Real Rigidities and Nominal Price Changes[*]

Peter J. Klenow[†]and Jonathan L. Willis[‡]

March 2006

Abstract

A large literature seeks to provide microfoundations of price setting for macro models. A challenge has been to develop a model in which monetary policy shocks have the highly persistent effects on real variables estimated by many studies. Nominal price stickiness has proved helpful but not sufficient without some form of "real rigidity" or "strategic complementarity." We embed a model with a real rigidity a la Kimball (1995), wherein consumers flee from relatively expensive products but do not flock to inexpensive ones. We estimate key model parameters using micro data from the U.S. CPI, which exhibit sizable movements in relative prices of substitute products. When we impose a significant degree of real rigidity, fitting the micro price facts requires very large idiosyncratic shocks and implies large movements in micro quantities.

1 Introduction

Many studies estimate that monetary policy shocks have persistent effects on real output – effects lasting well beyond a year. For a sampling of estimates, each using a different identification strategy, see Christiano, Eichenbaum and Evans (1999), Romer and Romer (2003), and Bernanke, Boivin and Eliasz (2004).

In terms of theoretical microfoundations, one way of obtaining real effects of nominal shocks is, of course, nominal price rigidity. In recent quantitative treatments, however, the real effects of nominal price stickiness do not last much longer than the average duration of a price. Chari, Kehoe and McGrattan (2000) examine time-dependent models in the spirit of Taylor (1980). Golosov and Lucas (2003) characterize a state-dependent model, i.e., a model with fixed "menu costs" of changing prices and endogenous timing of price changes. Dotsey, King and Wolman (1999) investigate a hybrid of Calvo's (1983) time-dependent model and a conventional state-dependent model in that menu cost shocks influence the timing of re-pricing.

The recent micro empirical literature, meanwhile, finds that nominal prices typically change at least once per year. Bils and Klenow (2004) and Klenow and Kryvtsov (2005) report that U.S. consumer prices change every six months or so, on average. Dhyne et al. (2005), surveying a spate of recent studies, conclude that Euro Area prices typically change around once per year. Similarly, Taylor (1999) summarized the earlier evidence as saying prices change once a year on average.

Putting the micro evidence together with the quantitative theory, nominal rigidity,

2

by itself, appears unable to generate the persistent non-neutrality seen in the aggregate data. This failure has rekindled interest in combining nominal rigidities with "real rigidities", i.e., ingredients that makes firms reluctant to change their relative prices. Ball and Romer (1990) emphasized the need for such real rigidities on top of nominal rigidities to generate realistic persistence. These real rigidities, sometimes also called "strategic complementarities," can be on the factor supply side or on the goods demand side. Examples on the factor supply side include the real wage rigidities modeled recently by Blanchard and Gali (2005), and the firm-specific inputs suggested by Rotemberg (1996) and pursued by Woodford (2003), Altig, Christiano, Eichenbaum and Linde (2005), and Gertler and Leahy (2005). Alternatively, real rigidities could be on the preference side, as proposed by Kimball (1995) and used recently by Eichenbaum and Fisher (2004) and Dotsey and King (2005).

In this paper we focus on the Kimball-style real rigidity because it is capable of generating an arbitrary amount of persistence. Under Kimball's preferences, the elasticity of substitution between a given variety and others is decreasing in the relative quantity consumed of the variety. Thus sellers face a price elasticity of demand that is increasing in their good's relative price. In contrast to the Dixit-Stiglitz world of a constant elasticity and a constant desired markup of price over marginal cost, in Kimball's world the desired markup is decreasing in the relative price. When a re-pricing firm faces a higher marginal cost, say due to higher wages in the wake of monetary stimulus, the firm will temper its price increase because of the endogenous drop in its desired markup. The

lack of coordination is critical in this story, as it means a re-pricing seller will be raising its relative price. Each round of re-pricing is more tentative under Kimball's preferences, so that it takes longer for a monetary shock to fully pass through to the average price level.

As Dotsey and King (2005) and Basu (2005) discuss, Kimball's specification creates a smoothed version of a "kink" in the demand curve facing a given firm. Consumers flee from individual items with high relative prices, but do not flock to individual items with low relative prices. The result is that profits decline more steeply around a relative price of one. This is what creates "rigidity" in the relative price a firm wants.

We investigate the compatibility of the Kimball real rigidity with patterns of nominal and relative price changes in the micro data collected by the U.S. Bureau of Labor Statistics (BLS) for the Consumer Price Index (CPI). In this data, nominal prices changes are much larger than needed to keep up with overall inflation, as stressed by Golosov and Lucas (2004). Given little synchronization, these large changes in nominal prices translate into big movements in relative prices, suggesting firms face important idiosyncratic shocks to their marginal cost and/or desired markup.

Embedding Kimball's real rigidity in an industry equilibrium model, we assess how large the idiosyncratic shocks must be in order to rationalize the observed changes in relative prices. Given the degree of real rigidity suggested by Kimball (1995) and used by Eichenbaum and Fisher (2004), we find that the model requires very large idiosyncratic shocks (on the order of 35% per month to item-specific productivity), with concomitant

4

changes in item-specific quantities (including prices that entirely eclipse demand in about 20% of simulated months).

The rest of the paper is organized as follows. In section 2 we write down an industry equilibrium model that combines Kimball preferences with firm pricing decisions in the face of fixed menu costs. In section 3 we briefly describe the CPI microdata. In section 4 we estimate the model's parameters under different assumed levels of Kimball's superelasticity (the elasticity of the elasticity of demand). In section 5 we conclude.

2 Model

The model is a variant of the standard monopolistic competition model, and it describes optimal pricing behavior within a particular sector of the economy.[1] To explore the role of real rigidities, we adapt the standard model to include a flexible variety aggregator a la Kimball (1995).

2.1 Consumers

A representative agent consumes goods and provides labor for production. The economy has S sectors with each sector containing n_s producers. The representative agent chooses consumption of goods across all sectors and all producers $\{\{C_{si}\}_{i=1}^{n_s}\}_{s=1}^{S}$. The agent derives utility from consumption of sectoral composite goods which are created by the costless aggregation of goods within each sector.

---

[1]See Blanchard and Fischer (1994).

Each sectoral composite good C_s is created using the following Kimball preferences:

$$\frac{1}{n_s}\sum_{i=1}^{n_s}\Upsilon\left(\frac{n_s C_{si}}{C_s}\right) = 1, \tag{1}$$

where $\Upsilon(1) = 1$, $\Upsilon' > 0$, and $\Upsilon'' < 0$. The Kimball formulation features an elasticity of substitution decreasing in x, the relative quantity consumed of the item. CES preferences, which are used in the standard model of monopolistic competition, are nested within this specification.[2] More generally, C_s may be defined only implicitly by (1).

The representative agent chooses consumption and labor (L) to maximize utility

$$U = \max_{\left\{\{C_{si}\}_{i=1}^{n_s}\right\}_{s=1}^{S}}\frac{1}{1-\sigma}\left(\prod_{s=1}^{S}\left(\frac{C_s}{\alpha_s}\right)^{\alpha_s}\right)^{1-\sigma} + \frac{\nu(1-L)}{1-\sigma}, \tag{2}$$

subject to equation (1) and the budget constraint

$$\sum_{s=1}^{S}\sum_{i=1}^{n_s}P_{si}C_{si} = \tilde{w}L + \Pi, \tag{3}$$

where α_s is the Cobb-Douglas preference parameter for the composite good of sector s, $\tilde{w}$ is the nominal wage rate, and Π are the agent's profits from producers. We assume that the sectoral shares sum to 1, $\sum_{s=1}^{S}\alpha_s = 1$. The assumption of costless aggregation of composite goods from each sector provides an additional constraint relating the cost

---

[2]The Kimball specification reduces to CES preferences if $\Upsilon(x) = x^{\frac{\theta-1}{\theta}}$, where θ is the elasticity of substitution between items.

of the sectoral composite to the cost of individual goods within that sector:

$$P_s C_s = \sum_{i=1}^{n_s} P_{si} C_{si},$$

(4)

where P_s is the price of the sectoral good.

Since there is no saving in this economy, all income will be spent on purchases of goods. Based on first order conditions for consumption, the representative agent will spend a constant share of nominal income on goods from sector s:

$$\alpha_s = \frac{P_s C_s}{PY}.$$

(5)

PY is the nominal value of all goods produced in the economy, which also satisfies $PY = \tilde{w}L + \Pi$.

We use the first order conditions to derive the following relative demand for good C_{sj}:

$$\frac{n_s C_{sj}}{C_s} = \Upsilon'^{-1} \left(\frac{P_{sj}}{P_s} \sum_{i=1}^{n_s} \frac{C_{si}}{C_s} \Upsilon' \left(\frac{n_s C_{si}}{C_s} \right) \right).$$

(6)

To simplify notation, we define the following:

$$\varphi(x) \equiv \Upsilon'^{-1}(x)$$

(7)

$$D_s \equiv \sum_{i=1}^{n_s} \frac{C_{si}}{C_s} \Upsilon' \left(\frac{n_s C_{si}}{C_s} \right).$$

(8)

7

Using these definitions, the relative demand function can now be expressed as

$$\frac{n_s C_{sj}}{C_s} = \varphi\left(\frac{P_{sj}}{P_s} D_s\right).$$

(9)

2.2 Producers

Each firm in sector s produces a differentiated good and is monopolistically competitive. Producers are assumed to meet all demand, implying that $Y_{si} = C_{si}$. Given the demand function for their goods, they set their price to maximize profits. To implement a price change, firms must pay a labor cost of $\tilde{\phi}$.

Contemporaneous profits, excluding the implementation cost, are

$$\tilde{\Pi}_{si} = P_{si} Y_{si} - \tilde{w} L_{si}.$$

(10)

The producer faces the demand function given by (9) and the following production function:

$$Y_{si} = Z_s A_{si} L_{si}^{\eta},$$

(11)

where Z_s is a sectoral productivity index, A_{si} is an idiosyncratic (to the firm) productivity index, and η parameterizes the returns to scale of production.

We normalize firm profits by the portion of (smoothly-growing) nominal GDP contributed by the average firm in sector s ($\alpha_s PY/n_s$) and substitute in the demand function

to arrive at

$$\Pi_{si} \;=\; \frac{P_{si}}{P_s}\varphi\left(\frac{P_{si}}{P_s}D_s\right) - \frac{n_s w}{\alpha_s}L_{si}, \tag{12}$$

where $w = \frac{\tilde{w}}{PY}$. Note that consumer utility maximization implies that α_s equals the nominal output share for sector s, $\alpha_s = \frac{P_s Y_s}{PY}$.

The amount of labor, L_{si}, required to meet demand is derived by substituting (9) into (11):

$$L_{si} = \frac{1}{n_s}A_{si}^{\frac{-1}{\eta}}\left(\frac{n_s^{\eta-1}Y_s}{Z_s}\right)^{\frac{1}{\eta}}\left(\varphi\left(\frac{P_{si}}{P_s}D_s\right)\right)^{\frac{1}{\eta}}. \tag{13}$$

Contemporaneous profits can then be expressed as

$$\Pi_{si} \;=\; \frac{P_{si}}{P_s}\varphi\left(\frac{P_{si}}{P_s}D_s\right) - wA_{si}^{\frac{-1}{\eta}}\zeta_s\left(\varphi\left(\frac{P_{si}}{P_s}D_s\right)\right)^{\frac{1}{\eta}}, \tag{14}$$

where $\zeta_s \equiv \frac{1}{\alpha_s}\left(\frac{n_s^{\eta-1}Y_s}{Z_s}\right)^{\frac{1}{\eta}}$ is what we will call the "sectoral output-productivity ratio."

2.3 Menu Costs

Firms choosing to adjust their price in a given period will be faced with an adjustment cost. To implement a price change, a firm in sector s must hire $\tilde{\phi}_s$ units of labor at the going wage $\tilde{w}$. Expressed relative to nominal output per firm in sector s, this adjustment

cost is $\phi_s \equiv \frac{n_s \tilde{w}}{\alpha_s PY} \tilde{\phi}_s$. Normalized profits, net of implementation costs, are therefore

$$\Pi_{si} = \frac{P_{si}}{P_s} \varphi \left(\frac{P_{si}}{P_s} D_s \right) - w A_{si}^{\frac{-1}{\eta}} \zeta_s \left(\varphi \left(\frac{P_{si}}{P_s} D_s \right) \right)^{\frac{1}{\eta}} - \phi_s. \qquad (15)$$

Note that these normalized profits will be stationary because they involve only relative prices and quantities, the idiosyncratic productivity process will be stationary, and the menu cost is expressed relative to average firm revenue in the sector.

2.4 Dynamic Optimization

Given the implementation cost of a price change, the firm solves a dynamic optimization problem to maximize profits. In each period the firm decides whether or not to adjust its price. If it decides to adjust, it pays an implementation cost and resets its price. If it does not adjust, its nominal price remains fixed, and its relative price, $p_{si} = \frac{P_{si}}{P_s}$, decreases at the rate of sectoral inflation. As noted, we assume that the nominal output of the economy, PY, is growing at a constant rate. Sectoral inflation rates, however, will be buffeted about by shocks to the sectoral technology index Z_s.

The state variables for the firm's optimization problem are the firm's relative price at the end of the previous period $(p_{si,-1})$, the growth rate of the sectoral technology index (g_{Z_s}), the sectoral inflation rate (π_s), the sectoral output-productivity ratio (ζ_s), the idiosyncratic productivity index (A_{si}), and the information set Ω used to form future expectations.

Given these state variables, $S = \{p_{si,-1}, g_{Z_s}, \pi_s, \zeta_s, A_{si}, \Omega\}$, the firm maximizes the

10

value function

$$V(S) = \max(V^C(S), V^{NC}(S)), \tag{16}$$

where $V^C(S)$ represents the firm's value if it changes its price and $V^{NC}(S)$ represents its value if it does not change its price. These value functions, in turn, are

$$V^C(S) = \max_{p_{si}} \left\{ p_{si}\varphi\left(p_{si}D_s\right) - wA_{si}^{\frac{-1}{\eta}}\zeta_s\left(\varphi\left(p_{si}D_s\right)\right)^{\frac{1}{\eta}} - \phi_s + E_{S'|S}\left[\beta V(S')\right] \right\}, \tag{17}$$

with $S' = \{p_{si}, g'_{Z_s}, \pi'_s, \zeta'_s, A'_{si}, \Omega'\}$, and

$$V^{NC}(S) = \frac{p_{si,-1}}{1+\pi_s}\varphi\left(\frac{p_{si,-1}}{1+\pi_s}D_s\right) - wA_{si}^{\frac{-1}{\eta}}\zeta_s\left(\varphi\left(\frac{p_{si,-1}}{1+\pi_s}D_s\right)\right)^{\frac{1}{\eta}} + E_{S'|S}\left[\beta V\left(S'\right)\right] \tag{18}$$

with $S' = \{\frac{p_{si,-1}}{1+\pi_s}, g'_{Z_s}, \pi'_s, \zeta'_s, A'_{si}, \Omega'\}$. The parameter β is the discount factor, appropriately adjusted for the rescaling of the problem in terms of profits relative to growing nominal GDP.

In order to solve this optimization problem, each firm must be able to form expectations over the state variables in the subsequent period. Based on the consumer's optimization problem, each sector's nominal output share is constant: $\alpha_s = \frac{P_s Y_s}{PY}$. The assumption that the nominal output of the economy is growing at a constant rate therefore implies that the sectoral level of nominal output is also growing at a constant rate. Since all firms know this constant growth rate, they only need to compute a forecast of one of the sectoral aggregates (inflation or real output growth) and then they can back

out the implied forecast of the other. Here we will describe forecasts of inflation.

In the spirit of Krusell and Smith (1998), we assume each firm forecasts next period's inflation using the following linear forecasting rule:

$$\pi_{s,t+1}^f = \alpha_0 + \alpha_1 \pi_{s,t} + \alpha_2 \log \zeta_{s,t} + \alpha_3 g_{Z_s,t} + \epsilon_{\pi_s,t}, \ \epsilon \sim N(0, \sigma_{\epsilon_{\pi_s}}^2), \tag{19}$$

where the residual is assumed to be orthogonal to the other right-hand-side variables. The "regressors" are all state variables in firms' information sets at time t. A firm's idiosyncratic shock is not included because the price setting behavior of a single firm should not affect the sectoral inflation rate. Because of the error term, firms are not simply using a point forecast for next period's inflation, but rather are taking into account the distribution of next period's inflation conditional on this period's observables.

Given their forecast for next period's inflation rate, firms can derive expectations for the sectoral output-productivity ratio, $\zeta_{s,t+1}$. First, they will back out the forecast of sectoral output growth:

$$g_{Y_s,t+1}^f = g_{PY} - \pi_{s,t+1}^f. \tag{20}$$

Second, they will compute the forecasted value of $\zeta_{s,t+1}$, shown here in log-levels,

$$\log \zeta_{s,t+1}^f = \log \zeta_{s,t} + \frac{1}{\eta} g_{Y_s,t+1}^f - \frac{1}{\eta} g_{Z_s,t+1}^f, \tag{21}$$

based on their forecast of $g_{Z_s,t+1}$.

Regarding the exogenous processes, we assume that the idiosyncratic productivity index follows a log-normal autoregressive process:

$$\log A_{si,t+1} = \rho_A \log A_{si,t} + \epsilon_{A,t+1}, \ \epsilon \sim N(0, \sigma_{\epsilon_A}^2). \tag{22}$$

We assume that the growth rate of the sectoral technology index follows a normal autoregressive process:

$$g_{Z_s,t+1} = \mu_{g_{Z_s}} + \rho_{g_{Z_s}} g_{Z_s,t} + \epsilon_{g_{Z_s},t+1}, \ \epsilon \sim N(0, \sigma_{\epsilon_{g_{Z_s}}}^2). \tag{23}$$

2.5 Modeling expectations of sectoral inflation and output growth

In order to compute expectations of sectoral inflation and output growth, we will set up a three-variable VAR(1) using the sectoral state variables:

$$\begin{bmatrix} \pi_{s,t+1}^f \\ \log \zeta_{s,t+1}^f \\ g_{Z_s,t+1} \end{bmatrix} = A_0 + A_1 \begin{bmatrix} \pi_{s,t} \\ \log \zeta_{s,t} \\ g_{Z_s,t} \end{bmatrix} + \xi_{t+1}. \tag{24}$$

We assume that ξ_{t+1} is not known until after all pricing decisions are made in period t.

With a little manipulation, we can convert (19), (20), (21), and (23) into the following

13

VAR system:

$$\pi^f_{s,t+1} = a_0 + a_1\pi_{s,t} + a_2\log\zeta_{s,t} + a_3 g_{Z_s,t} + \epsilon_{\pi,t+1} \tag{25}$$

$$\log\zeta^f_{s,t+1} = \frac{g_{PY} - a_0 - \mu_{g_{Z_s}}}{\eta} - \frac{a_1}{\eta}\pi_{s,t} + \frac{\eta - a_2}{\eta}\log\zeta_{s,t} - \frac{a_3 + \rho_{g_{Z_s}}}{\eta}g_{Z_s,t} \tag{26}$$

$$-\frac{1}{\eta}\epsilon_{\pi,t+1} - \frac{1}{\eta}\epsilon_{g_{Z_s},t+1}$$

$$g_{Z_s,t+1} = \mu_{g_{Z_s}} + \rho_{g_{Z_s}}g_{Z_s,t} + \epsilon_{g_{Z_s},t+1} \tag{27}$$

where

$$a_0 = (1 - a_1)\pi_{ss} - a_2\log(\alpha_s) - a_3\frac{\mu_{g_{Z_s}}}{1 - \rho_{g_{Z_s}}}. \tag{28}$$

2.6 Kimball aggregator

In order to explore the role of real rigidities, we have selected a flexible function for the aggregator $\Upsilon(x)$. Recall that x is the relative quantity consumed of an individual variety. Our function is parsimoniously governed by two parameters, $\{\bar\theta, \bar\varepsilon\}$:

$$\Upsilon(x) = 1 + (\bar\theta - 1)\exp\left(\frac{1}{\bar\varepsilon}\right)\bar\varepsilon^{\frac{\bar\theta}{\bar\varepsilon}-1}\left(\Gamma\left(\frac{\bar\theta}{\bar\varepsilon}, \frac{1}{\bar\varepsilon}\right) - \Gamma\left(\frac{\bar\theta}{\bar\varepsilon}, \frac{x^{\frac{\bar\varepsilon}{\bar\theta}}}{\bar\varepsilon}\right)\right) \tag{29}$$

where $\Gamma(u, z)$ is the incomplete gamma function

$$\Gamma(u, z) \equiv \int_z^\infty s^{u-1}e^{-s}ds. \tag{30}$$

This function is a generalization of the CES aggregator, Υ^{CES}. In the limit as $\bar{\varepsilon} \to 0$, then $\Upsilon \to \Upsilon^{CES}$.

The solution to the model derived above depends on the derivative and the inverse of the derivative of Υ:

$$\Upsilon'(x) = \frac{\bar{\theta} - 1}{\bar{\theta}} \exp\left(\frac{1 - x^{\frac{\bar{\varepsilon}}{\bar{\theta}}}}{\bar{\varepsilon}}\right) \tag{31}$$

$$\varphi(x) \equiv (\Upsilon')^{-1}(x) = \left(1 + \bar{\varepsilon}\ln\left(\frac{\bar{\theta} - 1}{\bar{\theta}x}\right)\right)^{\frac{\bar{\theta}}{\bar{\varepsilon}}}. \tag{32}$$

With our functional form for Υ^{CES}, the price elasticity of demand for a given variety can vary with the variety's relative price. The elasticity is

$$\theta(x) \equiv -\frac{\Upsilon'(x)}{x\Upsilon''(x)} = \bar{\theta}x^{-\frac{\bar{\varepsilon}}{\bar{\theta}}}, \tag{33}$$

where $x = \frac{n_s Y_{si}}{Y_s} = \varphi(\frac{P_{si}}{P_s}D_s)$. Recall that $D_s = \sum_{i=1}^{n_s} \Upsilon'\left(\frac{n_s Y_{si}}{Y_s}\right)\frac{Y_{si}}{Y_s}$. In the Dixit-Stiglitz case ($\bar{\varepsilon} \to 0$), the elasticity is constant and equal to $\bar{\theta}$.

This functional form also produces variation in the super-elasticity, or the rate of change of the elasticity. The super-elasticity is expressed as

$$\varepsilon(x) \equiv 1 - \frac{\Upsilon'(x)}{x\Upsilon''(x)} - \frac{\Upsilon'(x)\Upsilon'''(x)}{\Upsilon''(x)^2} = \bar{\varepsilon}x^{-\frac{\bar{\varepsilon}}{\bar{\theta}}}, \tag{34}$$

where, again, $x = \frac{nY_{si}}{Y_s} = \varphi(\frac{P_{si}}{P_s}D_s)$. Depending on the value of $\bar{\varepsilon}$, the super-elasticity can provide a strong incentive for a firm to keeps its price close to the average sectoral price. Note that $\bar{\theta}$ and $\bar{\varepsilon}$ are the values of the elasticity and the super-elasticity at any

symmetric equilibrium, i.e., whenever $\frac{nY_{si}}{Y_s} = 1 \; \forall \; i$.

The effects of the super-elasticity on demand for a given variety are illustrated in Figure 1. Compared to the Dixit-Stiglitz case of $\bar{\varepsilon} = 0$, the demand curve is less convex with $\bar{\varepsilon} > 0$. When $\bar{\varepsilon} = 5$ the demand curve is approximately linear, and with $\bar{\varepsilon} = 10$ it is ostensibly concave. Kimball's preferences create a smoothed version of a kinked demand curve, although for different reasons than in the traditional use of the term (other prices are held fixed here, so it does not hinge on asymmetric responses of competitor prices). As a firm's relative price rises above one, its demand is choked off more quickly than with CES. And as its relative price declines below one, its demand rises less rapidly than it does under CES. Unlike CES preferences, with concavity there is a finite "choke price" at which demand is zero. This will play an important role in our simulations, because it effectively offers a firm the option of selling no output if it should so desire in the face of comparatively low idiosyncratic productivity.

Figure 2 plots a firm's profits as its price moves away from the symmetric point, assuming common productivity and constant returns to scale. The higher the super-elasticity, the more concave the profit function. Profits decline more steeply away from one because price increases are penalized by plummeting demand and price decreases are not rewarded by soaring demand. As in Figure 1, the prices at which demand disappears entirely are clearly visible. The greater concavity drives home the "real rigidity" induced by Kimball's preferences. When idiosyncratic productivity shocks hit, firms will be less aggressive in passing these marginal cost shocks on to their relative prices. And when

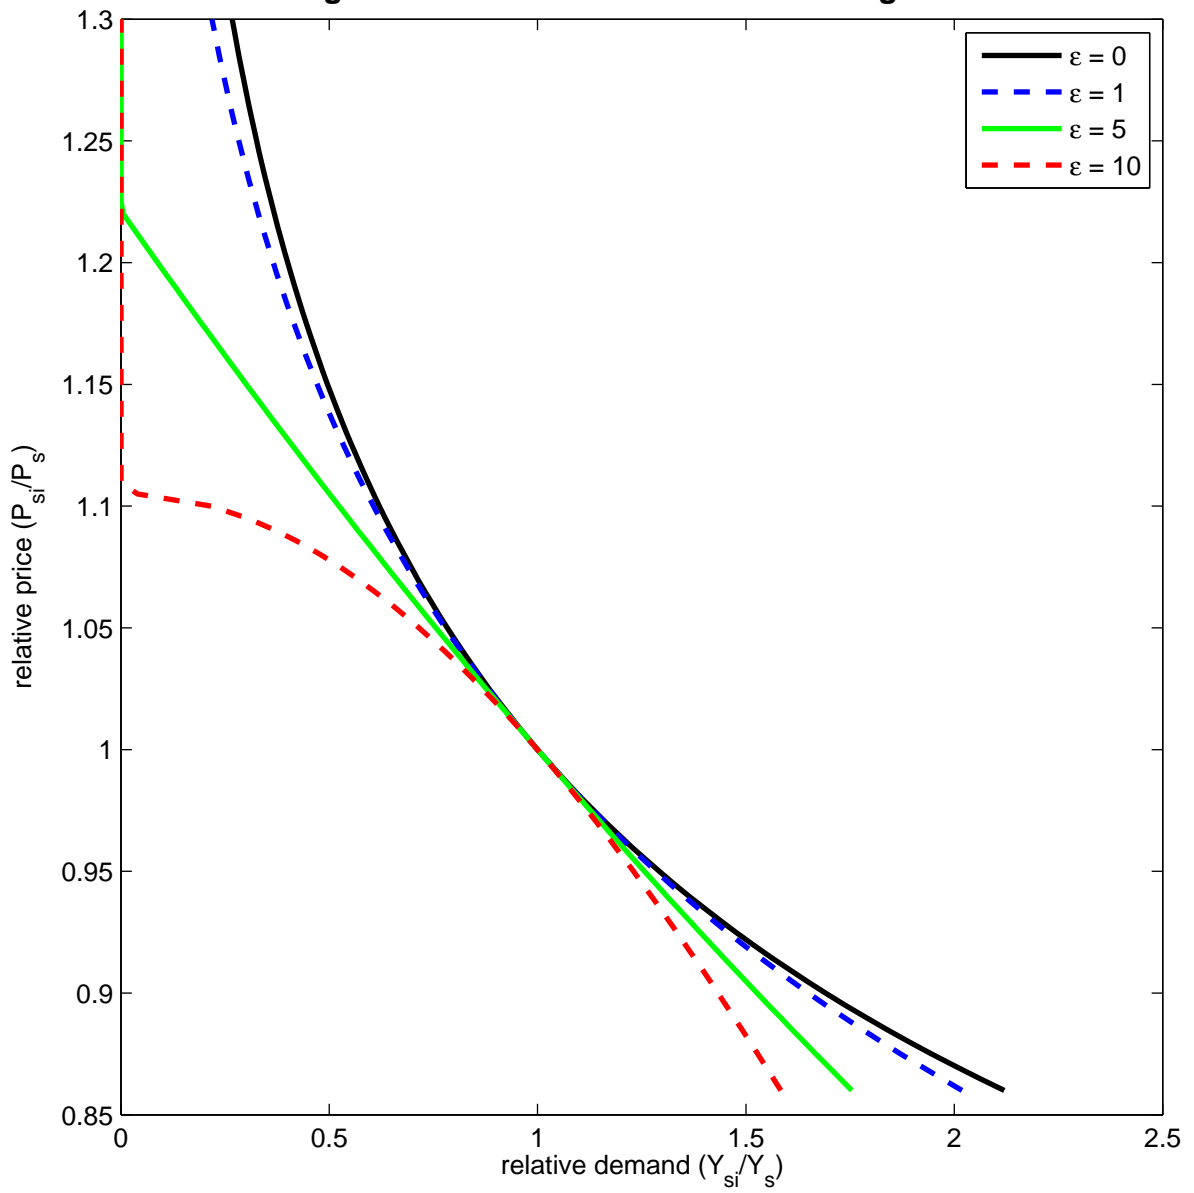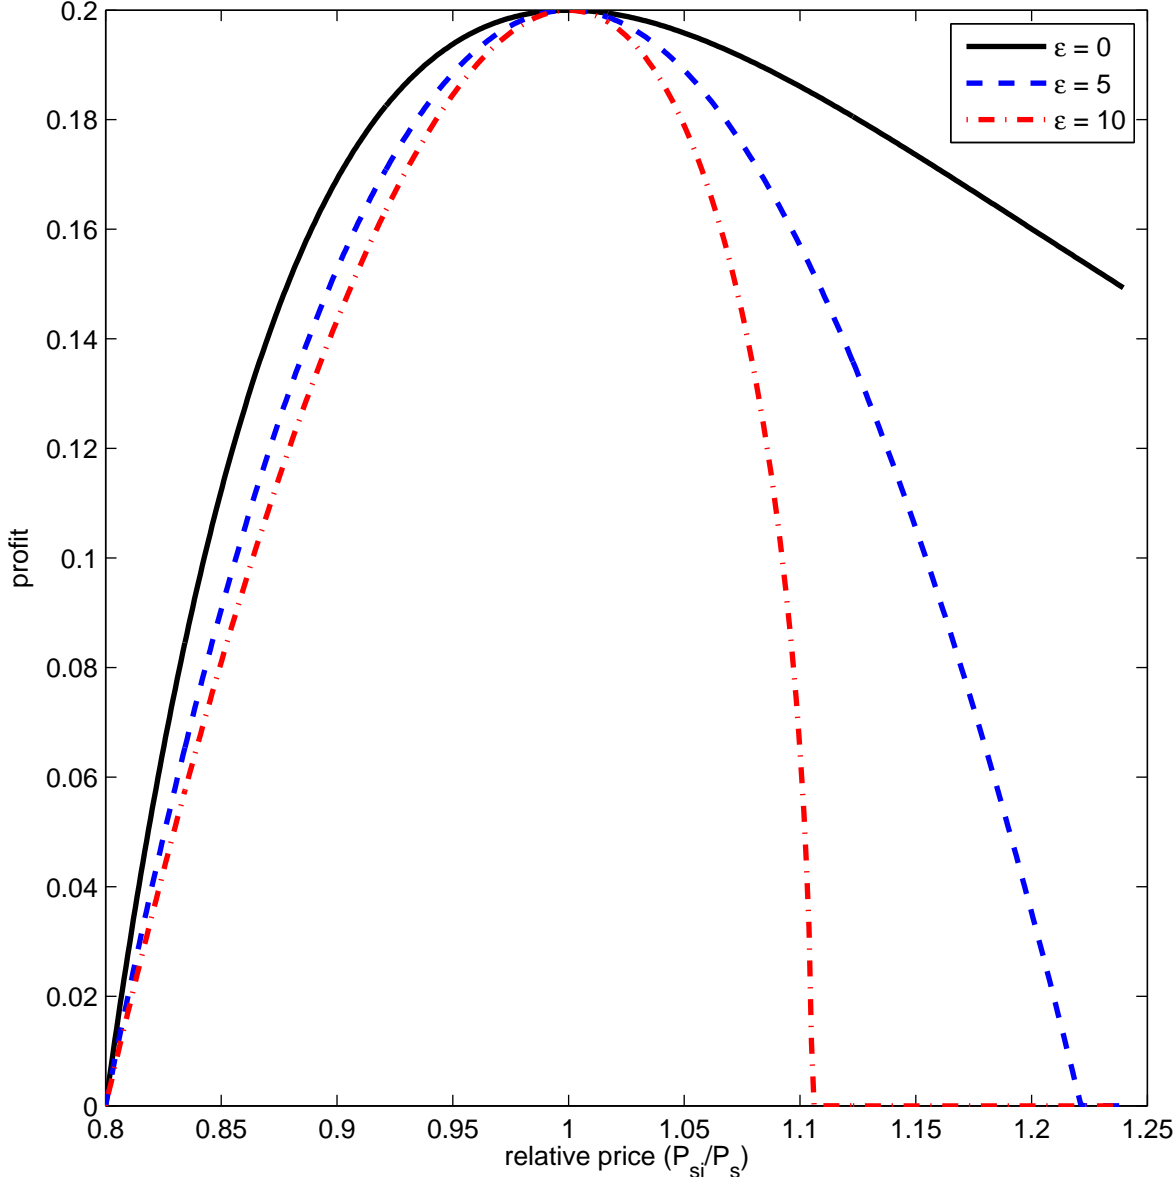**Figure 1: Demand function with real rigidities**

Figure 2: Profit function

common sectoral shocks hit, firm price responses will not be synchronized because of the idiosyncratic shocks. As a result, the "Kimball kink" will slow down the response to common shocks as well. How much so we will see in section 4 below. In the interim we will briefly describe how we will solve the model and the data we will use to discipline the model's predictions.

2.7   Model solution

Due to the presence of a discrete-choice decision in the optimization problem expressed in (16), the model is solved numerically using value function iteration. In this solution, all state variables are placed on discrete grids. The bounds of the relative price state are set wide enough to include all optimal pricing decisions, and prices are placed on the grid in increments of 0.5%. The autoregressive process for idiosyncratic productivity is transformed into a discrete-valued Markov chain following Tauchen (1986).[3]   The three-variable VAR for sectoral inflation, the sectoral output-productivity ratio, and the sectoral technology growth are similarly converted into a first-order Markov chain.[4] This conversion results in a transition matrix expressing the probability of observing any realization of future sectoral-level state variables as a function of the current state variables.

In addition to the parameters which we will estimate, we set several parameters based on the literature, U.S. data, or the steady state solution of the model. The growth rate

---

[3]The discrete grid for idiosyncratic productivity contains 7 points.

[4]The discrete grids for sectoral inflation, the sectoral output-productivity ratio, and the sectoral technology growth shock contain 11, 7, and 7 points, respectively.

of nominal output for the economy is set at 0.37 percent per month, which reflects average nominal GDP growth of 4.4 percent in the United States over 1988-2004. The monthly discount rate, β, is set at 0.996. The (normalized) wage is set at its symmetric equilibrium steady state value, $w = \frac{\bar{\theta}-1}{\theta}$.

Following Willis (2000), the inflation forecasting equation in (19) is used to compute a rational expectations equilibrium of the model. For a given specification of the structural parameters of the model along with the inflation forecasting parameters, $\Theta = \{\alpha_1, \alpha_2, \alpha_3\}$, the model is solved and the policy function is generated. A panel of 320 firms over 240 months is then simulated using the policy functions.[5]

Simulating data from the model requires an updating process to determine the evolution of the endogenous sectoral-level state variables. For tractability, we assume that the sectoral variable D_s, which is a function of relative output levels of firms within the sector, is held constant at its average value. The steady state value of this variable when $\bar{\epsilon} = 0$ is $D_s = \frac{\bar{\theta}-1}{\theta}$. More generally, D_s is concave in the dispersion of relative output and decreasing in $\bar{\epsilon}$. Since no closed form solution is available for D_s, its value is set equal to the average value of D_s computed using simulated data and equation (8).

The sectoral inflation rate and the sectoral output-productivity ratio ($\log \zeta_s$) are determined by the collective actions of firms in the simulation. When setting prices in the current period, firms know the current value of inflation and $\log \zeta_s$. To determine the current period inflation rate, which in turn determines the value of $\log \zeta_s$ using equations (20) and (21), we locate the grid point in the discretized inflation state space that most

---

[5]The size of the panel was chosen to match the size of the average sector in the BLS dataset.

closely matches the following two conditions from the model:[6]

$$1 \;=\; \frac{1}{n_s}\sum_{i=1}^{n_s}\Upsilon\left(\varphi\left(\frac{P_{si}}{P_s}D_s\right)\right) \tag{35}$$

$$1 \;=\; \frac{1}{n_s}\sum_{i=1}^{n_s}\frac{P_{si}}{P_s}\varphi\left(\frac{P_{si}}{P_s}D_s\right). \tag{36}$$

As a reminder, the first equation is the Kimball flexible variety aggregator. The second equation is the implicit definition of D_s given in (8) after substituting the demand function given in (9).

After simulating a panel of firm-months, we evaluate the forecasting rule used to form expectations for future inflation. Adding in the exogenous sectoral productivity growth shock, g_{Z_s}, an OLS regression of the linear forecasting rule in (19) is executed on the simulated data. The initial assumed values of the forecast parameters, Θ_0, are then compared to the OLS estimates, Θ_1. If these values differ, then the forecast parameters are updated based on Θ_1 and a new solution for the model is derived. This updating process continues until a fixed point is reached. This fixed-point solution represents a rational expectations equilibrium where the inflation forecasting rule assumed by firms matches up with the behavior of the simulated data.[7]

---

[6]Ideally, we would use these equations to endogenously determine D_s and π_s. We are working on an improved solution method that will allow us to include D_s as an additional state variable without a significant loss in computation speed. Currently, the assumption of a constant value for D_s does not appear to be too restrictive for the model. The standard deviation of D_s, computed using simulated data and equation (8), is only 0.002.

[7]Following Krusell and Smith (1998), we are currently exploring whether the inclusion of additional variables into the forecasting rule will lead to a significant improvement in the inflation forecast.

3 CPI Data

In its Commodities and Services Survey, the U.S. Bureau of Labor Statistics checks the prices of around 85,000 items a month in order to tabulate the U.S. CPI. An individual item refers to a product or service with specific attributes sold by a particular outlet in a given location. The Survey covers all goods and services other than shelter, or about 70% of the CPI based on BLS consumer expenditure weights. The CPI Research Database, maintained by the BLS Division of Price and Index Number Research, contains all prices in the Commodities and Services Survey from January 1988 to the present. [See Klenow and Kryvtsov (2005) for a more detailed description of the CPI Research Database.] We base our statistics on data through December 2004 for the three largest areas – New York, Los Angeles, and Chicago – for which all items are surveyed every month (as opposed to bi-monthly for most items in other areas). This sub-sample consists of about 14,000 prices a month.

The BLS identifies each collected price as either a "regular" price or a "sale" price (i.e., a temporarily low price that is labelled so in some way). Although sale prices may require implementation costs, we focus on regular prices because they exhibit smaller relative price changes. As we will report shortly, this will be a conservative approach. We also exclude all price changes coinciding with a change in the item surveyed, seasonal changeovers, and temporary stockouts. To minimize the importance of measurement error, we drop price changes that exceed 10 natural log points in absolute value. These price jumps constitute less than one-tenth of one percent of all price changes. Using a

lower threshold, such as five log points, has almost no effect on our tabulations.

In order to estimate the parameters in our model, we calculate five statistics from the CPI data. Several of these are related to statistics calculated by Klenow and Kryvtsov (2005) from earlier vintages of the CPI Research Database. Further, several of the statistics are similar to those used by Golosov and Lucas (2004) to calibrate their general equilibrium state-dependent pricing model (albeit, without any real rigidities).

The first moment we calculate is the average sectoral inflation rate over time. Let P_{sit} denote the price of item i in sector s in month t, and ω_{sit} the BLS weight on item i within category s in month t. The weights in sector s sum to ω_s^{93} in every month, the BLS consumption expenditure weight of category s in 1993 (which themselves sum to 1). We define the sectoral inflation rate in month t as

$$\pi_{st} = \sum_i \omega_{sit}[log(P_{sit}) - log(P_{sit-1})]/\omega_s^{93}.$$

For each of 67 sectors ("Expenditure Classes") in the BLS data, we calculate the mean of inflation across the 203 months from February 1988 to December of 2004, or $\pi_s = \sum_{t=1}^{203} \pi_{st}/203$. We then take the weighted average of these across sectors to arrive at 0.153% per month:

$$\overline{\pi} = \sum_s \omega_s^{93}\pi_s = 0.00153.$$

In similar fashion we calculate our second moment, the average (across sectors) of the standard deviation of sectoral inflation (across months). We first compute the standard

deviation of inflation across months for each sector, and then calculate the weighted mean of these sectoral standard deviations to be 1.02%:

$$\sigma_\pi = \sum_s \omega_s^{93} \sqrt{\sum_{t=1}^{203} (\pi_{st} - \pi_s)^2 / 202} = 0.0102.$$

Our third moment is the average fraction of items changing price from one month to the next. Let $I(\Delta P_{sit} \neq 0)$ be a price-change indicator for item i in sector s in month t. It takes on the value 1 if the item changed price from month $t-1$ to t, and 0 otherwise. Weighting items and sectors appropriately, this indicator averages 21.5% across items, sectors, and time:

$$\overline{I(\Delta P \neq 0)} = \sum_s \omega_s^{93} [\sum_{t=1}^{203} (\sum_i \omega_{sit} I(\Delta P_{sit} \neq 0)/\omega_s^{93})/203] = 0.215.$$

For the fourth and fifth moments, it is useful to define the sectoral price index $\overline{P}_{st} = exp(\sum_{k=2}^{t} \pi_{sk})$. The index is 1 at $t=1$ in each sector, and cumulates inflation going forward. We let p_{sit} denote the ratio of the price of item i to the sectoral price index, $p_{sit} = P_{sit}/\overline{P}_{st}$. This is the relative price of item i within sector s at time t. For each sector, we calculated the serial correlation and standard deviation of $log(p_{sit})$ across months with price changes. We took out item-specific means to deal with any discrepancy in units (e.g., sizes of cereal boxes). We then took the weighted mean of sector statistics to obtain a serial correlation of 0.318 and a standard deviation of 13.9%,

24

again across new prices:

$$\rho_p = \sum_s \omega_s^{93} \sqrt{\sum_i \omega_{si} \sum_{I_{sit}=1} (logp_{sit} - \overline{logp_{sit}})(logp_{sit-\tau_{sit}} - \overline{logp_{sit}})} = 0.318.$$

$$\sigma_p = \sum_s \omega_s^{93} \sqrt{\sum_i \omega_{si} \sum_{I_{sit}=1} (logp_{sit} - \overline{logp_{sit}})^2} = 0.139.$$

Here I_{sit} is shorthand for $I(\Delta P_{sit} \neq 0)$, and τ_{sit} is the age (duration) of the price at the time of its "death" in month t. Separately, note that the sectoral price index we defined above is *not* the same as that implied by the Kimball aggregator. We do not observe all of the prices in the market, and hence do not construct this ideal price index. When we simulate the model below, we will construct a simulation counterpart to what we calculated in the data.

In Table 1, nearby, we provide these moments. We compute bootstrapped standard errors by drawing "quotelines" (strings of prices for a given item) with replacement. As shown, the moments are estimated with great precision – not surprising given the 2.8 million micro datapoints on prices underlying them. If we had looked at posted prices rather than regular prices (i.e., omitted temporary price discounts), the main difference would be a higher standard deviation of new relative prices over time (19% rather than 14%). Including price changes involving product turnover, seasonal changeovers, or temporary stockouts would also have boosted the standard deviation. Finally, with more disaggregate BLS sectors, specifically 250 instead of 67, the standard deviation was virtually identical.

Table 1: BLS CPI Moments

$\bar{\pi}$	σ_π	$\overline{I(\Delta P \neq 0)}$	ρ_p	σ_p
0.00153	0.0102	0.215	0.318	0.139
(0.00001)	(0.0002)	(0.001)	(0.012)	(0.002)

4   Model Estimation and Simulation

We selected the statistics in Table 1 from many possible statistics because we think they represent key features of the data that the model should be able to mimic. We now use the statistics in Table 1 to estimate some of the structural parameters of the model. The parameters we will estimate are the standard deviation of innovations to sectoral productivity growth (σ_{g_z}), the autocorrelation coefficient for the idiosyncratic technology process (ρ_A), the standard deviation of innovations to the idiosyncratic technology component (σ_A), and the magnitude of the implementation cost (μ_ϕ). One additional parameter, the mean growth rate of sectoral productivity, is calibrated directly using the mean sectoral inflation rate. Although the other four parameters do not map one-to-one to data statistics, we have strong intuition for how they relate to each other. More volatile sectoral productivity growth, ceteris paribus, should boost the volatility of sectoral inflation and the frequency of price changes. Higher serial correlation of the idiosyncratic productivity term should increase the serial correlation and standard deviation of relative price movements. A bigger standard deviation of idiosyncratic in-

novations should increase the frequency of price changes and the size of relative price movements. Finally, a higher menu cost should, ceteris paribus, reduce the frequency of price changes.

We use the Simulated Method of Moments procedure to estimate these parameters. In this procedure, the BLS moments, $\Psi_{\text{BLS}} = \left\{ \sigma_\pi, \overline{I(\Delta P \neq 0)}, \rho_p, \sigma_p \right\}$ are matched up against the same moments computed from simulated data, $\Psi_{\text{sim}}(\Phi)$. The moments from the simulated data are functions of the structural parameters, $\Phi = \{\sigma_{g_Z}, \rho_A, \sigma_A, \mu_\phi\}$. The estimation involves finding the vector of structural parameters, Φ, that minimizes the weighted distance between BLS moments and simulated moments.

$$\min_\Phi \left(\Psi_{\text{BLS}} - \Psi_{\text{sim}}(\Phi)\right) W \left(\Psi_{BLS} - \Psi_{\text{sim}}(\Phi)\right)' \tag{37}$$

Note that W represents the weighting matrix.[8]

Before estimating, we fix the value of three parameters. One is the returns to scale parameter η. In this draft we set $\eta = 1$ in all simulations. In future drafts we will entertain $\eta < 1$ as an extreme version of firm-specific factor markets. We also set the elasticity of demand (evaluated at a relative price of 1) to $\bar{\theta} = 5$, which implies a markup of 25 percent in the case with no real rigidities. This is at the high end of most estimates in the IO literature, but lower than the value of 11 (10 percent markup) typically used

---

[8]As discussed in Gourieroux and Monfort (1996), the resulting estimator is consistent. To limit simulation error, we simulate a panel in the estimation procedure that is 10 times as long as the BLS dataset. Since there are no permanent differences across firms, this approach is the same as simulating 10 panels, each with a different initial distribution, and then taking the average of the moments across the 10 panels. We choose the former approach for computational simplicity.

in the macro literature. In future drafts we will consider higher and lower values to check the robustness of our findings. Finally and most crucially, we fix the value of the super-elasticity at the symmetric point, $\bar{\varepsilon}$. Initially we will set $\bar{\varepsilon} = 0$, the Dixit-Stiglitz case of a constant elasticity. But we will contrast this with the case of $\bar{\varepsilon} = 10$. This is the low value entertained by Eichenbaum and Fisher (2004), who also considered the value of 33 suggested by Kimball (1995). Ideally we would like to estimate the super-elasticity parameter, but even the elasticity parameter has proved hard to estimate in the literature. Because we have price data but no quantity data, the challenge would be even greater for us. We leave this to future versions.

In Table 2 below we present estimates of the four model parameters for the case when we impose $\bar{\varepsilon} = 0$. We call this our baseline case. Note that the idiosyncratic shock must be sizable (innovation standard deviation of around 12%) and somewhat persistent (serial correlation around 0.7) in order to match the persistence and volatility of item relative prices (across newly set prices). Also worth noting is the menu cost, which is estimated to be around 6% of average firm revenue *when spent*. The menu cost must be multiplied by the frequency of price changes to obtain the average expenditures on menu costs relative to average firm revenue. This comes to 1.4%, which is in the neighborhood of estimates by Levy, Bergen, Dutta and Venable (1997) and Zbaracki et al. (2004).

Table 2: Estimation with $\bar{\epsilon} = 0$

Parameter Estimates				Forecast Parameters			
σ_{gz}	ρ_A	σ_A	μ_ϕ	α_1	α_2	α_2	R^2
0.017	0.678	0.119	0.064	0.01	0.33	-0.05	0.14
(0.001)	(0.012)	(0.001)	(0.002)				

Figure 3 plots the model response of the sectoral price index to a negative shock to sectoral productivity. Impulse responses are computed by introducing a shock to sectoral productivity growth relative to the baseline simulation. The shock is applied iteratively to each possible period in the simulation, creating a series of impulse responses. Each impulse response is based on a 1-period deviation in the sectoral productivity growth rate relative to baseline. The variation across the responses arises from the model's nonlinearity. In a linear model the response is the same regardless of the starting distribution of relative prices.

In Figure 3, the solid line in the middle is the *average* response across simulations. The two dotted lines represent the boundaries for the region that contains the middle 68 percent of the impulse responses, which approximates one standard deviation. The width of the bands illustrates that the underlying distribution plays a large role in the response to sectoral shocks.

On average, prices ultimately rise about 1% in response to the shock. Our focus is on how long it takes to get there. The longer it takes, the greater the real output response in the meantime. In the absence of the real rigidity and in the presence of

Figure 3: Sectoral price response to 1 percent sectoral productivity decline (estimation with ε = 0)
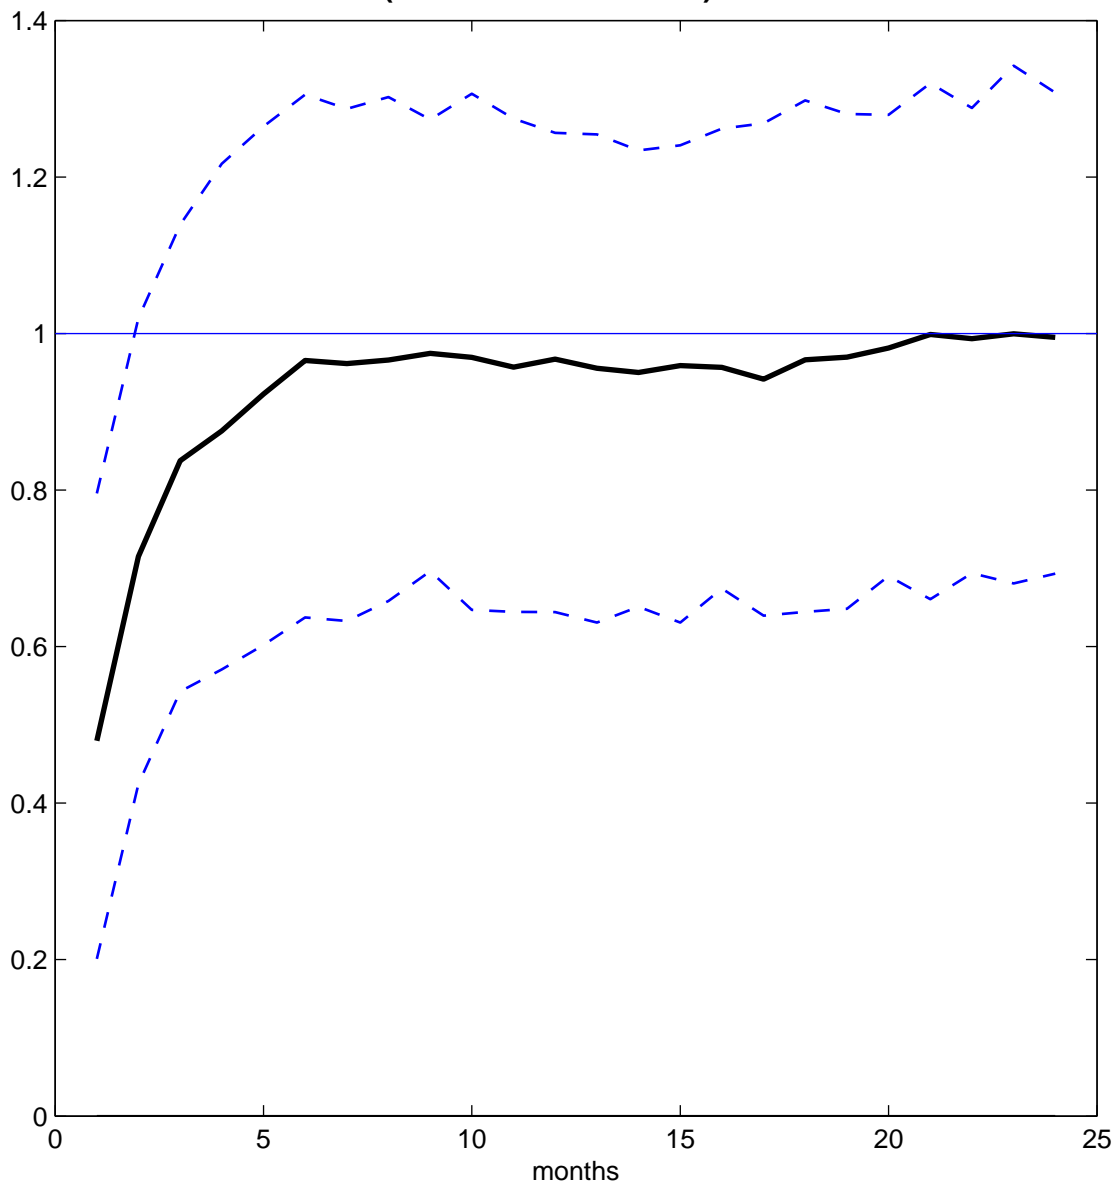
months

modest nominal rigidity (over 20% of items changing prices per month), the response is swift. The half-life is about one month, and prices almost fully respond after 6 months. Clearly, the baseline model does not generate as much persistence as sought to match structural VAR evidence of effects lasting well beyond a year.

We next simulate a model with the baseline parameter values from Table 2, except with $\bar{\varepsilon} = 10$ rather than zero. Table 3 compares the data moments to those in the baseline model and to those with $\bar{\varepsilon}$ set to 10. The results are predictable. Adding a real rigidity a la Kimball (1995) makes firms more reluctant to change prices (9% of the time, down from 22%), and makes relative prices more stable (serial correlation of 0.15 down from 0.31, and standard deviation of 4% down from 14%). With the Kimball kink, firms do not pass marginal cost shocks as fully onto their prices.

Table 3: Moments

	σ_π	$\overline{I(\Delta P \neq 0)}$	ρ_p	σ_p
BLS	0.010	0.215	0.318	0.139
$\bar{\varepsilon} = 0$	0.010	0.215	0.312	0.140
$\bar{\varepsilon} = 10$	0.005	0.089	0.154	0.038

Adding the real rigidity does prolong the response to a sectoral shock. Figure 4 displays a comparison of impulse responses from the two versions of the baseline estimation model. The addition of real rigidities ($\bar{\varepsilon} = 10$) to the model increases the half-life of the response to 5 months, and it now takes about 18 months for the full effect to be realized.

31

Figure 4: Sectoral price response to 1 percent sectoral productivity decline
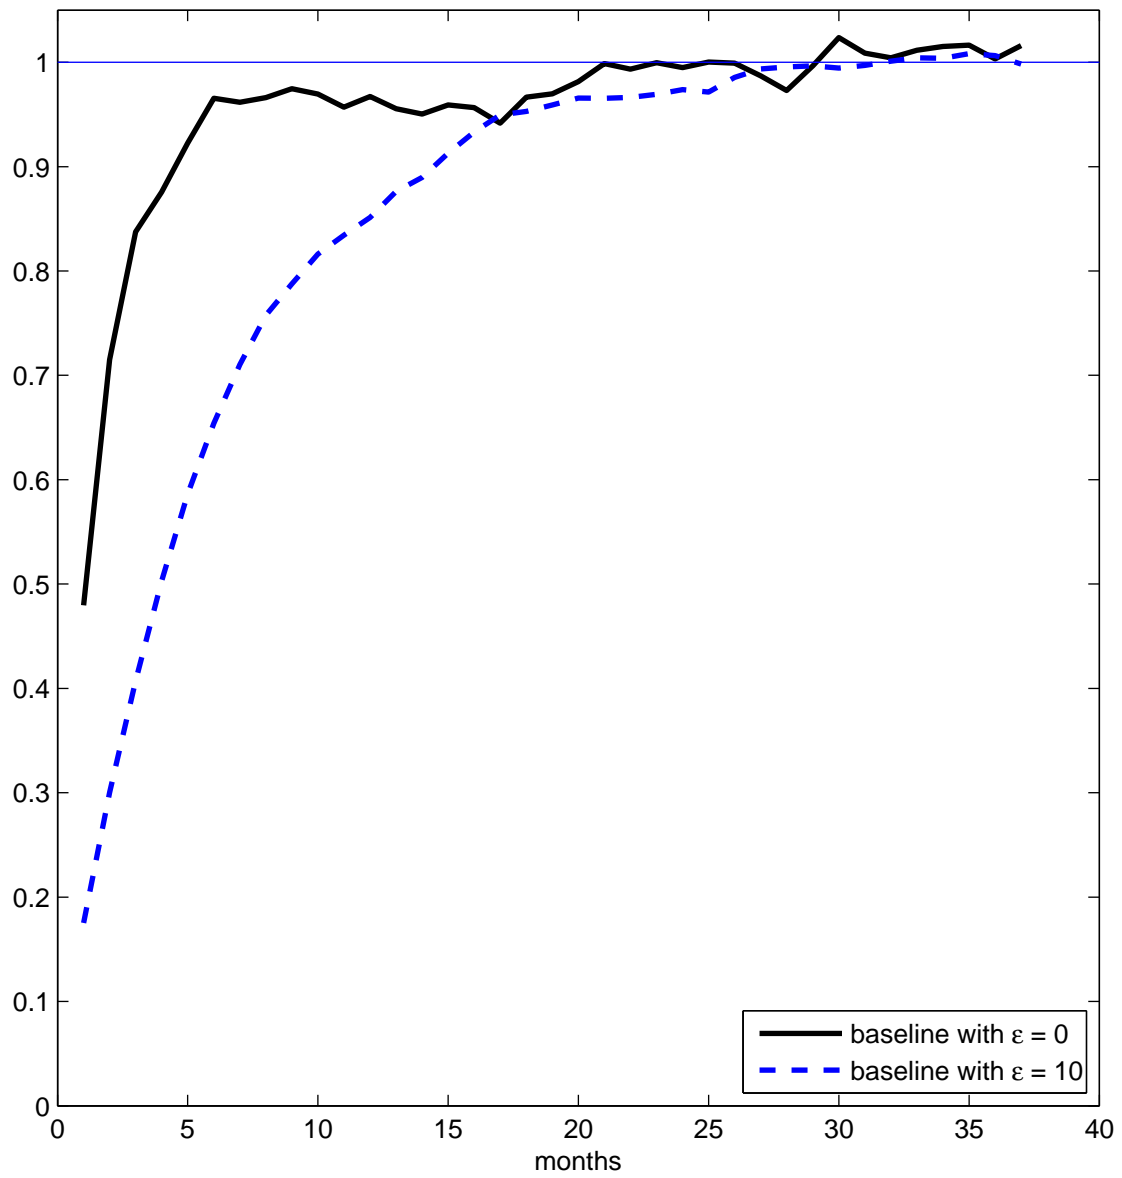
As shown in Table 3, however, adding the real rigidity pushes the model moments away from the data moments. We therefore re-estimate the model subject to $\bar{\varepsilon} = 10$. The resulting parameter estimates (which closely match model and data moments) are given in Table 4. Figure 5 provides the average impulse response and the one-standard-deviation bands. Figure 6 compares the baseline impulse response function to that with $\bar{\varepsilon} = 10$. With the real rigidity, the half-life is 4 months, and the full effect on prices not seen until about 5 months. While it was shown in Figure 4 that real rigidity of the degree considered here can considerably slow the response of prices, when the model with real rigidity is estimated to match the BLS moments, the response of prices occurs slightly more quickly.

Table 4: Parameter Estimates When $\bar{\epsilon} = 10$

Parameter Estimates				Forecast Parameters			
σ_{g_z}	ρ_A	σ_A	μ_ϕ	α_1	α_2	α_2	R^2
0.024	0.705	0.349	0.130	0.01	0.22	-0.04	0.16
(0.003)	(0.013)	(0.025)	(0.055)				

Comparing the parameter estimates in Tables 2 and 4, two important differences are evident. First, with the real rigidity the idiosyncratic shock must be very large – about 35%, compared to 12% without the real rigidity. This is the standard deviation of the *monthly* innovation to firm productivity. In future drafts we will try to compare this to evidence on firm-level productivity, but these shocks seem very large to us. They

Figure 5: Sectoral price response to 1 percent sectoral productivity decline (estimation with $\varepsilon = 10$)

months
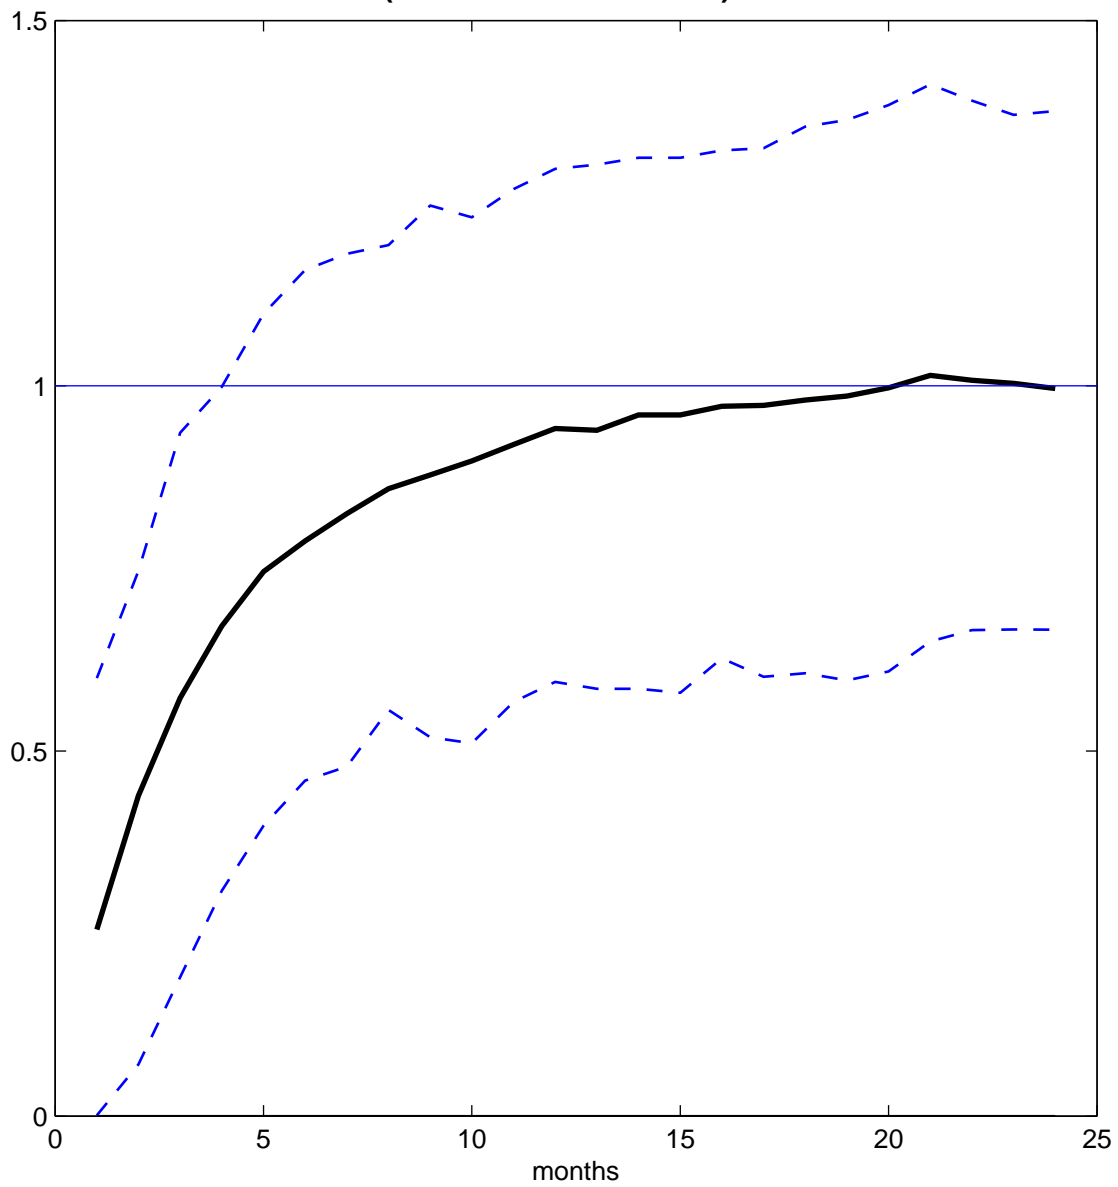
**Figure 6: Sectoral price response to 1 percent sectoral productivity decline**
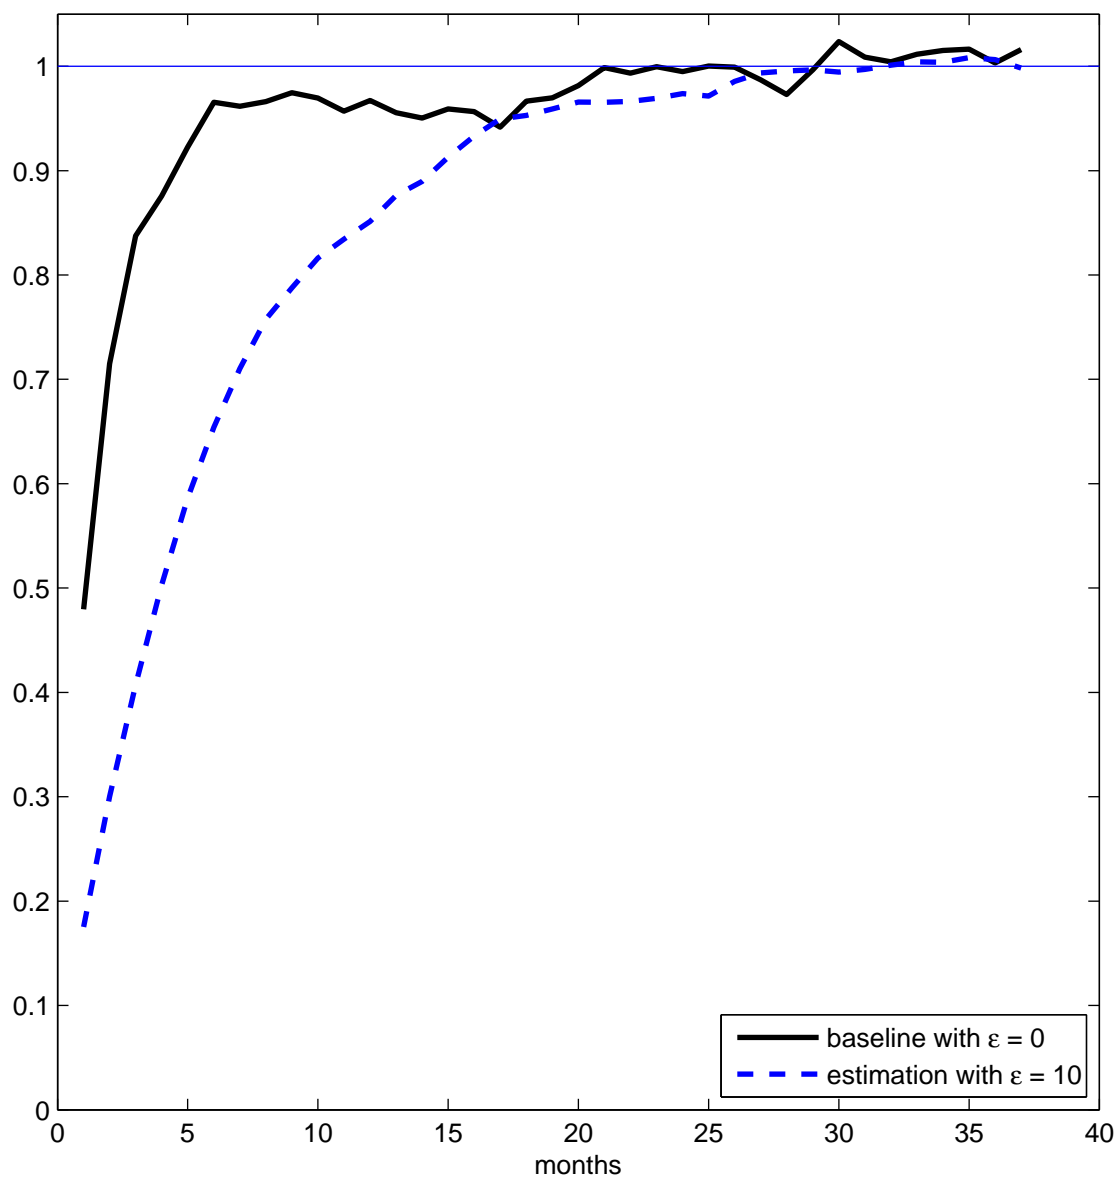
are necessary because, in the presence of the real rigidity, firms must face very large marginal cost shocks in order to change their relative prices as much as we observe in the CPI data. The second important change is to the size of menu costs. It doubles in size to 13% of average firm revenue in the presence of the real rigidity. Taking into account the frequency of price changes, menu costs absorb nearly 3% of average firm revenue. This is larger than estimated by the papers with direct evidence on menu costs (around 1%).

The models with and without the real rigidity also differ markedly in their implications for quantity movements. For the model with $\bar{\varepsilon} = 0$, Figure 7 plots simulated prices and quantities for 100 months for a single item/firm. Both the prices and quantities are relative to the industry aggregates. Given that supply (productivity) shocks drive price movements in the model, the price and quantity movements are in opposite directions. And given that demand is elastic ($\bar{\theta} = 5 \gg 1$), the quantities move significantly more, in percentage terms, than the prices do.

Figure 8 plots simulated prices and quantities when $\bar{\varepsilon} = 10$. Compared to when $\bar{\varepsilon} = 0$, quantities do not reach the same highs with $\bar{\varepsilon} = 10$. The real rigidity dampens the rise in quantity demanded when the price falls; relative quantities do not even reach twice the symmetric value, compared to over three times the symmetric value with no real rigidity. The flip side is that quantities fall more sharply with the real rigidity in response to relative price increases. Whereas quantities bottom out at half the symmetric level without the real rigidity, they frequently fall to zero in the presence of the real

Figure 7: Simulation of firm−level prices and output when ε = 0
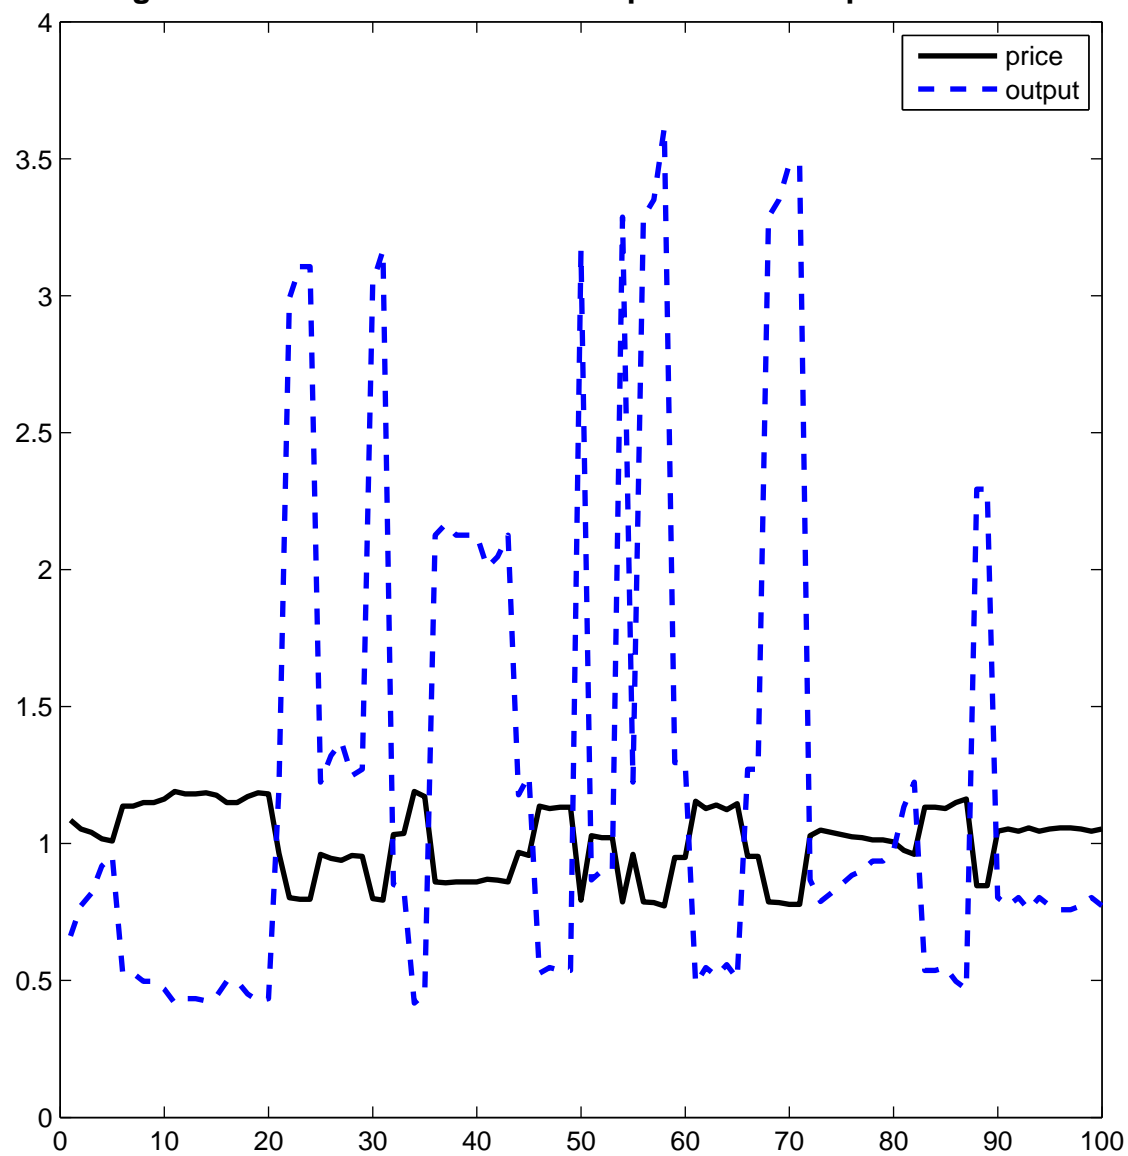
**Figure 8: Simulation of firm–level prices and output when ε = 10**
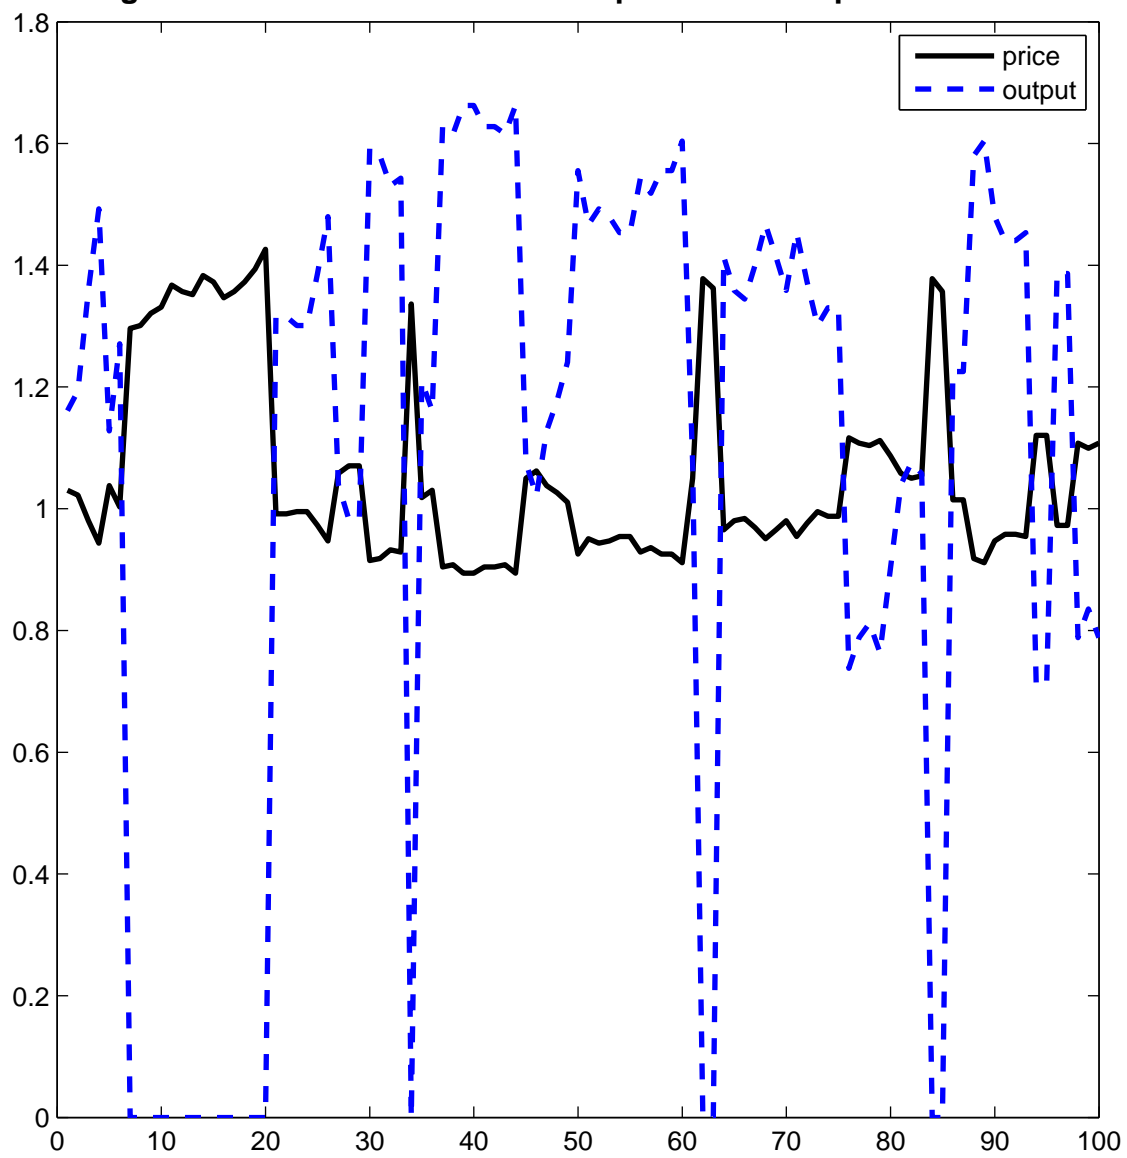
rigidity. Strong real rigidity induces concavity in the demand curve, as shown in Figure 1. So quantities hit zero at finite relative prices. Figure 8 demonstrates that this is not just a possibility, but a regular occurrence. Across many simulations, "total eclipse of demand" occurs in about 22% of months, close to the frequency in Figure 8. We find this implication implausible, but in the future we will investigate it systematically using scanner data from U.S. grocery stores.

We next look at the histogram of relative prices and relative quantities in the absence and presence of the Kimball real rigidity, respectively. Figures 9 and 10 are the histograms of relative prices (pooled across firm-months). With $\bar{\epsilon} = 10$, relative prices are bimodal. Firms keep their relative price close to 1 unless their marginal cost is so high that it is not profitable to sell, in which case they price themselves out of the market. [In the model firms do not have the option of simply stocking out temporarily. This occurs in about 5% of months in the CPI microdata. But no data on prices is available in such months. The CPI relative price variability applies to items in stock.]

Figures 11 and 12 are the histograms of relative quantities (again, pooled across firm-months). Quantities are more tightly distributed around 1 most of the time in the presence of the real rigidity. But the left tail of zeroes stands out relative to what happens without the real rigidity.

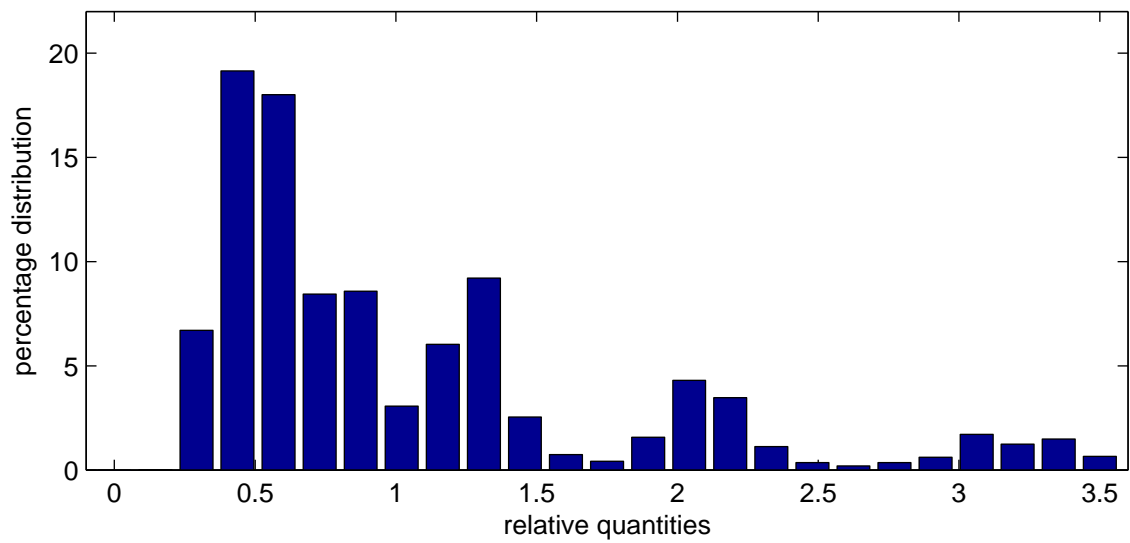**Figure 9: Histogram of prices, estimation with $\varepsilon = 0$**
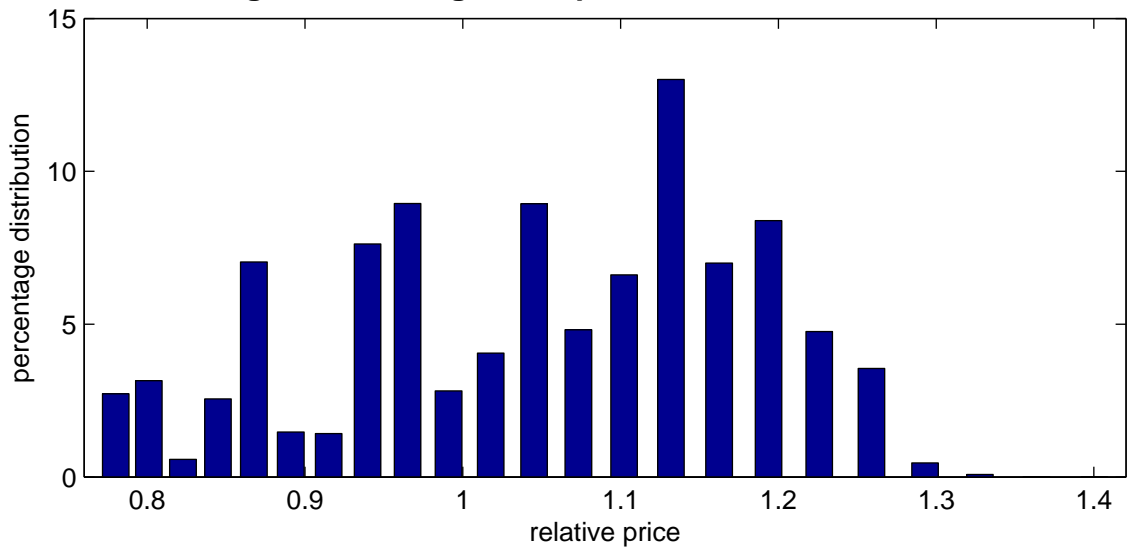
Figure 10: Histogram of prices, estimation with $\varepsilon = 10$
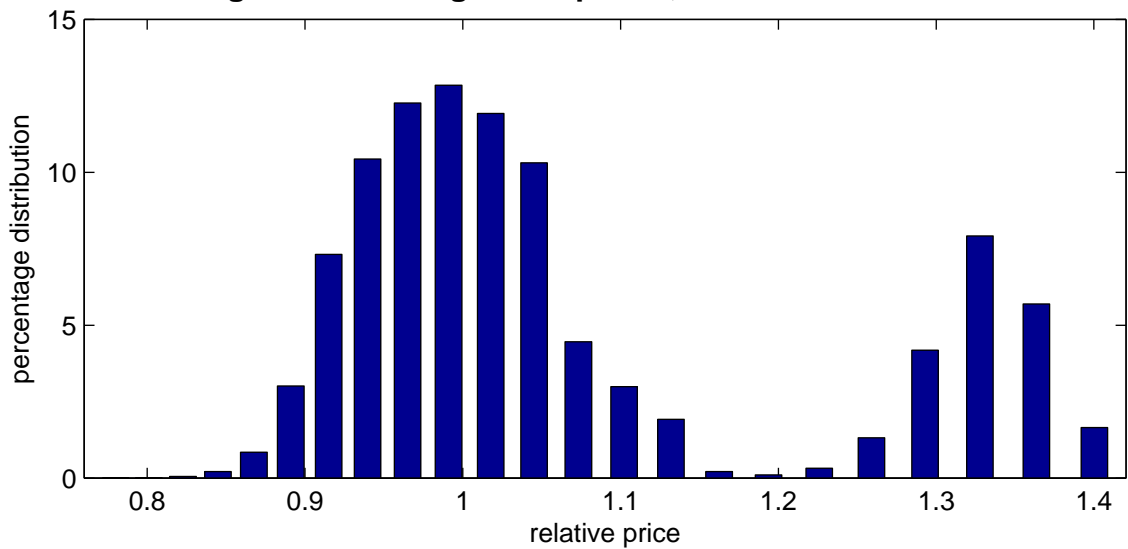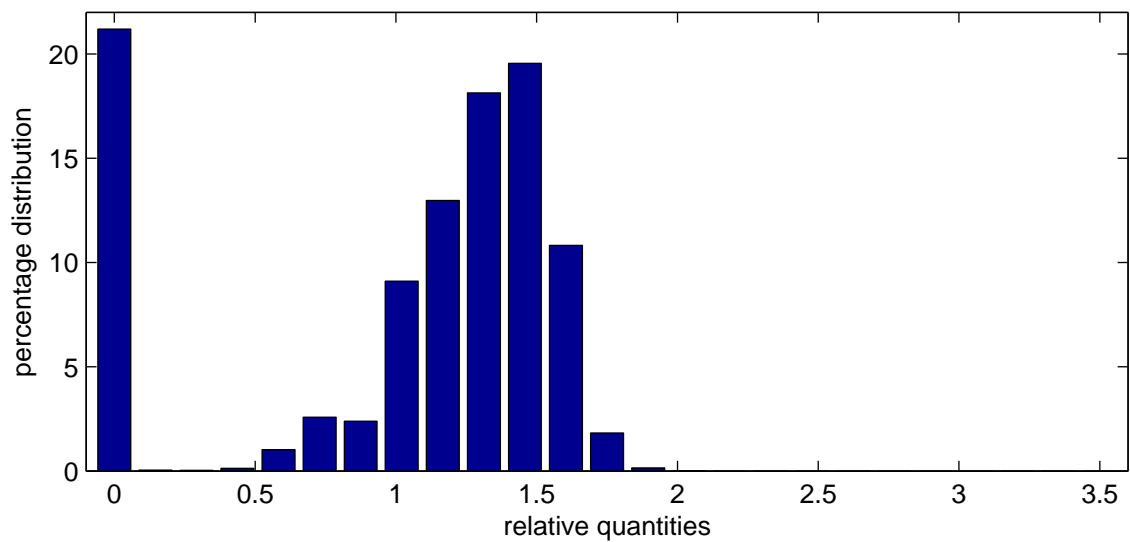
Figure 11: Histogram of quantities, estimation with $\varepsilon = 0$

Figure 12: Histogram of quantities, estimation with $\varepsilon = 10$

5 Conclusion

Research on monetary policy shocks seeks a model in which these shocks have real effects lasting beyond a year. Promising ingredients include real rigidities coupled with nominal rigidities. In this paper we explored the implications of Kimball's (1995) concave demand curve. Such a real rigidity makes firms highly averse to changing their relative prices, so that without coordination it takes a long time for aggregate shocks to fully work themselves into prices. But the micro evidence from the U.S. CPI displays large changes in relative prices. Reconciling this micro fact with the Kimball real rigidity, we find, requires big shocks to firm productivity (around 35% per month) and implies that firms frequently price themselves out of the market (about 22% of months). These properties appear extreme to us, but in future revisions we will compare them to micro evidence on productivity and quantities.

A Kimball rigidity need not always have extreme implications for shocks and quantities. Perhaps, at a sufficiently disaggregate level, price changes are highly synchronized. Synchronization, however, might undo the aggregate persistence that real rigidities were conceived to generate. (No need to dampen one's response to a common shock if competitor price changes are synchronized.) Nevertheless, we will investigate this hypothesis using scanner data from grocery stores.[9] Another possibility is that taste shocks hit demand for individual items in ways that affect relative prices but not relative quantities.

---

[9]Our preliminary investigation of publicly available Dominick's data suggests that our CPI-based calculations do not overstate relative price variability of close substitutes such as different brands of soft drink, paper towels, cereal, etc.

In the future we will consider idiosyncratic taste shocks as well. A third possibility is that the Kimball rigidity applies at a higher, rather than lower, level of aggregation. Consumers might have concave preferences over cereal vs. toothpaste, rather than over competing brands of cereal or toothpaste. We will consider this hypothesis too in future revisions. But, like synchronization in the presence of a low level real rigidity, a real rigidity at a high level of aggregation might not prolong real effects of nominal shocks. (No need to dampen one's response to a common shock if one has little effect on the category aggregate.)

In closing, we stress that we have focused on a single, albeit powerful, real rigidity: concave demand. The facts we document may or may not be relevant for evaluating some other real rigidities, such as specific factor markets or real wage rigidities. We leave that to future work.

References

Altig, David, Lawrence Christiano, Martin Eichenbaum, and Jesper Linde, "Firm-Specific Capital, Nominal Rigidities and the Business Cycle," NBER Working Paper 11034, 2005.

Ball, Laurence and David Romer, "Real Rigidities and the Non-Neutrality of Money," *The Review of Economic Studies*, April 1990, *57* (2), 183–204.

Basu, Susanto, "Comment on Implications of State-Dependent Pricing for Dynamic Macroeconomic Models," *Journal of Monetary Economics*, January 2005, *52* (1), 243–247.

Bernanke, Ben S., Jean Boivin, and Piotr Eliasz, "Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach," NBER Working Paper 10220, 2004.

Bils, Mark and Peter J. Klenow, "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, October 2004, *112*, 947–985.

Blanchard, Olivier and Jordi Gali, "Real Wage Rigidities and the New Keynesian Model," mimeo, MIT, 2005.

⸺ **and Stanley Fischer**, *Lectures on Macroeconoics*, 1994. Cambridge: MIT Press.

Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, September 1983, *12*, 383–398.

Chari, V.V., Patrick J. Kehoe, and Ellen R. McGrattan, "Sticky Price Models of the Business Cycle," *Econometrica*, September 2000, *68*, 1151–1180.

Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans, "Monetary Policy Shocks: What Have We Learned and to What End?," *Handbook of Macroeconomics*, volume 1A, Elsevier: New York, 1999. John B. Taylor and Michael Woodford, eds.

Dhyne, Emmanuel, Luis J. Alvarez, Herve Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffman, Nicole Jonker, Patrick Lunnermann, Fabio Rumler, and Jouko Vilmunen, "Price Setting in the Euro Area: Some Stylized Facts from Individual Consumer Price Data," European Central Bank Working Paper 524, 2005.

Dotsey, Michael and Robert King, "Implications of State-Dependent Pricing for Dynamic Macroeconomic Models," *Journal of Monetary Economics*, January 2005, *52* (1), 213–242.

─── , ─── , **and Alexander Wolman**, "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics*, May 1999, *114*, 655–690.

Eichenbaum, Martin and Jonas D. M. Fisher, "Evaluating the Calvo Model of Stickly Prices," NBER Working Paper 10617, 2004.

Gertler, Mark and John Leahy, "A Phillips Curve with an Ss Foundation," mimeo, New York University, 2005.

Golosov, Mikhail and Robert E. Lucas Jr., "Menu Costs and Phillips Curves," NBER Working Paper 10187, 2003.

Gourieroux, Christian and Alain Monfort, *Simulation-Based Econometric Methods*, Oxford University Press, 1996.

Kimball, Miles, "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, November 1995, *27* (4), 1241–1277.

Klenow, Peter J. and Oleksiy Kryvtsov, "State Dependent or Time Dependent Pricing: Does It Matter for Recent U.S. Inflation?," NBER Working Paper 11043, 2005.

Krusell, Per and Anthony Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, October 1998, *106* (5), 867–96.

Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable, "The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains," *Quarterly Journal of Economics*, 1997, *112* (3), 791–825.

Romer, David H. and Christina D. Romer, "A New Measure of Monetary Shocks: Derivation and Implications," NBER Working Paper 9866, 2003.

Rotemberg, Julio, "Prices, Output and Hours: An Empirical Analysis Based on a Sticky Price Model," *Journal of Monetary Economics*, 1996, *37*, 505–533.

Tauchen, George, "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 1986, *20*, 177–181.

Taylor, John B., "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, February 1980, *88*, 1–23.

——, "Staggered Price and Wage Setting in Macroeconomics," *Handbook of Macroeconomics*, volume 1B, Elsevier: New York, 1999. John B. Taylor and Michael Woodford, eds.

Willis, Jonathan, "General Equilibrium of a Monetary Model with State-Dependent Pricing," unpublished paper, Federal Reserve Bank of Kansas City, 2000.

Woodford, Michael, *Interest and Prices*, Princeton University Press: Princeton, NJ, 2003.

Zbaracki, Mark J., Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen, "Managerial and Customer Dimensions of Cost of Changing Prices: Direct Evidence from Industrial Marketsi," *Review of Economics and Statistics*, 2004, *86* (2), 514–533.