

Menu Costs and Phillips Curves

Mikhail Golosov

The University of Minnesota and the
Federal Reserve Bank of Minneapolis

Robert E. Lucas, Jr.*

The University of Chicago and the
Federal Reserve Bank of Minneapolis

May, 2003

Abstract

This paper develops a model of a monetary economy in which individual firms are subject to idiosyncratic productivity shocks as well as general inflation. Sellers can change price only by incurring a real “menu cost.” We calibrate this cost and the variance and autocorrelation of the idiosyncratic shock using a new U.S. data set of individual prices due to Klenow and Kryvtsov. The prediction of the calibrated model for the effects of high inflation on the frequency of price changes accords well with the Israeli evidence obtained by Lach and Tsiddon. The model is also used to conduct numerical experiments on the model’s response to credible and incredible disinflations and other shocks.

*We have benefitted from discussions with V.V. Chari, Larry Jones, Boyan Jovanovic, Patrick Kehoe, Timothy Kehoe, Bob King, Pete Klenow, Oleksiy Kryvtsov, Ellen McGrattan, Thomas Sargent, Nancy Stokey, Julia Thomas, and Aleh Tsyvinski.

1. Introduction

As an explanation for some of the apparent real effects of inflations, menu costs—real resources required to change nominal prices—have an evident appeal. Such costs are by definition a source of nominal “rigidity” or “stickiness,” and lump-sum costs have the added attraction that action will be taken infrequently and only in response to changes that are viewed as permanent. Thus menu cost models are consistent with the fact that even large disinflations have small real effects if credibly carried out.¹ Finally, and not least, menu costs are really *there*: The fact that many individual goods prices remain fixed for weeks or months in the face of continuously changing demand and supply conditions testifies conclusively to the existence of a fixed cost of repricing.

Granting these qualitative observations, it would be good to know how large the real effects of price stickiness induced by menu costs might be. Phillips curves estimated from aggregate data provide one source of information, but such estimates are notoriously unstable and can be given many different interpretations. In this paper, we calibrate a menu cost model using a new data set on prices, based on Bureau of Labor Statistics surveys, assembled and described by Bils and Klenow (2002) and Klenow and Kryvtsov (2003). This estimation makes use of both cross-section and time series evidence on the prices of narrowly-defined individual goods, and summary statistics on the frequency of individual price changes. The calibrated model is then used to carry out several numerical experiments involving the economy’s responses to specified changes in the overall inflation rate.

A distinctive feature of the model we use is the assumption that each seller is subject to two shocks that affect its pricing decision: an aggregate shock, reflecting economy-wide inflation, and an idiosyncratic shock, reflecting good-specific changes in

¹Sargent (1986).

technology or preferences. In such a world, an observed price change may be triggered by either kind of shock, or a mix of the two, and the pricing decisions of individual sellers will not be synchronized. Observations on aggregates must be interpreted as averages over a more complex microeconomic reality. In the next section we define an equilibrium in a model with this feature, in which inflation is defined by changes in an economy-wide nominal wage rate that follows an exogenously given random walk, with constant drift and variance.

The definition we propose is a little too general to permit us to calculate or characterize equilibria, and the rest of the paper examines a series of special cases and approximations. In Sections 3 and 4, we specialize to the case where the inflation rate is constant and the idiosyncratic shock is mean-reverting. In this version of the model, suitably normalized prices follow a continuous time Markov process that has an invariant distribution. We derive the Bellman equation for individual firms in this context. We describe a discrete-time model that can provide an arbitrarily good approximation to the continuous-time model, and for which solutions are inexpensive to compute. We use the approximate model to compute the optimal pricing behavior of firms.

This model predicts the fraction of prices that are changed per month, the key parameter in the widely-applied model of Calvo (1982). We use the model to examine the sensitivity of this fraction to changes in inflationary behavior. Then we compare the predictions of the model, as estimated from data from low-inflation U.S. economy, to Lach and Tsiddon's (1992) evidence on the frequency of price changes during the Israeli inflation of 1978-1982. The predictions of the model are surprisingly accurate.

Section 5 then describes the results of some simulations of disinflations, conducted using the model of Section 3. We compare credible and non-credible disinflations. Section 6 calculates some impulse-response functions. Section 7 re-introduces a stochastic shock to the inflation rate, and proposes an approximation to the firm's

pricing policy. This approximation is then used to examine the behavior of Phillips curves, in the sense of correlations between inflation rates and levels of production and employment. Estimates of the fraction of the variability in these variables that can be accounted for by monetary shocks and menu costs are also provided.

The effects of inflation shocks on aggregate production and employment, as we estimate them, are small and transient. For a credible disinflation—we use an immediate reduction in the rate of wage inflation from 15 to one percent per quarter—there are essentially no real effects, consistent with Sargent’s (1986) findings. But even when the disinflation is carried out in a *non-credible* way, the resulting decline in production never exceeds one percent of GDP. We summarize these and other findings in more detail in a concluding section.

The model we describe in the next section builds on the original formulation of the pricing problem of a single firm Sheshinski and Weiss (1977), and on the long literature of other papers that apply (S, s) -type inventory theory to pricing problems.² It has proved difficult to situate these pricing models into equilibrium models, but several precedents have been influential and valuable. Calvo (1983) leave the fixed cost in the background, assuming instead that a firm gets an option to re-price in any period with a fixed probability that does not depend on its own state. By varying this probability, any desired degree of price “stickiness” can be built into a model. Caplin and Spulber (1987) considered equilibrium in an economy in which the money supply always grows, the firm with the lowest price at any instant re-prices, and to the top of the price distribution. They show that a rectangular distribution of firms by their price levels is stationary, and that the distribution of prices relative to the

²Sheshinski and Weiss (1977) analyzed the pricing decision of an individual seller facing a deterministic trend in the desired price level. Versions of this problem, many of them stochastic, have been studied by Sheshinski and Weiss (1983), Frenkel and Jovanovic (1980), Chang (1999), Mankiw (1985), and Stokey (2002).

money supply is stationary and independent of the path of money. This was a striking example of any economy with an extremely low re-pricing rate (in Calvo's terms) that would appear to an observer of aggregates to have no price stickiness at all. Further developments can be found in Caplin and Leahy (1991) and elsewhere.

Lach and Tsiddon (1992) looked at individual price distributions in Israel, finding them not to be rectangular, and the changes not to be synchronized, even for firms at the same initial price. They suggested that a successful model would need to have idiosyncratic shocks as well as economy-wide shocks. Dotsey, King, and Wolman (1999) propose a monetary equilibrium model in which the synchronization of price changes is broken by a transient, random shock to the menu cost itself: Firms that draw high cost have an incentive to postpone re-pricing. They explore a number of issues numerically using a log-linear approximation. Further developments are described in Willis (2000) and Burstein (2002).

The model developed and applied below shares both its two-shock character and its use of numerical methods with Dotsey, King, and Wolman (1999) and their successors. We use an idiosyncratic productivity shock rather than a shock to menu costs. We close the model in a somewhat different way. Our calibration is entirely new, as are the numerical questions we pose and the answers we offer.

2. A Model of Monetary Equilibrium

The theory that we calibrate and simulate in this paper is a Bellman equation for a single, price-setting firm that hires labor at a given nominal wage, produces a consumption good with a stochastically varying technology, and sets product price subject to a menu cost of re-pricing. We situate our model of a firm in a model of a monetary economy in search equilibrium so as to be able to relate its predictions to aggregative evidence. In this economy, there is a continuum of infinitely lived

households, each of which consumes a continuum of goods. A Spence-Dixit-Stiglitz³ utility function is used to aggregate across goods to form current period utility. Each household also supplies labor on a competitive labor market. Firms hire labor, used to produce the consumption good and to re-set nominal prices for the good, and sell goods to consumers. Each firm produces only one of the continuum of consumption goods.

The economy is subject to two kinds of shocks: a monetary shock, which we summarize in the implied money wage rate W_t , and a firm-specific productivity shock V_t . The log of the nominal wage, $w_t = \log(W_t)$, is assumed to follow a Brownian motion with drift parameter a and variance σ_w^2 ,

$$dw = adt + \sigma_w dZ_w, \quad (2.1)$$

where Z_w denotes a standard Brownian motion with zero drift and unit variance. We think of the behavior of the money wage as ultimately being determined by the behavior of the money supply, but this connection will not be modeled and the only explicit role of dollars in the theory is as a unit of account. In the absence of the real menu costs associated with changing prices, the evolution of W_t would have no effect on resource allocation.

There are also firm-specific productivity shocks V_t , which is assumed to be independent across firms. We assume that $v_t = \log(V_t)$ follows a mean-reverting Ornstein-Uhlenbeck process:

$$dv = -\eta v dt + \sigma_v dZ_v, \quad \eta > 0. \quad (2.2)$$

where Z_v is a standard Brownian motion with zero drift and unit variance, distributed independently of Z_w .

Equilibrium prices and quantities will be modeled as stochastic processes, defined in terms of an initial joint distribution $\phi_0(P, V)$ of firms by their inherited price P

³Spence (1976) and Dixit-Stiglitz (1977).

and productivity level V and by the evolution of the exogenous processes V_t and W_t . Specifically, each firm chooses a pricing strategy $\{P_t\}$, a stochastic processes that takes the form of a right-continuous step function whose date t value depends on the histories of its own productivity shocks $\{V_s\}_0^t$ and of the nominal wage shocks $\{W_s\}_0^t$.⁴ The strategy $\{P_t\}$ that a firm chooses will also depend on its inherited initial price P_0 . When we want to emphasize this dependence we write $\{P_t(P_0)\}$.

For given firm behavior $\{P_t(P_0)\}$, the initial joint distribution of (P_0, V_0) , the initial wage W_0 , and the probabilities implied by (2.1) and (2.2), induce a family of probability measures ϕ_t for the prices facing consumers at all dates t . For any Borel set $A \subset \mathbf{R}^2$,

$$\phi_t(A) = \int \Pr\{(P_t(P_0), V_t)(\omega) \in A \mid V_0\} \phi_0(dP_0, dV_0). \quad (2.3)$$

There is also a capital market, in which consumers trade contingent claims to the monetary unit. We adopt the convention that

$$E \left[\int_0^\infty q_t Y_t dt \right]$$

is the value at date 0 of a dollar earnings stream $\{Y_t\}_0^\infty$, also a stochastic process defined in terms of W_t . Thus q_t must be multiplied by the appropriate probabilities to obtain the Arrow-Debreu prices.

We next describe the decision problem of consumers in this environment. At each date t , each household buys from every seller, each of which is characterized by a pair (P, V) , distributed according to the measure $\phi_t(P, V)$. The household chooses a buying strategy $\{C_t(\cdot)\}$, where $C_t(P)$ is the number of units of the consumption good that it buys from a seller who charges price P after the shock history $\{W_s\}_0^t$ has occurred. It also chooses a labor supply strategy $\{\ell_t\}$, where ℓ_t is the units of labor

⁴That is, the process $\{P_t\}$ is adapted to the filtration generated by $\{V_t, W_t\}$. (See, for example, Stokey (2002), ch. 1.)

supplied after the shock history $\{W_s\}_0^t$ has occurred. The values $C_t(P)$ and ℓ_t also depend on the initial price distribution $\phi_0(P, V)$.

For any buying strategy $C_t(P)$, define ζ_t as the Spence-Dixit-Stiglitz consumption aggregate

$$\zeta_t = \left[\int C_t(P)^{1-1/\varepsilon} \phi_t(dP, dV) \right]^{\varepsilon/(\varepsilon-1)}. \quad (2.4)$$

The corresponding current period utility is assumed to be

$$\frac{1}{1-\gamma} \zeta_t^{1-\gamma}.$$

Preferences over time are

$$E \left[\int_0^\infty e^{-\rho t} \left(\frac{1}{1-\gamma} \zeta_t^{1-\gamma} - \alpha \ell_t \right) dt \right]. \quad (2.5)$$

The operator $E(\cdot)$ is defined by the shock processes (2.1) and (2.2).

We write the consumer's budget constraint as

$$E \left[\int_0^\infty q_t \left(\int PC_t(P) \phi_t(dP, dV) - W_t \ell_t - \Pi_t \right) dt \right] \leq 0, \quad (2.6)$$

where Π_t is profit income, obtained from the household's holdings of a fully diversified portfolio of claims on the individual firms. The household chooses buying and labor supply strategies $\{C_t(\cdot)\}$ and $\{\ell_t\}$ so as to maximize (2.5), subject to (2.6), taking $\{q_t\}$, $\{W_t\}$, and $\{\phi_t\}$ as given.

At each date, a firm faces consumer demand $C_t(\cdot)$, a nominal wage rate W_t , and a stochastically determined productivity parameter (goods per hour worked) V_t . The firm enters the period with a price level P carried over from the past. If it leaves price unchanged, its current profit level is

$$C_t(P)(P - W_t/V_t).$$

If it chooses any price $Q \neq P$, its current profit level is

$$C_t(Q)(Q - W_t/V_t) - kW_t,$$

where the parameter k is the hours of labor needed to change price, the real menu cost.

A pricing strategy $\{P_t\}$ for the firm takes the form of an increasing sequence of stopping times $\{T_i\}_{i=1}^{\infty}$, dependent on the histories $\{V_t, W_t\}$ of shocks, and an associated sequence $\{P^i\}_{i=1}^{\infty}$ of prices, such that

$$P_t = P^i \quad \text{for } T_i \leq t < T_{i+1}.$$

Any such strategy yields the expected return

$$\begin{aligned} & E\left[\int_0^{T_1} q_t C_t(P_0)[P_0 - W_t/V_t]dt\right. \\ & \left. + \sum_{i=1}^{\infty} \int_{T_i}^{T_{i+1}} q_t C_t(P^i)[P^i - W_t/V_t]dt - k \sum_{i=1}^{\infty} q_{T_i} W_{T_i}\right]. \end{aligned} \quad (2.7)$$

Firms choose the price strategy $\{P_t\}$ to maximize (2.7). For a given initial price distribution ϕ_0 , the pricing strategy $\{P_t\}$, adopted by all sellers, defines the distributions $\{\phi_t\}$. A process $\{\Upsilon_t\}$ with the interpretation that $\Upsilon_t dt$ is the number of firms that re-price during the time interval $(t, t + dt)$, is also defined.

The market clearing condition for this economy is

$$\ell_t = \int \frac{C_t(P)}{V} \phi_t(dP, dV) + k\Upsilon_t. \quad (2.8)$$

The equality of goods consumed and goods produced is incorporated in (2.7). There is no money market to clear because money does not appear as a good in the model: It is just a unit of account.

The consumer side of the model is familiar. We will use the first-order conditions for consumers to reformulate the firms' problems, and then turn to the study of the latter. The first-order conditions for the consumer include

$$e^{-\rho t} \zeta_t^{-\gamma} \left[\int C_t(P)^{1-1/\varepsilon} \phi_t(dP, dV) \right]^{1/(\varepsilon-1)} C_t(P)^{-1/\varepsilon} = \lambda q_t P, \quad (2.9)$$

where the multiplier λ does not depend on time, and

$$e^{-\rho t} \alpha = \lambda q_t W_t. \quad (2.10)$$

Eliminating the multiplier between (2.9) and (2.10) and simplifying using (2.4) yields

$$\zeta_t^{-\gamma+1/\varepsilon} C_t(P)^{-1/\varepsilon} = \alpha \frac{P}{W_t}.$$

Solving for $C_t(P)$, we obtain the demand function facing each firm:

$$C_t(P) = \zeta_t^{1-\varepsilon\gamma} \left(\frac{\alpha P}{W_t} \right)^{-\varepsilon}. \quad (2.11)$$

Applying the natural normalization $q_0 = 1$ to (2.10), we obtain

$$q_t = e^{-\rho t} \frac{W_0}{W_t}. \quad (2.12)$$

Inserting the solutions (2.11) and (2.12) from the consumer's problem into the firm's objective function (2.7) then yields

$$\begin{aligned} & W_0 E \left[\int_0^{T_1} e^{-\rho t} \zeta_t^{1-\varepsilon\gamma} \left(\frac{\alpha P_0}{W_t} \right)^{-\varepsilon} \left[\frac{P_0}{W_t} - \frac{1}{V_t} \right] dt \right. \\ & \left. + \sum_{i=1}^{\infty} \int_{T_i}^{T_{i+1}} e^{-\rho t} \zeta_t^{1-\varepsilon\gamma} \left(\frac{\alpha P^i}{W_t} \right)^{-\varepsilon} \left[\frac{P^i}{W_t} - \frac{1}{V_t} \right] dt - k \sum_{i=1}^{\infty} e^{-\rho T_i} \right]. \end{aligned} \quad (2.13)$$

We sum up in the following

Definition. Given an initial wage W_0 and a joint price-shock distribution ϕ_0 , an *equilibrium* is a pricing strategy $\{P_t\}$ for firms, processes $\{\zeta_t\}$ and $\{\phi_t\}$ for the consumption aggregate and the price-shock distributions, and processes $\{\Upsilon_t\}$ and $\{\ell_t\}$ for the re-pricing frequenting and labor supply, such that (i) given $\{\zeta_t\}, \{P_t\}$ maximizes (2.13), (ii) given $\{P_t\}, \{\phi_t\}$ satisfies (2.3), (iii) $\{\zeta_t\}$ and $\{\phi_t\}$ satisfy (2.4), and such that $\{\Upsilon_t\}$ and $\{\ell_t\}$ satisfy (2.8), given $\{\phi_t\}$ and with $C_t(P)$ given by (2.11).

In our view it is a limitation of this definition of equilibrium that inflation is defined by the assumed exogenous behavior of nominal wages, and not indirectly by the behavior of the money supply. We will not remedy this deficiency here, but a brief digression on the issues involved in doing so may help in interpreting our results. Suppose we introduce money M_t into the theory, treating it as a given stochastic process, and motivate its use by adding real balances M_t/P_t as a second argument in consumers' utility function. Retaining additivity, the the instantaneous utility flow could then be written

$$\frac{1}{1-\gamma}\zeta_t^{1-\gamma} + u\left(\frac{M_t}{P_t}\right) - \alpha\ell_t, \quad (2.14)$$

say. An additional equilibrium condition would then be the equality of the marginal rate of substitution between real balances and leisure to their relative price, or

$$u'\left(\frac{M_t}{P_t}\right)\frac{1}{\alpha P_t} = \frac{R_t}{W_t}, \quad (2.15)$$

where R_t is the instantaneous nominal interest rate. The other equilibrium conditions are not affected. The nominal interest rate in turn is essentially minus the rate of change of the log of the price q_t , given in (2.12), and is itself a stochastic process. This familiar Fisherian connection between the *level* of real balances demanded and the *rate of change* of the nominal wage means that, in general, nominal wages and prices will not be proportional to the money supply along an equilibrium path.

But proportionality is a possibility. Suppose, for example, that u is the log function, so (2.15) becomes

$$\frac{W_t}{\alpha M_t} = R_t.$$

Suppose further that the growth rate of money is constant at a rate a . If the nominal wage grows at the same rate, the instantaneous interest rate will be constant at $\rho + a$ and the wage path

$$W_t = \alpha(\rho + a)M_t$$

will be an equilibrium.

The results we report for constant wage growth in Sections 3 and 4 could be interpreted in this way. But (2.14) is not a particularly attractive way to model the services of cash holdings, since it does not relate the stock of money held to spending flows, and in any case the restriction to constant inflation rates rules out cases of substantive interest that we want to discuss in Sections 5-7. Our view is that the main theme of this paper is the relation of inflation and changes in the rate of inflation to real output and employment—the Phillips curve—and not the analysis of the motives for holding money and the monetary sources of inflation.

Most previous work has simplified by eliminating or avoiding the idiosyncratic shocks, $\{V_t\}$ in our set-up, and focusing on aggregate inflation shocks only. We will initially go in the opposite direction, treating the special case in which the variance σ_w^2 of the wage process is zero, so that the drift parameter a is simply the constant rate of wage inflation. In this situation, we will seek an invariant joint distribution ϕ for prices and idiosyncratic shocks. In the next two sections, we formulate, calibrate, and study a Bellman equation for this case of a stationary equilibrium with constant inflation.

3. Stationary Equilibrium with Constant Inflation

Let the inflation rate be constant at a , so that

$$W_t = W_0 e^{at}. \tag{3.1}$$

Denote the value of the objective (2.13) of a firm in a stationary equilibrium that has an inherited nominal price P and exogenous shocks with current values (V, W) by $W\Phi(P, V, W)$. In a stationary equilibrium, the consumption aggregate ζ_t will be constant at a value ζ (say) and Φ will satisfy the Bellman equation

$$\Phi(P, V, W) = \max_T E \left[\int_0^T e^{-\rho t} \zeta^{1-\varepsilon\gamma} \left(\frac{\alpha P_0}{W_t} \right)^{-\varepsilon} \left[\frac{P_0}{W_t} - \frac{1}{V_t} \right] dt \right]$$

$$+e^{-\rho T} \cdot \max_Q [\Phi(Q, V_T, W_T) - k]. \quad (3.2)$$

It is convenient to restate (3.2) in terms of the logs: $(P, V, W) = (e^p, e^v, e^w)$ and to define a new state variable x_t by

$$x_t = p_t - w_t. \quad (3.3)$$

Then between price changes the motion of x_t is described by

$$dx = -adt \quad (3.4)$$

and the motion of v follows (2.2). Then we seek a solution to (3.2) of the form $\Phi(P, V, W) = \Phi(e^p, e^v, e^w) = \varphi(p - w, v) = \varphi(x, v)$, where the function φ satisfies

$$\varphi(x, v) = \max_T E \left[\int_0^T e^{-\rho t} \Pi(x_t, v_t) dt + e^{-\rho T} \max_{x'} [\varphi(x', v(T)) - k] \right], \quad (3.5)$$

where

$$\Pi(x, v) = \zeta^{1-\varepsilon\gamma} (\alpha e^x)^{-\varepsilon} [e^x - e^{-v}]. \quad (3.6)$$

Note that the pricing policy x that would maximize the profit function Π if there were no menu cost would be

$$x^* = \log\left(\frac{\varepsilon}{\varepsilon - 1}\right) - v. \quad (3.7)$$

We will study the problem (3.5) using a discrete-time and state approximation to (2.2)—a Markov chain—following Kushner and Dupuis (2001). The approximate process $v(t)$ is assumed to take values on the finite grid $V = \{-\bar{v}, \dots, -2h, -h, 0, h, 2h, \dots, \bar{v}\}$, so V is the state space of the Markov chain. For a given grid size h , we specify a time interval Δt and a transition function π on V ,

$$\pi(v', v) = \Pr\{v(t + \Delta t) = v' \mid v(t) = v\},$$

in terms of the defining parameters (η, σ_v) of the original continuous-time process in such a way that in the limit as $h \rightarrow 0$ the conditional, local means and variances of the original and approximate processes match. See the Appendix for details.

In the calculations reported below, we fix the grid size h and the set V on which the shocks values v lie. Analogous to the definition of the set V , we define a set

$$X = \{-\bar{x}, \dots, -2h, -h, 0, h, 2h, \dots, \bar{x}\}$$

of possible x values, so that states (x, v) will be elements of $S = X \times V$. Then we approximate the continuous problem (3.5) with the discrete problem

$$\begin{aligned} \psi(x, v) = \max\{ & \Pi(x, v)\Delta t + e^{-r\Delta t} \sum_{v'} \pi(v' | v)\psi(xe^{a\Delta t}, v'), \\ & \max_{x'}[\Pi(x', v)\Delta t + e^{-r\Delta t} \sum_{v'} \pi(v' | v)\psi(x'e^{a\Delta t}, v')] - k\}. \end{aligned} \quad (3.8)$$

Under the assumptions we have imposed there is a unique function ψ on S that satisfies (3.8).

To see the dynamic behavior implied by (3.8), it is useful to define some auxiliary functions. Let

$$\Omega(v) = \max_{x'}[\Pi(x', v)\Delta t + e^{-r\Delta t} \sum_{v'} \pi(v' | v)\psi(x'e^{a\Delta t}, v')],$$

so $\Omega(v)$ is the value of discounted profit if a one-time, costless repricing can occur in state v . Let $g(v)$ be the price that attains the maximum. Then define the “inaction region” $D(v)$ as

$$D(v) = \{x \in X : \psi(x, v) > \Omega(v) - k\}$$

(so that $D(v)$ is the set of $x = p - w$ values at which it does not pay to re-price when the productivity shock is v). In our application, this region will be an open interval:

$$D(v) = (L(v), U(v)).$$

The policy function for (3.8) is thus the function $f : S \rightarrow X$ defined by

$$\begin{aligned} f(x, v) &= x \text{ if } x \in D(v) \\ f(x, v) &= g(v) \text{ if } x \notin D(v) \end{aligned}$$

The (log) price in state (x, v) is $p = w + f(x, v)$.

Figures 1 and 2 plot the three functions g, U and L , for quarterly inflation rates of one and 15 percent. (These and other figures in this section are based on the baseline parameter values given in Table 1 in Section 4.) These functions are decreasing: high productivity shocks imply price decreases. Note that the inaction intervals $D(v)$ are wider for low v values: Getting prices “right” is more important when productivity shocks—and hence quantities sold—are high. Note also the widening of the inaction intervals as inflation increases. With high inflation, prices are changed often just to keep up with general inflation, and pricing rules do not need to be as sensitive to idiosyncratic productivity shocks. Higher inflation also induces a shift in the target prices toward the upper boundary of the inaction region.

INSERT FIGURES 1 AND 2

Figures 3 and 4 illustrate simulated sample paths for the ideal price given in (3.7), reflecting changes in relative productivity only. Figure 3 also plots the optimal price path for the cost parameter $k = 0.004$ and the constant inflation rate $a = .01$. Figure 4 is the same in all respects, except that the cost parameter is decreased to $k = 0.0004$. One can see the way the step function for price tracks the ideal price, and the fact that as the fixed cost decreases prices are changed more often and the tracking gets better.

INSERT FIGURES 3, 4, AND 3a

Figure 3a, taken from Chevalier *et al.* (2000), shows the time series of actual prices for Triscuits, based on scanner data from a Chicago supermarket chain. It is instructive to compare this time series to those in our Figures 3 and 4. On Figure 3a, temporary sales are evident in the many times the Triscuit price is reduced for a short time and returned to exactly the former price soon thereafter. Such patterns are of course common to many price series. To obtain a good match between theory

and data, then, sales must either be removed from the data or added to the model. As discussed below, we took the first course, made easier by the fact that the BLS flags observations that they regard as sale prices. For us, then, the data have sale observations removed. Do this with your eye on Figure 3a, and then compare it to Figures 3 and 4.

The joint process (x, v) , where x is controlled in the way we have just described, has a unique ergodic set E , given by

$$E = \{(x, v) \in S : x \in \bar{D}(v)\},$$

(where the overbar means closure) and there is a unique invariant distribution $\mu(x, v)$ on E . It is impossible to pass the boundaries of the inaction region because whenever the uncontrolled process would do so, price is revised to prevent this from occurring. Figure 5 plots the implied density of the pricing error $x_t - x^*$ for the two quarterly inflation rates one percent and 15 percent that were used in Figures 1 and 2.

INSERT FIGURE 5

4. Data, Calibration, and a Test

Our basic model lacks many features that a business cycle model needs—it has no capital and no aggregate shocks—but we drew on that literature for the preference parameters ρ , γ , α , and ε . We used the annual discount rate $\rho = .04$, the risk aversion parameter $\gamma = 2$, the elasticity of substitution parameter $\varepsilon = 7$, and the disutility of labor $\alpha = 0.9$. These ρ and γ values are conventional. The value of ε is related to the degree of monopoly power firms have: With $\varepsilon = 7$, firm profits are about 15 percent of GDP. The elasticity of substitution implies that a firm’s mark-up—defined as the percent by which price exceeds marginal cost—is about 16 percent. Estimates of mark-ups typically fall in the 10-20 percent range, implying values of ε in the

6-10 percent range.⁵ Our results are not sensitive to changes in ε within that range. We interpreted our linear labor disutility as indivisible labor with lotteries, following Hansen (1985). The value $\alpha = 0.9$ implies an average labor level of 0.85.

For the menu cost parameter k , the drift parameter a , and the two parameters σ_v^2 and η that characterize the idiosyncratic productivity shocks, we used a new data set on individual prices due to Klenow and Kryvtsov (2003). This price data set is based on the BLS survey, and contains about 80,000 time series of individual price quotes in 88 geographical locations. The series are either monthly or bimonthly, depending on the location, for the years 1988-1997. The individual price quotes comprise 123 narrowly defined goods categories. The data set also provides the positive weights ω_i , $\sum_i \omega_i = 1$, that are used to form the consumer price index from the individual prices (the ‘‘CPI weights’’). We used the prices and weights for the New York metropolitan area to calibrate the parameters $(a, \sigma_m^2, \eta, \sigma_v^2)$ and the fixed cost k of the model described in the last two sections.

For calibrating the model under the assumption of a deterministic trend, we imagine that the Klenow-Kryvtsov prices are generated by our model in the following way. Prices for goods $i = 1, \dots, N$ are collected at the time intervals Δt . Denote these prices by $\{p_i(t)\}$, $i = 1, \dots, N$. For each good i , let

$$p_i(t) = x_i(t) + w(t),$$

where for each t , $w(t)$ grows at the constant rate a , $v_i(t)$ is generated by (2.2), and equilibrium prices are determined as described in the last two sections.⁶

We identify sample averages from the Klenow-Kryvtsov data with means taken with respect to the invariant distribution $\mu(x, v)$. If a seller begins a period in state (x, v) , he immediately moves (re-prices) to $(f(x, v), v)$, so if F is any function on S

⁵See, for example, Rotemberg and Woodford (1995) and Basu and Fernald (1997).

⁶The BLS identifies some recorded prices as sale prices. We have replaced such observations with the recorded price for the same good in the period before the sale. See Figure 3a.

its observed average value will be

$$\sum_{x,v} F(f(x, v), v) \mu(x, v). \quad (4.1)$$

We first estimate the parameters (a, σ_w^2) of the stochastic inflation process defined in (2.1), and then turn to the calibration of (η, σ_v^2, k) .

The weighted average price of goods (CPI) at time t is

$$\begin{aligned} \bar{p}(t) &= \sum_i \omega_i p_i(t). \\ &= \sum_i \omega_i [x_i(t) + w(t)] \\ &\cong \sum_{x,v} x \mu(x, v) + w(t) \end{aligned}$$

where the second equality is a prediction of the theory, and the third replaces a sample average with the population mean. Then since the first term on the right is constant with respect to time, the model implies that

$$\bar{p}(t + \Delta t) - \bar{p}(t) = w(t + \Delta t) - w(t). \quad (4.2)$$

Equation (4.2) suggests using the average and the variance of the differences of the observed CPI series $\{\bar{p}(t)\}$ as estimates of $a\Delta t$ and $\sigma_w^2 \Delta t$ respectively (though for now we set $\sigma_w^2 = 0$)

The numbers in the first column of the first two rows of the table below report these estimates. For the deterministic inflation version of the model, we set the mean inflation rate equal to one percent and the standard deviation equal to zero, as indicated in column 2. (Both means and standard deviations are expressed as percentages.) The last three columns of the table indicate variations in the model's predictions under different choices for the parameters (η, σ_v^2, k) . The first two rows are not affected by these variations.

TABLE 1 : CALIBRATED PARAMETER VALUES

Moment	Data	Model	$\eta = 0.15$	$\sigma_v^2 = .0075$	$k = .003$
Quarterly inflation rate	.009	.01	.01	.01	.01
S.d. of inflation	.0374	0	0	0	0
Frequency of change	.219	.231	.227	.267	.261
Mean price increase	.095	.075	.078	.084	.071
S.d. of new prices	.094	.1	.133	.124	.101

To calibrate the three parameters (η, σ_v^2, k) , we calculate three additional sample moments that intuition suggests will convey information. The results are given in the last three rows of the table. The first is the frequency of price change: the average over all months in the data of the fraction of prices that were changed in that month. As shown on the table, this fraction is .219. Second, we calculated the average log price increase over all prices that increased from any date to the next date: .095 in the data. Finally, from among all prices that were increased, we calculated the standard deviation of the new prices. To do that we calculated log deviations from the average, $z_i(t) = p_i(t) - \bar{p}(t)$, for each good i and then computed the standard deviation of $z_i(t)$ over time for each good i . Then we averaged over goods i . This yielded the number .094.

For any values of $(a, 0, \eta, \sigma_v^2, k)$ we can calculate the corresponding moments predicted by the theory, under the assumption that the probability distribution of (x, v) is the invariant distribution $\mu(x, v)$ and that prices are given by $h(x, v)$. The fraction of firms that do not change price in the time interval Δt is just the probability that the invariant distribution assigns to the set $D(v)$:

$$\sum_{v, x \in D(v)} \mu(x, v).$$

The fraction of firms that change per month is then

$$1 - \left(\sum_{v,x \in D(v)} \mu(x, v) \right)^{1/(3\Delta t)}.$$

This value, .231, shown in the appropriate row of column 2 of the table, is calculated with $(a, \sigma_w^2, \eta, \sigma_v^2, k)$ set equal to (.01, 0, .25, .005, .004). Columns 3-5 of the table indicate how the calculated moments change as (η, σ_v^2, k) are changed one at a time from these benchmark values. That is, column 3 shows the computed statistics when the parameter vector $(a, \sigma_w^2, \eta, \sigma_v^2, k) = (.01, 0, .25, .005, .004)$ is replaced by $(.01, 0, .15, .005, .004)$. Thus the table shows that the frequency of price changes is insensitive to changes in the rate of mean reversion in the idiosyncratic shock, that it increases with the variance of these shocks, and that it decreases with increases in the menu cost.

There are many studies that try to estimate or calibrate menu costs for particular products. For example, Levy *et al.* (1997) estimate that the cost of changing prices in supermarkets is about 0.7 percent of firms' revenue. In our baseline model with $k = 0.004$, menu costs are about 0.25 percent of revenues and 1.9 percent of profits. The labor required to adjust prices is equal to 0.5 percent of overall employment.

We solved the model, calibrated as just described, for quarterly inflation rates a ranging from 0 to 20 percent, calculated the invariant distribution μ in each case, and calculated the fraction of firms that change price each month in this stationary equilibrium. For comparison, we carried out the same calculations for the deterministic Sheshinski-Weiss case where the variance of the idiosyncratic shocks is set equal to zero. The results are shown in Figure 6. Also shown are two empirical observations: the inflation-repricing pair (0.9,21.9) from the Klenow-Kryvtsov data, and the pair (15,41) taken from the Lach-Tsiddon study of the 1978-82 Israeli inflation. The first pair lies exactly on the upper curve, reflecting the fact that we used the Klenow-Kryvtsov data to calibrate our model. The model so calibrated fits exactly the Israeli

inflation, too. This is a genuine test of the theory, carried out with evidence that is far out of sample.

Figure 6 also confirms the necessity of including idiosyncratic shocks if the model is to fit the evidence from low inflation economies. As inflation rates are reduced, a lot of “price stickiness” remains in the data. Of course, the evidence does not bear on our interpretation of the idiosyncratic shocks as productivity differences, as opposed to shifts in preferences, responses to inventory build-ups, or other factors.

INSERT FIGURE 6

The quarterly re-pricing rates ranging from 22 to 41 percent shown on the figure may be compared to Calvo’s (1983) assumption that the re-pricing probability is constant, independent of the firm’s situation. When this fraction is endogenously determined, as in our theory, it varies with the inflation rate, but we were surprised at how small a change in the re-pricing probability is needed to fit the large differences in inflation rates shown on the figure. The most important differences are in the two theory’s implications for a cross-section of firms, as shown in our Figures 1 and 2. In our theory, re-pricing is more likely when sales are large and price is more important. In the Calvo theory, the firm’s situation does not effect the re-pricing probability.

5. Disinflations: Credible and Non-credible

We used the model with deterministic wage growth to study the effects of two kinds of disinflations: reductions in the given rate of wage growth. Each experiment considers an economy that is initially in the stationary equilibrium described in Section 3, with a constant rate of inflation. In both experiments, there is a one-time, unanticipated, and permanent reduction in the growth rate of nominal wages from a (say) to $a' < a$. The first experiment involves a fully credible, rational expectations equilibrium disinflation. The second describes a non-credible disinflation, in which

firms continue to expect the old inflation rate to be resumed.

A disinflation of either kind will take the economy out of its stationary equilibrium. Let $\zeta(a)$ denote the constant value of the consumption aggregate defined in (2.4) under the original policy. We first construct an equilibrium response in which a new stationary distribution is attained, with the constant value $\zeta(a')$ corresponding to the new inflation rate, and in which the time path $\{\zeta_t\}$, $\zeta_0 = \zeta(a)$ and $\zeta_t \rightarrow \zeta(a')$, induced by the shock is perfectly foreseen by firms.

It is easiest to describe this construction in terms of the discrete approximation (3.8). We set a limit n on the number of transition periods, and begin with an assumed finite sequence $\zeta^n = (\zeta_1, \zeta_2, \dots, \zeta_n)$ of values of the consumption aggregate. Then we define the sequence $\{\psi_i(x, v, \zeta^n)\}_{i=1}^n$ of value functions recursively by

$$\psi_n(x, v, \zeta^n) = \psi(x, v), \quad (5.1)$$

where $\psi(x, v)$ is the solution to (3.8) at the new stationary equilibrium with ζ_t constant at $\zeta(a')$, and

$$\begin{aligned} \psi_i(x, v, \zeta^n) = \max\{ & \Pi(x, v, \zeta_i)\Delta t + e^{-r\Delta t} \sum_{v'} \pi(v' | v) \psi_{i+1}(x e^{a\Delta t}, v', \zeta^n), \\ & \max_{x'} [\Pi(x', v, \zeta_i)\Delta t + e^{-r\Delta t} \sum_{v'} \pi(v' | v) \psi_{i+1}(x' e^{a\Delta t}, v', \zeta^n)] - k \} \end{aligned} \quad (5.2)$$

for $i = 1, 2, \dots, n - 1$. Let $\{f_i(x, v, \zeta^n)\}_{i=1}^n$ be the sequence of policy functions corresponding to the value functions $\{\psi_i(x, v, \zeta^n)\}_{i=1}^n$ so defined. For given behavior ζ^n of the consumption aggregate, these functions can be calculated by the usual backward induction.

The pricing behavior $\{f_i(x, v, \zeta^n)\}_{i=1}^n$ in turn implies a sequence $\{\phi_i(P, V, \zeta^n)\}_{i=1}^n$ of joint distributions of prices and productivity shocks, taking the original invariant distribution ϕ as the initial condition. Individual firm sales are given by (2.11), and then new values of the consumption aggregate by (2.4):

$$(\Gamma\zeta)_i = \left[\int \zeta_i^{(1-\varepsilon\gamma)(1-1/\varepsilon)} \left(\frac{\alpha P}{W_i} \right)^{1-\varepsilon} \phi_i(dP, dV) \right]^{\varepsilon/(\varepsilon-1)}. \quad (5.3)$$

The construction described in equations (5.1)-(5.3) thus defines a function Γ taking an n -vector ζ into $\Gamma\zeta$.

In our calculations we used the policy functions from the stationary equilibrium with wage growth equal to a' to generate $\zeta(a')$, and then iterated using Γ until a fixed point was found. This procedure requires a choice of the length n of the transition period. We chose n large enough that the last few terms of the fixed point ζ^n were close to the value $\zeta(a')$ associated with the new stationary equilibrium. The resulting description of the transition is thus a rational expectations equilibrium in which agents have perfect foresight about the evolution of aggregate variables.

Figure 7a plots the response of two variables—real GDP and employment—following a one-time permanent and fully credible decrease in the quarterly growth rate of nominal wages from 15 to 1 percent. Real GDP y_t is defined simply as the sum of production of the different goods,

$$y_t = \int C_t(P)\phi_t(dP, dV).$$

Employment is the equilibrium path of the variable ℓ_t , the labor used to produce goods *and* the labor used in re-pricing. Both variables are normalized so that the initial levels are one. One can see that neither variable declines by as much as one half of one percent, and this very modest response is completed within a quarter of the initial shock.

Figure 7b plots the response of the inflation rate and the re-pricing rate to this same credible disinflation. In constructing the figure, the price *level* is the fixed-weight index P_t defined by

$$P_t y_t = \int P C_t(P)\phi_t(dP, dV),$$

and the inflation rate is its quarterly rate of growth. The initial re-pricing rate is about twice the long-run, post-disinflation level (see Figure 7b) and except for an initial uptick, the figure shows a small adjustment that is over in less than a month.

The inflation rate follows an erratic path from 15 to 1. The pricing bounds are narrower with the new, lower rate of wage growth (see Figures 1, 2 and 5) so as soon as the new rate is announced, many firms find themselves with a price level that is outside the bounds. These firms immediately re-price, some with increases and others with decreases. These changes almost cancel, and produce the very low initial inflation rate in the first instant. After that, inflation rises to above 3 percent, and gradually declines to its new equilibrium level of one. This adjustment, too, is over in a quarter.

INSERT FIGURES 7a AND 7b

Sargent's (1986) analysis contrasts the effects of a disinflation that is credible with the effects of one that is not. There is only one way for a disinflation to be credible, but there are many non-credible ways for it to occur. Figures 8a and 8b represent one possibility. They describe the results of the following experiment. Consider the economy in the stationary equilibrium with a constant 15 percent quarterly inflation rate. Exactly as in the previous example, there is a permanent decrease in the growth rate of nominal wages from 15 percent to one percent. This decrease is completely non-credible, in the sense that all the agents in the economy believe that wages will continue to grow at 15 percent and treat one percent growth as a zero probability event. Over time they repeatedly observe the one percent growth in W but continue to believe that the next instant the growth rate will revert to its old 15 percent level. It would be more realistic to assume that agents would use Bayes rule to decide on the probability of W reverting to the high inflationary level, but we will show that even with this extreme assumption real responses are very small.

The computational strategy for this non-credible case is more complicated than the one we used to calculate the equilibrium under a credible disinflation, because beliefs about the consumption aggregator do not coincide with the time path that actually occurs. Let the original value of the consumption aggregate be $\zeta(a)$, as in

the credible case, and let the distribution of the firms be the stationary distribution $\phi_1(x, v)$ corresponding to a . Let the inflation rate be a' for one period. At this point, everyone believes that inflation will revert to a immediately and remain there, so that the original distribution $\phi(x, v)$ and the original value of the consumption aggregate be $\zeta(a)$ will eventually be restored. We assume that this is expected to occur in n periods, and define an operator Γ with fixed point ζ^n exactly as in the perfect foresight equilibrium described above. This procedure gives everyone's action in the first period, corresponding to the actual value a' and their *belief* that wage growth will immediately revert to its original value a .

But this belief is not accurate. In the second period, wage growth is a' again and everyone believes it will revert to a in period 3. But the situation differs from period 1 because a new distribution $\phi_2(x, v) \neq \phi_1(x, v)$ has been established. We need to recalculate a new time path ζ^n for the consumption aggregate, which will give us everyone's behavior in period 2. So we proceed period by period, calculating a perfect foresight path from each date on, using only the first action to describe what firms actually do, and discarding the rest of the sequence.

INSERT FIGURES 8a AND 8b

Comparing Figures 8 to the credible disinflation shown in Figures 7, we can see that as expected, the real effects are larger in the non-credible case: The ultimate decline in GDP is 0.9 percent rather than 0.3 percent. Even so, the effect is hardly supportive of the conventional view that the welfare costs of non-credible disinflations are high: The GDP response remains less than one percent, and the employment response is only slightly larger. The effects shown on Figures 8 are very persistent, but of course this is only because we have modeled firms' beliefs as being unrealistically persistent.

We think the reason for the surprisingly unimportant role played by credibility

in this model is something like the following. Since firms can choose the frequency of their price changes, they increase it when inflation is high. Each time they adjust their prices, they expect to adjust them again on average within two months. When disinflation is not credible, the new prices are higher than the ones firms choose with a credible disinflation, and this reduces output. On the other hand, firms wait longer than in a credible disinflation to adjust their prices. That means that in the non-credible disinflation, more firms have prices that are suboptimally low, and they produce higher outputs. By comparing Figures 7b and 8b, one can see that the re-pricing rate is lower under non-credible than credible disinflation. These effects work in opposite directions: The first effect leads to a decline in the output while the second one cause it to increase. The resulting overall effect on output is negative but small.

6. An Impulse-Response Function

In the next section we will describe approximate behavior of the economy in a two-shock equilibrium, in which firms are subject both to idiosyncratic and aggregate shocks. It is instructive first, however, to study the dynamics of the model with deterministic wage behavior that is subjected to a single, one-time shock to the nominal wage level. The thought experiment we conduct in this section subjects an economy in the stationary equilibrium nominal wage rate growth is constant at a to an unanticipated jump in the nominal wages from W to $(1+h)W$, after that the wage growth resumes its original level a .

We calculate the response to this shock using the iterative method for calculating rational expectations equilibrium that we applied to the credible disinflation case in Section 5. Figures 9a and 9b plot the impulse-response functions calculated in this way when a equals one percent per quarter and $h = 0.0125$.

INSERT FIGURE 9a AND 9b

First, note that the initial response in output is less than the size of the monetary shock. Since aggregate output is

$$\begin{aligned}
y_t &= \int C_t(P)\phi_t(dP, dV) \\
&= \alpha^{-1/\gamma} \left(\int \left(\frac{P}{W}\right)^{1-\varepsilon} \phi_t(dP, dV) \right)^{(1-\gamma\varepsilon)/(\gamma\varepsilon-\gamma)} \times \int \left(\frac{P}{W}\right)^{-\varepsilon} \phi_t(dP, dV)
\end{aligned}$$

the increase of W to $(1+h)W$ can increase total output by at most $(1+h)^{(1-\gamma\varepsilon)/\gamma} \times (1+h)^\varepsilon = (1+h)^{1/\gamma}$.

The increase in W leads to a temporary increase in the number of the firms changing their prices. This effect is over very quickly, occurring right after the jump in wages, after which the frequency of price changes reverts to its steady state level. The effect on the real output lasts longer, but it also declines to zero by the middle of the first quarter. The reason for such a fast decline is the fact that even if firms do not react to the aggregate shock alone, many of them re-price due to idiosyncratic shocks. Once a firm decides to re-price for any reason, it will take the higher level of nominal wages into account when choosing a new price.

This result can be contrasted with the previous menu cost literature. In the absence of idiosyncratic shocks, the log-linear approximation of our firms' problem would be equivalent to the set-up of Caplin and Spulber (1987). Their result that aggregate shocks are completely neutral would then hold in ours: In a stationary equilibrium the distribution of the relative prices of the firms would be uniform, and a $\kappa\%$ increase in W would cause $\kappa\%$ of the firms to adjust their prices. The resulting distribution of the relative prices is the same as the stationary distribution, and so total output would remain unchanged. The presence of idiosyncratic shocks introduces more complicated distributions of the relative prices, so in our case the shock to nominal wages leads to a real responses.

Caplin and Leahy (1991) study a slightly different economy where the level of nominal wages follow Brownian motion. Their result suggest that response of the

output to the nominal shocks may be very persistent: It is not worth paying the menu cost to adjust the nominal price if the aggregate shock is small. We show that with the reasonably calibrated idiosyncratic shocks, this effect will be very transient. Even though the firms do not react to the aggregate shocks *per se*, they respond to idiosyncratic shocks and take the higher level of nominal wages into account when they choose a new price.

7. Approximations to a Two-Shock Equilibrium

In Section 2 we defined an equilibrium for an economy subject to two independent shock process: the nominal wage process (2.1) and the idiosyncratic productivity process (2.2). All of the results reported in Sections 3-6, though, are based on special cases where the variance σ_w^2 of the wage shock is set equal to zero. The only inflation shocks we have studied are one-time impulses or permanent changes in the drift parameter.

The only barrier to formulating and studying a Bellman equation for the firm's pricing problem in the case where both aggregate and idiosyncratic shocks operate is the appearance in the firm's objective function (2.13) of the consumption aggregate ζ_t . We have analyzed situations where this aggregate is constant or varies deterministically, but if $\sigma_w^2 > 0$, ζ_t is a stochastic process that depends in a complicated way on the evolution of the distributions $\{\phi_t\}$ of all firms' prices. We do not know how to calculate equilibria in these circumstances. Instead, we will study a Bellman equation based on (2.13) but with ζ_t replaced with a suitable constant. The hope is to obtain an approximation to equilibrium pricing policies that is useful when σ_w^2 is not too large.

The Bellman equation suited for this purpose is the variant of (3.5) described in

the Appendix:

$$\varphi(x, v) = \max_T E \left[\int_0^T e^{-\rho t} \Pi(x_t, v_t) dt + e^{-\rho T} \max_{x'} [\varphi(x', v(T)) - k] \right], \quad (7.1)$$

where (see (3.6))

$$\Pi(x, v) = \zeta^{1-\varepsilon\gamma} \alpha^{-\varepsilon} e^{-\varepsilon x} [e^x - e^{-v}]. \quad (7.2)$$

The processes (x, v) are assumed to follow

$$dx = -adt + \sigma_w dZ_w, \quad (7.3)$$

and

$$dv = -\eta v dt + \sigma_v dZ_v. \quad (7.4)$$

The corresponding discrete approximation, analogous to (3.8), is

$$\begin{aligned} \psi(x, v) = \max \{ & \Pi(x, v) \Delta t + e^{-r\Delta t} \sum_{x', v'} \pi(x', v' | x, v) \psi(x', v'), \\ & \max_{\xi} [\Pi(\xi, v) \Delta t + e^{-r\Delta t} \sum_{x', v'} \pi(x', v' | \xi, v) \psi(x', v')] - k \}. \end{aligned} \quad (7.5)$$

The transition probabilities $\pi(x', v' | x, v)$ for the finite-state Markov chain that we use to approximate the x and v processes are described in the Appendix.

As in Section 3, we let $f(x, v)$ be the policy function for (7.5). Then the transition function for the controlled (x, v) process will be $\pi(x', v' | f(x, v), v)$. We assume that this process has a unique invariant distribution $\mu(x, v)$. The value and policy functions and the invariant distribution μ will all depend on the value chosen for ζ . We then obtain the equilibrium (or pseudo-equilibrium) value of ζ as the solution to

$$\zeta = \left(\zeta^{1-\varepsilon\gamma} \alpha^{-\varepsilon} \right) \left[\int e^{(1-\varepsilon)x} \mu(dx, dv; \zeta) \right]^{\varepsilon/(\varepsilon-1)}, \quad (7.6)$$

using (2.4), (2.10), and (3.3). We calculated this solution iteratively. The policy function computed in this way is the policy of a firm that correctly observes the mean

level of ζ_t but ignores all the fluctuations about this level. We propose this function as an approximation to the true behavior of the firms in a two-shock equilibrium.

To get some idea about the likely accuracy of this approximation, we recalculated the impulse-response functions displayed in Figures 9a and 9b in Section 6 (which display a rational expectations equilibrium in which ζ_t varies over time) using the constant ζ approximation just described. Figure 10 shows the results for real GDP and employment. Evidently, the approximation works very well for the effects of a one-time shock. We take this as an indication that it will be highly accurate for stochastic shocks as well.

INSERT FIGURE 10

We conduct several thought-experiments using this approximation. First, we will study the effect of the volatility of inflation on the volatility of the real output. We simulate 40 quarters of data. Since the model is in continuous time (actually, about 40 discrete periods per quarter) and the economic data come in discrete intervals, we aggregate the output of the model into quarterly values by taking the means of the relevant variables over the quarter. Thus, the level ν_T^Q at quarter T of any function of time $\nu(t)$ is defined as

$$\nu_T^Q = \int_T^{T+1} \nu(t) dt.$$

For our simulations we chose $\sigma_w^2 = 0.0014$ which corresponds to the .037 standard deviation of quarterly inflation in the Kryvtsov and Klenow data set. The standard deviation of the log level of output is equal to .0017 in our simulation. The standard deviation of actual U.S. quarterly consumption for the same period (1989 to 1998) around linear trend is equal to .015. Thus monetary fluctuations in this model can account for only about 10 percent of the observed fluctuations in output. This estimate is consistent with estimates from other sources.⁷

⁷For example, see Lucas's (2003) survey.

In the second experiment we regress the log level of real output on the log difference of the nominal wages, using the simulated series generated by the model:

$$\ln(y_t^Q) = \alpha + \beta[\ln(w_t^Q) - \ln(w_{t-1}^Q)].$$

In this regression, we obtain the estimate $\beta = 0.04$ with the standard error 0.007. Thus an increase in nominal wage rates leads to an increase in real output, as in standard Phillips curve regressions, but the effect is very small. This conclusion is not sensitive to different specifications of the parameters (a, σ_w) .

8. Conclusions

We have constructed a model of a monetary economy in which re-pricing of goods is subject to a menu cost, and studied the behavior of this economy numerically. The model is distinguished from its many predecessors by the presence of idiosyncratic shocks along with general inflation. We used a data set on individual U.S. prices recently compiled by Kryvtsov and Klenow to calibrate the menu cost and the variance and autocorrelation of the idiosyncratic shocks. We conducted several experiments with the model.

A key prediction of any menu cost model is that the fraction of firms that re-price in a given time interval will increase with increases in the inflation rate. We simulated our model at inflation rates varying from zero to 20 percent per quarter. The results, shown on Figure 6, trace out a curve that passes through the inflation rate-re-pricing rate pair observed by Kryvtsov and Klenow in the low-inflation U.S. economy of the 1990s. The curve also fits the inflation rate-re-pricing rate observed by Lach and Tsiddon from the high inflation Israeli economy of 1978-82. We note that a model without idiosyncratic shocks could not fit both observations.

The model also predicts that credible disinflations will have smaller real effects than incredible ones: Menu costs make firms reluctant to re-price in response to

transient shocks. We simulated a credible (perfect foresight) reduction in the rate of wage inflation from 15 percent to one percent, and that re-ran the simulation of a disinflation that was non-credible in a specified way. The real effects were large in the second case by a factor of three, but in both simulations this large and sudden disinflation resulted in output and employment declines of one percent or less.

These results all refer to a special case in which inflation is deterministic. We also solved an approximation to a more realistic two-shock model. With a realistic inflation variance, this model can account for perhaps one-tenth of observed variance of U.S. real consumption about trend. A Phillips curve estimated from data generated by the model implies that a one percentage point reduction in inflation slope will depress production by 0.04 of one percent.

In summary, the model we propose fits microeconomic evidence on U.S. pricing behavior well, and does a surprisingly good job of accounting for behavior differences between countries with very different inflation rates. It does not appear to be consistent with large real effects of monetary instability.

Appendix

The construction of approximating Markov chains for the one shock model of Sections 3-5 and the two-shock model of Section 6 is based on Kushner and Dupuit (2001). This appendix provides the details, based on the two-shock model of Section 6. For the most part, the specialization to the one-shock case is obvious.

In the calculations described below, we fix the grid size h and define the state space $S = X \times V$ as we did in Section 3. Then we approximate the continuous problem (7.1) with the discrete problem (7.5), repeated as

$$\begin{aligned} \psi(x, v) &= \max\{\Pi(x, v)\Delta t + e^{-r\Delta t} \sum_{x', v'} \pi(x', v' | x, v)\psi(x', v'), \\ &\max_{\xi}[\Pi(\xi, v)\Delta t + e^{-r\Delta t} \sum_{x', v'} \pi(x', v' | \xi, v)\psi(x', v')] - k\}, \end{aligned} \quad (\text{A.1})$$

where π is a transition function defined on $S \times S$ that we define in a moment. The time interval Δt is related to the grid size and other parameters by

$$\Delta t = \frac{h^2}{Q}, \quad (\text{A.2})$$

where

$$Q = \sigma_w^2 + ah + \sigma_v^2 + \eta\bar{v}h. \quad (\text{A.3})$$

We assume that in a given time interval Δt at most one of the variables x and v changes.⁸ Provided that neither x nor v is at its upper or lower bound, we assume that if x changes, it moves either to $x + h$ or to $x - h$; if v changes, it moves either to $v + h$ or to $v - h$. The final possibility is that neither of the variables changes and the

⁸This means that the Markov chains approximating $v(t)$ and $x(t)$ will not be independent for $h > 0$, even though the continuous processes are. But independence will hold in the limit, as $h \rightarrow 0$.

state remains at (x, v) . The probability of all other transitions is zero. Away from the boundaries of S , the five non-zero transition probabilities will then be defined by

$$\pi(x+h, v, x, v) = \frac{\sigma_w^2/2}{Q}, \quad (\text{A.4})$$

$$\pi(x-h, v, x, v) = \frac{\sigma_w^2/2 + ah}{Q}, \quad (\text{A.5})$$

$$\pi(x, v+h, x, v) = \frac{\sigma_v^2/2}{Q} \quad (\text{if } v \geq 0), \quad (\text{A.6})$$

$$\pi(x, v-h, x, v) = \frac{\sigma_v^2/2 + \eta vh}{Q} \quad (\text{if } v \geq 0), \quad (\text{A.7})$$

and

$$\pi(x, v, x, v) = 1 - \frac{\sigma_w^2 + \sigma_v^2 + \eta vh + ah}{Q} \quad (\text{if } v \geq 0). \quad (\text{A.8})$$

(The $v(t)$ process is symmetric about zero, so the adaptations of (A.6) and (A.7) for the case $v < 0$ are obvious.) Transitions at the boundaries are handled by assuming that if, for example, x hits its upper bound \bar{x} then x goes one step down to $\bar{x} - h$ with probability $\pi(\bar{x} - h, v, \bar{x}, v)$ and stays at \bar{x} with probability $\pi(\bar{x} + h, v, \bar{x}, v) + \pi(\bar{x}, v, \bar{x}, v)$ as given by the formulas (A.4) and (A.8). It is evident that the five probabilities (A.4)-(A.8) add to one, and that the probabilities (A.4)-(A.7) are positive. That (A.8) is non-negative follows from the fact that $|v| \leq \bar{v}$.

The first and second moments of the Markov chain we have just defined, conditional on the current state (x, v) (assumed not be a boundary point of S) are readily calculated from (A.4)-(A.8). They are

$$\begin{aligned} \frac{E\{x(t + \Delta t) \mid x(t) = x, v(t) = v\} - x}{\Delta t} &= -a, \\ \frac{E\{v(t + \Delta t) \mid x, v\} - v}{\Delta t} &= -\eta v \quad (\text{if } v \geq 0), \\ \frac{Var\{x(t + \Delta t) \mid x, v\}}{\Delta t} &= \sigma_m^2 + ah - a\Delta t, \\ \frac{Var\{v(t + \Delta t) \mid x, v\}}{\Delta t} &= \sigma_v^2 + h|\eta v| - (\eta v)^2 \Delta t, \end{aligned}$$

$$\frac{\text{Cov}\{x(t + \Delta t), v(t + \Delta t) \mid x, v\}}{\Delta t} = -\eta va \Delta t \quad (\text{if } v \geq 0).$$

From (A.2) and (A.3)

$$\frac{\Delta t}{h} = \frac{h}{\sigma_m^2 + ah + \sigma_v^2 + \eta \bar{v} h} \rightarrow 0 \quad \text{as } h \rightarrow 0.$$

This is the sense in which the conditional, local moments of the approximating chain approximate the conditional, local moments of the continuous time $(x(t), v(t))$ process defined by (7.3) and (7.4). See Kushner and Dupuis (2001), Chapter 9, for a proof that this approximation converges in distribution to the continuous time diffusion process when $h \rightarrow 0$.

REFERENCES

- [1] Basu, Susanto, and John G. Fernald. 1997. "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy*, 105: 249-283.
- [2] Bils, Mark and Peter J. Klenow. 2002. "Some Evidence on the Importance of Sticky Prices". NBER Working Paper #9069.
- [3] Burstein, Ariel. 2002. "Inflation and Output Dynamics with State Dependent Pricing Decisions." Working paper.
- [4] Calvo, Guillermo A. 1983. "Staggered Prices in a Utility Maximizing Framework." *Journal of Monetary Economics*, 12: 383-398.
- [5] Caplin, Andrew S., and Daniel F. Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics*, 102: 703-726.
- [6] Caplin, Andrew S., and John Leahy. 1991. "State Dependent Pricing and the Dynamics of Money and Output." *Quarterly Journal of Economics*, 106: 683-708.

- [7] Chang, Fwu-Ranq. 1999. "Homogeneity and the Transactions Demand for Money." *Journal of Money, Credit, and Banking*, 31: 720-30
- [8] Chari, V.V., Patrick J. Kehoe, and Ellen R. McGrattan. 2000. "Sticky Price Models of the Business Cycle." *Econometrica*, 68: 1151-1180.
- [9] Chevalier, Judith A., Anil K. Kashyap, and Peter E. Rossi. 2000. "Why don't prices rise during periods of peak demand? Evidence from scanner data." NBER Working Paper #7981.
- [10] Dixit, Avinash, and Joseph E. Stiglitz. 1977. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review*, 67: 297-308.
- [11] Dotsey, Michael, Robert G. King, and Alexander L. Wolman. 1999. "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output." *Quarterly Journal of Economics*, 114: 655-690.
- [12] Frenkel, Jacob A., and Boyan Jovanovic. 1980. "On Transactions and Precautionary Demand for Money." *Quarterly Journal of Economics*, 95: 25-43.
- [13] Klenow, Peter J. and Oleksiy Kryvtsov. 2003. "State-Dependent Pricing and the Variance of U.S. Inflation, 1988-2002. Federal Reserve Bank of Minneapolis working paper.
- [14] Kushner, Harold J., and Paul Dupuis. 2001. *Numerical Methods for Stochastic Control Problems In Continuous Time*. Second Edition. New York: Springer-Verlag.
- [15] Lach, Saul, and Daniel Tsiddon. 1992. "The Behavior of Prices and Inflation: An Empirical Analysis of Disaggregated Price Data." *Journal of Political Economy*, 100: 349-389.

- [16] Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable. 1997. "The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains." *Quarterly Journal of Economics*, 113: 791-825.
- [17] Lucas, Robert E., Jr. 2003. "Macroeconomic Priorities." *American Economic Review*, 93: 1-14.
- [18] Mankiw, N. Gregory. 1985. "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly." *Quarterly Journal of Economics*, 100: 529-538.
- [19] Rotemberg, Julio J., and Michael Woodford. 1995. "Dynamic General Equilibrium Models with Imperfect Competition." In Thomas Cooley, ed., *Frontiers of Business Cycle Research*. Princeton: Princeton University Press.
- [20] Sargent, Thomas J. 1986. *Rational Expectations and Inflation*. New York: Harper & Row.
- [21] Sheshinski, Eytan, and Yoram Weiss. 1977. "Inflation and Costs of Price Adjustment." *Review of Economic Studies*, 54: 287-303.
- [22] Sheshinski, Eytan, and Yoram Weiss. 1983. "Optimum Pricing Policy under Stochastic Inflation." *Review of Economic Studies*, 60: 513-529.
- [23] Spence, Michael E. 1976. "Product Selection, Fixed Costs, and Monopolistic Competition." *Review of Economic Studies*, 43: 217-235.
- [24] Stokey, Nancy L. 2002. *Brownian Models in Economics*. Manuscript.
- [25] Willis, Jonathan L. 2000. "General Equilibrium of a Monetary Model with State-Dependent Pricing." Working paper.

FIGURE 1: PRICING BOUNDS FOR 1% QUARTERLY INFLATION




FIGURE 2: PRICING BOUNDS FOR 15% QUARTERLY INFLATION




FIGURE 3: SAMPLE PATH OF ACTUAL AND IDEAL PRICE LEVEL $k=0.004$





FIGURE 3A: PRICE OF TRISCUIT 9.5 oz IN
DOMINICK'S FINER FOODS SUPERMARKET IN CHICAGO

Figure 1



Source: Chevalier, Kashyap and Rossi (2000)

FIGURE 4: SAMPLE PATH OF ACTUAL AND IDEAL PRICE LEVEL $k=0.0004$




FIGURE 5: DENSITIES OF DEVIATIONS FROM TARGET PRICE




FIGURE 6: FRACTION OF PRICES CHANGED EACH QUARTER

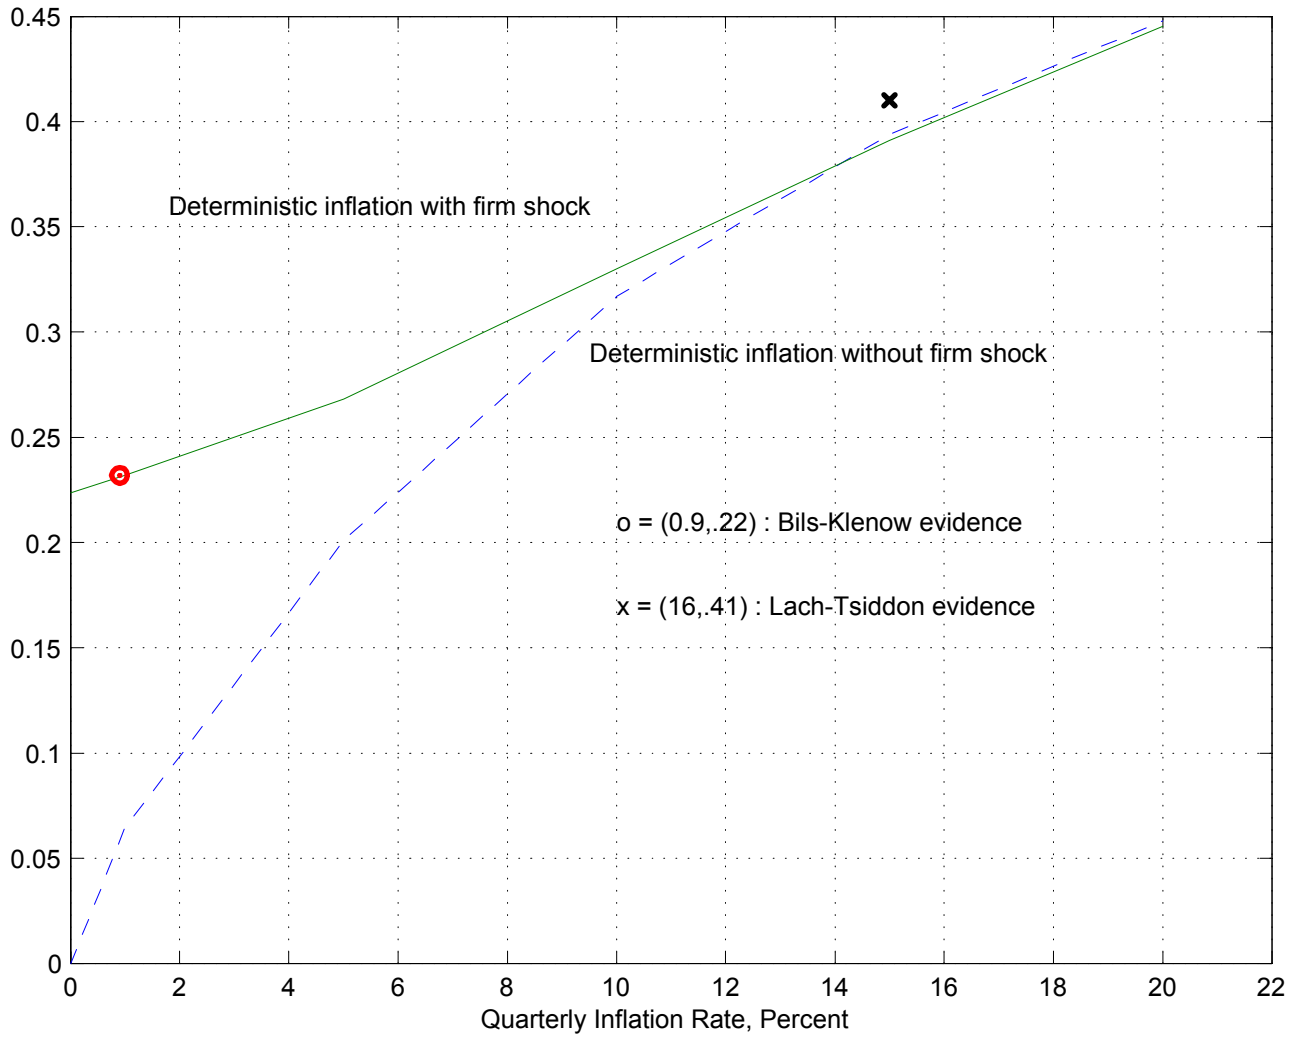


FIGURE 7a: RESPONSES TO A CREDIBLE DISINFLATION




FIGURE 7b: RESPONSES TO A CREDIBLE DISINFLATION




FIGURE 8a: RESPONSES TO A NON-CREDIBLE DISINFLATION




FIGURE 8b: RESPONSES TO A NON-CREDIBLE DISINFLATION




FIGURE 9a: RESPONSES TO A TRANSIENT WAGE INCREASE




FIGURE 9b: RESPONSES TO A TRANSIENT WAGE INCREASE




FIGURE 10: APPROXIMATE AND EXACT IMPULSE RESPONSE FUNCTIONS

