

NBER WORKING PAPER SERIES

A PHILLIPS CURVE WITH AN SS FOUNDATION

Mark Gertler
John Leahy

Working Paper 11971
<http://www.nber.org/papers/w11971>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2006

We would like to thank Andrew Caplin, Jordi Gali, Mikhail Golosov, Per Krusell, Robert Lucas, Lars Svensson, Andrea Tambalotti, and Michael Woodford. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Mark Gertler and John Leahy. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Phillips Curve with an Ss Foundation
Mark Gertler and John Leahy
NBER Working Paper No. 11971
January 2006, Revised March 2006
JEL No. E1, E3

ABSTRACT

We develop an analytically tractable Phillips curve based on state-dependent pricing. We differ from the existing literature by considering a local approximation around a zero inflation steady state and introducing idiosyncratic shocks. The resulting Phillips curve is a simple variation of the conventional time-dependent Calvo formulation, but with some important differences. First, the model is able to match the micro evidence on both the magnitude and timing of price adjustments. Second, holding constant the frequency of price adjustment, our state-dependent model exhibits greater flexibility in the aggregate price level than does the time-dependent model. On the other hand, with real rigidities present, our state-dependent pricing framework can exhibit considerable nominal stickiness, of the same order by a conventional time-dependent model.

Mark Gertler
Department of Economics
New York University
269 Mercer Street, 7th Floor
New York, NY 10003
and NBER
mark.gertler@nyu.edu

John Leahy
Department of Economics
New York University
269 Mercer Street
New York, NY 10003
and NBER
john.leahy@nyu.edu

1. 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 the future. In addition to its forward-looking nature, this relationship also differs from the traditional Phillips curve in that all the coefficients are explicit functions of the primitives of the model.

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 where 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 which are, in general, difficult to aggregate.² For this reason, the time-dependent approach has proven 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; KK) 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 KK provide support for these notions by showing that a conventional state-dependent pricing model (Dotsey, King and Wolman, 1999; DKW) and a conventional time-dependent model (Calvo, 1983) yield very similar dynamics when calibrated to recent U.S. data.

A interesting recent paper by Golosov and Lucas (2003; GL) challenges this rationalization. The authors first note that to reconcile the evidence on the large size of individual firm price adjustments in the KK data with the low US 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 (due to 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 within a state-dependent

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

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

model with idiosyncratic productivity shocks, an exogenous shock to the money supply has a much stronger effect on the price level and a much weaker effect on real output than it does within 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 non-neutrality of money resulting from time dependence on its head.

Because pricing behavior in their model is very complex, GL 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, for example, that to obtain an empirically reasonable degree of nominal stickiness within a time-dependent price 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 idiosyncratic shocks and real rigidities. Accordingly, 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 DKW state-dependent framework, we focus on a local approximation around the steady state. We differ from DKW by introducing idiosyncratic shocks, as in GL. We differ from GL, 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.) On the other hand, our framework does capture the “selection” effect of state-dependent pricing: those farthest away from target tend to adjust more frequently, a feature that need 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 2 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

³Ball 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.

shocks, but we differ from GL by assuming that at any moment in time, there is a spatial nature to the idiosyncratic shock; i.e., at any moment only a subset of the economy is hit by the turbulence from idiosyncratic shocks. In addition, following Danziger (1999), we restrict the distribution of idiosyncratic shocks to be uniform. As we show, the latter restriction greatly simplifies the aggregation of Ss policies, while the former permits the resulting Phillips curve to be as flexible in parametric form as the conventional New Keynesian Phillips curve.

In section 3 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. We then turn to the key theoretical result that makes the Ss problem tractable. Given the uniform distribution of idiosyncratic shocks, a firm adjusting at time t can ignore the future states of the world where an idiosyncratic shock hits, up to a second order. Put differently, up to a second order, the firm’s continuation value conditional on an idiosyncratic shock at $t + 1$ is independent of its price at t . As we show, this “simplification” theorem makes the state-dependent pricing problem as easy to solve as the conventional time-dependent pricing problem. We then proceed to derive an approximate analytical solution, which includes deriving loglinear expressions for both the target price and the set of Ss bands.

In section 4 we characterize the complete model and present a log-linear approximation about the steady state. 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. As we discuss in section 5, a distinctive feature of our Phillips curve is that the key primitive parameter that enters the slope coefficient on the real activity measure (typically real marginal cost) is the Poisson arrival process for the idiosyncratic shock, as opposed to the measure of the degree of price rigidity that enters the standard formulation. The reason for this difference is that in within our framework, the frequency of price adjustment is endogenous and cannot be taken as a model primitive. As we show, further, because the frequency of the idiosyncratic shock will in general exceed the frequency of price adjustment, our state-dependent Phillips curve will exhibit greater price flexibility than the corresponding time-dependent relation. As in GL, the selection effect is at work: Firms that receive an idiosyncratic shock but do not adjust have a price that is already close to the target. On the other hand, because we can allow for real rigidities, our state-dependent framework is nonetheless capable of delivering considerable nominal stickiness.

In section 6 we calibrate the model to match the KK evidence on the frequency and absolute magnitude of price adjustments and also evidence on the costs of price adjustment. We then show that the framework 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 7.

2. Model: Environment

We begin with a conventional New Keynesian model. The basic features of the standard model include monopolistic competition, money, and nominal price stickiness. Also, for convenience, there are only consumption goods. To this familiar baseline framework we add three features: 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 modelling issues. The first is the need to conserve on state variables. In the most general state-dependent pricing model, the entire distribution of prices relative to the optimum 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 short cuts. Our strategy will be to make restrictions on the distribution of idiosyncratic shocks to simplify the distributional dynamics.⁴ We borrow our distributional assumption from 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 is able to solve. We differ from Danziger by allowing for a more flexible parameterization of the model in which significant nominal inertia and hence a significant first order effect of money on output is possible. Though we cannot solve for an exact solution, we can obtain an approximate analytical solution by consideration 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)).

The second modelling issue arises from the discontinuities and non-differentiabilities associated with in Ss adjustment. This issue potentially complicates finding a loglinear approximation of the model, since Taylor's theorem does not apply to functions that are not differentiable. Fortunately, as we discuss, this technical problem is applicable to only a small percentage of firms that happen to lie near the Ss bands and have not faced an idiosyncratic shock in the recent past. Since the bands tend to be wide relative to the aggregate shock, only these firms are potentially motivated to adjust their prices in response to the aggregate shock. 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 cost prior to the decision whether to adjust. This assumption, together with the assumption that aggregate shocks are small relative idiosyncratic shocks, guarantees that firms only consider adjusting when idiosyncratic shocks hit.⁵ It leads to smooth behavior of firms as they approach the Ss bands, eliminating any complications to linearizing our model. In an appendix we confirm that this decision

⁴Caplin and Spulber (1987), Caplin and Leahy (1991, 1997), and Benabou (1988) also make distributional assumptions that reduce the state space. DKW make assumptions that limit the number of prices observed in the economy to a finite number. Willis (2002) follows Krussel and Smith (1988) and approximates the distribution by a finite number of moments. GL 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.

⁵Note that a firm may or may not adjust in the wake of an idiosyncratic shock. Thus the price adjustment frequency is endogenous and not simply tied to the frequency of idiosyncratic shocks. Indeed, in general the former will be smaller than the latter.

cost need only be very tiny and that setting the decision cost to zero has only a minor effect on the dynamics of the model.

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

2.1. 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 only supply labor on the island that they live. There is perfect consumption insurance across islands, and any firm profits are redistributed lump sum to households.

Time is discrete and indexed by t . Let C_t be consumption; M_t 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+1}^n the nominal interest rate from t to $t + 1$. Then the objective for a representative household on island z is given by:

$$\max E_t \sum \beta^i \left\{ \log \left[C_{t+i} \cdot \left(\frac{M_{t+i}}{P_{t+i}} \right)^\nu \right] - \frac{1}{1+\varphi} N_{z,t+i}^{1+\varphi} \right\} \quad (2.1)$$

subject to budget constraint:

$$C_t = \frac{W_{z,t}}{P_t} N_{z,t+i} + \Gamma_{z,t} - \frac{M_t - M_{t-1}}{P_t} - \frac{(1/R_t^n)B_t - B_{t-1}}{P_t} \quad (2.2)$$

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:

$$\frac{W_{z,t}}{P_t} = \frac{N_{z,t+i}^\varphi}{(1/C_t)} \quad (2.3)$$

$$E_t \left\{ \beta \frac{C_t}{C_{t+1}} R_t^n \frac{P_t}{P_{t+1}} \right\} = 1 \quad (2.4)$$

$$\frac{M_t}{P_t} = \nu C_t \frac{R_t^n}{R_t^n - 1} \quad (2.5)$$

2.2. Final goods firms

Production occurs in two stages. Monopolistically competitive intermediate 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 $P_{z,t}^j$ be the associated nominal price. The production function for final goods is the following CES aggregate of intermediate goods:

$$Y_t = \left[\int_0^1 \int_0^1 (Y_{z,t}^j)^{\frac{\varepsilon-1}{\varepsilon}} dj dz \right]^{\frac{\varepsilon}{\varepsilon-1}} \quad (2.6)$$

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

$$Y_{z,t}^j = \left(\frac{P_{z,t}^j}{P_t} \right)^{-\varepsilon} Y_t \quad (2.7)$$

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

$$P_t = \left[\int_0^1 \int_0^1 [P_{z,t}^j]^{1-\varepsilon} dj dz \right]^{\frac{1}{1-\varepsilon}} \quad (2.8)$$

2.3. Intermediate goods firms

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

$$Y_{z,t}^j = X_{z,t}^j \cdot N_{z,t}^j \quad (2.9)$$

Here $X_{z,t}^j$ 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 multiplicative i.i.d. productivity shocks. These shocks follow a compound Poisson process. The arrival of the shock on island z at date t (i.e., whether firms are subject to a draw from random productivity variable at t) is perfectly correlated across all firms on the island and independent of shocks to other islands. The realization of the draw, however, is uncorrelated across firms on the island. Specifically, let $\xi_{z,t}^j$ denote the shock to firm j on island z at date t ,

$$X_{z,t}^j = \begin{cases} X_{z,t-1}^j e^{\xi_{z,t}^j} & \text{if a productivity shock occurs} \\ X_{z,t-1}^j & \text{otherwise} \end{cases} \quad (2.10)$$

We assume that the productivity shock arrives with probability $1 - \alpha$ and that conditional on arrival the random variable $\xi_{z,t}^j$ is distributed uniformly with density $1/\phi$ such that

$$E\{e^{(\varepsilon-1)\xi}\} = 1$$

As will become obvious, this normalization ensures that the expected multiplicative impact of the shock on the firm's discounted profits is unity. In addition, we assume that the support of $\xi_{z,t}^j$ is sufficiently large relative to the steady state Ss band. Let ω denote the absolute value of the log deviation of the each band from the target price in steady state. We assume:

$$\phi > 4\omega \tag{2.11}$$

As illustrated by Danziger (1999), the uniform distribution satisfying (2.11) introduces considerable tractability to the general Ss problem. The distribution of prices following an idiosyncratic shock has a simple form: uniform within the adjustment triggers and a mass at the target. As we show, this feature makes possible a reasonably simple approximation of the solution to the decision problem. It also simplifies the steady state equilibrium, as well as the local approximation around the steady state.^{6,7}

Each producer faces a fixed cost of adjusting price. We assume that this cost takes the form $b_t(X_{z,t}^j)^{\varepsilon-1}$ where

$$b_t = \begin{cases} b & \text{if } P_t^j \neq P_{t-1}^j \\ 0 & \text{otherwise} \end{cases} \tag{2.12}$$

We scale the adjustment cost by the factor $(X_{z,t}^j)^{\varepsilon-1}$ to keep the firm's decision problem homogenous as it size varies.⁸ This adjustment cost is in units of the final consumption good.

The fixed cost b 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. As noted earlier, Ss policies lead naturally to non-differentiabilities that make linearization difficult. In our case, this problem arises with firms close to the boundaries of the range of inaction that do not receive an idiosyncratic shock. For these firms an aggregate shock in one direction will cause some of these firms to adjust, whereas a shock in the other direction will take all of them deeper into the inaction region. There is therefore a kink in the response of these firms to the aggregate state.

Since the aggregate shock will be small, this non-differentiability arises among a small number of

⁶The steady state gap between the upper and lower adjustment triggers will be 2ω . 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 inequality implies that this condition is met in a neighborhood of the steady state.

⁷Note that our shock process is only conditionally uniform. With our Poisson assumption, the ex ante distribution of shocks combines a mass at zero with the wide uniform distribution.

⁸If the economy were growing we would also have to normalize the cost of price adjustment by the real wage and aggregate output.

firms. Moreover, since these firms lie near the boundary of the inaction region, they are essentially indifferent between adjustment and non-adjustment. Our solution to the differentiability problem is to add a small decision cost so that these firms never choose to adjust in response to an aggregate shock. 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 the state of the economy and to also organize this information to contemplate a price adjustment, they must pay a small decision cost, $d \cdot (X_{z,t}^j)^{\varepsilon-1}$. If a firm elects to pay the decision cost, it then can decide whether to adjust price. If it chooses to adjust, then it also incurs the fixed cost $b(X_{z,t}^j)^{\varepsilon-1}$. We make assumptions that guarantee that a firm will only elect to pay the decision cost if idiosyncratic turbulence hits. In particular, we assume that

$$0 < d < \frac{\varepsilon - 1}{1 - \alpha\beta} \left(\frac{1}{12}\phi^2 - \omega^2 + \frac{4}{3}\frac{\omega^3}{\phi} \right).$$

As the appendix shows, the upper bound ensures that the firm wishes to pay the decision cost whenever it observes the arrival of the idiosyncratic shock.⁹ The lower bound ensures that there exists a neighborhood of the steady state in which firms choose not to pay d whenever the idiosyncratic shock does not arrive.

We emphasize that the main justification for introducing the decision cost is that it solves a particular technical problem that arises in the log-linearization, and by doing so yields a simple solution to a complex aggregation problem. In the appendix we present a numerical solution to a version of the model that omits the decision cost and show that it yields dynamics nearly identical to those of our baseline model that includes the decision costs. Intuitively, since we will be considering only small perturbations about steady state, the decision cost only affects a small number of firms near the Ss bands, and only affects the aggregate dynamics through the effect that these firms have on the price level. For the remainder of the paper we only need to keep the implication that a firm contemplates price adjustment if and only if it observes the arrival of the idiosyncratic shock.

2.4. The monetary shock

We close the model with an exogenous stationary process for either the money supply or the nominal interest rate. We present the exact form of this process when we discuss the complete linearized model.

⁹To calculate the amount that a firm is willing to pay to find out the state, consider a firm whose price is at the frictionless optimum and receives an idiosyncratic productivity shock. We will see below that the loss from non-adjustment is approximately equal to $(\varepsilon - 1)\xi^2$. We integrate this loss over the region of price adjustment $[-\phi/2, -\omega] \cup [\omega, \phi/2]$ using the density $1/\phi$, and subtract the cost of price adjustment. $1 - \alpha\beta$ comes from taking the present value until the next price adjustment. All other initial conditions create greater gains to learning.

3. The Firm's Optimal Pricing Decision

Given the fixed cost of price adjustment, the solution to the firm's optimal policy will involve an Ss-style of price adjustment. Specifically, there will be a range of inaction, where the gain in discounted earnings from adjusting is not sufficient to cover the fixed cost. The optimal policy will involve an upper trigger, a lower trigger, and a target price. The firm adjusts when its price either reaches or moves beyond either of the trigger prices.

In this section we first characterize the firm's objective function. We argue that based on a plausible assumption about the size of the adjustment cost b , 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. With this simplified objective, we characterize both the steady state and a log-linear approximation of the decision rules about the steady state.

Before continuing with the firm's problem it is useful to state the following lemma which is proved in an appendix.

Lemma 3.1. *Given the assumption on the decision cost d , there exists a neighborhood of the non-stochastic steady state in which the firm pays the decision cost if and only if the firm experiences an idiosyncratic shock.*

The intuition behind the lemma is straightforward. At the beginning of the period the firm only initiates the price adjustment process if the gains justify the total cost $(b + d)X^{\varepsilon-1}$.¹⁰ After paying the cost d , the firm adjusts if the gains outweigh the menu cost $bX^{\varepsilon-1}$. Hence the range of inertia in the first phase is greater than the range of inertia in the second. The firm only initiates the price adjustment process if there is an event big enough that there is a good chance it is taken outside the larger bands. Given the assumption that the idiosyncratic shocks are much larger than the aggregate shocks, the firm reacts to the former but not the latter. The full problem of the firm is solved in the appendix.

In the next section, we appeal to the lemma and impose that the firm adjusts its price only in response to the idiosyncratic shock. This simplifies the exposition by eliminating the decision cost d from the discussion.

3.1. The Firm's Objective

Real profits net of adjustment costs, $\tilde{\Pi}_{t+i}^j$, are given by

¹⁰Zbaracki et al. (2004) find that the managerial costs of information gathering and decision making d are much larger than the physical costs of adjusting prices b . Fabiani et al. (2004) find that firms in the Euro area review their prices more often than they change them.

$$\tilde{\Pi}_{z,t+i}^j = \left(\frac{P_{z,t+i}^j}{P_{t+i}} - \frac{W_{z,t+i}}{P_{t+i}X_{z,t+i}} \right) Y_{t+i} - b_{t+i}(X_{z,t+i}^j)^{\varepsilon-1} \quad (3.1)$$

where b_{t+i} is defined by equation (2.12). Note that from cost minimization, the firm's real marginal cost is $W_{z,t+i}/X_{z,t+i}$.

It is convenient to define the “normalized” price, $Q_{z,t}^j$, which is the price, $P_{z,t}^j$, normalized by multiplicative impact of the idiosyncratic productivity shock on the firm's marginal cost ($1/X_{z,t}^j$):

$$Q_{z,t}^j = \frac{P_{z,t}^j}{1/X_{z,t}^j} = P_{z,t}^j X_{z,t}^j$$

There are two advantages of working with the normalized price. First, assuming that the firm's desired markup is stationary, $Q_{z,t}^j$ is stationary. In contrast $P_{z,t}^j$ is nonstationary since $X_{z,t}^j$ is nonstationary. Second, all firms that reset their price in period t will wind up choosing the same normalized price, which simplifies the aggregation. Since idiosyncratic productivity differs across firms, firms will not choose the same absolute price. Note that because $Q_{z,t}^j$ depends on $X_{z,t}^j$, it may change even if the firm keeps its nominal price constant.

Restating period profits in terms of the normalized price and making use of the demand function the firm faces (equation (2.7)) yields

$$\tilde{\Pi}_{z,t+i} = X_{z,t+i}^{\varepsilon-1} \left[P_{t+i}^{\varepsilon-1} \left(Q_{z,t+i}^j \right)^{-\varepsilon} \left(Q_{z,t+i}^j - W_{z,t+i} \right) Y_{t+i} - b_{t+i} \right] \equiv X_{z,t+i}^{\varepsilon-1} \Pi_{t+i} \quad (3.2)$$

At this point we drop the j, z subscripts. It is useful to define the variable $A_{t+i} = P_{t+i}^{\varepsilon-1} Y_{t+i}$ and $\Lambda_{t,t+i} = U'(C_{t+i})/U'(C_t)$. We define the firm's value function as the maximized stream of discounted net profits, as follows:

$$\begin{aligned} V(Q_{t-1}e^{\xi_t}, \Omega_t) &= \max E_t \sum_{i=0}^{\infty} \beta^i \Lambda_{t,t+i} \tilde{\Pi}_{t+i} \\ &= \max X_t^{\varepsilon-1} E_t \sum_{i=0}^{\infty} \beta^i \Lambda_{t,t+i} \left(\frac{X_{t+i}}{X_t} \right)^{\varepsilon-1} \Pi_{t+i} \end{aligned} \quad (3.3)$$

V depends on the normalized price inherited from the previous period. We write this price as $Q_{t-1}e^{\xi_t}$ with the understanding that $\xi_t = 0$ when there is no idiosyncratic productivity shock. V also depends on Ω_t which summarizes the aggregate state of the economy, which depends on the current values of C_t, Y_t, W_t , and P_t , as well as their future evolution.

Given that gross profits and adjustment costs are homogeneous in $X_t^{\varepsilon-1}$, it is convenient to define

the normalized value function $v(\cdot)$:

$$V(Q_{t-1}e^{\xi_t}, \Omega_t) = X_t^{\varepsilon-1} \cdot v(Q_{t-1}e^{\xi_t}, \Omega_t) \quad (3.4)$$

with

$$v(Q_{t-1}e^{\xi_t}, \Omega_t) = \max E_t \sum_{i=0}^{\infty} \left(\frac{X_{t+i}}{X_t} \right)^{(\varepsilon-1)} \beta^i \Lambda_{t,t+i} [A_{t+i} Q_{t+i}^{-\varepsilon} (Q_{t+i} - W_{t+i}) - b_{t+i}]$$

To express the normalized value function in a recursive form, let $\bar{v}(Q_t, \Omega_t)$ denote the value after price adjustment. Then

$$v(Q_{t-1}e^{\xi_t}, \Omega_t) = \max \left\{ \bar{v}(Q_{t-1}e^{\xi_t}, \Omega_t), \max_{Q_t} \bar{v}(Q_t, \Omega_t) - b \right\} \quad (3.5)$$

and

$$\begin{aligned} \bar{v}(Q_t, \Omega_t) &= A_t Q_t^{-\varepsilon} (Q_t - W_t) \\ &+ \beta E_t \Lambda_{t,t+1} \left\{ \alpha \bar{v}(Q_t, \Omega_{t+1}) + (1 - \alpha) e^{(\varepsilon-1)\xi_{t+1}} v(Q_t e^{\xi_{t+1}}, \Omega_{t+1}) \right\}. \end{aligned} \quad (3.6)$$

Given all our assumptions, what complicates the firm's problem, in general, is that it must take account of the continuation value conditional on an idiosyncratic shock, $E_t \left\{ \Lambda_{t,t+1} e^{(\varepsilon-1)\xi_{t+1}} v(\cdot) \right\}$. Absent this consideration, the choice of the target price at time t would just involve taking into account discounted profits in states where the firm's price remains fixed at its period t target. In this respect, the choice of the target is no more difficult than in the conventional time-dependent framework. The choice of the triggers also simplifies.

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

3.2. Approximate Value Function

It is convenient to define the target and trigger in logarithmic terms. Let $\ln Q_t^*$ denote the natural log of the target (normalized) price and let $\ln Q_t^L$ and $\ln Q_t^H$ be the natural logs of the upper and lower triggers. Under the Ss policy, the firm adjusts to $\ln Q_t^*$ if $\ln(Q_t) \notin [\ln Q_t^L, \ln Q_t^H]$. Our goal now is to derive an approximate value function that leads a tractable (approximate) solution to the decision problem

We begin by assuming that the fixed cost of price adjustment b is second order. As is well known, doing so implies that the range of inaction $[\ln Q_t^L, \ln Q_t^H]$ is first order (e.g. Mankiw (1985), Akerlof and Yellen (1985)). This in turn implies that it is reasonable to restrict attention to a second order

approximation of profits.¹¹

There is an important additional implication of our “small” b assumption: 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 Q_t up to a second order.

Proposition 3.2. *Suppose (a) b is second order (implying $\ln Q_t^L - \ln Q_t^*$ and $\ln Q_t^H - \ln Q_t^*$ are first-order), and (b) $\phi > 2(\ln Q_t^H - \ln Q_t^L)$, then the expected value at date t of an optimal policy after an idiosyncratic shock at date $t + 1$, $E\{e^{(\varepsilon-1)\xi_{t+1}}v(Q_t e^{\xi_{t+1}}, \Omega_{t+1})\}$, is independent of the current value of Q_t to a second order. In particular, the firm can treat its objective as*

$$\bar{v}_n(Q_t, \Omega_t) = \Pi_t + \alpha\beta E_t \{\Lambda_{t,t+1} v(Q_t, \Omega_{t+1})\} + \mathcal{O}^3$$

The main insight of the proposition is that in future states where the idiosyncratic shock will hit, history will be erased. The subsequent continuation value $E_t \left\{ \Lambda_{t,t+1} e^{(\varepsilon-1)\xi_{t+1}} v(Q_t e^{\xi_{t+1}}, \Omega_{t+1}) \right\}$ is irrelevant to current pricing decision to a second order. This proposition has the flavor of an envelope theorem.

In the appendix we provide a formal proof of the proposition. Here we present the intuition, which follows from Figure 1. Consider a firm with log normalized price equal to $\ln Q_t$ that receives an idiosyncratic shock in period $t + 1$. The shock leaves the firm log-uniformly distributed between $\ln Q_t - \phi/2$ and $\ln Q_t + \phi/2$. Now in period $t + 1$ the firm follows a pricing strategy characterized by the triplet $\{\log Q_{t+1}^L, \log Q_{t+1}^*, \log Q_{t+1}^H\}$. Given this policy, $\ln Q_{t+1}$ will be uniformly distributed over $(\ln Q_{t+1}^L, \ln Q_{t+1}^H)$ if the firm does not adjust (the dark gray region in Figure 1). If the firm does adjust (the light gray regions of Figure 1), then $\ln Q_{t+1} = \ln Q_{t+1}^*$. Since the triplet $\{\log Q_{t+1}^L, \log Q_{t+1}^*, \log Q_{t+1}^H\}$ is independent of Q_t (it depends only on the state at $t + 1$), it follows that the distribution of $Q_t e^{\xi_{t+1}}$ and hence $v(Q_t e^{\xi_{t+1}}, \Omega_{t+1})$ is independent of Q_t . This can be seen from Figure 1: A shift in $\ln Q_t$, shifts the entire distribution of $\ln Q_t + \xi_{t+1}$. This does not affect the distribution after adjustment, only the states in which the firm adjusts up and the states in which the firm adjusts down.

Accordingly the only way that Q_t could possibly affect $E_t \left\{ \Lambda_{t,t+1} e^{(\varepsilon-1)\xi_{t+1}} v(Q_t e^{\xi_{t+1}}, \Omega_{t+1}) \right\}$, is by affecting the correlation between $\Lambda_{t,t+1} e^{(\varepsilon-1)\xi_{t+1}}$ and $v(Q_t e^{\xi_{t+1}}, \Omega_{t+1})$. However, given our restrictions on b , this correlation is second order and its dependence on dependence on Q_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 Q_t . The second is that b is second order, which makes the correlation between the decision to change price and Q_t third order.

¹¹The terms first order and second order refer to rates of convergence relative to some scaling variable. In this case that variable is \sqrt{b} .

3.3. Approximate Optimal Pricing Policy

Armed with the preceding proposition, we now take a second-order approximation of the profit function about the frictionless optimal price. Let Q_t^o be the optimal normalized price in the frictionless optimum (i.e. the optimum with no adjustment costs.), then

$$\Pi_t = \chi_1 Y_t (W_t/P_t)^{1-\varepsilon} - \chi_2 Y_t (W_t/P_t)^{1-\varepsilon} (\ln Q_t - \ln Q_t^o)^2 + \mathcal{O}^3 \quad (3.7)$$

where χ_1 and χ_2 are constants, with $\chi_2 = \frac{1}{2} \frac{\varepsilon^{1-\varepsilon}}{(\varepsilon-1)^{-\varepsilon}}$, and \mathcal{O}^3 collects terms that are third order. Because we are approximating the profit function about the frictionless optimal price, the first-order term is zero.¹² Given the elasticity of demand, Q_t^o is simply a markup over the nominal wage:

$$Q_t^o = \mu W_t \quad (3.8)$$

with $\mu = \varepsilon/(\varepsilon - 1)$.

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 firms only consider adjustment only following the idiosyncratic shock, it follows that Q remains fixed in all states in which the idiosyncratic shock does not hit. We can therefore write \bar{v} as:

$$\begin{aligned} \bar{v}(Q_t, \Omega_t) &= E_t \sum_{i=0}^{\infty} (\alpha\beta)^i \Lambda_{t,t+i} \left[\chi_1 Y_{t+i} (W_{t+i}/P_{t+i})^{1-\varepsilon} - \chi_2 Y_{t+i} (W_{t+i}/P_{t+i})^{1-\varepsilon} (\ln Q_t - \ln Q_{t+i}^o)^2 \right] \\ &\quad + [\text{terms independent of } Q_t] + \mathcal{O}^3 \end{aligned} \quad (3.9)$$

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 term in this expression is dated t since the normalized price remains constant in these states.

It is now straightforward to derive the optimality conditions for the target and the two triggers. The first-order necessary condition for the target Q^* is given by:

$$E_t \sum_{i=0}^{\infty} (\alpha\beta)^i \Lambda_{t,t+i} \left[Y_{t+i} (W_{t+i}/P_{t+i})^{1-\varepsilon} (\ln Q_t^* - \ln Q_{t+i}^o) \right] = 0 \quad (3.10)$$

The triggers in turn are given by a value matching condition that equates the gain from not adjusting to the gain from adjustment, net the adjustment cost: For $J = H, L$:

$$\bar{v}(Q_t^J, \Omega_t) = \bar{v}(Q_t^*, \Omega_t) - b \quad (3.11)$$

¹²Given that the first-order term is zero implies that our linear-quadratic model will not be subject to the problems described in Woodford (2002).

Given our quadratic approximation, the value matching condition can be restated as:

$$\begin{aligned} & E_t \sum_{i=0}^{\infty} (\alpha\beta)^i \left[\chi_2 Y_{t+i} (W_{t+i}/P_{t+i})^{1-\varepsilon} (\ln Q_t^J - \ln Q_{t+i}^o)^2 \right] \\ &= E_t \sum_{i=0}^{\infty} (\alpha\beta)^i \left[\chi_2 Y_{t+i} (W_{t+i}/P_{t+i})^{1-\varepsilon} (\ln Q_t^* - \ln Q_{t+i}^o)^2 \right] + b \end{aligned} \quad (3.12)$$

Since we are interested in a local approximation about the steady state, we now analyze the non-stochastic steady state as a necessary first step.

3.4. Non-stochastic Steady State

We first set the aggregate shocks at their respective means. The only disturbance in the steady state is the idiosyncratic productivity shock, which washes out in the aggregate. We will use bars above variables to indicate steady state values.

It is straightforward to derive the optimal steady state target and adjustment triggers. Given that P_t and W_t are fixed, it follows from the first-order conditions (3.8) and (3.10) that the steady state target price, \bar{Q}^* , is a constant equal to the steady state frictionless optimal price \bar{Q}^o , as follows:

$$\bar{Q}^* = \bar{Q}^o = \mu \bar{W} \quad (3.13)$$

The steady state triggers are pinned down by the value matching condition with P_t and W_t at their respective steady state means:

$$\bar{v}(\bar{Q}^J, \bar{\Omega}) = \bar{v}(\bar{Q}^*, \bar{\Omega}) - b \quad (3.14)$$

for $J = H, L$. Given the quadratic approximation of $\bar{v}_N(\cdot)$ and given $\bar{Q}^* = \bar{Q}^o$, we can write:

$$\sum_{i=0}^{\infty} (\alpha\beta)^i \left[\chi_2 \bar{Y} (\bar{W}/\bar{P})^{1-\varepsilon} (\ln \bar{Q}^J - \ln \bar{Q}^o)^2 \right] = b \quad (3.15)$$

The solution to this quadratic equation yields two steady state triggers:

$$\begin{aligned} \ln \bar{Q}^H &= \ln \bar{Q}^o + \sqrt{(1-\alpha\beta) \frac{b}{\chi_2 \bar{Y} (\bar{W}/\bar{P})^{1-\varepsilon}}} = \ln \bar{Q}^o + \sqrt{2 \frac{1-\alpha\beta}{\varepsilon-1} \frac{b}{\bar{Y}}} \\ \ln \bar{Q}^L &= \ln \bar{Q}^o - \sqrt{(1-\alpha\beta) \frac{b}{\chi_2 \bar{Y} (\bar{W}/\bar{P})^{1-\varepsilon}}} = \ln \bar{Q}^o - \sqrt{2 \frac{1-\alpha\beta}{\varepsilon-1} \frac{b}{\bar{Y}}} \end{aligned} \quad (3.16)$$

The second equality follows from noting that $\bar{W} = \mu \bar{P}$ and substituting for χ_2 and μ in terms of ε . Note that since b is second order, the steady state bands $\ln \bar{Q}^H - \ln \bar{Q}^*$ and $\ln \bar{Q}^* - \ln \bar{Q}^L$ are first order and symmetric, as we maintained earlier.

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.

Finally, the steady state probability of price adjustment conditional on an idiosyncratic shock is $1 - \frac{\ln \bar{Q}^H + \ln \bar{Q}^L}{\phi}$.¹³ The unconditional probability of price adjustment, then is simply the product of the probability of an idiosyncratic shock times the probability of adjusting conditional on this shock: $(1 - \alpha) \cdot \left(1 - \frac{\ln \bar{Q}^H + \ln \bar{Q}^L}{\phi}\right) \equiv 1 - \theta$. In the Calvo model, this frequency of price adjustment is taken as a primitive. Here, the primitive is the frequency of idiosyncratic shocks. The frequency of price adjustment depends on the likelihood of adjustment conditional on an idiosyncratic shock, where the latter depends on the primitives of the model, including the distribution of idiosyncratic shocks, the fixed costs, and so on.

The average time a price is fixed is simply the inverse of the frequency of price adjustment and is given by $1/(1 - \theta)$. Note that, in general, the average time a price is fixed exceeds the average amount of time in between idiosyncratic shocks, $\frac{1}{1 - \alpha}$:

$$\frac{1}{(1 - \theta)} = \frac{1}{(1 - \alpha) \cdot \left(1 - \frac{\ln \bar{Q}^H + \ln \bar{Q}^L}{\phi}\right)} > \frac{1}{1 - \alpha} \quad (3.17)$$

This difference occurs because firms may choose to keep their prices fixed in the event of a shock. Again, however, those firms that do not adjust in this instance will have their price within a first order of the target price.

3.5. Aggregate Shocks and Local Dynamics

We now consider (small) aggregate shocks to the steady state. Let lowercase letters represent log deviations from steady state values, so that $q_t^* = \ln Q_t^* - \ln \bar{Q}_t^*$ and let $w_t = \ln W_t - \ln \bar{W}$. Log-linearizing (3.10) about the steady state values of P , W , Q^0 and Q^* :

$$\begin{aligned} q_t^* &= (1 - \beta\alpha) E_t \sum_{i=0}^{\infty} (\beta\alpha)^i q_{t+i}^o + \mathcal{O}^2 \\ &= (1 - \beta\alpha) E_t \sum_{i=0}^{\infty} (\beta\alpha)^i w_{t+i} + \mathcal{O}^2 \end{aligned} \quad (3.18)$$

since $q_t^o = w_t$. As in the pure time-dependent model, the target depends on a discounted stream of future values of nominal marginal cost. In the time-dependent framework, however, future marginal cost in each period is weighted by the probability that the price remains fixed. In our state-dependent

¹³Since both $\bar{q}^H + \bar{q}^L$ and ϕ may be first order $(\ln \bar{Q}^H + \ln \bar{Q}^L) / \phi$ need not approach zero as the b approaches zero.

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 the price has stayed fixed.¹⁴

We next consider the local dynamics for the optimal triggers. Log-linearizing (3.12) about the steady state values of P , W , Q^0 and Q^* and using the definition (3.18) yields

$$q_t^H = q_t^* + (1 - \beta\alpha) \frac{\ln \bar{Q}^H - \ln \bar{Q}^o}{2} E_t \sum_{i=0}^{\infty} (\alpha\beta)^i [y_{t+i} + (\varepsilon - 1)(w_{t+i} - p_{t+i}) - c_{t+i} + c_t] \quad (3.19)$$

$$q_t^L = q_t^* - (1 - \beta\alpha) \frac{\ln \bar{Q}^o - \ln \bar{Q}^L}{2} E_t \sum_{i=0}^{\infty} (\alpha\beta)^i [y_{t+i} + (\varepsilon - 1)(w_{t+i} - p_{t+i}) - c_{t+i} + c_t] \quad (3.20)$$

Note that the width of the bands $q_t^H - q_t^L$ may fluctuate. However, given the quadratic profit function, they widen and contract symmetrically. Thus, given the uniform distribution, the average price within bands is simply the target: $\frac{1}{2}(q_t^H + q_t^L) = q_t^*$. This result will prove useful when we next consider the local dynamics of the price level.¹⁵

3.6. Price Index

At this point we reintroduce the j and z indexes. We may express the price index in terms of normalized prices, as follows:

$$\begin{aligned} P_t &= \left(\int \int P_{z,t}^j 1^{-\varepsilon} dj dz \right)^{\frac{1}{1-\varepsilon}} \\ &= \left(\int \int Q_{z,t}^j 1^{-\varepsilon} X_{z,t}^j \varepsilon^{-1} dj dz \right)^{\frac{1}{1-\varepsilon}} \end{aligned} \quad (3.21)$$

Given that Q is equal to \bar{Q}^o to a first order, the right-hand side is equal to

$$\left(\int \int Q_{z,t}^j 1^{-\varepsilon} X_{z,t}^j \varepsilon^{-1} dj dz \right) = \ln \bar{P} + \int \int q_{z,t}^j dj dz + \mathcal{O}^2$$

It follows that

$$p_t = \int \int q_{z,t}^j dj dz \quad (3.22)$$

¹⁴Note that we have taken two distinct approximations: a second-order approximation of the firm's value function and a linearization of the first order conditions. The first approximation depends on the value of b . The second is in the deviation of the aggregate variables from steady state. One can think of this as taking these two limits in succession. First, choose b small enough such that the firms' objectives are approximately quadratic. Second, choose an aggregate forcing process that is small enough that it does not trigger further adjustment.

¹⁵The 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 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 Q .

Now consider an island z that received an idiosyncratic shock at date $t - i$. Firms that adjusted at that date set their price equal to q_{t-i}^* . Those that did not adjust remained uniformly distributed on (q_{t-i}^L, q_{t-i}^H) . Given (3.19) and (3.20), the average price of the non-adjusters is also q_{t-i}^* . Hence the average price on island z is q_{t-i}^* . Given that the arrival rate of the shock is $1 - \alpha$, it follows that

$$p_t = (1 - \alpha) \sum_{i=0}^{\infty} \alpha^i q_{t-i}^* \quad (3.23)$$

which may be expressed as

$$p_t = (1 - \alpha)q_t^* + \alpha p_{t-1} \quad (3.24)$$

4. 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 “state-dependent” Phillips curve and then turn our attention to the rest of the model.

Manipulation of (3.18) yields the optimal reset price q_t^* as the following discounted stream of future nominal wages.

$$q_t^* = (1 - \beta\alpha)E_t \sum_{i=0}^{\infty} (\beta\alpha)^i w_{z,t+i} \quad (4.1)$$

Note that q_t^* depends on the island-specific wage $w_{z,t+i}$. As a step toward aggregation, we would like to derive this relation in terms of the economy-wide average wage, w_{t+i} .

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

$$w_{z,t+i} - p_{t+i} = \varphi n_{z,t+i} + c_{t+i} \quad (4.2)$$

Averaging over this condition yields $w_{t+i} - p_{t+i} = \varphi n_{t+i} + c_{t+i}$, implying the following relation between the island z relative wage and the relative employment levels:

$$w_{z,t+i} - w_{t+i} = \varphi(n_{z,t+i} - n_{t+i}) \quad (4.3)$$

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 :

$$w_{z,t+i} = w_{t+i} - \varphi \varepsilon (q_t^* - p_{t+i}) \quad (4.4)$$

Notice that $w_{z,t+i}$ depends inversely on 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 non-adjusting firms.¹⁶ This strategic complementarity, in turn, contributes to the overall stickiness in the movement of prices. Combining (4.4) with (4.1) yields

$$q_t^* = (1 - \beta\alpha) [\Psi(w_t - p_t) + p_t] + \beta\alpha E_t q_{t+1}^* \quad (4.5)$$

with

$$\Psi = \frac{1}{1 + \varphi\varepsilon} \quad (4.6)$$

In equilibrium, the real wages of adjusting firms, $w_{z,t} - p_t$ moves less than one for one with the aggregate real wage, implying similarly sluggish movement in the target price q_t^* . In this respect, the strategic complementarity measured inversely by the coefficient Ψ , dampens the adjustment of prices. With economy-wide labor markets, Ψ equals unity, implying $w_{z,t}$ simply is equal to w_t .

We are now in a position to present the Phillips curve. Let $\pi_t = p_t - p_{t-1}$ denote inflation. Combining the equation for the target price (4.5) with the price index (3.24) yields

$$\pi_t = \lambda(w_t - p_t) + \beta E_t \pi_{t+1} \quad (4.7)$$

and

$$\lambda = \frac{(1 - \alpha)(1 - \beta\alpha)}{\alpha} \Psi \quad (4.8)$$

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. In particular, for the time-dependent Phillips curve the slope coefficient, λ_{td} is given by

$$\lambda_{td} = \frac{(1 - \theta)(1 - \beta\theta)}{\theta} \Psi \quad (4.9)$$

where θ is the probability of no price adjustment, which is a primitive in the time-dependent framework.

Because in general $\alpha < \theta$ (since a fraction of firms that receive an idiosyncratic shock may not adjust), λ is greater than λ_{td} . The implication is that inflation is more sensitive to movements in real marginal cost in the state-dependent framework relative to the time-dependent model. This of course is just another way of saying that state dependence leads to greater price flexibility relative to time

¹⁶For a menu of alternative ways to introduce real rigidities, see Kimball (1995) and Woodford (2003).

dependence. The intuition for this outcome follows directly from GL, which in turn comes from Caplin and Spulber. Frequent idiosyncratic shocks give firms the option of also adjusting to aggregate shocks. Unlike the time-dependent case, firms not adjusting are those that are already close to the target. Thus in general, the state-dependent formulation will yield greater flexibility than does the time dependent. How much difference this makes, however, will depend upon the values of the primitive parameters of the model.

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.

$$w_t - p_t = \kappa(y_t - y_t^*) \tag{4.10}$$

where $\kappa = (\sigma + \varphi)$ is the elasticity of marginal cost and y_t^* is the natural (flexible price equilibrium) level of output. Combining equations (4.7) and (4.10) then yields the Phillips curve in terms of the output gap:

$$\pi_t = \lambda\kappa(y_t - y_t^*) + \beta E_t \pi_{t+1} \tag{4.11}$$

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:

$$y_t = -(r_t^n - E_t \pi_{t+1}) + E_t y_{t+1} \tag{4.12}$$

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

$$m_t - p_t = y_t - \zeta r_t^n \tag{4.13}$$

Equations (4.11), (4.12) and (4.13) determine the equilibrium aggregate dynamics, conditional on a monetary policy rule (and given the definition, $p_t = \pi_t + p_{t-1}$). To illustrate the model dynamics in a way that sharpens the focus on the issue of how state versus time dependence affects the degree of nominal stickiness, we close the model with a simple money growth rule:

$$m_t - m_{t-1} = \eta_t^m \tag{4.14}$$

where η_t^m is a mean zero i.i.d. exogenous shock to the evolution of the money stock.

Both recent empirical work and conventional wisdom hold that central banks use the nominal interest rate rather than a monetary aggregate as the instrument of monetary policy. We accordingly also explore closing the model with an interest rate rule of the following form:

$$r_t^n = \vartheta \pi_{t-1} + \rho r_{t-1}^n + \eta_t^{rn} \quad (4.15)$$

where η_t^{rn} is a mean zero i.i.d. exogenous shock to the evolution of the money stock

5. Properties of the Model

Before proceeding to some numerical exercises with the model, we first characterize some general properties. We begin by noting that our model is consistent with Klenow and Kryvtsov's (2003) evidence on the decomposition of inflation. In particular, these authors show that for the recent low inflation decade in the U.S.: (i) the proportion of firms that adjust their prices has fairly constant and that (ii) the variation in the inflation has driven almost entirely by variation in the size of price adjustment, not by variation in the frequency of price adjustment. The local approximation of our model has this property, even though it allows for state-dependent pricing. In particular, within the local approximation, the proportion of firms adjusting their price in any given period is practically constant. In particular, the fraction of firms that keep their prices fixed is equal to

$$\theta_t \approx \bar{\theta} + \frac{1 - \alpha}{\phi} (q_t^H - q_t^L) \quad (5.1)$$

In general the fluctuations in q_t^H and q_t^L will be quite small (see equations (3.19) and (3.20)), implying that movements in the fraction of firms that adjust $1 - \theta_t$ is quite small. Moreover, these movements have no effect on inflation, which depends only on the target q^* not the triggers q_t^H and q_t^L . Fluctuations in the rate of inflation are thus entirely explained by variation in the size of price adjustments, as in the standard time-dependent framework.

Second, we note that our state-dependent formulation of inflation is quite flexible: At one extreme the model can generate the kind of complete flexibility suggested by Caplin and Spulber. At the other, it can perfectly mimic the degree of nominal stickiness in the pure time-dependent Calvo model.

When the Ss bands are small relative to ϕ , most firms adjust in response to an idiosyncratic shock and $\alpha \sim \theta$. 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.

In the other extreme, when $\alpha = 1$ our model behaves exactly like a flexible price model: The slope coefficient on marginal cost in the Ss Phillips curve goes to infinity. In this case, the idiosyncratic productivity shock hits each firm each period. According to (3.18), $q_t^* = q_t^o$, and according to (3.23), the price index is equal to q_t^o as well. The economy is always at its frictionless optimum. Money is neutral. Neutrality holds in spite of the fact that a fraction θ of firms do not adjust their prices in each period.¹⁷

¹⁷Note that Danziger does not find neutrality in his model even though he assumes that $N = 1$. The reason is that

What is the source of this neutrality? It is instructive to analyze it both from the perspective of a firm and from the economy as a whole. Consider 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 mark up 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 be of 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 rises in the next period. A time-dependent firm would find that its price is too low. The state-dependent firm shifts the set of 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 Caplin and Spulber, 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 where 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, monetary policy will affect the distribution of markups. How important these effects are depends on the model calibration. We turn to this issue next.

6. 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 and the coefficient of relative risk aversion σ at 1.0. 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

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.

φ) at 1.0, which is a reasonable intermediate range value in the literature. Finally, we set the interest elasticity of money demand $\nu = .1$.

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 relative to average steady state firm output, b/\bar{Y} . Note first that the steady state Ss band $\omega = \ln \bar{Q} - \ln \bar{Q}^o$ is a function of these parameters: From equation (3.16): $\omega = \sqrt{\frac{1-\alpha\beta}{\varepsilon-1} \frac{b}{\bar{Y}}}$. The average size of price adjustment will depend on ω and the range of the idiosyncratic productivity shock (which depends on ϕ). In turn, the frequency of price adjustment $1 - \theta$ equals $(1 - \alpha)(1 - \frac{2\omega}{\phi})$, where $1 - \frac{2\omega}{\phi}$ is the probability of adjustment conditional on an aggregate shock. We can then proceed to derive a system of relations that pin down the triplet $(\alpha, \phi, b/\bar{Y})$, using evidence on: (i) the frequency of price adjustment; (ii) the absolute size of price adjustments; and (iii); the costs of price adjustment.

Klenow and Kryvtsov (2003) report that the median time a price is fixed is slightly over four months. Accordingly we fix the average frequency of price adjustment at $(1 - \alpha)(1 - \frac{2\omega}{\phi}) = 0.4$, which yields a median duration of prices of 1.36 quarters or 4.1 months.¹⁸ Next we require that the model match Klenow and Kryvtsov's evidence that the average absolute size of price adjustments is about 8.0 percent: This implies $\frac{\phi}{4} + \frac{\omega}{2} = .08$. Finally, we set the steady state resources devoted to price adjustment equal to 0.4 percent of revenue, based on the evidence in Zbaracki et al. (2004).¹⁹ This implies $b/\bar{Y} = .004/[(1 - \alpha)(1 - \frac{2\omega}{\phi})]$.

Table 1 shows the values of $\alpha, \phi, b/\bar{Y}$, and ω implied by this parameterization, as well as $1 - \frac{2\omega}{\phi}$, the probability of price adjustment conditional on the idiosyncratic shock. Note 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.²⁰

Table 2 shows implied value of λ , the slope the coefficient on marginal cost in Phillips curve given by equation (4.7). 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. The respective

¹⁸As Cogley and Sbordone (2005) note, given adjustment is binomial random variable, the time until the next adjustment can be approximated as an continuous time exponential random variable, implying a median waiting time equal to $-\ln(2)/\ln(\theta)$. Note that the median waiting time is less than the mean duration of prices ($1/0.4 = 2.5$ quarters, roughly seven and a half months) since the exponential distribution implies that some prices may not change for a very long time.

¹⁹Zbaracki et al. (2004) quantify the physical costs of price adjustment for a large manufacturing firm. 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. GL, on the other hand, find that a value of 0.24 allows their model to best match certain properties of the data. Our quantitative results are robust to either of these alternative values.

²⁰It is often noted that the size of price adjustment that we observe in the data is not constant, as it would be in a continuous time Ss model with a stationary and continuous forcing process. We see both large and small price adjustments in the data. This model is consistent with this observation. The largest price adjustment, $\phi/2 + \omega \sim 16\%$, occurs when a firm at the edge of the band receives a shock of size $\phi/2$. The smallest adjustment is $\omega \sim 3\%$. In the current parameterization the former is five times the latter. The reason that this model generates both large and small adjustment is that the shock process is discontinuous.

formulation for the two cases are given by.

$$\begin{aligned}\lambda &= \frac{(1-\alpha)(1-\beta\alpha)}{\alpha(1+\varphi\varepsilon)} \\ \lambda_{td} &= \frac{(1-\theta)(1-\beta\theta)}{\theta(1+\varphi\varepsilon)}\end{aligned}\tag{6.1}$$

where in each case, the term $1/(1+\varphi\varepsilon)$ reflects the influence of the strategic complementarity stemming from local labor markets²¹. For comparison, we also report slope coefficients for the case where real rigidities are absent (global labor markets). In this case $\varphi = 0$. Let a $\bar{\lambda}$ denote the case without real rigidities. The respective slope coefficients are then:

$$\begin{aligned}\bar{\lambda} &= \frac{(1-\alpha)(1-\beta\alpha)}{\alpha} \\ \bar{\lambda}_{td} &= \frac{(1-\theta)(1-\beta\theta)}{\theta}\end{aligned}\tag{6.2}$$

For the case with local labor markets ($\varphi = 1$), the slope coefficient is .053 for our Ss model, while for the Calvo model the parameter shrinks to about .023. λ exceeds λ_{td} by a factor of two, indicating greater nominal flexibility with state dependence. However, the absolute difference is small. Further, λ lies within the range of estimates reported by Galí 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.

We next consider the response of the model economy to an unanticipated monetary shock. We begin with the case where monetary policy is governed by the simple money growth rule defined in equation (4.14) and where real rigidities are absent. This scenario corresponds closest to the policy experiment considered by Golosov and Lucas (2005). Figure 2 illustrates the response of the model economy to a permanent 1 percent decrease in the money stock. The solid line is the response of our state-dependent pricing model, while 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 three 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 GL, the state-dependent model also exhibits greater nominal flexibility than does 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.

Figure 2 clearly illustrates the effect of the absence of complementarities. The optimal target price

²¹Introducing other complementarities such as firm-specific capital (Woodford 2005) or a chain of production (Basu 1995) may reduce this parameter further.

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.

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 ten 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 q^* . The initial response of the target price is only half the size of the money shock and it takes about 10 quarters for q^* to adjust to steady state. In this case even after the majority of firms have adjusted their prices, the economy will not have returned to 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.

The simple exogenous process for monetary policy given by equation (4.14) is useful for illustrating the degree of nominal inertia in the model, but is not an empirically reasonable characterization of central bank behavior. A vast recent literature suggests that simple interest rate rules in the spirit of equation (4.15) provide a better characterization of central bank policy for industrialized economies than do money growth rules. We accordingly reconsider money shock experiment, replacing equation (4.14) with equation (4.15) in the model economy. In line with recent estimates (e.g. Clarida, Gali and Gertler (2000)), we set the coefficient on inflation at 1.5 and the interest rate smoothing parameter at 0.7.

Figure 4 reports the impulse responses of the model economy to an unanticipated 100 basis point increase in the short-term interest rate, for the case with local labor markets. Note that for the Ss model there is a sharp decline in real output that lasts nearly a year. There is relatively little movement in the price level, a reflection that under the interest rate policy the money supply reverts at least part of the way back to trend. As always, the Calvo model exhibits more nominal stickiness, but the relative differences in the behavior of output and inflation are small in percentage terms. More significantly, however, with complementarities present, the Ss model is capable of generating considerable nominal stickiness.

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. On the other hand, it is straightforward to add a number of features (e.g., habit formation, investment with delays and adjustment costs, and so on) that have proved useful in improving the empirical performance of such models.

7. 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 (2005), we find that for a given frequency of price adjustment, the Ss model exhibits greater nominal flexibility than a corresponding time-dependent framework, due to a selection effect: Firms farthest away from target adjust in the Ss model, while 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.

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 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.

²²For 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.

8. Appendix: Proofs

Proof of the Lemma 3.1: We prove the lemma in two parts. In the first part, we show that there exists a neighborhood of the steady state in which the firm would never choose to pay the decision cost if the idiosyncratic shock did not hit. In the second part, we show that the firm would always choose to pay the decision cost should the idiosyncratic shock hit.

The first part of the proof is constructive. Fix a date t and a firm j . We suppose that all firms in the economy are following the strategy of paying the decision cost if and only if they are hit by the idiosyncratic shock, and when they do pay the decision cost, they adjust according to the policy $\{q_s^L, q_s^*, q_s^H\}$. We suppose that at t , firm j does not receive the idiosyncratic shock. We show that j chooses not to pay the decision cost.

Suppose that the firm last experienced an idiosyncratic shock at date $s < t$. By assumption, the firm paid the decision cost at that date and as a consequence learned the state of the world. Let $\Omega_{s,t}$ denote the firm's information at date t based on information collected at date s . $\Omega_{s,t}$ contains all of the aggregate variables, as well as the idiosyncratic shock at date s . By assumption, no idiosyncratic shock has hit the firm since s and the firm has not adjusted its price; therefore the firm knows its current normalized price and it is Q_s .

Let $v(Q_s, \Omega_{s,t})$ denote the optimal policy given the normalized price Q_s and the information $\Omega_{s,t}$. We describe the firm's decision in three steps. The first step is whether or not to pay the decision cost. We have

$$v(Q_s, \Omega_{s,t}) = \max\{E[\hat{v}(Q_s, \Omega_{t,t})|\Omega_{s,t}] - d, E[\Pi_t(Q_s) + \beta\Lambda_{t,t+1}v(Q_s, \Omega_{s,t+1})|\Omega_{s,t}]\}$$

If the firm pays the decision cost d , then it receives the expected value of $\hat{v}(Q_t, \Omega_{t,t})$, where $\hat{v}(Q_s, \Omega_{t,t})$ is the value after having paid the decision cost of receiving the updated information $\Omega_{t,t}$ and having the option to alter its price from Q_s . This expectation is taken with respect to the current information $\Omega_{s,t}$. If the firm does not pay the decision cost, then it cannot alter its price, it receives the expected profit $\Pi_t(Q_s)$ defined in equation (3.2) and the discounted value of $v(Q_t, \Omega_{s,t+1})$, where here the information is one period older. Again the expectation is taken with respect to $\Omega_{s,t}$.

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

$$\hat{v}(Q_s, \Omega_{t,t}) = \max\{\bar{v}(Q_s, \Omega_{t,t}), \max_Q \bar{v}(Q, \Omega_{t,t}) - b\}$$

where \bar{v} (as in the text) is the value after having made the pricing decision. As the final step, we define \bar{v} :

$$\bar{v}(Q, \Omega_{t,t}) = \Pi_t + \beta E\{\Lambda_{t,t+1}v(Q, \Omega_{t,t+1})|\Omega_{t,t}\}$$

Now, by assumption, $\ln Q_s \in [q_s^L, q_s^H]$, since the firm last paid the decision cost in period s . Recall the firm cannot adjust without paying d . The first thing to note is that in steady state j would never

pay the decision cost at t . The reason is that in steady state $q_j \in [\bar{q}^L, \bar{q}^H]$ which is within b of the optimum. Paying the decision cost would only make sense if the firm wanted to adjust its price, but then the total cost of adjustment would be $b + d > b$. Paying the extra d would only make sense if there were some extra value to paying the decision cost beyond the option of price adjustment. In steady state, this is only the case in the event of the idiosyncratic shock.

Now suppose that the aggregate shocks are small enough that they keep the firm in a neighborhood of the steady state $\ln Q_s \in [q_s^L, q_s^H] \subset [\bar{q}^L - \delta_1, \bar{q}^H + \delta_1]$ with probability $1 - \delta_2$ for some $\delta_1, \delta_2 > 0$ and all s . Note the loss to non-adjustment in period t for a firm with normalized price Q_s relative to a firm at q_t^L can be no worse than $\chi_2 Y_t (W_t/P_t)^{1-\varepsilon} [(q_t^L - \delta_1 - q_t^*)^2 - (q_t^L - q_t^*)^2]$ (with probability $1 - \delta_2$). As the aggregate shock shrinks to zero, δ_1 and δ_2 approach zero as well, and the difference between these losses also approaches zero. Since the loss to non-adjustment at q_t^L is b , there exists a neighborhood of the steady state in which the loss to non-adjustment is less than $b + d$. This establishes the first part of the proposition.

To establish the second part, consider the steady state and a firm with the optimal price. It is easy to show that any other firm would have a strictly greater desire to pay the decision cost in steady state. Arguments similar to the first part show that outside steady state, the cost to not paying the decision cost will be in a neighborhood of the steady state cost. Suppose that this firm receives an idiosyncratic shock. The firm compares the cost of inaction until the next idiosyncratic shock to the benefit of learning the state and acting optimally. Should the latter exceed the former by more than the decision cost, the firm will choose the latter. This difference is equal to

$$\frac{2(\varepsilon - 1)\bar{Y}}{(1 - \alpha\beta)\phi} \int_{\omega}^{\phi/2} [q^2 - \omega^2] dq$$

If the firm learns that the idiosyncratic shock is less than ω , it will not adjust and there is no difference between the two options. If the idiosyncratic shock is $q > \omega$, the firm that adjusts receives an added benefit of $\frac{(\varepsilon-1)\bar{Y}}{(1-\alpha\beta)} [q^2 - \omega^2]$ over the period until the next price adjustment. This benefit is integrated with respect to the density $1/\phi$ to the maximum shock $\phi/2$. The 2 comes from the symmetry of the problem: whatever happens at the upper trigger, happens at the lower trigger. This expression simplifies to

$$\frac{(\varepsilon - 1)\bar{Y}}{(1 - \alpha\beta)} \left[\frac{\phi}{12} - \omega^2 + \frac{4\omega^3}{3\phi} \right]$$

Therefore if the decision cost is less than the bound stated in the Lemma, the firm would choose to pay the decision cost in the event of an idiosyncratic shock. This establishes the second part of the argument and completes the proof of the lemma. QED

Proof of Proposition 3.2: Suppose that the firm has a current level of Q_t such that $\ln Q_t \in [\ln Q_t^L, \ln Q_t^H]$. We are interested in the expected value of an optimal policy conditional on an idiosyn-

cratic productivity shock in period $t + 1$. Also let Q_{t+1}^* denote the optimal choice of Q_{t+1} in the event of adjustment.

Consider $E\{\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}}v(Q_t e^{\xi_{t+1}}, P_{t+1}, W_{t+1})\}$ over the states of the world in which the idiosyncratic shock hits. Given the assumption on ϕ , $\xi^H > q_{t+1}^H - \ln Q_t$ and $\xi^L < q_{t+1}^L - \ln Q_t$:

$$\begin{aligned} & E\{\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}}v(Q_t e^{\xi_{t+1}}, P_{t+1}, W_{t+1})\} \\ = & E \left\{ \Lambda_{t,t+1} \frac{1}{\phi} \int_{\ln Q_{t+1}^H - \ln Q_t}^{\xi^H} e^{(\varepsilon-1)\xi_{t+1}} \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1}) d\xi_{t+1} \right. \\ & + \Lambda_{t,t+1} \frac{1}{\phi} \int_{\ln Q_{t+1}^L - \ln Q_t}^{\ln Q_{t+1}^H - \ln Q_t} e^{(\varepsilon-1)\xi_{t+1}} \bar{v}_n(Q_t e^{\xi_{t+1}}, P_{t+1}, W_{t+1}) d\xi_{t+1} \\ & \left. + \Lambda_{t,t+1} \frac{1}{\phi} \int_{\xi^L}^{\ln Q_{t+1}^L - \ln Q_t} e^{(\varepsilon-1)\xi_{t+1}} \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1}) d\xi_{t+1} \right\} \end{aligned}$$

Rearranging yields:

$$\begin{aligned} & E\{\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}} v(Q_t e^{\xi_{t+1}}, P_{t+1}, W_{t+1})\} \\ = & E \left\{ \Lambda_{t,t+1} \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1}) \right. \\ & \left. + \Lambda_{t,t+1} \frac{1}{\phi} \int_{\ln Q_{t+1}^L - \ln Q_t}^{\ln Q_{t+1}^H - \ln Q_t} e^{(\varepsilon-1)\xi_{t+1}} \left[\bar{v}(Q_t e^{\xi_{t+1}}, P_{t+1}, W_{t+1}) - \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1}) \right] d\xi_{t+1} \right\} \end{aligned}$$

A change of variable, $\Phi_{t+1} = \ln Q_t + \xi_{t+1}$, gives:

$$\begin{aligned} & E\{\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}} \Lambda_{t,t+1}(Q_{t+1}^*, P_{t+1}, W_{t+1})\} \tag{A1} \\ = & E \left\{ \Lambda_{t,t+1} \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1}) \right. \\ & \left. + \Lambda_{t,t+1} \frac{1}{\phi} \int_{\ln Q_{t+1}^L}^{\ln Q_{t+1}^H} e^{(\varepsilon-1)(\Phi_{t+1} - \ln Q_t)} \left[\bar{v}(e^{\Phi_{t+1}}, P_{t+1}, W_{t+1}) - \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1}) \right] d\Phi_{t+1} \right\} \end{aligned}$$

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

$$\begin{aligned} & \Lambda_{t,t+1} e^{(\varepsilon-1)(\Phi_{t+1} - \ln Q_t)} \left[\bar{v}(e^{\Phi_{t+1}}, P_{t+1}, W_{t+1}) - \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1}) \right] \\ = & \Lambda_{t,t+1} e^{(\varepsilon-1)(\ln Q_{t+1}^* - \ln Q_t^*)} \left[\bar{v}(e^{\Phi_{t+1}}, P_{t+1}, W_{t+1}) - \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1}) \right] + \mathcal{O}^3 \end{aligned}$$

Further, the assumption that $\phi > 2(\ln Q_t^H - \ln Q_t^L)$ and that the range of integration is first order,

$$\begin{aligned} & E\{\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}}v(Q_t e^{\xi_{t+1}}, P_{t+1}, W_{t+1})\} \\ = & E\left\{\Lambda_{t,t+1}\bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1})\right. \\ & \left.+ \Lambda_{t,t+1}\frac{1}{\phi}\int_{q_{t+1}^L}^{q_{t+1}^H}e^{(\varepsilon-1)(\ln Q_{t+1}^* - \ln Q_t^*)}[\bar{v}(e^{\Phi_{t+1}}, P_{t+1}, W_{t+1}) - \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1})]d\Phi_{t+1}\right\} + \mathcal{O}^3 \end{aligned}$$

where the integral is second or third order (it is second order if ϕ is first order and third order if ϕ is fixed independently of b). It follows that $E\{e^{(\varepsilon-1)\xi_{t+1}}v(Q_t e^{\xi_{t+1}}, P_{t+1}, W_{t+1})\}$ is independent of Q_t to a second order. QED

9. Appendix: 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) the idiosyncratic shocks are dispersed enough that firms always expect that both price increases and decreases are possible; (2) The Ss bands are small enough that the second-order approximation of the profit function is valid and the continuation value following an idiosyncratic shock is third order; (3) the aggregate shocks are sufficiently small that the log-linearization is accurate. In this appendix, we evaluate the first of two assumptions. The third 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. To summarize, the model holds up well.

9.1. Idiosyncratic shocks

The wide distribution of idiosyncratic shocks greatly simplifies the solution to the model. In the limit, as the variance of the aggregate shock shrinks to zero, the condition $\phi > 4\omega$ assumed in equation (2.11) is sufficient to ensure that some firms raise their prices while others lower their prices. Away from steady state, we also need the range of the idiosyncratic shock to be wide enough to compensate for movements in the bands caused by aggregate shocks. Fortunately, we have a lot of leeway on this dimension. In our calibration, $\phi = .254$, whereas $4\omega = .132$. In the simulations the effect of a monetary shock on the position of the bands is essentially the same as the effect of a shock on the desired price. Since this latter effect is considerably less than ω , there is ample room for the idiosyncratic shock to compensate for movements in the bands.

9.2. 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 steady-state value for the bands. The ratio is .21. Hence it is not obvious that the third-order terms are small.²³ Below, we solve a non-linear version of the firm's problem and find that these terms do affect the steady state position of the bands, but have a negligible effect on the dynamics of the price level relative to steady state.

9.3. 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 (8.1) as follows

$$E\{\Lambda_{t,t+1}e^{(\varepsilon-1)\xi_{t+1}}v(Q_t e^{\xi_{t+1}}, P_{t+1}, W_{t+1})\} \\ = C + E_t \left\{ \Lambda_{t,t+1} \frac{1}{\phi} \int_{\ln Q_{t+1}^L}^{\ln Q_{t+1}^H} e^{(\varepsilon-1)(\Phi_{t+1})} [\bar{v}(e^{\Phi_{t+1}}, P_{t+1}, W_{t+1}) - \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1})] d\Phi_{t+1} \right\} e^{-(\varepsilon-1)\ln Q_t}$$

where C and the coefficient on $e^{-(\varepsilon-1)\ln Q_t}$ are independent of Q_t . Taking a second-order approximation with respect to $\ln Q_t$ about the frictionless optimal price, the coefficient on $(\ln Q - \ln Q^*)^2$ is

$$\frac{(\varepsilon-1)^2}{2} E_t \left\{ \Lambda_{t,t+1} \frac{1}{\phi} \int_{\ln Q_{t+1}^L}^{\ln Q_{t+1}^H} e^{(\varepsilon-1)(\Phi_{t+1})} [\bar{v}(e^{\Phi_{t+1}}, P_{t+1}, W_{t+1}) - \bar{v}(Q_{t+1}^*, P_{t+1}, W_{t+1})] d\Phi_{t+1} \right\} e^{-(\varepsilon-1)\ln Q_t^*}$$

Now the term in square 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 $.13\bar{Y}$. To get an idea of how large is the effect that we have omitted from our approximation of the value function, we need to multiply this by $\beta(1-\alpha)/(1-\alpha\beta)$ in order to account for all of the times that this term enters the present value calculation (The term appears in period $t+i$ with probability $(1-\alpha)\alpha^{i-1}$). The resulting coefficient is $.129\bar{Y}$. This should be compared with the coefficient on $(\ln Q - \ln Q^*)^2$ 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.

²³Devereux and Siu (2004) argue in another context that these third order terms may be quantitatively important.

9.4. The decision cost

The introduction of the decision cost d solves a particular technical problem in the linearization. Ss models have threshold rules that make them difficult to linearize. Without the decision cost, a firm with Q_t in the neighborhood of Q_t^H 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 non-differentiability of the price index in the neighborhood of the steady state: shocks in one direction may trigger adjustment, while shocks in the other direction may not. When the idiosyncratic productivity shock hits, this non-differentiability does not matter, the idiosyncratic shock smooths it out. We introduce the decision cost to eliminate this non-differentiability in other states of the world.²⁴

In the model, the decision cost only affects the calculation of the bands q^H and q^L , and these bands were only used in the calculation of the price index. We did not need the decision cost for Proposition 3.1, 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^* . 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. In this experiment, we set the decision cost equal to zero. We consider a 1 percent reduction in the money supply.²⁵ We assume that the aggregate variables follow the paths predicted by our model. We calculate the optimal non-linear adjustment policies, given perfect foresight of the response of the aggregate variables and the quadratic approximation of the profit function (3.7). We then calculate a new price index using equation from (3.22) by assuming that the distribution of prices begins in steady state and then iterating the constructed optimal policies. We can judge the effect of the decision cost by comparing this price index to that predicted by our model.

Figure 5 presents this comparison. The two price paths are very close. Eliminating the decision cost has very little effect on the price level.

To see why this is so, consider the period of the shock. The $\ln q$ are initially distributed uniformly over $[-0.30, 0.30]$. After the shock the barriers fall to 0.24 and -0.33 . Without the decision cost those left in the interval $[0.24, 0.30]$ want to adjust. With the decision cost they do not. How large an effect does this have? Recall that only α of the firms originally in this interval remain. A fraction $\alpha(q_0^H - q_1^H)/\phi$, or 1 percent, of firms are affected by the decision cost. These firms charge $q_1^* = .005$ rather than the average of q_0^H and q_1^H . This implies an average price change of 2.7 percent. The price

²⁴DKW eliminate these non-differentiabilities by introducing idiosyncratic cost shocks. We cannot do this, since the distribution of prices for firms who last received the idiosyncratic shock at date t but chose not to adjust, would no longer be uniform between $\ln Q_t^L$ and $\ln Q_t^H$. It would be a convolution of this uniform distribution and the additional idiosyncratic shock.

²⁵The Ss bands contract when the money supply falls. Hence reductions in the money supply lead to more adjustment than do increases. A reduction in the money supply is therefore a stronger test of the effect of the decision cost.

levels therefore differ by $.01 \cdot .027 \sim .0003$. The difference is an order of magnitude smaller than the change in the price level which is just under 0.3 percent. It is even smaller in subsequent periods since the policies move by less.

Larger monetary shocks reproduce Figure 5 on a larger scale. The two impulse response 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 $\ln q_0^*$ may find themselves outside of the bands. This has a big effect on the performance of the model.

As one final check on the model, we repeated the experiment in Figure 5, but this time we replaced the quadratic approximation of the profit function with the exact profit functions (3.1). Figure 6 shows the resulting impulse response. It is very close to the linear model. This does not mean that the third order terms do not have any effect. They do impact the position of the bands relative to the optimal price. In the linear model, $\omega = 3.3$ percent. With the exact profit function, the upper band is equal to 3.77 percent, whereas the lower band is equal to 2.97 percent. These differences have only a minor effect on the dynamics of the model, because the dynamics depend on the movements in the bands not the position of the bands.

References

- [1] Akerlof, George, and Janet Yellen (1985), "A Near-Rational Model of the Business Cycle with Wage and Price Inertia," *Quarterly Journal of Economics* **100**, 823-838.
- [2] Ball, Laurence, and David Romer (1990), "Real Rigidities and the Non-neutrality of Money," *Review of Economic Studies* **57**, 183-203.
- [3] Benabou, Roland (1988), "Search, Price Setting and Inflation," *Review of Economic Studies* **55**, 353-376.
- [4] Bilts, Mark, and Peter Klenow (2002), "Some Evidence on the Importance of Sticky Prices," NBER Working Paper No. 9069.
- [5] Basu, Susanto (1995), Intermediate Goods and Business Cycles: Implications for Productivity and Welfare," *American Economic Review* **85**, 512-531
- [6] Caballero, Ricardo, and Eduardo Engel (1991), "Dynamic (S, s) Economies," *Econometrica* **59**, 1659-86.
- [7] Calvo, Guillermo (1983), "Staggard Prices in a Utility-maximizing Framework," *Journal of Monetary Economics* **12**, 383-398.
- [8] Caplin, Andrew, and Daniel Spulber (1987), "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics* **102**, 703-725.
- [9] Caplin, Andrew, and John Leahy (1991), "State-Dependent Pricing and the Dynamics of Money and Output," *Quarterly Journal of Economics* **106**, 683-708.
- [10] Caplin, Andrew, and John Leahy (1997), "Aggregation and Optimization with State-Dependent Pricing," *Econometrica* **65**, 601-623.
- [11] Christiano, Larry, Martin Eichenbaum, and Charles Evans (2005), "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy* **113**, 1-45.
- [12] Clarida, Richard, Jordi Gali, and Mark Gertler (2000), "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics* **115**, 147-180.
- [13] Cogley, Timothy, and Argia Sbordone (2004), "A Search of a Stable Phillips Curve," Arizona State University working paper.
- [14] Danziger, Lief (1999), "A Dynamic Economy with Costly Price Adjustments," *American Economic Review* **89**, 878-901.

- [15] Devereux, Michel, and Henry Siu (2004), "State Dependent Pricing and Business Cycle Asymmetries," University of British Columbia Department of Economics Working Paper.
- [16] Dotsey, Michael, Robert King, Alex Wolman (1999), "State Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics*, **104**, 655-690.
- [17] Dutta, S., M. Bergen, D. Levy, and R. Venable (1999), "Menu Costs, Posted Prices, and Multi-product Retailers," *Journal of Money, Credit and Banking* **31**, 683-703.
- [18] Eichenbaum, Martin, and Jonas, Fisher (2004), "Evaluating the Calvo Model of Sticky Prices," NBER Working Paper No. 10617.
- [19] Fabiani, S., et al. (2004), "The Pricing Behavior of Firms in the Euro Area: New Survey Evidence," European Central Bank working paper.
- [20] Fischer, Stanley (1977), "Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule," *Journal of Political Economy* **85**, 191-205
- [21] Gali, Jordi, and Mark Gertler (1999), "Inflation Dynamics: A Structural Econometric Approach," *Journal of Monetary Economics* **44**, 195-222.
- [22] Gali, Jordi, Mark Gertler, and David Lopez-Salido (2001), "European Inflation Dynamics," *European Economic Review* **47**, 1237-1270.
- [23] Golosov, Mikhail, and Robert Lucas (2003), "Menu Costs and Phillips Curves," NBER Working Paper No. 10187.
- [24] Kimball, Miles (1995), "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking* **27**, 1247-1277.
- [25] Klenow, Peter, and Oleksiy Kryvtsov (2005) "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" NBER Working Paper No. 11043.
- [26] Krusell, Per, and Anthony Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy* **106**, 867-96.
- [27] Levy, D., M. Bergen, S. Dutta, and R. Venable (1997), "The Magnituded of Menu Cost: Direct Evidence from a Large U.S. Supermarket Chain," *Quarterly Journal of Economics* **112**, 791-825.
- [28] Mankiw, N. Gregory (1985), "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly," *Quarterly Journal of Economics*, **100**, 529-539.
- [29] Sbordone, Argia (2002), "Prices and Unit Labor Costs: A New Test of Price Stickiness," *Journal of Monetary Economics* **49**, 265-292.

- [30] Willis, Jonathan (2002), “General Equilibrium of a Monetary Model with State-Dependent Pricing,” mimeo, Federal Reserve Bank of Kansas City.
- [31] Woodford, Michael (2002) “Inflation Stabilization and Welfare,” *Contributions to Macroeconomics* **2**. (<http://www.bepress.com/bejm/contributions/vol2/iss1/art1>)
- [32] Woodford, Michael (2003), *Interest and Prices*, Princeton: Princeton University Press.
- [33] Zbaracki, Mark, Mark Ritson, Daniel Levy, and Shantanu Dutta (2004), “Managarial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets, *Review of Economics and Statistics* **86**, 514-533.

Table 1: Imputed values of the parameters

θ	α	$1 - \frac{2\omega}{\phi}$	ω	ϕ	b/\bar{Y}
.6	.4594	.7400	.0330	.2540	.0100

Table 2: The effect of complementarities on the coefficient on marginal cost

φ	λ	λ_{td}
0	.642	.271
1	.053	.023




Figure 1: A Firm's Response to a Productivity Shock





Figure 2: Response to a Money shock (no strategic complementarity)





Figure 3: Response to a Money Shock (with complementarities)



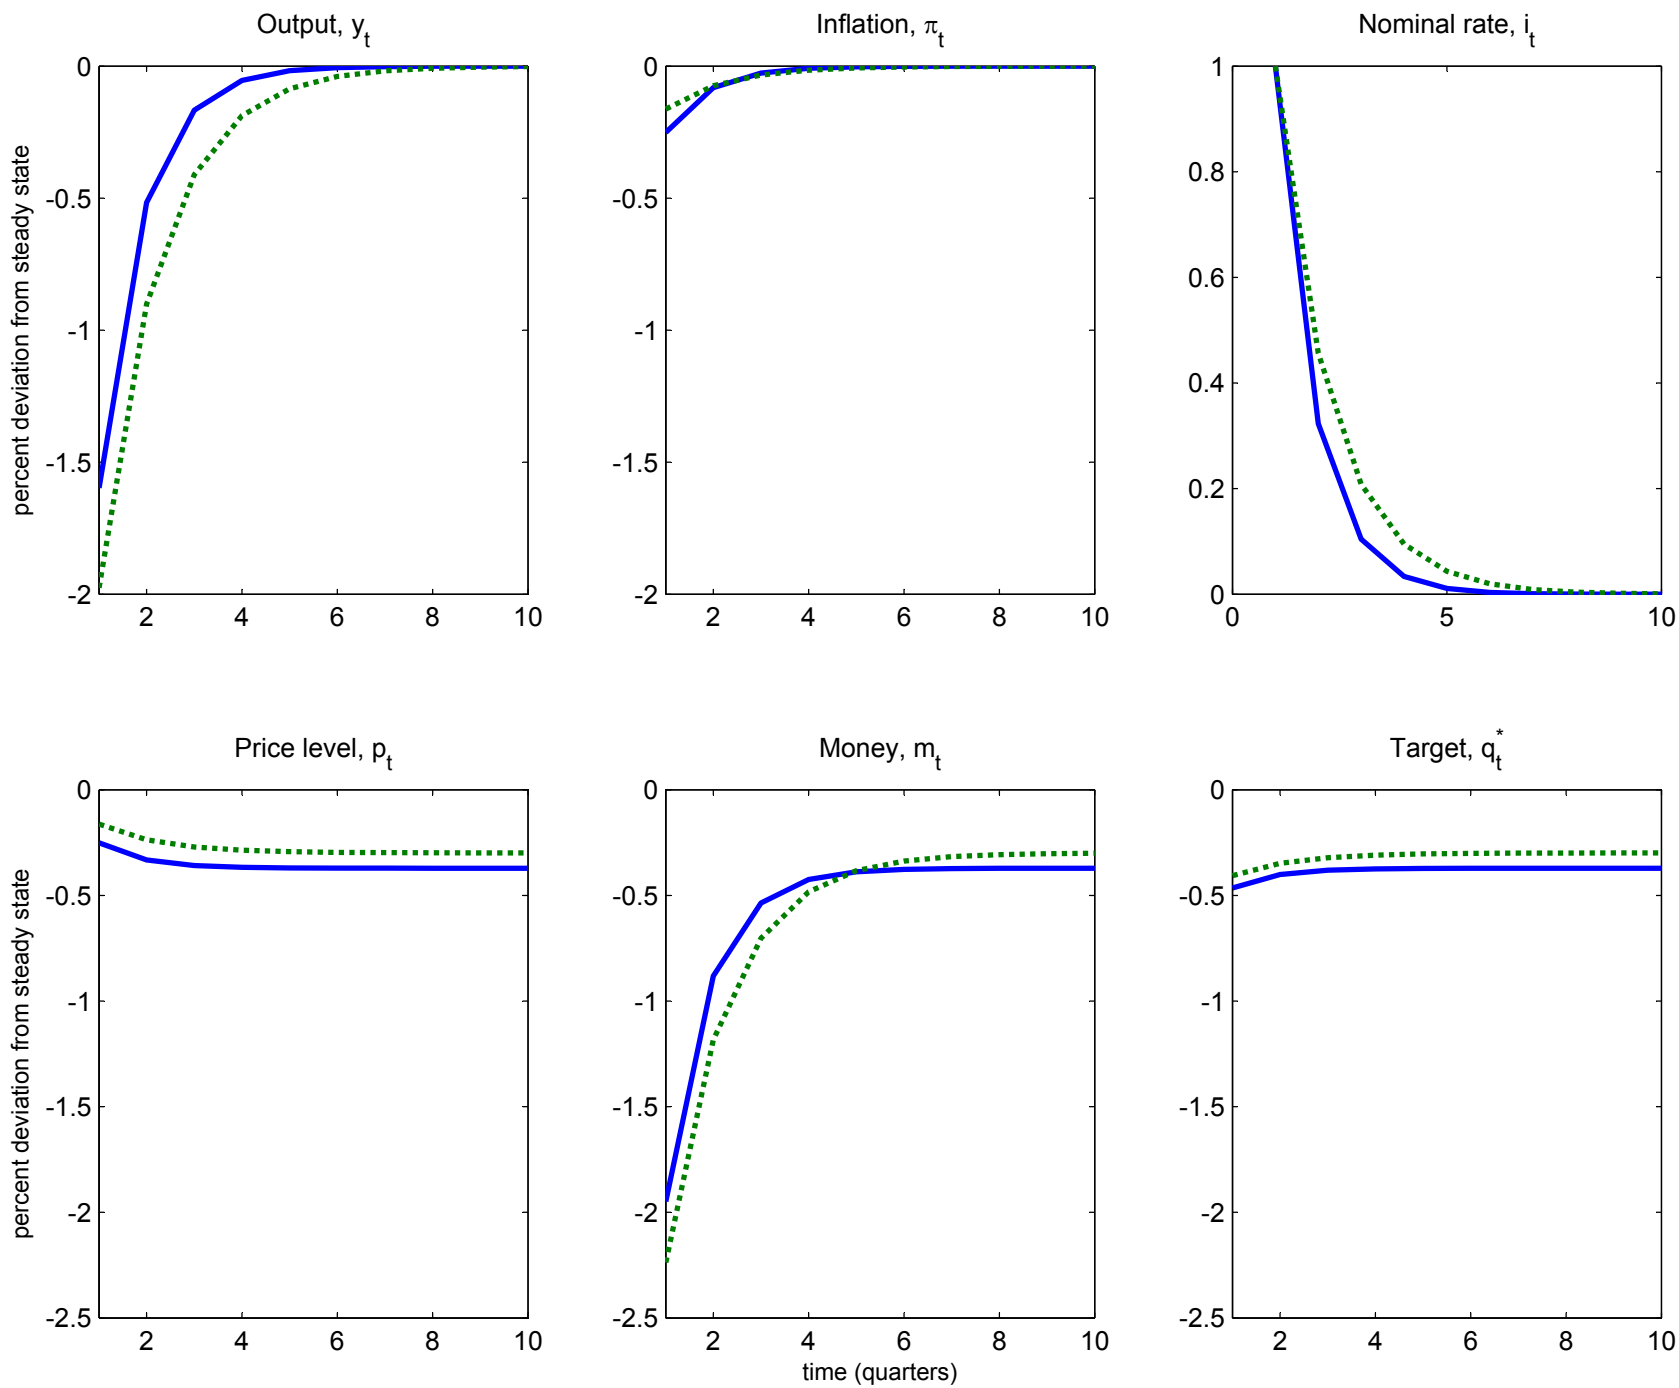


Figure 4: Response to an interest rate shock



Figure 5: The response of the price level to a 1% money shock with and without the decision cost




Figure 6: Comparing the response of the price level to a 1% money shock in the linear model to the response of firms exact profit functions and no decision costs

