When

[24] Assuming that the firm needed to pay d to learn aggregate information would further reduce the decision cost. Removing the real rigidity would raise $\underline{d}$ to about $.02\bar{Y}$.

[25] Devereux and Siu (2004) argue in another context that these third-order terms may be quantitatively important.

we take a second-order approximation with respect to $\ln \mu_t$, the coefficient on $\hat{\mu}^2$ is

$$\frac{(\varepsilon-1)^2}{2} E_t\left\{\Lambda_{t,t+1}\frac{1}{\phi}\int_{\ln \mu^L_{t+1}}^{\ln \mu^H_{t+1}} \exp\left[(\varepsilon-1)\Phi_{t+1}\right]\left[\bar{v}(\exp(\Phi_{t+1}),\Omega_{t+1})\right.\right.$$
$$\left.\left. -\bar{v}(\mu^*_{t+1},\Omega_{t+1})\right]d\Phi_{t+1}\right\}\exp\left[-(\varepsilon-1)\ln\bar{\mu}_t\right].$$

Now the term in brackets is bounded above by b. We can get some idea of how large this coefficient is by replacing the term in brackets by b and evaluating it at the steady-state values of the other variables. The result is $.239\bar{Y}$. To get an idea of how large the effect is that we have omitted from our approximation of the value function, we need to multiply this by $\beta\tau(1-\alpha)/(1-\alpha\beta)$ in order to account for all the times this term enters the present value calculation (the term appears in period $t+i$ with probability $\tau[1-\alpha]\alpha^{i-1}$). The resulting coefficient is $.227\bar{Y}$. This should be compared with the coefficient on $\hat{\mu}^2$ that we include in our approximation. This coefficient is

$$\frac{1}{2}\chi_2\left(\frac{W}{P}\right)^{\varepsilon-1}\bar{Y} = \frac{\varepsilon-1}{2}\bar{Y}.$$

Given the parameters in our calibration, this is equal to $5\bar{Y}$. It follows that the omitted terms are indeed small relative to the terms that we include. For our parameterization the coefficient on the continuation value following the idiosyncratic shock is an order of magnitude smaller than the coefficient on profits prior to the shock.

The Importance of the Decision Cost

The introduction of the decision cost d solves a particular technical problem in the linearization. An *Ss* model has threshold rules that make the model difficult to linearize. Without the decision cost, a firm with Q_t in the neighborhood of Q^H_t will want to adjust in period $t+1$ if Q^H falls and not adjust if Q^H rises. Since the price index is equal to the average of $\ln Q$, this creates a nondifferentiability of the price index: shocks in one direction may trigger adjustment, whereas shocks in the other direction may not. When the idiosyncratic productivity shock hits, this nondifferentiability does not matter, and the idiosyncratic shock smooths it out. We introduce the decision cost to eliminate this nondifferentiability in other states of the world.[26]

In the model, the decision cost affects only the calculation of the bands Q^H and Q^L, and these bands were used only in the calculation of the price index. We did not need the decision cost for proposition 1 since that proposition considered only states of the world in which the idiosyncratic shock arrived. We did not need the decision cost for the calculation of the optimal target Q^*.

[26] Dotsey et al. (1999) eliminate these nondifferentiabilities by introducing idiosyncratic cost shocks. We cannot do this since the distribution of prices for firms that last received the idiosyncratic shock at date t but chose not to adjust would no longer be uniform between $\ln Q^L_t$ and $\ln Q^H_t$. It would be a convolution of this uniform distribution and the additional idiosyncratic shock.

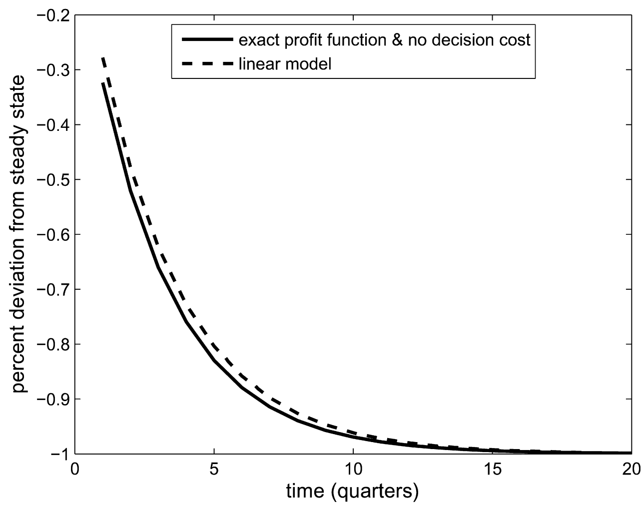

FIG. C1.—Comparing the response of the price level to a 1 percent money shock in the linear model to the response with exact profit functions and no decision cost.

Since the bands are wide relative to the aggregate fluctuations, firms at the target rarely reach the bands before they receive another idiosyncratic shock.

We can get some idea of how the decision cost affects the aggregate dynamics by considering how firms might want to adjust if we removed the decision cost. We consider an experiment similar to that of Golosov and Lucas (2007). In this experiment, we set the decision cost equal to zero. We consider a 1 percent reduction in the money supply and solve for the perfect foresight dynamics by iterating on the model. We begin by assuming that the aggregate variables follow the paths predicted by our linear model. We then calculate the optimal adjustment policies, given perfect foresight of these aggregate variables. We then use these policies to construct a new price index. Here we assume that prices are initially distributed according to the stationary distribution that would result if the aggregate shocks were all zero. We resolve the linear model taking this price as given and use these aggregate variables as inputs into the next stage. We iterate to convergence.

We perform this experiment using the quadratic approximation of the profit function and the exact profit function. In both cases, the resulting equilibrium price paths are close to those of our linear model. Figure C1 presents the comparison for the case of the exact profit function.

We can get some intuition for why the decision cost does not matter by considering the period after the shock. Suppose that we normalize the model such that initially $\ln Q_0^* = 0$. The stationary distribution of $\ln Q$ has a mass of firms massed at zero equal to $(1 - \tau) + \tau[1 - (2\omega/\phi)]$ and the remainder distributed uniformly with a density of τ/ϕ, just under 4 in our parameterization, over the interval $[-0.036, 0.048]$. Note that since we are using the exact profit function, the bands are asymmetric. In the period after the shock, the barriers fall to -0.041 and 0.043. Without the decision cost, those left in the interval $[0.043, 0.048]$ want to adjust. With the decision cost, they do not. How large an effect does this have? Recall that only α, just under one-half, of the firms originally in this interval remain. Therefore, only a small fraction, $\alpha(\ln Q_0^H - \ln Q_1^H)/\phi$ or just over 0.8 percent, of firms are affected by the decision cost. If they are allowed to adjust, these firms charge $\ln Q_1^* = -.005$ rather than the average of $\ln Q_0^H$ and $\ln Q_1^H$. This implies an average price change of just over 5 percent. Firms affected by the decision cost therefore contribute $-0.008 \times 0.05 \sim -0.0004$ to the decline in the price level. The difference is small relative to the decline in the price level itself, which is 0.28 percent in the linear model. The difference is even smaller in subsequent periods since the policies move by less. It is also smaller in response to a positive shock to the money supply, since the skewness of the *Ss* bands implies that the affected firms adjust by only 4 percent.

Larger monetary shocks reproduce figure C1 on a larger scale. The two impulse responses look exactly the same; only the vertical axis changes. Once the money shock exceeds 5 percent, the model begins to break down. Firms initially at Q^* may find themselves outside of the bands. This has a big effect on the performance of the model.

References

Akerlof, George, and Janet Yellen. 1985. "A Near-Rational Model of the Business Cycle with Wage and Price Inertia." *Q.J.E.* 100 (suppl.): 823–38.

Ball, Laurence, and David Romer. 1990. "Real Rigidities and the Non-neutrality of Money." *Rev. Econ. Studies* 57 (April): 183–203.

Basu, Susanto. 1995. "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare." *A.E.R.* 85 (June): 512–31.

Benabou, Roland. 1988. "Search, Price Setting and Inflation." *Rev. Econ. Studies* 55 (July): 353–76.

Caballero, Ricardo, and Eduardo Engel. 1991. "Dynamic (*S*, *s*) Economies." *Econometrica* 59 (November): 1659–86.

Calvo, Guillermo. 1983. "Staggered Prices in a Utility-Maximizing Framework." *J. Monetary Econ.* 12 (September): 383–98.

Caplin, Andrew, and John Leahy. 1991. "State-Dependent Pricing and the Dynamics of Money and Output." *Q.J.E.* 106 (August): 683–708.

———. 1997. "Aggregation and Optimization with State-Dependent Pricing." *Econometrica* 65 (May): 601–23.

Caplin, Andrew, and Daniel Spulber. 1987. "Menu Costs and the Neutrality of Money." *Q.J.E.* 102 (November): 703–25.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005."Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *J.P.E.* 113 (February): 1–45.

Clarida, Richard, Jordi Gali, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Q.J.E.* 115 (February): 147–80.

Cogley, Timothy, and Argia Sbordone. Forthcoming. "A Search of a Stable Phillips Curve." *A.E.R.*

Danziger, Lief. 1999. "A Dynamic Economy with Costly Price Adjustments." *A.E.R.* 89 (September): 878–901.

Devereux, Michael, and Henry Siu. 2004. "State Dependent Pricing and Business Cycle Asymmetries." Working paper, Dept. Econ., Univ. British Columbia.

Dotsey, Michael, and Robert King. 2005. "Implications of State-Dependent Pricing for Dynamic Macroeconomic Models." *J. Monetary Econ.* 52 (January): 213–42.

Dotsey, Michael, Robert King, and Alex Wolman. 1999. "State Dependent Pricing and the General Equilibrium Dynamics of Money and Output." *Q.J.E.* 104 (May): 655–90.

Eichenbaum, Martin, and Jonas Fisher. 2004. "Evaluating the Calvo Model of Sticky Prices." Working Paper no. 10617, NBER, Cambridge, MA.

Fabiani, Silvia, et al. 2004. "The Pricing Behavior of Firms in the Euro Area: New Survey Evidence." Working paper, European Central Bank, Frankfurt.

Gali, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Approach." *J. Monetary Econ.* 44 (October): 195–222.

Gali, Jordi, Mark Gertler, and David Lopez-Salido. 2001. "European Inflation Dynamics." *European Econ. Rev.* 45 (June): 1237–70.

Golosov, Mikhail, and Robert E. Lucas Jr. 2007. "Menu Costs and Phillips Curves." *J.P.E.* 115 (April): 171–99.

Kimball, Miles. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model." *J. Money, Credit and Banking* 27 (November): 1247–77.

Klenow, Peter, and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Q.J.E.* 123 (August).

Krusell, Per, and Anthony Smith. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *J.P.E.* 106 (October): 867–96.

Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable. 1997. "The Magnitude of Menu Cost: Direct Evidence from a Large U.S. Supermarket Chain." *Q.J.E.* 112 (August): 791–825.

Mankiw, N. Gregory. 1985. "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly." *Q.J.E.* 100 (May): 529–39.

Midrigan, Virgiliu. 2006. "Menu Costs, Multi-product Firms, and Aggregate Fluctuations." Working paper, Ohio State Univ.

Nakamura, Emi, and Jon Steinsson. 2007. "Five Facts about Prices: A Reevaluation of Menu Cost Models." Working paper, Harvard Univ.

Sbordone, Argia. 2002. "Prices and Unit Labor Costs: A New Test of Price Stickiness." *J. Monetary Econ.* 49 (March): 265–92.

Willis, Jonathan. 2002. "General Equilibrium of a Monetary Model with State-Dependent Pricing." Working paper, Fed. Reserve Bank Kansas City.

Woodford, Michael. 2002. "Inflation Stabilization and Welfare." *Contributions to Macroeconomics* 2 (1). http://www.bepress.com/bejm/contributions/vol2/iss1/art1.

———. 2003. *Interest and Prices.* Princeton, NJ: Princeton Univ. Press.

Zbaracki, Mark, Mark Ritson, Daniel Levy, and Shantanu Dutta. 2004. "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets." *Rev. Econ. and Statis.* 86 (May): 514–33.