Menu costs and asymmetric price adjustment

Tore Ellingsen [a,1], Richard Friberg [a,1,*], John Hassler [b,2]

[a] *Department of Economics, Stockholm School of Economics, Box 6501, SE–113 83, Stockholm, Sweden*
[b] *Institute for International Economic Studies, Stockholm University, SE–106 91, Stockholm, Sweden*

ARTICLE INFO

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 generate observed frequencies of price changes, the model can account for several aggregate 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 latter asymmetry is more pronounced when input prices are less volatile, as documented by Pelzman (2000).

1. Introduction

The inflationary surge after the COVID-19 pandemic and Russia's invasion of Ukraine has rekindled interest in how prices respond to input costs. In many markets, prices are quick to increase when input costs rise but slow to fall when input costs drop. Such "rockets and feathers"-pricing has been documented for e.g. gasoline retailers (e.g., Bacon, 1991; Neumark and Sharpe, 1992; Arbatskaya and Baye, 2004; Asane-Otoo and Dannemann, 2022), and for many agricultural products (von Cramon-Taubadel and Goodwin, 2021). Indeed, prices respond faster to cost increases than to cost decreases in a large majority of Peltzman (2000)'s broad sample of more than two hundred different product markets; see also Nakamura and Steinsson (2008).

The literature in the field of industrial organization has focused on two explanations for such asymmetries: tacit collusion and search costs. The tacit collusion narrative was developed in Borenstein et al. (1997), if costs fall, firms "automatically" come closer to a collusive markup by waiting to adjust. Although the argument still awaits full formalization, tacit collusion has served as a motivation for much of the subsequent empirical and experimental work (Bayer and Ke, 2018; Bulutay et al., 2021). The other prominent explanation for price asymmetry is associated with asymmetric search intensity by consumers (Yang and Ye, 2008; Tappata, 2009) – in these models a low search intensity when costs are high makes consumers relatively unresponsive to price decreases, whereas the opposite holds when costs are low.

Here, we pursue an orthogonal reason for response asymmetry that relies on the combination of an asymmetric profit function and menu costs of changing price.[1] Our theoretical contribution is most easily explained by pinpointing the main departures from

---

* Corresponding author.
  E-mail addresses: tore.ellingsen@hhs.se (T. Ellingsen), richard.friberg@hhs.se (R. Friberg), john.hassler@iies.su.se (J. Hassler).
[1] Also affiliated with Norwegian School of Economics and CEPR.
[2] Also affiliated with CEPR.
[1] For evidence on the connection between the magnitude of menu costs and the frequency of price adjustment, see for example Levy et al. (1997) and Anderson et al. (2015). The former study exploits item-pricing legislation that raise the menu cost in the relevant jurisdictions, while the latter study exploits the seller's policy to price multiple varieties of the same good in a uniform manner.

previous menu cost models. 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. One essential difference is that instead of Barro's linear demand our firm faces a demand with constant price elasticity. The constant elastic demand implies that the static profit function is not symmetric around the optimal static price, and the profit function drops off more rapidly for too low prices than for too high prices.[2] Intuitively, starting from a static optimal price, a shock that gives a higher cost will then be associated with a stronger incentive to raise price than the incentive to lower price for a fall in costs of the same magnitude. The shape of the profit function, together with menu costs, will determine the optimal size, frequency and asymmetry of price adjustments in our dynamic model.

Early work in the menu cost literature eschewed the otherwise popular constant elasticity case because the interaction between the fixed menu cost and the non-linear demand function generates complicated solutions for the firm's pricing problem. [3] Alvarez and Lippi (2014) show that one can build a tractable menu-cost model with constant demand elasticity by considering a second-order approximation to the profit function instead of the exact profit function. For many purposes, this simplified model suffices, but because the second-order approximation is symmetric, their model eliminates entirely the adjustment asymmetry that we seek to investigate.

However, Danziger (1999) shows that constant elasticity of demand can be compatible with tractability if the standard assumption of constant menu costs is replaced with the assumption that price-adjustment costs are proportional to the frictionless profit flow. This, then, is our key assumption here.[4]

We do not claim that Danziger's menu cost assumption is entirely realistic, or even that it is always more realistic than Barro's fixed-cost assumption. There are almost certainly some scale economies in price-setting.[5] Indeed, if one considers merely the narrowly defined physical costs of changing prices in retail outlets, the fixed-cost assumption looks increasingly accurate, as the need to put price tags on individual items steadily diminishes.[6]

However, there are at least two sets of reasons why the costs of changing prices are not independent of volume, even among retailers. First, selling large volumes is related both to selling many variants and to being part of retail chains. In both cases policies of uniform pricing are well documented – firms set the same price for e.g. different flavors of a product[7] and for different store locations across the country, even if local demand conditions vary substantially (DellaVigna and Genzkow, 2019). Articles documenting uniform pricing typically conclude that the practice is puzzling for economists, but if we take such a policy as given, then it suggests that determining whether to update a price that trades off considerations across many different locations or varieties is more complex and time consuming than simply updating one price in one location. This implies that menu costs, broadly interpreted, are increasing in the number of varieties and locations.[8] Consistent with such reasoning, Anderson et al. (2015) use data from a US grocery retailer to document that the larger the number of varieties of a product, the lower is the probability that price will increase following a cost shock.

Second, price changes typically involve communication costs that go well beyond the change of price tags. Price increases cause disgruntled customers, as customers tend to suspect that higher prices involve higher mark-ups; see Eyster and Madarasz (2021) and the references therein. Sellers should therefore (be prepared to) justify the price increases that are caused by higher costs. Price decreases, on the other hand, create a potential for additional sales, but only to the customers that pay attention to them. Hence, sellers should again (be prepared to) inform the customers. If the benefits from these communication activities scale with profits, it seems plausible that the costs of optimally chosen communication do so, too. Danziger's variable menu costs may thus offer a reasonable approximation to communication costs.[9]

In the context of retailing, we are not aware of any empirical study that has carefully measured the importance of communication costs associated with price changes. However, in a detailed case study of a large industrial manufacturer, Zbaracki et al. (2004) document price adjustment costs that are quite consistent with the above storyline. Physical menu costs are small, and the costs of informing and negotiating with customers constitute by far the largest price adjustment cost.

Why is Danziger's formulation so tractable? Due to the profit-proportional menu cost assumption, the inaction band has constant logarithmic width, and in the no-drift case we can analytically characterize the ergodic distribution of price deviations. Moreover,

---

[2] The fact that constant demand elasticity generates larger static losses from negative price deviations than from positive ones has previously been pointed out by e.g. Kuran (1983).

[3] The early theoretical work on the inflation/output trade–off that did assume constant price elasticity thus limited its attention to non–stochastic models and to the case of very small menu costs; see for example Bénabou and Konieczny (1994).

[4] For a somewhat related application of Danziger's framework, see Gertler and Leahy (2008). Since their analysis restricts attention to a local approximation around a steady state of zero inflation, asymmetric price adjustment still does not occur.

[5] In the context of Dominick's, a large US retail food chain with 93 stores in the greater Chicago area, Sayag et al. (2024) document that small price changes were more likely when sales volumes were large already in the 1990s, suggesting that there were significant scale economies in retail price adjustment.

[6] For studies of these costs, which comprises store employers taking time off from other valuable tasks, see Levy et al. (1998) and Dutta et al. (1999).

[7] See e.g. Anderson et al. (2015) for a retail example or closely related examinations of uniform pricing for movies (Orbach and Einav (2007)) and rental cars (Cho and Rust (2010)).

[8] Relatedly we may note that in differentiated product oligopolies pass-through depends on the set of products that firms control and its interaction with menu costs are far from trivial, see e.g. Nakamura and Zerom (2010), Friberg and Romahn (2018).

[9] It is far from obvious that the benefits associated with providing information will be symmetric for cost increases and cost decreases, but the costs associated with optimal information provision could still be similar enough for the symmetric model to be adequate.
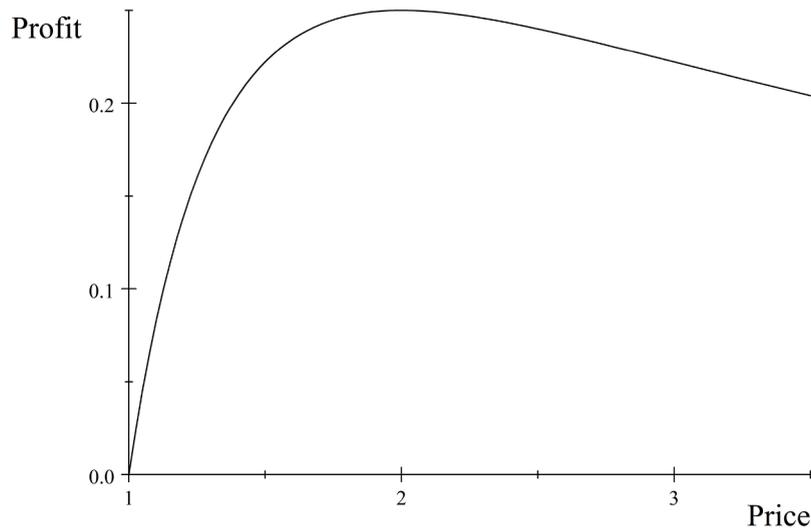
Fig. 1. Profits flow as a function of price with constant unitary marginal cost and a constant demand elasticity of -2.

even in the positive drift case the numerical algorithms converge quickly. Under realistic parameter values the model produces asymmetric price responses in the same ballpark as the evidence of Peltzman (2000) and Nakamura and Steinsson (2008); see also Bils and Klenow (2004), Klenow and Kryvtsov (2008), and Gautier et al. (2024). Our model reproduces three patterns in these data: price increases are more frequent, smaller, and faster than price decreases.

Our analysis also speaks to the business cycle analysis that relies on menu costs as a source of price rigidity; see e.g. Golosov and Lucas (2007), Gertler and Leahy (2008), Midrigan (2011), Alvarez and Lippi (2014), Karadi and Reiff (2019), and Alvarez and Lippi (2022). Key questions in this literature concern the real effects of monetary policy, how they depend on the nature of shocks, and whether predictions from models relying on time- vs. state-dependent adjustment costs differ. To the best of our knowledge, the kind of asymmetries that we analyze have not been highlighted in this literature. We return to this connection in our conclusions.

Finally, our work relates to the immense empirical literature on the pass-through into consumer prices of cost shifters such as exchange rates (Burstein and Gopinath, 2014) and commodity prices (Hong and Li, 2017). Some work in this literature documents asymmetries in the pass-through of price increases and decreases; see, for instance, Benzarti (2024) for a survey of asymmetries in the pass-through of increases and decreases of taxes. However, the workhorse theoretical models of pass-through are typically static and do not feature asymmetric price adjustment (see, e.g. Weyl and Fabinger, 2013). A final point to note in relation to the static pass-through literature is that the case of constant elastic demand and constant marginal cost is known to imply that the desired markup is constant and hence the pass-through elasticity equals one. Our results provide a stark example of how dynamics and menu costs alter the patterns of pass-through relative to this static and frictionless benchmark.

Our paper is organized as follows. Section 2 contains a simple exposition of the firm's static pricing problem. In Section 3, we go on to analyze the fully dynamic and stochastic pricing problem, illustrating how the pricing pattern asymmetries vary with parameters such as the elasticity of demand and input price volatility. Section 4 in turn analyzes adjustment lag asymmetries. 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

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

While input price variation is going to be the major source of changes in prices and quantities, the losses associated with an inappropriate choice of output price p, can be studied without explicit reference to ψ. For notational simplicity, we therefore suppress the argument ψ in this section.

Fig. 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 Fig. 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. The overall asymmetry is intuitive in this simple case. As the price falls towards the marginal cost, profits approach zero. Any price above the optimal will, on the other hand, leave the firm with positive profits.

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) and in particular Bénabou 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

$$3D''(p^* + x) + p^* D'''(p^* + x) - C'''(D(p^*))(D'(p^*))^3$$
$$-3C''(D(p^*))D''(p^*)D'(p^*) - C'(D(p^*))D'''(p^*) \qquad > 0.$$

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.

Remark 1: *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 excessively low price.*

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 optimal 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 isoelastic 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[10]

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

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

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

---

[10] Adding a real shock to demand would be straightforward to analyze. In particular, if $\theta = 0$ and the real demand shock is multiplicative, it only adds noise to real profits, without changing optimal pricing.

Inserting (4) into (3) and using (2), it takes a few manipulations to write the real frictionless profit as

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

Let us now turn to the dynamic pricing problem. 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

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

With positive 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. As noted in the Introduction, we seek conditions that entail an optimal inaction band of constant logarithmic width, and the key is to assume that the real menu cost is a constant fraction of Π_f,

$$M_t = m\Pi_f\left(\psi_t\right). \tag{6}$$

The rough intuition is clear: When output and profits are high the gain from adjusting the price is high. Thus, inaction bands can be constant only if the menu cost is positively related to output.

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 the current cost ψ. Our central analytical result is summarized as follows.

Proposition 1. *If the real menu cost is a constant fraction $m > 0$ of the frictionless profit Π_f and if parameters satisfy the inequality*

$$E\left(e^{-\frac{v-1}{1+v\theta}\Delta\ln\psi}\right)/(1+r) < 1, \tag{7}$$

then the monopolist's dynamic pricing problem has a unique solution and it entails inaction bands with constant logarithmic width.

We leave the proof of existence and uniqueness for the Appendix, but provide the main steps of the characterization (how the constant inaction band hinges on the menu cost function) here.

By definition, under constant inaction bands, the optimal policy can be completely characterized by three numbers, μ_u, μ_l, μ_r, and the decision rule:

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

To derive conditions such that the inaction band is constant, define the normalized (output) price at time t,

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

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

$$\Delta\ln\mu_{t+1} \equiv \ln\mu_{t+1} - \ln\mu_t = \Delta\ln p_t - (\pi + \tilde{\pi}_{t+1}) - \frac{\Delta\ln\psi_{t+1}}{1+v\theta}, \tag{8}$$

and $\ln\mu_t$ is a random walk with drift

$$E_{\Delta\mu} = -\pi - \frac{\gamma - \pi}{1+v\theta}$$

whenever the price is unchanged. Assuming that real cost shocks and nominal demand shocks are independent, the variance of $\Delta\ln\mu_{t+1}$, denoted $\sigma_{\Delta\mu}^2$, is a weighted sum of their respective variances, i.e., $\frac{\sigma_{\Delta\psi}^2}{1+v\theta} + \sigma_{\tilde{\pi}}^2$.[11]

---

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

We denote the change in $\ln \mu_{t+1}$ when the output price is constant 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, 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,

$$\Pi\left(\mu_t p_f\left(\psi_t, P_t\right); \psi_t, P_t\right) = f\left(\mu_t\right)\psi_t^{\frac{1-\nu}{1+\nu\theta}}, \tag{9}$$

where

$$f\left(\mu_t\right) \equiv \left(\frac{\nu}{\nu - 1}\right)^{\frac{1-\nu}{1+\nu\theta}} \left(\mu_t^{1-\nu} - \mu_t^{-\nu(1+\theta)} \frac{\nu - 1}{\nu(1 + \theta)}\right).$$

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

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

Using this conjecture, denoting

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

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

$$V\left(\mu_t\right) = \max \left\{f\left(\mu_t\right) + E\left(V\left(\mu_t e^{\Delta \ln \mu_{t+1}}\right)\tilde{R}\right), \tag{10}\right.$$

$$\left.\max_{\mu'}\left\{f\left(\mu'\right) - \frac{M_t}{\psi_t^{\frac{1-\nu}{1+\nu\theta}}} + E\left(V\left(\mu' e^{\Delta \ln \mu_{t+1}}\right)\tilde{R}\right)\right\}\right\}.$$

To get rid of the dependence on ψ_t and P_t, we must assume that the menu cost M_t is proportional to $\psi_t^{\frac{1-\nu}{1+\nu\theta}}$. Using (2) and (4) it follows immediately that this is equivalent to assuming that M_t is proportional to $q_f^{(2+\theta\nu)/(1+\theta\nu)}$, where q_f is the "frictionless quantity" $D(p_f)$. Equivalently, it is proportional to the frictionless profit, as claimed.

In Eq. (10), $\tilde{R}$ 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 $E\tilde{R} < 1$, Eq. (10) is a contraction mapping with a unique solution that can be found by standard numerical methods.[12] 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 (10) until it converges. 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 $(\ln \mu_u - \ln \mu_r)/(\ln \mu_r - \ln \mu_l)$.

Observe that the condition $E\tilde{R} < 1$ is violated if $\frac{\nu-1}{1+\nu\theta}$ is large and the volatility of the real cost shock is high, because the possibility of a series of negative cost shocks in combination with highly elastic demand makes expected discounted profits unbounded. As we show in the Online Appendix, the inequality (7) ensures that discounted profits are finite.[13]

4. Average price dynamics

We now consider price adjustments across firms, or equivalently, across time. When the distribution of price deviations is stationary, the average response to a small cost shock is given by the share of firms changing their price times the magnitude of the price change of those who change. To compute the average responses, we must therefore derive the stationary distribution of price deviations in addition to the optimal policy given by the numbers $\mu_u, \mu_l,$ and μ_r.

To do this, 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.[14] Thus, we assume

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

---

[12] See the Online Appendix for details on the numerical implementation.

[13] The stochastic discount factor $\tilde{R}$ inherits the profit function's convexity in ψ_{t+1}. A simple way to illustrate that convexity is to note that when $\theta = 0$ the second derivative of frictionless profits 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. Therefore, even if $E\frac{\psi_{t+1}}{\psi_t} = 1$ and $r > 0$, $\tilde{R}$ may be larger than unity in expectation.

[14] In the calibration, we set the period length to 1/1000 years.

where

$$\varepsilon = \left(E_{\Delta\mu}^2 + \sigma_{\Delta\mu}^2\right)^{\frac{1}{2}},$$

$$\delta = \frac{1}{2}\frac{-E_{\Delta\mu}}{\varepsilon},$$

and $E_{\Delta\mu}$ and $\sigma_{\Delta\mu}^2$ are the drift and the variance per period of the normalized price deviation when no price changes are undertaken.

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

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

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. We can then easily find the unconditional (long-run) share of firms with deviation μ_s, denoted $d(\mu_s)$ for any μ_s in the domain $s \in [\mu_{-L}, \mu_U]$. To do this, we note that for $s \neq 0, U, L$, the share of firms with normalized price μ_s must be those who in the previous period had a relative price μ_{s-1} and moved upward (happens with probability) $\left(\frac{1}{2} - \delta\right)$ plus those with relative price μ_{s+1} who moved downward (happens with probability) $\left(\frac{1}{2} + \delta\right)$. Therefore, for $s \in \{-L+1, \ldots -1, 1, \ldots U-1\}$

$$d(\mu_s) = (1/2 - \delta)d(\mu_{s-1}) + (1/2 + \delta)d(\mu_{s+1}). \tag{11}$$

Solving this linear difference equation is straightforward (see the Online Appendix).

4.1. Average price dynamics without drift

When $\delta = 0$, there is a closed form solution to the distribution of normalized prices given by

$$d(\mu_s) = \begin{cases} a + \tilde{b}_u s \text{ for } s \in \{0, \ldots, U\}; \\ a + \tilde{b}_l s \text{ for } s \in \{-L, \ldots, -1\}, \end{cases} \tag{12}$$

with

$$a = \frac{2}{2 + U + L}, \tilde{b}_l = \frac{a}{1 + L}, \tilde{b}_u = \frac{-a}{1 + U}. \tag{13}$$

Lemma 1. 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 $(\mu_r - \mu_l)/(\mu_u - \mu_r)$.

Proof: Observe that

$$d(\mu_u)/d(\mu_l) = \frac{a + \tilde{b}_u U}{a + \tilde{b}_l(-L)} = \frac{1 + L}{1 + U}.$$

As the period length gets small, L and U increase and $d(\mu_u)/d(\mu_l)$ approaches L/U, which is equal to $(\mu_r - \mu_l)/(\mu_u - \mu_r)$ given our assumptions of the binomial shock.

Lemma 1 generalizes a result by Tsiddon (1993), who considered the case of symmetric bands. Since the shares are linear with a common intercept a, 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 Fig. 2, we depict the result graphically. The area of the two shaded triangles to the left and right in the figure, represents the share of firms for which 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. Intuitively, for a given magnitude of shock, firms are more likely to be in a section of the inaction band where such a shock triggers them to raise price than to lower price. Given the asymmetry of the inaction band around the optimal return point, the resulting price increase is lower than the corresponding price decrease.

When there is drift in the frictionless price, the argument no longer applies, because drift implies a departure from the linear ergodic distribution of price deviations.

4.2. Average price dynamics with drift

With drift, the general solution to (11) is given by

$$d(\mu_s) = \begin{cases} \tilde{a}_u + \tilde{b}_u \rho_1^s \text{ for } s \in \{0, U\}; \\ \tilde{a}_l + \tilde{b}_l \rho_1^s \text{ for } s \in \{-L, 0\}, \end{cases} \tag{14}$$

with ρ_1 given by $(1 - 2\delta)/(1 + 2\delta)$. The system is linear in the four coefficients $\{a_l, b_l, a_u, b_u\}$ and is again straightforward to solve (see the Online Appendix). With inflationary drift, the average immediate response to a positive shock must be larger than the average
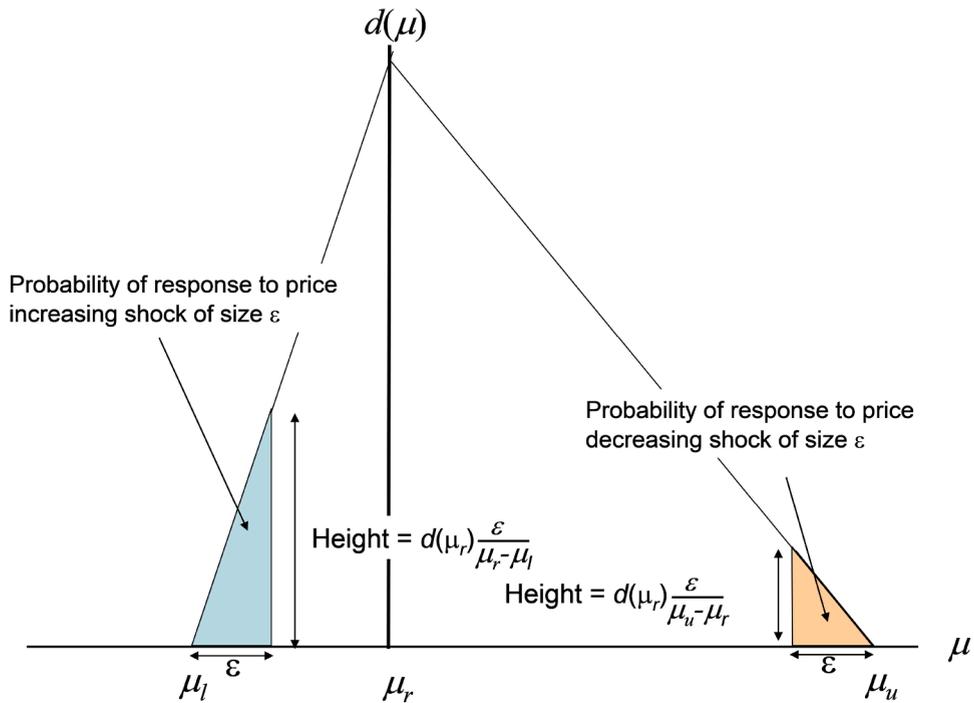
Fig. 2. Ergodic density with no drift.

response to a deflationary shock of equal size. To understand why, consider the slope of the distribution function, i.e., $(d(\mu_{s+1}) - d(\mu_s))/\varepsilon$, which is given by $(\rho_1 - 1)\tilde{b}_i\rho_1^s/\varepsilon$ for $i \in \{l, u\}$ while the density is $\tilde{a}_i + \tilde{b}_i\rho_1^s$. Since $\rho_1 - 1 < 0$ when $\delta > 0$ (inflationary drift), it follows that the *slope* of the distribution function falls in s whenever the distribution itself increases ($s < 0$) and vice versa for $s > 0$. Thus, when there is inflationary drift, the ergodic distribution is no longer piecewise linear as in Fig. 2, but first concave and then convex as in Fig. 3. 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. Compared to the case where there was no drift (Fig. 2), firms are now in the case depicted in Fig. 3 even more likely to find themselves in a part of the inaction band where a shock would trigger them to raise price.

5. Can the model match the asymmetries documented by Peltzman (and others)?

Having developed the model we now proceed to a quantitative evaluation of whether it is able to provide asymmetries in line with those empirically observed. Indeed, for empirically reasonable parameter values we show that our model generates results that are broadly consistent with three key pricing asymmetries that have been established in broad datasets: price increases are more frequent, smaller, and faster than price decreases.

As noted, Peltzman (2000)'s study of asymmetric price responses in more than 200 markets is a key reference in establishing that asymmetric speed of price adjustment is not an anomaly confined to gasoline retail markets but rather a stylized fact. We choose to let the patterns that he finds guide our calibrations. We thus rely on Peltzman (2000) for the adjustment lag asymmetries. To capture other possible asymmetries in pricing in the U.S. for broadly the same time period we rely on microdata evidence from the U.S. Bureau of Labor Statistics, as analyzed by Nakamura and Steinsson (2008).[15]

In our *baseline* calibration we set $\pi = 0.03$ per year and $\sigma_{\bar{\pi}}^2 = 0$. Peltzman (2000) computes the standard deviation of input prices, $\sigma_{\Delta\mu}$, and finds the standard deviation of monthly input price changes to be 4.1 % for the producer goods sample. As a first pass, we thus set $\sigma_{\Delta\psi}^2 = 0.041$ per month. We set the drift in input prices to zero and the real interest rate relevant to businesses at $r = 0.07$ per year. We set the demand elasticity to $v = 4$; for a discussion and references, see Nakamura and Steinsson (2008, p. 1447). The parameter θ, capturing the extent of returns to scale, is somewhat more difficult to calibrate. In the long run, it is reasonable and standard practice to assume constant returns to scale, while decreasing returns to scale may apply in the short run. Our model does not

---

[15] Klenow and Kryvtsov (2008), building on prior work by Bils and Klenow (2004), also estimate adjustment pattern asymmetries using BLS data. However, they only study consumer prices. Moreover, Nakamura and Steinsson go further when it comes to cleaning regular price changes from the effects of sales and product substitutions. Since the shocks we consider follow a random walk, it is only the regular price changes that can be explained by our model.
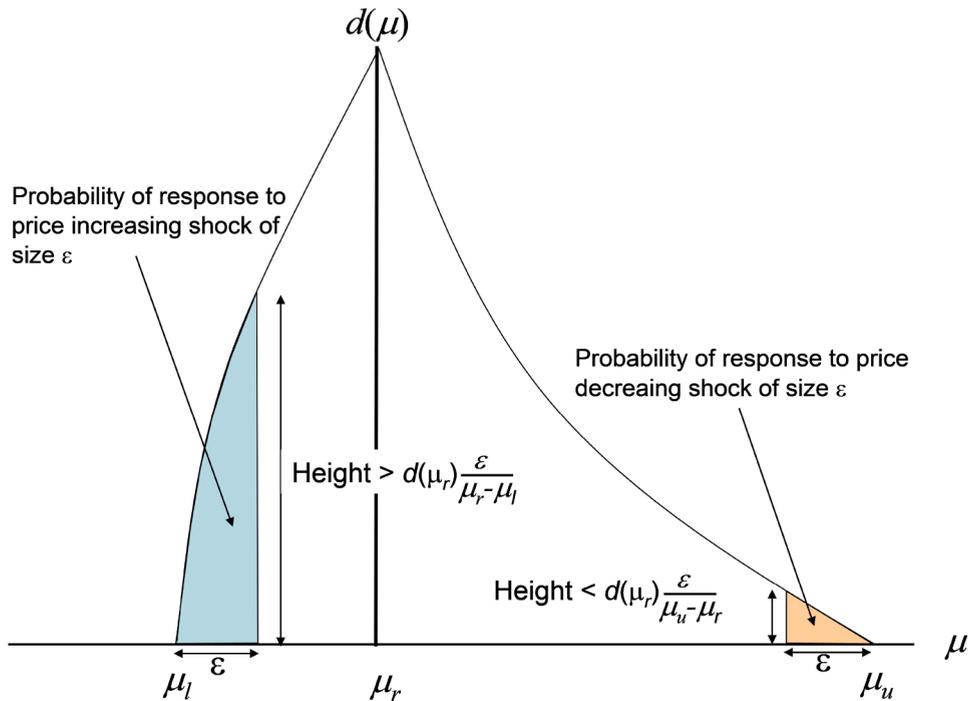
Fig. 3. Ergodic density with inflationary drift.

make a distinction between long and short-run production functions, but it seems appropriate not to deviate to much from constant returns to scale. We thus set $\theta = 0.1$ as a baseline case.

The median implied duration for a price is 8.9 months during 1998–2005 in Nakamura and Steinsson (2008, Table II) and we set the menu cost to match this duration, implying $m = 0.0018$, close to two tenths of a percent of *yearly frictionless profits*.

Baseline calibration

The optimal policy, under the baseline calibration is to allow prices to be 5.73 % below the static profit maximizing price and 6.33 % above before adjusting. The return point is close to the static profit maximizing price, at 0.12 % above. Note that our assumptions imply a constant inaction region and constant return point expressed in terms of the frictionless price. The price increase is thus $5.73 + 0.12 = 5.85$% and the price decrease $6.33 - 0.12 = 6.21$% of the static profit maximizing price. Define the adjustment size asymmetry as the ratio of the size of price increases to the size of price decreases. In the baseline calibration the size asymmetry is 0.942 so that price increases are around 6 % smaller than price decreases. The corresponding number, across all sectors, is 0.695 in Nakamura and Steinsson (2008).

The ratio of the frequencies of price increases to price decreases is 2.19 in our calibration and the average across all sectors is 1.841 in Nakamura and Steinsson (2008, Table II). Thus, in both our simulation and in the data, price increases are about twice as common as price decreases.

For the key asymmetry of interest, asymmetric responses to cost shocks, we can not compare to Nakamura and Steinsson (2008) and instead turn to Peltzman (2000). He estimates the pass-through of cost increases and decreases. Define the adjustment lag asymmetry as the ratio of the average impact response of the price level to cost increasing and cost reducing shocks respectively. In our baseline calibration this lag asymmetry is 2.15. In other words, the increase in price when costs rise is twice as large as the decrease in price when costs fall. Peltzman (2000, Tables 4 and 5) documents a corresponding lag asymmetry of 1.850 for consumer prices and 2.895 for producer prices. We conclude that, with plausible parameter values, the model is able to provide a rough quantitative match of the three dimensions of asymmetric price adjustment that have been documented.

In the following we keep the same parameter values as in the baseline but change parameters of interest one by one to understand how results change if key parameters change.

Lower menu cost

In alternative calibrations that use menu cost that are half as large ($m = 0.0009$) the asymmetries are still pronounced. Lower menu costs imply lower price changes and, when adjusted, prices rise by 4.86 % and decrease by 5.13 %.[16] The adjustment size asymmetry is thus similar to the baseline and is now 0.947, i.e., price increases are around 5 % smaller than price decreases. Price increases are slightly less than twice as common as decreases, the ratio between the frequencies is 1.95. The adjustment lag asymmetry is 2.07, i.e.

---

[16] Price is up to 5.19 % above the profit maximizing price down to 4.80 % below the profit maximizing price before it changes. When prices are adjusted, they are set very close to the profit maximizing price (0.06 % above).

the average price increase in response to a positive cost shock is more than twice as large as the price decrease in response to falling costs. Lower menu costs imply more frequent adjustments and the average time between price changes is 0.546 years.

Lower variance of input prices

Now use the baseline values but lower the input price volatility to $\sigma^2_{\Delta\psi} = 0.03$. The lower volatility narrows the inaction region compared to the baseline. The optimal policy now allows price to be 3.92 % below the static profit maximizing price and 4.69 % above before adjusting. The return point is 0.49 % above the static profit maximizing price and thus substantially higher than in the baseline case. The higher return point reflects the increased importance of trend inflation relative to cost shocks. Prices now increase by 5.16 % and decrease by 5.22 % when adjusted. The resulting adjustment size asymmetry is thus less pronounced at 0.988. In contrast, the other asymmetries are strengthened. The frequency of price increases to price decreases is 3.28. The lag asymmetry is now 3.21, i.e. the rise in price in response to a cost increase is around 3 times the fall in price following a cost decrease. This is close to what Peltzman (2000) reports for producer prices (2.895). The main point, however, is that lower input price volatility is associated with more lag asymmetry, which is precisely the pattern found by Peltzman (2000). The average time between price adjustment is 0.99 years.

More convex costs

Our last calibration explores a more convex cost function and sets $\theta = 0.5$. We keep the other parameters as in the baseline. There are two direct effects on the optimal policy. First, more convex costs leads to tighter inaction regions, since in particular low output prices (high sales) lead to high marginal costs. Second, as seen in (8) a higher θ also makes price deviations less responsive to input price shocks. Again, the relative importance of the drift term then increases. The average price increase is now 3.99 % and the average decrease 3.81 %.[17] This implies an adjustment size asymmetry of 1.047, i.e., price increases are (in contrast to data) 4.7 % larger than price decreases. Despite the more narrow band, the average time span between price adjustments is 1.12 years. The larger importance of the inflationary drift implies a high ratio of the frequency of price increases to price decreases of 8.28 and the lag asymmetry is very pronounced at 7.58.

6. Final remarks

We have shown that the lag asymmetry, that has been the focus of much empirical work in industrial organization, can be generated within a simple framework that also generates two other asymmetries (size and frequency of price increases and decreases) that have been documented in the menu cost literature.

Clearly, this is not to say that tacit collusion or search costs are unimportant for explaining asymmetric price adjustment. While our simple monopoly model does surprisingly well in reproducing common patterns in the data one size will clearly not fit all. The kind of setup that we examine, monopolistic firms facing non-trivial menu costs, is for instance less likely to capture some key aspects of competition in gasoline retailing, where price changes often occur at a daily frequency and strategic interaction is key; see e.g. Byrne and De Roos (2019) or Assad et al. (2024).

To explain our mechanism as simply as possible, we have confined attention to the case of a single-product monopoly. As a result, there are some features of the data that we have neglected to address, notably the existence of small price changes. Following Midrigan (2011), the conventional response has been to extend the model to multi-product sellers and assume that there are economies of scope in adjusting prices; for influential subsequent models with this feature, see Alvarez and Lippi (2014) and Alvarez et al. (2016). These models rely on local approximations that eliminate price adjustment asymmetries.[18] We think a natural next step would be to introduce Danziger's menu cost assumption into these settings to see if they still remain tractable without the approximations. If so, can the more general model explain why the price adjustment asymmetries are particularly pronounced for small price changes, as observed by Chen et al. (2008) and Gautier et al. (2024), among others?

Another relevant question that such a richer macro-model might address is how much larger the welfare cost of business cycles becomes if one takes into account the feature that firms are more tolerant of prices that are too high than of prices that are too low. Presumably, the answer to this question is relevant for analyses of the optimal inflation target; with higher trend inflation the need for downward adjustment diminishes.

Finally, our analysis serves as yet another reminder of the general insight that tiny adjustment costs can have a large impact on behavior under conditions of ongoing uncertainty (Mankiw, 1985; Dixit, 1991). Seen from this perspective, it is perhaps not too surprising that the third-derivative of the profit function has a first-order impact on pricing.

CRediT authorship contribution statement

Tore Ellingsen: Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization; **Richard Friberg:** Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization; **John Hassler:** Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization.

---

[17] The lower trigger is now 3.53 % below the profit maximizing price while the upper trigger and the return point are 4.27 % and 0.46 % above, respectively.

[18] They also add other features, such as fat-tailed shocks (Alvarez and Lippi, 2014) and random menu costs (Alvarez et al., 2016).

Data availability

No data was used for the research described in the article.

Declaration of competing interest

None.

Acknowledgments

Supplementary material

Supplementary material associated with this article can be found in the online version at 10.1016/j.ijindorg.2026.103259

References

Alvarez, F., Le Bihan, H., Lippi, F., 2016. The real effects of monetary shocks in sticky price models: a sufficient statistic approach. Am. Econ. Rev. 106 (10), 2817–2851.

Alvarez, F., Lippi, F., 2014. Price setting with menu costs for multiproduct firms. Econometrica 82 (1), 89–135.

Alvarez, F., Lippi, F., 2022. The analytic theory of a monetary shock. Econometrica 90 (4), 1655–1680.

Anderson, E., Jaimovich, N., Simester, D., 2015. Price stickiness: empirical evidence of the menu cost channel. Rev. Econ. Stat. 97 (4), 813–826.

Arbatskaya, M., Baye, M.R., 2004. Are prices 'sticky'online? market structure effects and asymmetric responses to cost shocks in online mortgage markets. Int. J. Ind. Org. 22 (10), 1443–1462.

Asane-Otoo, E., Dannemann, B.C., 2022. Rockets and feathers revisited: asymmetric retail gasoline pricing in the era of market transparency. Energy J. 43 (6), 103–122.

Assad, S., Clark, R., Ershov, D., Xu, L., 2024. Algorithmic pricing and competition: empirical evidence from the german retail gasoline market. J. Polit. Econ. 132 (3), 723–771.

Bacon, R.W., 1991. Rockets and feathers: the asymmetric speed of adjustment of UK retail gasoline prices to cost changes. Energy Econ. 13 (3), 211–218.

Barro, R.J., 1972. A theory of monopolistic price adjustment. Rev. Econ. Stud. 39 (1), 17–26.

Bayer, R.-C., Ke, C., 2018. What causes rockets and feathers? an experimental investigation. J. Econ. Behav. Org. 153, 223–237.

Benzarti, Y., 2024. Tax incidence anomalies. Natl. Bureau Econ. Res. WP No 32819 .

Bils, M., Klenow, P.J., 2004. Some evidence on the importance of sticky prices. J. Polit. Econ. 112 (5), 947–985.

Borenstein, S., Cameron, A.C., Gilbert, R., 1997. Do gasoline prices respond asymmetrically to crude oil price changes? Q. J. Econ. 112 (1), 305–339.

Bulutay, M., Hales, D., Julius, P., Tasch, W., 2021. Imperfect tacit collusion and asymmetric price transmission. J. Econ. Behav. Org. 192, 584–599.

Burstein, A., Gopinath, G., 2014. International prices and exchange rates. In: Handbook of International Economics. Elsevier. Vol. 4, pp. 391–451.

Byrne, D.P., De Roos, N., 2019. Learning to coordinate: a study in retail gasoline. Am. Econ. Rev. 109 (2), 591–619.

Bénabou, R., Konieczny, J.D., 1994. On inflation and output with costly price changes: a simple unifying result. Am. Econ. Rev. 84 (1), 290–297.

Chen, H.A., Levy, D., Ray, S., Bergen, M., 2008. Asymmetric price adjustment in the small. J. Monet. Econ. 55 (4), 728–737.

Cho, S., Rust, J., 2010. The flat rental puzzle. Rev. Econ. Stud. 77 (2), 560–594.

von Cramon-Taubadel, S., Goodwin, B.K., 2021. Price transmission in agricultural markets. Annu. Rev. Resour. Econ. 13 (1), 65–84.

Danziger, L., 1999. A dynamic economy with costly price adjustments. Am. Econ. Rev. 89 (4), 878–901.

DellaVigna, S., Genzkow, M., 2019. Uniform pricing in u.s. retail chains. Q. J. Econ. 134 (4), 2011–2084.

Dixit, A., 1991. Analytical approximations in models of hysteresis. Rev. Econ. Stud. 58 (1), 141–151.

Dutta, S., Bergen, M., Levy, D., Venable, R., 1999. Menu costs, posted prices, and multiproduct retailers. J. Money Credit Bank. 31, 683–703.

Ellingsen, T., Friberg, R., Hassler, J., 2006. Menu costs and asymmetric price adjustment. CEPR Discussion Paper No 5749.

Eyster, E., Madarasz, Kristof, M.P., 2021. Pricing under fairness concerns. J. Eur. Econ. Assoc. 19 (3), 1853–1898.

Friberg, R., Romahn, A., 2018. Pass-through by multi-product firms. Int. J. Econ. Bus. 25 (2), 265–295.

Gautier, E., Conflitti, C., Faber, R.P., Fabo, B., Fadejeva, L., Jouvanceau, V., Menz, J.-O., Messner, T., Petroulas, P., Roldan-Blanco, P., et al., 2024. New facts on consumer price rigidity in the euro area. Am. Econ. J.: Macroecon. 16 (4), 386–431.

Gertler, M., Leahy, J., 2008. A Phillips curve with an Ss foundation. J. Polit. Econ. 116 (3), 533–572.

Golosov, M., Lucas, R.E., 2007. Menu costs and Phillips curves. J. Polit. Econ. 115 (2), 171–199.

Hong, G.H., Li, N., 2017. Market structure and cost pass-through in retail. Rev. Econ. Stat. 99 (1), 151–166.

Karadi, P., Reiff, A., 2019. Menu costs, aggregate fluctuations, and large shocks. Am. Econ. J.: Macroecon. 11 (3), 111–146.

Klenow, P.J., Kryvtsov, O., 2008. State-dependent or time-dependent pricing: does it matter for recent US inflation? Q. J. Econ. 123 (3), 863–904.

Kuran, T., 1983. Asymmetric price rigidity and inflationary bias. Am. Econ. Rev. 73 (3), 373–382.

Levy, D., Bergen, M., Dutta, S., Venable, R., 1997. The magnitude of menu costs: direct evidence from large US supermarket chains. Q. J. Econ. 112 (3), 791–824.

Levy, D., Bergen, M., Venable, R., 1998. Price adjustment at multiproduct retailers. Manager. Decis. Econ. 19 (1), 81–120.

Mankiw, N.G., 1985. Small menu costs and large business cycles: a macroeconomic model of monopoly. Q. J. Econ. 100 (2), 529–538.

Midrigan, V., 2011. Menu costs, multiproduct firms, and aggregate fluctuations. Econometrica 79 (4), 1139–1180.

Nakamura, E., Steinsson, J., 2008. Five facts about prices: a reevaluation of menu cost models. Q. J. Econ. 123 (4), 1415–1464.

Nakamura, E., Zerom, D., 2010. Accounting for incomplete pass-through. Rev. Econ. Stud. 77 (3), 1192–1230.

Neumark, D., Sharpe, S.A., 1992. Market structure and the nature of price rigidity: evidence from the market for consumer deposits. Q. J. Econ. 107 (2), 657–680.

Orbach, B.Y., Einav, L., 2007. Uniform prices for differentiated goods: the case of the movie-theater industry. Int. Rev. Law Econ. 27 (2), 129–153.

Peltzman, S., 2000. Prices rise faster than they fall. J. Polit. Econ. 108 (3), 466–502.

Sayag, D., Snir, A., Levy, D., 2024. Small price changes, sales volume, and menu cost. arXiv preprint arXiv:2403.07166.

Tappata, M., 2009. Rockets and feathers: understanding asymmetric pricing. Rand J. Econ. 40 (4), 673–687.

Tsiddon, D., 1993. The (mis) behaviour of the aggregate price level. Rev. Econ. Stud. 60 (4), 889–902.

Weyl, E.G., Fabinger, M., 2013. Pass-through as an economic tool: principles of incidence under imperfect competition. J. Polit. Econ. 121 (3), 528–583.

Yang, H., Ye, L., 2008. Search with learning: understanding asymmetric price adjustments. Rand J. Econ. 39 (2), 547–564.

Zbaracki, M.J., Ritson, M., Levy, D., Dutta, S., Bergen, M., 2004. Managerial and customer costs of price adjustment: direct evidence from industrial markets. Rev. Econ. Stat. 86 (2), 514–533.