Menu Costs and Asymmetric Price Adjustment

Tore Ellingsen*
Stockholm School of Economics and CEPR

Richard Friberg**
Stockholm School of Economics and CEPR

John Hassler***
Stockholm University and CEPR

June 22, 2006

Abstract

We study optimal price setting by a monopolist in an infinite horizon model with stochastic costs, moderate inflation, and costly price adjustment. For realistic parameters, chosen to replicate observed frequencies of price changes, the model fits numerically several empirical regularities. In particular, price reductions are larger but less frequent than price increases, and prices respond considerably faster to cost increases than to cost decreases. The associated kink in the steady state short-run Phillips curve implies that the output loss associated with a small negative inflation surprise is about twice as large as the output gain associated with a small positive inflation surprise.

Keywords: Asymmetric price adjustment, downward rigidity, menu costs, Phillips curve.

JEL Codes: D42, E31, E32.

*Department of Economics, Stockholm School of Economics, Box 6501, SE–113 83 Stockholm, Sweden. Email: tore.ellingsen@hhs.se.

**Department of Economics, Stockholm School of Economics, Box 6501, SE–113 83 Stockholm, Sweden. Email: richard.friberg@hhs.se.

***Institute for International Economic Studies, Stockholm University, SE–106 91 Stockholm.

Financial support from Torsten and Ragnar Söderbergs Stiftelse (Ellingsen) and the Hedelius travel grant (Friberg) is gratefully acknowledged. Part of this research was carried out during Friberg's visit to Princeton University, for whose hospitality he is deeply grateful. We thank Alan Blinder, Mark Gertler, Sam Pelzman and Elu von Thadden for helpful discussions and comments. The paper has also benefited from the comments of seminar participants at the Carlos III, the CIE Industrial Organization workshop in Århus 2005, IIES, Research Institute of Industrial Economics in Stockholm, Stockholm School of Economics, Stockholm University, Toulouse University, Uppsala University, and the 2004 ENTER jamboree (Barcelona). Remaining errors are ours.

1 Introduction

> "The odds are better than two to one that the price of a good will react faster to an increase in the price of an important input than to a decrease. This asymmetry is fairly labeled a "stylized fact." This fact poses a challenge to theory. The theory of markets is surely a bedrock of economics. But the evidence in this paper suggests that the theory is wrong, at least insofar as an asymmetric response to costs is not its general implication."

Sam Peltzman (2000, p 439), having studied how prices respond to costs in more than two hundred product markets, thus describes what appears to be an embarrassing failure of price theory. Asymmetric price rigidities have earlier been identified in a number of industry studies,[1] but after Peltzman's paper it is no longer possible to argue that asymmetric price adjustment is an exception. Asymmetry is the rule.

The asymmetric price rigidity implies a sizeable kink in the steady state short-run Phillips curve. The findings of Peltzman (2000) imply that the short-run output loss associated with a small negative inflation surprise is about twice as large as the output gain associated with a small positive inflation surprise of similar absolute size. Clearly, this is of general interest to macroeconomists and monetary policy makers.

Several macroeconomists, notably Tsiddon (1993) and Ball and Mankiw (1994), have proposed that asymmetric price adjustment to nominal shocks could be due to a combination of price adjustment (menu) costs and a drift in the desired nominal price. The drift could be caused by inflation or trends in demand or input costs. The argument is plain: Price adjustment costs imply that firms adjust prices infrequently. If inflation (drift) is high relative to other shocks, it is often not worth paying the adjustment cost in order to alter an excessively high nominal price, since future drift increases the desired nominal price automatically. An excessively low nominal price will be altered more often, because the expectation of future drift only exacerbates the need for upward price adjustment.

Building on the insights of Tsiddon and Ball and Mankiw, we investigate whether menu costs of empirically reasonable sizes can quantitatively explain the observed price adjustment patterns. Before explaining our model, let us point out three crucial challenges that it faces. First, there is an *adjustment lag puzzle*: In order to explain Peltzman's findings for the U. S. in the 1982-1996 period, the model needs to produce sizeable asymmetries in reaction time between a cost shock and the output price change even if inflation is low. Second, there is an *adjustment size puzzle*: Price reductions ought to be less frequent but larger than price increases, as seems to be a robust rule in economies with moderate inflation.[2] Previous models, that explain the adjustment lag puzzle with menu costs and inflation, predict the opposite to this, large price increases and small reductions.[3] Our model thus needs to disable or counterbalance the mechanism that creates larger price increases than price decreases in previous theoretical work. Finally, the average asymmetry of the price rigidity should be a decreasing function of the variance of the cost shocks, as found by Peltzman (2000). A calibrated version of our model tackles all three challenges while simultaneously matching the empirical frequencies of price adjustment (about two each year) with empirically reasonable menu costs (less than 1 % of yearly revenues).

[1] See for instance Borenstein, Cameron and Gilbert, (1997), and Borenstein and Shepard, (2002), on gasoline markets or Neumark and Sharpe, (1992), on bank deposit rates.

[2] For the U. S. Klenow and Kryvtsov (2005) find that about 58 percent of all regular retail price changes are increases. The fact that regular price decreases are relatively larger is evident from their Figure 4. Very similar findings for Europe are reported by Dhyne et al. (2005).

[3] Asymmetric price adjustment can also be related to strategic interaction, issues that we abstract from. Indeed, in a highly stylized setting Maskin and Tirole (1988) show how strategic interaction between firms producing an homogeneous good can lead to asymmetric price adjustment ("Edgeworth Cycles"). Here, a number of small price decreases are followed by a large price increase. Noel (2006) documents this pattern of price adjustment on the Toronto retail gasoline market. The evidence reported in footnote 2 however points to that a more typical pattern is that price increases are smaller and more frequent than price decreases.

Following the lead of Barro (1972), we analyze the behavior of an infinitely lived monopolistically competitive firm facing independent and identically distributed random shocks and having to pay a menu cost to change its output price. Two essential differences are that our firm faces a demand with constant price elasticity, instead of Barro's linear demand, and that we allow anticipated trends in costs and in the price level.[4] The primary role of the constant price elasticity assumption is that it generates an asymmetric cost of price misalignments.[5] With linear demand and constant marginal costs, our model would tend to reproduce Ball and Mankiw's falsified prediction of relatively large price increases when inflation is positive. Linear demand and costs entails symmetric losses from positive and negative price deviations, whereas convex demand (and costs) can entail relatively greater losses from negative price deviations. Since it takes a smaller negative deviation before the price is adjusted upwards, upward price changes are then going to be both more frequent and smaller, at least for low levels of inflation. The reduced form assumption of a quadratic loss functions, common to Tsiddon (1993), Ball and Mankiw (1994), Caplin and Leahy (1997), and to much other macroeconomic literature, essentially amounts to assuming linear demand and constant marginal costs, thereby implying that price increases must be larger than price decreases at positive rates of inflation.

Presumably, the reason why few other papers in the menu cost literature have considered the constant elasticity case is that it is known to be complicated.[6] However, Danziger (1999) showed that several complications can be avoided by reformulating the problem. One of Danziger's innovations was to replace the standard assumption of constant menu costs with the assumption that menu costs are scale-dependent. In the famous restaurant example, the standard assumption corresponds to the polar case in which the menu is posted on the wall for all customers to see. If this assumption were true, we should observe that successful firms should change their prices more frequently as they grow larger. To our knowledge, no evidence has been presented in favor of the constant menu cost assumption; instead it is justified on the basis of tractability. In Danziger's model (and in ours), constant menu costs would impose substantial computational difficulties, because inaction bands become state-dependent. Instead, Danziger assumed that price adjustment costs constitute a constant fraction of demand at the new price - thereby making inaction bands constant. In the restaurant example, Danziger's specification corresponds roughly to a case in which there are menus on each table. A successful restaurant expands by increasing the number of tables and thus also the number of menus to print when prices are changed. Apart from technical convenience, we think that Danziger's assumption provides the most natural benchmark - it depicts the case in which the frequency of price changes is independent of firm size.

Under our menu cost assumption, the inaction band has constant logarithmic width, and we can analytically characterize the ergodic distribution of price deviations in the no–drift case. Moreover, the numerical algorithms converge quite quickly even when inflation is positive. In this way, our work complements that of Golosov and Lucas (2003) whose model also assumes constant price elasticity, but assumes that menu cost are constant.[7] In order to solve the model despite non-constant inaction bands, Golosov and Lucas assume that the stochastic process is mean-reverting.[8]

[4] Sheshinsky and Weiss (1977) were the first to study deterministic trends in menu costs models. Early papers combining trends and idiosynchratic shocks include Frenkel and Jovanovic (1980) and Sheshinsky and Weiss (1983).

[5] The fact that constant demand elasticity generates larger static losses from negative price deviations than from positive ones has previously been pointed out by Kuran (1983, 1986), Naish (1986) and Konieczny (1990) among others. In some of these papers, the associated greater frequency of price increases or the associated higher average price level is interpreted as asymmetric *average* price rigidity. That is misleading, however. The average price response is the product of the probability of a price change and the size of that change, and a small probability of a large change may yield the same average as a large probability of a small change. This insight echoes Caplin and Spulber's (1987) demonstration that the existence of menu costs and firm level price rigidity need not lead aggregate price rigidity.

[6] Previous theoretical work on the inflation/output trade–off that assumes constant price elasticity has thus limited its attention to non–stochastic models and to the case of very small menu costs; see for example Bènabou and Konieczny (1994) and the references therein.

[7] Golosov and Lucas do not focus on pricing asymmetries.

[8] Mean-reversion could be a more accurate description of many input markets than is our random walk assumption, but the

Our paper is mainly positive, but it may have normative implications for monetary policy. In a recent paper that also (essentially) utilizes Danziger's menu cost formulation, Gertler and Leahy (2005) show that it is possible to generate an analytically tractable Phillips curve in a general equilibrium model with state-dependent pricing. Gertler and Leahy (2005) consider a quadratic approximation to the profit function about a zero drift steady state. The combination of quadratic losses and no drift generates symmetric inaction bands. That is, the cost of tractability in Gertler and Leahy's model is the neglect of asymmetric responses. Our analysis is complementary. We sacrifice an analysis of general equilibrium effects in order to gain precision with respect to the firms' pricing problem. Our findings suggest that the effects of short-run monetary policy are generically asymmetric when third-order effects are taken account of, and that the short-run Phillips curve may have a sizeable kink. In steady state, the output loss associated with unexpectedly low inflation could well be twice as large as the output gain associated with unexpectedly large inflation.

Our paper is laid out in the following way. We start out in Section 2 with a simple exposition of the firm's static pricing problem. In Section 3, we then go on to analyze the fully dynamic and stochastic pricing problem, illustrating the magnitude of asymmetric price rigidity in realistic numerical examples. Section 4 analyzes aggregation and Section 5 concludes.

2 The Static Price Adjustment Problem

Consider a firm producing a single output using a single input. Production costs are denoted $C(q, \psi)$, where q is the output quantity and ψ is the (random) input price. Demand for the firm's output is given by $D(p)$, where p denotes the output price. The firm's profit is

$$\Pi = pD(p) - C(D(p)).$$

Figure 1 displays the profit function in the standard case of constant marginal costs $C'(q)$ and constant elasticity of demand, ν. The figure is drawn for $C'(q) = 1$ and $\nu = 2$.

As Figure 1 shows, the firm's loss from setting too high a price is smaller than the loss from setting too low a price. The difference is small when deviations from the optimal price are small, but can be large when the deviations grow. In an environment with non-trending price level shocks, such an asymmetric profit function tends to generate more and smaller price increases than price decreases. Previous work by Kuran (1983, 1986), Naish (1986), Konieczny (1990) and in particular Benabou and Konieczny (1994) has linked the skewness of the profit function to the underlying cost and demand functions: Let p^* be the optimal static price, and assume that demand is decreasing in p and costs are increasing in output. Compare the profit from a price that is x dollars too high, $\Pi(p^* + x)$, to the profit from a price that is x dollars too low, $\Pi(p^* - x)$. Performing Taylor expansions and neglecting terms of order 4 and higher, the difference is approximately $\Pi'''(p^*)x^3/3$. This difference is positive, as desired, if $\Pi'''(p^*) > 0$, or equivalently if

$$\begin{aligned} &3D''(p^* + x) + p^* D'''(p^* + x) - C'''(D(p^*))(D'(p^*))^3 \\ &-3C''(D(p^*))D''(p^*)D'(p^*) - C'(D(p^*))D'''(p^*) \qquad > 0. \end{aligned}$$

Clearly, the expression is zero if both C and D are linear and if D is linear and C is quadratic. Strict convexity of demand and/or cost functions make the first and the penultimate term positive. Note that the asymmetry caused by convexity can be overturned if $D'''(p^*)$ and/or $C'''(D(p^*))$ are negative. However, for the convex demand and cost functions that are normally used, it is straightforward to check that this case does not arise.

We conclude; if the elasticity of demand is constant and the cost function is a polynomial with positive coefficients, then the profit associated with an excessively high price is larger than that associated with an

issue is not settled. For example, Schwartz and Smith (2000) find that crude-oil prices tend to revert to long-run equilibrium levels which themselves change randomly over time.

excessively low price if the deviations from the profit maximizing price is of the same absolute magnitude in the two cases. As we will show below, these third-order effects, that are disregarded under the often used practice of using a second-order approximation of the profit function, are of quantitative importance for individual and aggregate pricing behavior with menu costs.

If the nominal price deviates from the desired level due to cost shocks rather than price level shocks (as we shall assume below) the above approximation remains valid for the case in which the profit maximizing mark-up is constant. An upward cost shock then generates the same absolute price deviation as the corresponding downward cost shock. For example, an upward cost change generates a larger need to change the price than does a downward cost change if the firm faces iso-elastic demand and constant marginal costs.

3 Dynamic pricing

A static analysis of the costs of mispricing is an incomplete guide to how firms will make costly price adjustments in anticipation of random shocks to production costs. We need to endogenize the decision *when to* adjust prices to analyze the dynamic relation between cost shocks and price responses. Obviously, this requires a dynamic model. We start with an analysis of the dynamic pricing problem of an individual firm, identifying circumstances under which negative and positive price adjustments are of different size. In the second subsection, we analyze asymmetry in the immediate average price response, a measure of the adjustment lag asymmetry.

3.1 The absolute size of price adjustments

Let output be produced by a single input, and let the production function take the form

$$q = (1+\theta)\, x^{\frac{1}{1+\theta}},$$

where q and x are output and input quantities respectively and θ measures the degree of decreasing returns to scale. The nominal input price is stochastic and is denoted Ψ_t. Thus, the nominal cost function is also stochastic, with the cost at time t being

$$C_t = \frac{\Psi_t}{1+\theta} q_t^{1+\theta}. \tag{1}$$

We assume that cost shocks are multiplicative. The possibility of input price inflation and/or productivity growth is captured by including a constant drift rate γ. In other words, Ψ_t follows a geometric random walk with drift, written

$$\ln \Psi_{t+1} = \ln \Psi_t + \gamma + \tilde{\gamma}_{t+1},$$

where $\tilde{\gamma}_{t+1}$ is i.i.d. normal with zero mean.

Let demand be given by the constant elasticity function

$$D(p_t) = \left(\frac{p_t}{P_t}\right)^{-\nu}, \tag{2}$$

where p_t is the firm's current price and P_t is the current price level. We allow the price level (nominal demand) to move stochastically, that is[9]

$$\ln P_{t+1} = \ln P_t + \pi + \tilde{\pi}_{t+1}.$$

[9] Adding a real shock to demand would not be diificult to analyze. In particular, if $\theta = 0$ and the real demand shock is multiplicative, it should only add noice to real profits, without changing optimal pricing.

where π is the constant expected inflation rate and $\tilde{\pi}_{t+1}$ is the inflation surprise between t and $t+1$, assumed to be i.i.d over time with zero expectation.

Given the real cost $\psi_t \equiv \frac{\Psi_t}{P_t}$ and the chosen price p_t, the firm's current real profit is

$$\Pi_t = \Pi\left(\frac{p_t}{P_t};\psi_t\right) = \frac{p_t}{P_t}D(p_t) - \frac{\psi_t}{1+\theta}D(p_t)^{1+\theta}, \tag{3}$$

which is maximized by the frictionless price

$$p_f(\psi_t, P_t) \equiv \arg\max_{p_t} \Pi\left(\frac{p_t}{P_t};\Psi_t\right) = P_t\left[\left(\frac{\nu}{\nu-1}\right)\psi_t\right]^{\frac{1}{1+\nu\theta}}. \tag{4}$$

Two things should be noted here. First, under decreasing returns to scale ($\theta > 0$), the frictionless price is a concave function of input prices where the degree of concavity increases in $\nu\theta$. By Jensen's inequality, this implies that the expected change in output prices is lower than the expected change in input prices. Second, when $\theta = 0$, the frictionless profit is a convex function of ψ, implying that higher volatility of input prices increases expected profits.[10]

With menu costs, the optimal pricing policy is to allow some upward and downward deviations from the frictionless price. The maximum deviations that should be allowed under the optimal policy generally depend on all state variables, i.e., costs, the price level and on the cost of changing the price if this is non-constant. As already note, we will focus on a special case entailing an optimal inaction band of constant logarithmic width. In this case, we make assumptions such that the maximum percentage deviation allowed above and below the frictionless price is constant over time. As we will show, constant bandwidth requires a menu cost that is decreasing in the cost variable ψ_t. The intuition is clear: When marginal costs are high, frictionless profits are low and the gain from adjusting the price is small. Thus, inaction bands can be constant only if the menu cost is positively related to frictionless profits. Given a constant demand elasticity, this is implies that menu costs are increasing in the output quantity.

If the optimal inaction band has constant logarithmic width, the optimal policy can be completely characterized by three numbers, μ_u, μ_l, μ_r, and the decision rule:

- If $p_t \geq e^{\mu_u} p_f(\psi_t, P_t)$, adjust the price downwards to $e^{\mu_r} p_f(\psi_t, P_t)$.
- If $p_t \leq e^{\mu_l} p_f(\psi_t, P_t)$, adjust the price upwards to $e^{\mu_r} p_f(\psi_t, P_t)$.
- Otherwise, keep the price constant.

Let us now derive conditions such that the inaction band is constant. In the beginning of each period, the firm observes the current cost parameter and decides whether to change the price or not. If the firm changes the price, a menu cost M_t is paid and the new price immediately applies. The Bellman equation for this problem is

$$\begin{aligned} W\left(\frac{p_t}{P_t},\psi_t\right) &= \max\left\{\Pi(p_t,\psi_t,P_t) + \frac{E\left(W\left(\frac{p_t}{P_{t+1}},\psi_{t+1}\right)\right)}{1+r},\right. \\ &\qquad \left.\max_{\acute{p}}\left\{\Pi(\acute{p}) - M_t + \frac{E\left(W\left(\frac{\acute{p}}{P_{t+1}},\psi_{t+1}\right)\right)}{1+r}\right\}\right\}. \end{aligned} \tag{5}$$

[10] The second derivative of frictionless profits when $\theta = 0$ is $\psi^{-(1+\nu)}\nu^{1-\nu}(\nu-1)^{\nu}$. When input prices are low and ν large. This can be a very large number.

Define the normalized (output) price at time t,

$$\mu_t \equiv p_t/p_f(\psi_t, P_t),$$

and note that $\mu_t - 1$ is the relative misalignment of the actual price from the frictionless target. In order for the inaction band to be constant, the cost ψ_t and price level P_t should not appear in the Bellman equation. Thus, we seek assumptions such that μ is the only relevant state variable.

The change in log of the real cost can be written

$$\Delta \ln \psi_{t+1} \equiv \ln \psi_{t+1} - \ln \psi_t = \gamma - \pi + \tilde{\gamma}_{t+1} - \tilde{\pi}_{t+1}.$$

The normalized price hence moves according to the equation

$$\Delta \ln \mu_{t+1} \equiv \ln \mu_{t+1} - \ln \mu_t = \Delta \ln p_t - (\pi + \tilde{\pi}_{t+1}) - \frac{\Delta \ln \psi_{t+1}}{1+\nu\theta}, \tag{6}$$

and $\ln \mu_t$ is a random walk a drift given by $-\pi - \frac{\gamma-\pi}{1+\nu\theta}$ whenever no price changes are undertaken. Assuming that real cost shocks and nominal demand shocks shocks are independent, the variance of $\Delta \ln \mu_{t+1}$, denoted $\sigma^2_{\Delta\mu}$, is a weighted sum of their respective variances, i.e., $\frac{\sigma^2_{\Delta\psi}}{1+\nu\theta} + \sigma^2_{\tilde{\pi}}$.[11]

We denote the change in $\ln \mu_{t+1}$ when no output price changes are undertaken by

$$\Delta \ln \tilde{\mu}_{t+1} \equiv -(\pi + \tilde{\pi}_{t+1}) - \frac{\Delta \ln \psi_{t+1}}{1+\nu\theta},$$

noting that realized inflation, $(\pi + \tilde{\pi}_{t+1})$, and real input price increases, $\Delta \ln \psi_{t+1}$, reduce the normalized price, i.e., making the price misalignment more negative.

After a few manipulations, we can write current real profit as a separable function of the real input cost ψ_t and the normalized price μ_t,

$$\Pi\left(\mu_t p_f(\psi_t, P_t); \psi_t, P_t\right) = f(\mu_t)\, \psi_t^{\frac{1-\nu}{1+\nu\theta}}. \tag{7}$$

For a properly defined menu cost, we tentatively conjecture that the value function can be similarly separated,

$$W\left(\frac{p_t}{P_t}, \psi_t\right) = V(\mu_t)\, \psi_t^{\frac{1-\nu}{1+\nu\theta}}.$$

Using this conjecture, denoting

$$\tilde{R}\left(\frac{\psi_{t+1}}{\psi_t}\right) = \left(\frac{\psi_{t+1}}{\psi_t}\right)^{\frac{1-\nu}{1+\nu\theta}} \frac{1}{1+r}$$

and dividing both sides of the Bellman-equation by $\psi_t^{\frac{1-\nu}{1+\nu\theta}}$ yields

$$\begin{aligned} V(\mu_t) &= \max\left\{ f(\mu_t) + E\left(V\left(\mu_t e^{\Delta \ln \mu_{t+1}}\right)\tilde{R}\right), \right. \\ & \left. \max_{\mu'} \left\{ f(\mu') - \frac{M_t}{\psi_t^{\frac{1-\nu}{1+\nu\theta}}} + E\left(V\left(\mu' e^{\Delta \ln \mu_{t+1}}\right)\tilde{R}\right)\right\}\right\}. \end{aligned} \tag{8}$$

To get rid of the dependence on ψ_t and P_t, we finally assume that the menu cost M_t is proportional to $\psi_t^{\frac{1-\nu}{1+\nu\theta}}$.[12]

[11] It would be straightforward to relax the independence assumption.

[12] A similar assumption is made in Gertler and Leahy (2005).

Since the real frictionless profit is

$$\Pi_f(\psi_t) = \frac{1+\nu\theta}{\nu(1+\theta)} \left(\frac{\nu}{\nu-1}\right)^{\frac{1-\nu}{1+\nu\theta}} \psi_t^{\frac{1-\nu}{1+\nu\theta}},$$

another way to state the condition is that the real menu cost is a constant fraction of the real frictionless profit,

$$M_t = m\Pi_f(\psi_t). \tag{9}$$

In the parameterization below, we will express the frictionless profit in yearly streams, so that m denotes the menu cost in terms of the yearly frictionless profit given current ψ. Despite its ad hoc nature, this assumption might be an improvement on the standard assumption of constant menu costs. Evidence for U. S. retailers, reported by Levy et al. (2005), indicates that price adjustment costs are quite rapidly increasing in store traffic, a regularity that our model allows.[13]

In equation (8), $\tilde{R}$ acts as a stochastic discount factor. The first term in $\tilde{R}$ is the growth rate of profits (and the value function) given μ, and the second is the standard discount factor. If $E\tilde{R} < 1$, equation (8) is a contraction mapping with a unique solution that can be found by standard numerical methods.[14] The optimal policy therefore is time-invariant and can be characterized by the three numbers μ_u, μ_l, μ_r as described above. To find these we iterate on a discretized version of (8) until it converges.[15] The resulting values for μ_u, μ_l, μ_r determine the relative size of price adjustments. Accordingly, the degree of asymmetry in price adjustments can be measured by $(\ln\mu_u - \ln\mu_r)/(\ln\mu_r - \ln\mu_l)$.

The stochastic discount factor $\tilde{R}$ inherits the profit function's convexity in ψ_{t+1}. Therefore, if $E\frac{\psi_{t+1}}{\psi_t} = 1$ and $r > 0$, the stochastic discount factor may be larger than unity in expectation. If so, the value function is explosive. The reason is that even if input price increases and decreases are equally likely and of the same (log) size, convexity of the profit function implies that expected profits are increasing. If the rate of expected growth is larger than r, the value function is not well defined.

3.2 Asymmetries in the band of inaction when there is no drift

We now turn to a quantitative investigation of the inaction band when there is no drift in the frictionless price ($\pi = \gamma = 0$). Our main result in this section is that the inaction band is asymmetric. Since it takes less of an upward cost change to induce a price change, and firms aim for a constant markup, price increases will be smaller and more frequent than price decreases. The direction of the asymmetry is the same in all our simulations (as long as there is no drift in the frictionless price), thus confirming that the logic from the static setup carries over to dynamic price setting with menu costs. The strength of the asymmetry is determined by the elasticities of demand and costs, the size of menu costs as well as the volatility of shocks.

Two preliminary targets for our numerical simulations of the inaction band is to match the average frequency of price changes in the United States over the last decade, as reported by Bils and Klenow (2004), and to match the relative frequency of (ordinary) price increases, as reported by Klenow and Kryvtsov (2005). In other words, we want to generate roughly two ordinary price changes every year out of which about 55-60% are increases.

As discussed in the previous section, we need $E\tilde{R} < 1$ to have a well-defined value function. This condition is not satisfied if $\frac{\nu-1}{1+\nu\theta}$ is large and the volatility of the real cost shock is high, the reason being

[13] To be precise, the model implies that M is proportional to $q_f^{(2+\theta\nu)/(1+\theta\nu)}$, where q_f is the "frictionless quantity" that is associated with the price p_f. It would arguably be more natural to have the cost depend on the actual quantity sold, rather than the hypothetical frictionless quantity. However, such an assumption complicates the analysis considerably. In this case, if the the firm decides to change its price today, it must take into account that the price it sets affects the menu cost paid *next time* an adjustment is undertaken since it affects the volume sold at that point in time.

[14] See Appendix 1.

[15] See Appendix 3 for details.

that the possibility of a series of negative cost shocks is combination with highly elastic demand makes expected discounted profits unbounded. While this might be an unrealistic feature of the long-run behavior of the model, we believe it has negligible impacts on its implications for the short-run pricing behavior of the firm. In what follows, we will assume that the shock is nominal, implying that we can analyze also the case of high demand elasticities.[16] For values of $\nu < 4$, we have compared the results under real versus nominal shocks finding them being almost identical. An inspection of (8) sheds light on this. As we see, nominal and real shocks both shifts the normalized price, and when $\theta = 0$, by the same amount. The only difference is that a real shock also affects the discount factor $\tilde{R}$. When $\theta > 0$, a real shock has a smaller impact on the normalized price but is otherwise very similar to a nominal shock. The problem of controlling the price-misalignment is almost identical under the two assumptions about the shock. We are therefore confident that focusing on nominal shocks is innocuous.

In our first numerical exercise, we set the standard deviation of input prices to 4% per month and the interest rate to 10% per year.[17] The input price volatility of 4% corresponds exactly to the empirical volatility found by Peltzman (2000). That number is definitely on the high side, however, since it includes temporary price changes as well as permanent ones. We start with a conservative base-line case with $\nu = 2$, $\theta = 0$ and $m = .1\%$, where in particular the menu costs are very low. The associated value function $V(\mu)$ depicted in the upper left panel of Figure 5 and the characteristics of the optimal policy is given in Table 1. The value function is asymmetric, $(\ln \mu_u - \ln \mu_r)/(\ln \mu_r - \ln \mu_l)$ is 1.034, but the degree of asymmetry is low under this parameterization. While qualitatively right, our conservative benchmark does not match quantitatively the relatively large price declines in the data. Simulation shows that the firm on average waits 6.7 months between price changes, which is about right, but that price decreases are almost as likely as increases.

Given our analysis of the static case, one might conjecture that, *ceteris paribus*, a higher demand elasticity should make the loss of a misaligned price more asymmetric and thus causing more asymmetric price adjustments, i.e., a larger value of $(\mu_u - \mu_r)/(\mu_r - \mu_l)$. This conjecture is premature. The reason is that a higher elasticity entails a narrower inaction band, since the cost of being far from the frictionless price increases. The reduction in the band-width tends to reduce the asymmetry, since around $\mu = 1$ the profit function is well approximated by a quadratic second order Taylor approximation. Increasing the demand elasticity ν to 5, reduces the band considerably, without producing any significant asymmetry, as seen in the upper middle panel. In this case, the firm allows price deviation of only around 5.5%, corresponding to 2 months on average between price changes, and the degree of asymmetry is negligible. The value function is depicted in the upper right panel of Figure 5.

If we now increase the menu costs, the band widens and becomes more asymmetric, as seen in the lower left panel (still with $\nu = 5$). In this case, we have set $m = 1\%$. The asymmetry is non-negligible but not large $(\ln \mu_u - \ln \mu_r)/(\ln \mu_r - \ln \mu_l) = 1.13$, implying that price increases are slightly more likely (53 vs. 47%) and the average duration between price changes is 6.6 months. With $\nu = 10$ and a menu cost of 4% of the yearly frictionless profit, $(\ln \mu_u - \ln \mu_r)/(\ln \mu_r - \ln \mu_l) = 1.25$, and the average duration between price changes is 6.5 months. Of all price changes, 56% are price increases, in line with the data reported above. While a menu cost of 4% the frictionless profit may seem on the high side, observe that in the example it corresponds to only 0.4% of the value of sales. Thus, the yearly menu cost is $0.4\% \cdot 12/6.5 = 0.74\%$ of sales, which is well in line with empirical estimates. For example, Levy et al. (1997) find that total menu costs for large U.S. supermarket chains are around 0.7% of yearly sales.

If we simultaneously allow large demand elasticities, increasing marginal costs and high menu costs, the

[16] Of course, we could also have assumed that there is a real and a nominal component and that the real component is sufficiently small to imply

$$\frac{e^{\frac{1-\nu}{1+\nu\theta}\Delta \ln \psi_{t+1}}}{1+r} < 1.$$

[17] We solve the numerical problem where the state space is multiples of 0.05% around 1. This also implies that the volatility of the shocks has to be chosen from a discrete set. In the simulations, the actual standard deviation of the shocks is 4.11%.

asymmetry can be even more sizeable. For example, with $\nu = 10$, $\theta = 3$ and $m = 7.5\%$, $(\ln \mu_u - \ln \mu_r)/(\ln \mu_r - \ln \mu_l) = 1.52$. In this case, price increases are 50% more common than price decreases (60% of price changes are increases) and the average duration between price changes is close to six months. The maximum negative deviation is in this case 5.8% and the maximum positive is 14%.

Table 1 also shows that when the price is changed, it is set above the frictionless price $\ln \mu_r > 0$. Adjusting to a relatively high price partly compensates for the fact that too low prices are more costly than are too high prices. The compensation is rising in demand and cost elasticities because the asymmetry of the profit function increases in these parameters.

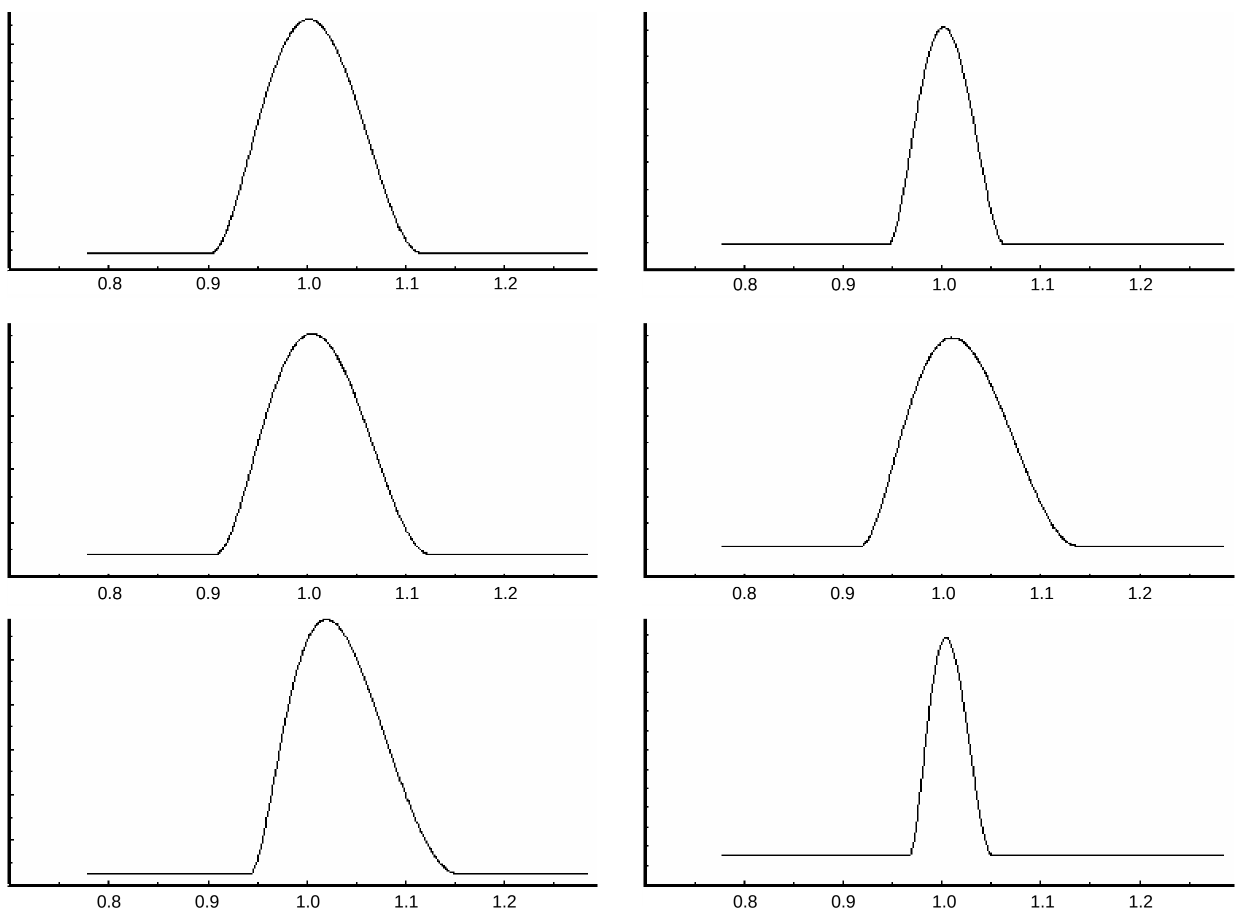

Figure 1: $V(\mu)$ for the six cases described in Table 1.

Before analyzing the effect of drift, we want to show that input price volatility is quantitatively important. With little uncertainty, the option value of waiting to adjust the price is low. This leads to a narrow inaction range, and consequently, little asymmetry. To see this, take the case where we found the most extreme asymmetry, i.e., $\nu = 10, \theta = 3$ and $m = 7.5\%$. If we reduce the volatility of the price shock to 1%, the band shrinks to a total of 8%, reducing the asymmetry from 1.52 to 1.23. Furthermore, the relatively narrow band, average duration between periods is too long, close to 21 months. Reducing the menu cost to get more reasonable duration between price changes would reduce the asymmetry further.

	$\ln \mu_l$	$\ln \mu_r$	$\ln \mu_u$	$\frac{\ln \mu_u - \ln \mu_r}{\ln \mu_r - \ln \mu_l}$	Average inaction	Price increases
Case 1	−0.1000	0.0015	0.1065	1.034	6.66 months	51.3%
Case 2	−0.0540	0.0025	0.0605	1.027	2.05 months	50.6%
Case 3	−0.0940	0.0040	0.1145	1.128	6.62 months	53.0%
Case 4	−0.0840	0.0095	0.1260	1.246	6.54 months	55.9%
Case 5	−0.0575	0.0210	0.1400	1.516	5.87 months	59.7%
Case 6	−0.033	0.004	0.0495	1.230	20.8 months	54.3%

Parameters	ν	θ	m	$\sigma_{\Delta\mu}$	Drift
Case 1	2	0	0.1%	4%	0
Case 2	5	0	0.1%	4%	0
Case 3	5	0	1%	4%	0
Case 4	10	0	4%	4%	0
Case 5	10	3	7.5%	4%	0
Case 6	10	3	7.5%	1%	0

Standard deviation of shocks is measured per month. The interest rate is $r = 10\%$ per year.

TABLE 1: Optimal pricing with no drift.

3.3 Adding drift

The numerical results above have been derived under the assumption of no drift. Adding negative drift to the price deviation - by introducing positive inflation of the price level or upward drift in costs - moves the peak of the value function to the right, reinforcing the tendency to set the price above the frictionless optimum. Since the price deviation is expected to fall over time, the price should be set higher when it is changed. Although price increases tend to be somewhat larger as a result, moderate levels of inflation have only a tiny impact on the degree of asymmetry, as comparison of Table 2 to Table 1 indicates.[18]

	$\ln \mu_l$	$\ln \mu_r$	$\ln \mu_u$	$\frac{\ln \mu_u - \ln \mu_r}{\ln \mu_r - \ln \mu_l}$	Average inaction	Price increases
Case 7	−0.0980	0.0040	0.1090	1.029	6.64 months	58.7%
Case 8	−0.0940	0.0150	0.1155	0.922	6.40 months	74.1%
Case 9	−0.0565	0.0225	0.1440	1.538	5.91 months	68.6%
Case 10	−0.0545	0.0330	0.1555	1.400	5.74 months	80.8%

Parameters	ν	θ	m	$\sigma_{\Delta\mu}$	Drift
Case 7	2	0	0.1%	4%	3%
Case 8	2	0	0.1%	4%	10%
Case 9	10	3	7.5%	4%	3%
Case 10	10	3	7.5%	4 %	10%

[18]Taking the moderate benchmark, with $\nu = 2, \theta = 0$ and $m = 0.1\%$ and adding a drift of 3% reduces the asymmetry from 1.034 to 1.029. Higher levels of inflation decreases the asymmetry more noticeably. With a drift of 10%, the return point is set 1.6% above the frictionless price, and price increases are somewhat larger than price decreases (11% vs. 10% of the frictionless price). However, with somewhat higher demand and/or cost elasticities, even ten percent inflation is not enough to make price increases larger than price decreases.

Standard deviation of shocks is measured per month. The interest rate is $r = 10\%$ per year.

Table 2: Optimal pricing with drift

Basically, our numerical experiments suggest that if one is willing to set the elasticity of demand equal to 5 or higher, the model offers a quantitatively good resolution the adjustment size puzzle. (Golosov and Lucas, 2003, use a demand elasticity of 7.)

It now remains to investigate whether the model also resolves the adjustment lag puzzle.

4 Average price dynamics

When the distribution of price deviations is stationary, the average response to a small cost shock is given by the probability that the price will change multiplied by the magnitude of the price change. To compute the average responses, the hardest step is to derive the stationary distribution of price deviations.

It is helpful to model the stochastic process of individual shocks as being discrete and binomial; when the period length becomes small (infinitesimal), this is without loss of generality.[19] Thus, we assume

$$\Delta \ln \tilde{\mu}_{t+1} = \begin{cases} \varepsilon \text{with probability } 1/2 + \delta; \\ -\varepsilon \text{with probability } 1/2 - \delta, \end{cases}$$

where

$$\begin{aligned} \varepsilon &= \left(\pi^2 + \sigma_\pi^2\right)^{\frac{1}{2}}, \\ \delta &= \frac{1}{2}\frac{\pi}{\varepsilon}. \end{aligned} \tag{10}$$

Given the optimally chosen triggers μ_l, μ_u and the return point μ_r, we now note that μ, will take a discrete number of values

$$\mu \in \mathcal{M} \equiv \{\mu_r + s\varepsilon, s \in \{-L, -L+1, ..., 0, 1, 2, ...U\}\},$$

where L is the largest number of accumulated positive price shocks accepted before adjustment, i.e., $\mu_r - (L+1)\varepsilon < \mu_l \leq \mu_r - L\varepsilon$ and similarly for U. Let us abuse notation by letting $\mu_s \equiv \mu_r + s\varepsilon$ so that the subscript s denotes the number of steps μ is away from the return point μ_r. Letting $d(\mu_s)$ denote the unconditional density of firms with deviation μ_s, we have that for $s \neq 0, U, L$,

$$d(\mu_s) = (1/2 - \delta)\, d\left(\mu_{s-1}\right) + (1/2 + \delta)\, d\left(\mu_{s+1}\right).$$

The roots of this homogeneous difference equation are $\rho_0 = 1$ and $\rho_1 = \frac{1-2\delta}{1+2\delta}$. Consequently, its general solution is

$$d(\mu_s) = \tilde{a}_i + \tilde{b}_i \rho_1^s, \tag{11}$$

when $\delta \neq 0$ and

$$d(\mu_s) = \tilde{a}_i + \tilde{b}_i s, \tag{12}$$

otherwise. The constants $\tilde{a}_i$ and $\tilde{b}_i$ are determined from end-point conditions and the fact the the density should sum to unity. Observe that the solution in the lower range μ_L to μ_0 is generally different from the solution in the upper range μ_0 to μ_U. Denote the constants in the two ranges by $\{a_l, b_l\}$ and $\{a_u, b_u\}$ respectively.

[19] As in the previous section, we set the period lengt to 1/1000 years.

To find end-conditions we first use the fact that the flows to the end points μ_u and μ_l are given by

$$\begin{aligned} d\left(\mu_U\right) &= (1/2-\delta) d\left(\mu_{U-1}\right), \\ d\left(\mu_{-L}\right) &= (1/2+\delta) d\left(\mu_{-L+1}\right). \end{aligned} \tag{13}$$

Second, we note that the inflow to the interior positive and negative regions equal the outflows, giving the equations

$$(1/2-\delta) d\left(\mu_0\right)=(1/2+\delta) d\left(\mu_1\right)+(1/2-\delta) d\left(\mu_U\right), \tag{14}$$

and

$$(1/2+\delta) d\left(\mu_0\right)=(1/2-\delta) d\left(\mu_{-1}\right)+(1/2+\delta) d\left(\mu_{-L}\right). \tag{15}$$

Finally, normalizing the total mass of firms to unity, we have

$$\sum_{s=-L}^{U} d\left(\mu_s\right)=1. \tag{16}$$

4.1 Average price dynamics without drift

Without drift, the problem has a solution on closed form. To find it, set $s=0$ in equation (11) to have $a_l+b_l \cdot 0=d\left(\mu_0\right)=a_u+b_u \cdot 0$, or equivalently

$$a_l=a_u=a. \tag{17}$$

Setting $\delta=0$ in (14) and (15), inserting the solution from (11) in them and in (16) provides a linear three-equation system in three unknowns. The system's solution is

$$\begin{aligned} a &= \frac{2}{2+U+L}, \\ b_l &= \frac{2}{(1+L)(2+U+L)}, \\ b_u &= -\frac{2}{(1+U)(2+U+L)}. \end{aligned}$$

Using (13), we conclude that

$$\frac{d(\mu_U)}{d(\mu_{-L})}=\frac{d(\mu_{U-1})}{d(\mu_{-(L-1)})}=\frac{a+b_u(U-1)}{a-b_l(L-1)}=\frac{1+L}{1+U},$$

immediately implying that when $\delta=0$ and the period length t converges to zero, the ratio of the end point densities $d(\mu_u)/d(\mu_l)$ converges to $(\mu_r-\mu_l)/(\mu_u-\mu_r)$. This generalizes a result by Tsiddon (1993), who considered the case of symmetric bands. Since the densities are linear with a common intercept c_0, the relation holds not only for the end point densities but also for the densities of any pair of equally large intervals starting at the end points. This should come as no surprise: The ratio of price decreases to price increases equals $(\mu_u-\mu_r)/(\mu_r-\mu_l)$ and in an environment without drift, the expected size of the price response to an upward cost shock, $d(\mu_l)(\mu_r-\mu_l)$, must equal the expected size of the price response to a downward cost shock, $d(\mu_u)(\mu_u-\mu_r)$.

In Figure 6, we depict the result graphically. The area of the two shaded triangles to the left and right in the figure, represents the probability that a positive and negative shock, respectively, induce a price adjustment. Clearly, the bases of the triangles are equal at ε, while the heights are inversely proportional to

$\mu_r - \mu_l$ and $\mu_u - \mu_r$ respectively. Thus, the probability of an adjustment is inversely related to the size of the adjustment and the product of the probability of adjustment and its size is thus the same for negative and positive adjustments.

When there is drift in the frictionless price, the argument no longer applies, because a sequence of shocks without drift implies a departure from the ergodic distribution of price deviations. – Only when the shocks reflect the underlying drift will the distribution of price deviations be the same.

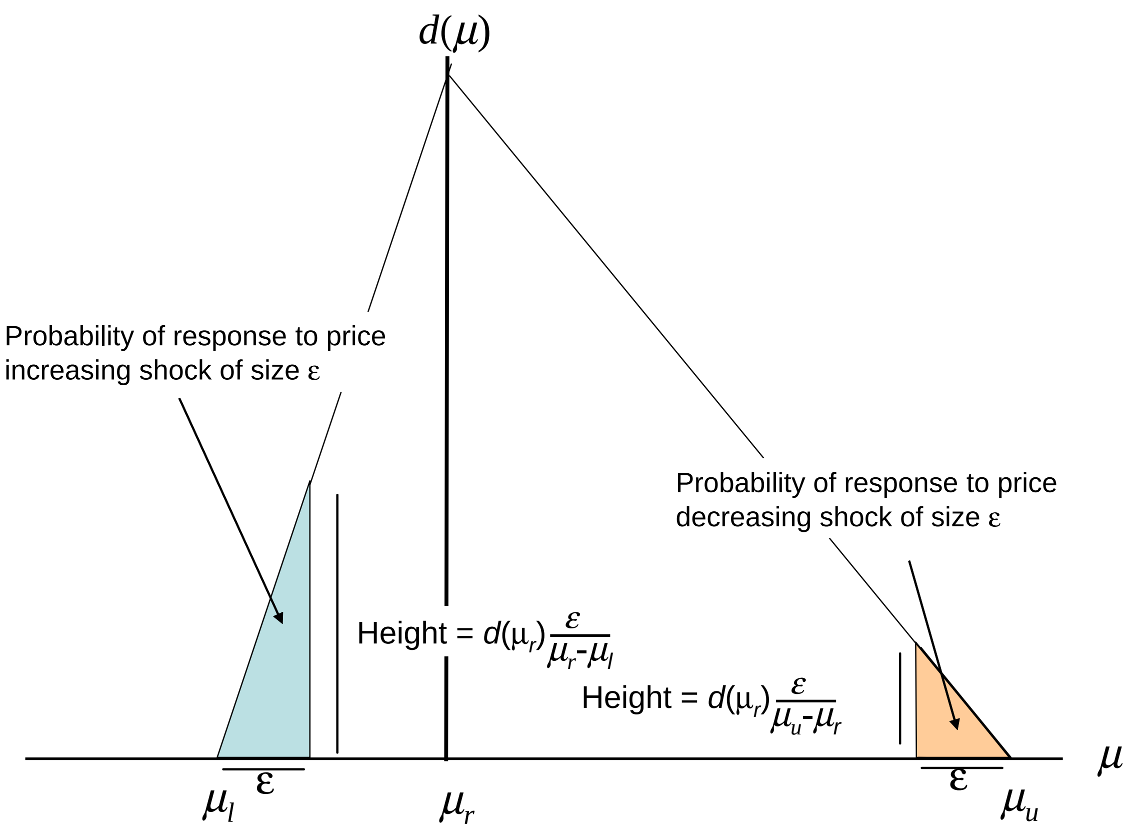

Figure 2: Upward and downward average adjustements are equal.

4.2 Average price dynamics with drift

With drift, the general solution is given by (11). All the equations (13), (14), (15) and (16) remain but (17) is replaced by

$$a_l + b_l = a_u + b_u. \tag{18}$$

We thus have the following four equations in the equally many unknowns:

$$\begin{aligned}
a_u + b_u &= a_l + b_l \\
(1/2-\delta)(a_u+b_u) &= (1/2+\delta)(a_u + b_u\rho_1) + (1/2-\delta)^2\left(a_u + b_u\rho_1^{U-1}\right) \\
(1/2+\delta)(a_l+b_l) &= (1/2-\delta)\left(a_l + b_l\rho_1^{-1}\right) + (1/2+\delta)^2\left(a_l + b_l\rho_1^{-L+1}\right) \\
1 &= (1/2+\delta)\left(a_l + b_l\rho_1^{-L+1}\right) + \sum_{s=-L+1}^{-1}\left(a_l + b_l\rho_1^s\right) \\
&\quad + \sum_{s=0}^{U-1}\left(a_u + b_u\rho_1^s\right) + (1/2-\delta)\left(a_u + b_u\rho_1^{U-1}\right).
\end{aligned}$$

The system is linear in the four coefficients $\{a_l, b_l, a_u, b_u\}$ and is straightforward to solve. With inflationary drift, the average immediate response to a positive shock must be larger than the average response to a deflationary shock. To get the idea, take the derivative of the ergodic density.[20] This yields,

$$\begin{aligned}
\frac{\partial\left(\tilde{a}_i + \tilde{b}_i\rho_1^s\right)}{\partial s} &= b_i\rho_1^s \ln\rho_1, \\
\frac{\partial^2\left(\tilde{a}_i + \tilde{b}_i\rho_1^s\right)}{\partial s^2} &= b_i\rho_1^s\left(\ln\rho_1\right)^2.
\end{aligned}$$

Since $\rho_1 \in (0,1)$ when there is inflationary drift, the first and the second derivative must be of different signs. Under inflationary drift, the ergodic distribution is no longer piecewise linear as in Figure 6, but first concave and then convex as in Figure 7. The fact that the density is concave below μ_r and convex above implies that the probability of an upward adjustment after a shock of size ε times the adjustment size must be strictly larger than the corresponding response to a negative shock.

4.3 Numerical adjustment lags

Let us finally compute the adjustment lag asymmetry as a function of the model's parameters. In the conservative base-line case of moderate elasticities, with $\theta = 0, \nu = 2$, $\sigma_{\Delta\mu} = 4\%$ per month, $m = .1\%$ and a drift of 3% per year we have seen that the adjustment size asymmetry is negligible. Under the discretization used in the previous section, $\delta = 0.003423$ and $\varepsilon = 0.00438188$.Thus,

$$\begin{aligned}
\mu_L &= \frac{0.004 + 0.098}{0.00438188} = 23.28, \\
\mu_U &= \frac{0.1090 - 0.004}{0.00438188} = 23.96,
\end{aligned}$$

so we set $L = 24$ and $U = 24$. Despite the negligible adjustment size asymmetry, the existence of drift implies that the average price responds more quickly to positive than to negative price shocks, in line with Tsiddon (1993). Specifically, the ratio of the immediate price impact for positive and negative price shocks (the shortest possible adjustment lag asymmetry) is

$$\frac{d(\mu_{-L})L}{d(\mu_U)U} = \frac{\left(a_l + b_l\rho_1^{-L}\right)L}{\left(a_u + b_u\rho_1^{U}\right)U} = 1.39.$$

[20] The density is a step-function under the discretization above, so it is not formally differentiable; the argument can be made precise by transforming the problem to continuous time.

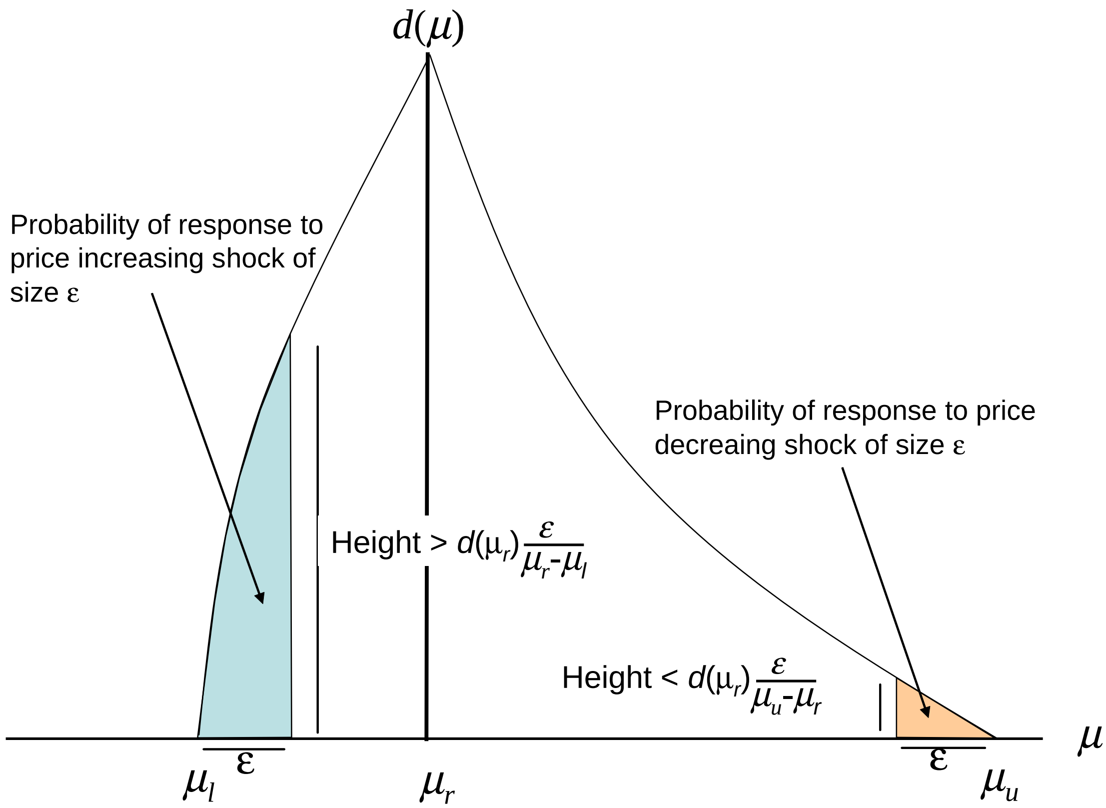

Figure 3: Upward adjustments are larger than downward adjustements.

In the case of high adjustment size asymmetry, i.e., with $\nu = 10, \theta = 3$, $m = 7.5\%$ we have $\mu_L = 18.0$ and $\mu_U = 27.7$, and we therefore set $L = 18$ and $U = 28$. Now, as expected, the adjustment lag asymmetry is somewhat smaller,

$$\frac{d(\mu_{-L})L}{d(\mu_U)U} = 1.35.$$

To produce larger adjustment lag asymmetries in line with Peltzman's (2000) findings (an asymmetry of almost 2 for producer prices and 3 for consumer prices), we need a larger δ. This can be achieved through larger drift π, but as seen from (10) we may as well increase δ by reducing the volatility ε. A rationale for reducing the volatility is that periodic sales tend to account for a large fraction of price variation. The price changes associated with sales are also substantially larger on average than regular price changes. For example, in the consumer price data studied by Klenow and Kryvtsov (2005) less than 20% of the variance in prices is due to regular price changes.[21] Since our model is only concerned with the effect of permanent

[21] Klenow and Kryvtsov (2005), Table 1, reveals that about 20% of all price changes are are due to temporary sales (the fraction of products that change price during a typical month is 0.293 and the fraction that change the ordinary price is 0.233).

input price changes, the empirical volatility of 4% is an overestimate. Consider therefore a reduction in volatility to the more reasonable level of 2%, letting the drift remain at 3% per year.[22] As discussed above, we need to reduce menu costs to produce a reasonable duration between price adjustments. We therefore set $m = 3\%$. In this case, we find $\mu_L = 20.4$ and $\mu_U = 24.1$ and $\delta = 0.0068459$. Adjustment size asymmetry is $24.1/20.4 = 1.18$. Setting $L = 20$ and $U = 24$, we finally compute the adjustment lag asymmetry as

$$\frac{d(\mu_{-L})L}{d(\mu_U)U} = 1.82.$$

The duration between price changes is five months and the share of price increases is 69%. Except a somewhat too high relative frequency of price increases, this example thus hits all the numerical targets.

Hence, our model passes at least the first, admittedly rudimentary, quantitative hurdle.

5 Final remarks

We have shown that a combination of small menu costs and mild inflation suffices to generate asymmetric price rigidity of the magnitude observed empirically. Moreover, the model replicates the observed negative relationship between input price volatility and output price asymmetry, which is just about the only effect that is statistically significant in Peltzman's (2000) regressions on the determinants of downward price rigidity.[23] Unlike earlier work, the model is also consistent with the fact that ordinary price reductions tend to be larger than ordinary price increases.[24]

Our results suggest that the short-run Phillips curve ought to be steeper for increases in the price level than for decreases, raising yet again the old question of whether downward price rigidity is important for optimal monetary policy. A kink in the Phillips curve certainly has qualitative implications for the optimal Taylor rule; see Dolado et al.(2005). However, as noted by Gertler and Leahy (2005), the quantitative effects of monetary policy in a menu cost model depend importantly on mechanisms which are neglected in our model, for instance strategic complementarities. We therefore choose not to undertake a numerical analysis of monetary policy here.

The average price change is 13.3% and the average ordinary price change is 8.5%. It follows that the average price change in connection to a sale is about 32.5%. (The equation is $0.8 \cdot 8.5\% + 0.2 \cdot x\% = 13.3\%$.) The relative contribution of ordinary price changes to the variance in prices is therefore something like $0.8 \cdot 8.5^2/(0.8 \cdot 8.5^2 + 0.2 \cdot 32.5^2) \approx 0.17$. This expression would be correct if all ordinary (and temporary) price changes were of the same size; an exact computation requires information about the variances of ordinary and temporary price changes.

[22] Here we set the gridstep to 0.025%.

[23] The model also predicts a positive relationship between the rate of input price drift and the magnitude of output price asymmetry, a pattern that Peltzman did *not* find to be statistically significant. However, as Peltzman used an ex post measure of drift, we think it is likely that some of the negative effect of high volatility is being picked up by the drift term, biasing the coefficient downwards.

[24] On a related note, the mechanisms that we investigate here can also entail asymmetric responses to currency appreciations and depreciations in international trade. Flodén and Wilander (2006) examine the interaction between the currency in which prices are set (exporter's currency or importers' currency) and the size and frequency of price adjustments.

References

Ball, Laurence and N. Gregory Mankiw (1994), Asymmetric Price Adjustment and Economic Fluctuations, *Economic Journal* 104, 247-261.

Barro, Robert J. (1972), A Theory of Monopolistic Price Adjustment, *Review of Economic Studies* 39, 17–26.

Bènabou Roland and Jerzy D. Konieczny (1994), On Inflation and Output with Costly Price Changes: A Simple Unifying Result, *American Economic Review* 84, 290–297.

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

Borenstein, Severin A., Colin Cameron and Richard Gilbert (1997), Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes? *Quarterly Journal of Economics* 112, 305-339.

Borenstein, Severin, A. and Andrea Shepard (2002), Sticky Prices, Inventories, and Market Power in Retail Gasoline Markets, *Rand Journal of Economics* 33, 116–139.

Caplin, Andrew S. and Daniel F.Spulber (1987), Menu Costs and the Neutrality of Money, *Quarterly Journal of Economics*, 102, 703-25.

Caplin, Andrew S. and John Leahy (1997), Aggregation and Optimization with State-Dependent Pricing, *Econometrica* 65, 601-625.

Danziger, Leif (1999), A Dynamic Economy with Costly Price Adjustments, *American Economic Review* 89, 878-901.

Dhyne, Emmanuel, Luis J. Álvarez, Hervé Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffmann, Nicole Jonker, Patrick Lünnemann, Fabio Rumler and Jouko Vilmunen (2005), *Price setting in the euro area: some stylized facts from individual consumer price data,* European Central Bank, WP nr 524.

Frenkel, Jacob A and Boyan Jovanovic (1980), On Transactions and Precautionary Demand for Money, *The Quarterly Journal of Economics* 95, 25-43.

Flodén, Martin and Fredrik Wilander (2006), State Dependent Pricing, Invoicing Currency and Exchange Rate Pass-Through, forthcoming, *Journal of International Economics.*

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

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

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

Konieczny, Jerzy.D. (1990), Inflation, Output and Labour Productivity when Prices are Changed Infrequently *Economica* 57, 201–218.

Kuran, Timur (1983), Asymmetric Price Rigidity and Inflationary Bias, *American Economic Review* 73, 373-382.

Kuran, Timur (1986), Price Adjustment Costs, Anticipated Inflation, and Output, *Quarterly Journal of Economics* 101, 407-418.

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* 112, 791-825.

Levy, Daniel, Georg Müller, Shantanu Dutta, and Mark Bergen, (2005), Holiday Price Rigidity and Cost of Price Adjustment, mimeo. Bar-Ilan University.

Maskin, Eric and Jean Tirole, (1988), A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles, *Econometrica,* 56:3, May, 571-99

Naish, Howard F. (1986): Price Adjustment Costs and the Output–Inflation Trade–Off, *Economica* 53, 219–230.

Neumark, David and Steven A. Sharpe (1992), Market Structure and the Nature of Price Rigidity: Evidence from the Market for Consumer Deposits, *Quarterly Journal of Economics* 107, 657-680.

Noel, Michael D. (2006). "Edgeworth Price Cycles: Firm Interaction in the Toronto Retail Gasoline Market", *Journal of Industrial Economics*, forthcoming.

Peltzman, Sam (2000), Prices Rise Faster Than They Fall, *Journal of Political Economy* 108, 466-502.

Rotemberg, Julio and Garth Saloner (1987), The Relative Rigidity of Monopoly Pricing, *American Economic Review* 77, 917-926.

Schwartz, Eduardo S. and James E. Smith (2000), Short-Term Variations and Long-Term Dynamics in Commodity Prices, *Management Science* 46, 893-911.

Sheshinski, Eytan and Yoram Weiss (1977), Inflation and Costs of Price Adjustment, *Review of Economic Studies* 44, 287-303.

Sheshinski, Eytan and Yoram Weiss (1983), Optimum Pricing Policy under Stochastic Inflation, *Review of Economic Studies* 50, 513-29.

Tsiddon, Daniel (1993), The (Mis)Behaviour of the Aggregate Price Level, *Review of Economic Studies* 60, 889-902.

6 Appendix

We shall first prove that the dynamic pricing problem with constant inaction bands has a unique solution as claimed in the text. We then illustrate how to solve the model when the price adjustment costs induce variable inaction bands.

6.1 Existence and Uniqueness

We will show that provided r is sufficiently high, $V(\mu)$ is bounded and satisfies discounting implying that (8) is a contraction mapping with a unique fixed point that can be found by iteration.

The real frictionless profit is

$$\Pi_f(\psi_t) = \frac{1+\nu\theta}{\nu(1+\theta)}\left(\frac{\nu}{\nu-1}\right)^{-\frac{\nu-1}{1+\nu\theta}} \psi_t^{-\frac{\nu-1}{1+\nu\theta}}.$$

We denote by $W\left(\frac{p}{P},\psi;m\right)$ and $V(\mu_t;m)$ the value function and the normalized value function for a given m. To see that V is bounded for sufficiently high r, we note that for $m>0$, we have

$$\begin{aligned}
V(\mu_t;m) &= \frac{W\left(\frac{p}{P},\psi;m\right)}{\psi_t^{-\frac{\nu-1}{1+\nu\theta}}} < \frac{W\left(\frac{p}{P},\psi;0\right)}{\psi_t^{-\frac{\nu-1}{1+\nu\theta}}} \\
&= \frac{1+\nu\theta}{\nu(1+\theta)}\left(\frac{\nu}{\nu-1}\right)^{-\frac{\nu-1}{1+\nu\theta}} E_t\sum_{s=0}^{\infty}(1+r)^{-s}\left(\frac{\psi_{t+s}}{\psi_t}\right)^{-\frac{\nu-1}{1+\nu\theta}} \\
&= \frac{1+\nu\theta}{\nu(1+\theta)}\left(\frac{\nu}{\nu-1}\right)^{-\frac{\nu-1}{1+\nu\theta}} \sum_{s=0}^{\infty}\left(\frac{E\left(e^{-\frac{\nu-1}{1+\nu\theta}\Delta\ln\psi}\right)}{1+r}\right)^{s}
\end{aligned}$$

where we used the fact that $\frac{\psi_{t+1}}{\psi_t}$ is i.i.d. for the second equality. Clearly, provided $E\left(e^{-\frac{\nu-1}{1+\nu\theta}\Delta\ln\psi}\right)/(1+r)<1$, the value function $V(\mu_t;m)$ is bounded from above. However, due to the convexity of the profit function, $r>0$ and $E\left(\frac{\psi_{t+1}}{\psi_t}\right)\geq 1$ is is not sufficient. Furthermore, $V(\mu)\geq \max V(\hat{\mu})-m$, since $\max V(\hat{\mu})$ can always be achieved by paying m, implying that V is bounded from below. Finally, equation (8) satisfies discounting if $E\left(e^{-\frac{\nu-1}{1+\nu\theta}\Delta\ln\psi}\right)/(1+r)<1$.

Recalling that $\frac{\psi_{t+1}}{\psi_t} = \exp\left(\gamma+\tilde{\gamma}_{t+1}-\pi-\tilde{\pi}_{t+1}\right)$, we can go further if $\tilde{\gamma}_{t+1}-\tilde{\pi}_{t+1}$ is log-normal with variance σ_ψ^2. Then

$$E\left(\frac{\psi_{t+1}}{\psi_t}\right)^{-\frac{\nu-1}{1+\nu\theta}} = e^{-\frac{\nu-1}{1+\nu\theta}(\gamma-\pi)+\frac{\left(\frac{\nu-1}{1+\nu\theta}\right)^2\sigma_\psi^2}{2}}.$$

As we see, a high variance increases $E\left(\frac{\psi_{t+1}}{\psi_t}\right)^{-\frac{\nu-1}{1+\nu\theta}}$, in particular when ν is large and θ is small.

6.2 One-sided costs

When the menu costs is one-sided, i.e., there is a cost associated with changing the price in one direction but not the other, the optimal policy collapses to one of two parameters – a trigger and a return point.

Specifically, if the menu costs applies for price increases $\mu_u = \mu_r$, i.e., $U = 0$ and if it applies for price reductions, instead $\mu_l = \mu_r$.

Following the same procedure as above, we recall that the difference equation for $d(\mu_s)$ satisfies $d(\mu_s) = a_u + a_l s$ for the interior range. At the trigger μ_l (the case when downward adjustments are costly is exactly symmetric to this case), we have

$$d(\mu_{-L}) = \frac{d(\mu_{-L+1})}{2}$$

We also know that the inflow to the share of firms with negative deviations equals the outflow. That is

$$\frac{d(\mu_0)}{2} = \frac{d(\mu_{-1})}{2} + \frac{d(\mu_{-L})}{2} = \frac{d(\mu_{-1})}{2} + \frac{d(\mu_{-L+1})}{4}$$

Finally,

$$\sum_{s=-L}^{0} d(\mu_s) = 1.$$

Therefore,

$$\begin{aligned} \frac{a_u}{2} &= \frac{a_u - a_l}{2} + \frac{a_u - (L-1)\,a_l}{4}, \\ \sum_{s=-(L-1)}^{0} (a_u + a_l s) + \frac{(a_u - (L-1)\,a_l)}{2} &= 1, \end{aligned}$$

with the solution

$$\begin{aligned} a_u &= \frac{2}{L+2}, \\ a_l &= \frac{2}{(L+2)(1+L)}, \end{aligned}$$

which is a special case of the solution above with $U = 0$. The ratio of the densities are given by

$$\begin{aligned} \frac{d(\mu_L)}{d(\mu_0)} &= \frac{\frac{a_u - a_l(L-1)}{2}}{a_u} = \frac{\frac{\frac{2}{L+2} - \frac{2}{(L+2)(1+L)}(L-1)}{2}}{\frac{2}{L+2}} \\ &= \frac{1}{1+L}. \end{aligned}$$

Clearly, the effect of a cost increase of ε if the distribution of normalized prices is the ergodic one, is that a share $\frac{2}{(L+2)(1+L)}$ of the firms increase their prices each by a fraction $(L+1)\,\varepsilon$. In the case of a cost reduction, instead a share $\frac{2}{L+2}$ of the firms reduce their prices by a fraction ε. The product is in both cases $\frac{2}{2+L}\varepsilon$.

6.3 Numerical implementation.

To solve (8) numerically, we first assume that the nominal and real shocks take two values each, $\pm\varepsilon_n$ and $\pm\varepsilon_r$. The probability of a positive shock $1/2 + \delta_r$ and $1/2 + \delta_n$ and of a negative $1/2 - \delta_r$ and $1/2 - \delta_n$. We assume the shocks are independent, although this could easily be relaxed.

The drift in real costs, satisfies

$$2\delta_r \varepsilon_r = \gamma - \pi,$$

and the variance is

$$\left(\left(\frac{1}{2}+\delta_r\right)(\varepsilon_r-2\delta_r\varepsilon_r)^2+\left(\frac{1}{2}-\delta_r\right)(-\varepsilon_r-2\delta_r\varepsilon_r)^2\right)=\varepsilon_r^2\left(1-4\delta_r^2\right)=\sigma_{\tilde{\psi}}^2,$$

implying

$$\begin{aligned}\varepsilon_r &= \left((\gamma-\pi)^2+\sigma_{\psi}^2\right)^{\frac{1}{2}}\\ \delta_r &= \frac{\gamma-\pi}{2\varepsilon_r}.\end{aligned}$$

For the nominal shock, we similarly have

$$\begin{aligned}\varepsilon_n &= \left(\pi^2+\sigma_{\pi}^2\right)^{\frac{1}{2}},\\ \delta_n &= \frac{1}{2}\frac{\pi}{\varepsilon_n}.\end{aligned}$$

We discretize the state space so that μ must be chosen from the set $\exp(s*0.0005)$, where s is any positive or negative integer. That is, prices must be chosen in steps of 0.05% around the current frictionless price. The possible combination of shocks are now, $\varepsilon_n+\varepsilon_r, \varepsilon_n-\varepsilon_r, -\varepsilon_n+\varepsilon_r$ and $-\varepsilon_n+\varepsilon_r$. The effect they have on μ are given by

$$\Delta\ln\mu_{t+1}=\begin{cases}-\varepsilon_n-\frac{\varepsilon_r}{1+\nu\theta}\text{with probability }\left(1/2+\delta_r\right)\left(1/2+\delta_n\right);\\ \varepsilon_n-\frac{\varepsilon_r}{1+\nu\theta}\text{with probability }\left(1/2-\delta_r\right)\left(1/2+\delta_n\right);\\ -\varepsilon_n+\frac{\varepsilon_r}{1+\nu\theta}\text{with probability }\left(1/2+\delta_r\right)\left(1/2-\delta_n\right);\\ \varepsilon_n+\frac{\varepsilon_r}{1+\nu\theta}\text{with probability }\left(1/2-\delta_r\right)\left(1/2-\delta_n\right).\end{cases}$$

The size of these shocks are finally rounded to the nearest integer times the grid size. We then iterate on (8) until $\max_{\mu}\left\|\frac{V_s(\mu)-V_{s+1}(\mu)}{f(1)}\right\|<\frac{1}{1000}$.