

Relative Prices as Aggregate Supply Shocks with Trend Inflation

David Demery
Nigel W. Duck

Discussion Paper No. 05/570

Revised July 2007

Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN

Relative Prices as Aggregate Supply Shocks with Trend Inflation

David Demery
Nigel W. Duck

JEL codes: E30 E31

Key Words: inflation, menu-cost, commodity price variance, commodity price skewness.

We would like to thank Jon Temple, the editor and two anonymous referees of this journal for very helpful comments on an earlier version. Remaining errors are our own.

David Demery is *Reader in Economics, University of Bristol* (E-mail: david.demery@bris.ac.uk). Nigel W. Duck is *Reader in Economics, University of Bristol* (E-mail: n.w.duck@bris.ac.uk).

Abstract

Ball and Mankiw (1995) use a static menu-cost model to explain the historical behaviour of the first and higher moments of commodity price changes in US producer prices. We show that when appropriately modified for a world of positive trend inflation and forward-looking behaviour by firms, the menu-cost model predicts a much weaker (possibly zero) correlation between the mean and the skewness of price changes than that found in the data.

The empirical correlation between the first and second moments of the distribution of price changes has long been recognized. Vining and Elwertowski (1976) and Parks (1978) provide evidence of this phenomenon in the US; Domberger (1987) provides similar evidence for the UK.¹ Ball and Mankiw (1995) - henceforth BM - propose a menu-cost model to explain not only this correlation but also the correlation between inflation and the third moment of the distribution, the skewness of commodity price changes. Indeed the distinctive prediction of their model is that there should be a strong positive correlation over time between inflation and skewness.

The central idea in BM's model is that if, because of menu-costs, firms do not re-set prices when faced with small shocks to their desired price, a skewed distribution of such shocks will affect the mean of actual price changes, the inflation rate: a positively-skewed distribution will raise it; a negatively-skewed distribution will lower it.² They go on to present evidence based on annual post-WW2 data on US producer prices, which, they argue, shows that the observed relationships between the three moments of the distribution of price changes are broadly in line with their model; in particular skewness and inflation are strongly positively correlated. From this they argue more generally that "the empirical validation of our theory provides evidence for menu-cost models of price adjustment" and they add, "menu-cost models were developed to explain monetary non-neutrality. They gain credibility from their ability to fit the facts regarding inflation and relative-price change" (BM, p191).

However, BM's model incorporates two questionable assumptions: first that there is no trend inflation; and second that firms are not forward-looking in their price-setting behaviour. Recognizing the possible limitations of such a static model, BM suggest that “future research might consider a fully dynamic version” (p.172). In this paper we develop such a model, and use it to assess the robustness of BM's interpretation.

The principal implication of our dynamic version is that the distinctive prediction of BM's model is noticeably weakened in the presence of positive trend inflation. In fact, when calibrated to mimic the main features of US producer price behaviour over the period 1948-2003, our model predicts only a weak, even zero, correlation between inflation and skewness, and therefore cannot explain the observed strong positive correlation. A version calibrated using a subset of low trend-inflation periods predicts a somewhat stronger correlation between skewness and inflation, though the predicted co-movements are still weaker than those observed.

Writing of the observed correlation between the first and second moments of the distribution of price changes Fischer (1982, p171) commented that “the nature of the association and its causation remain unclear”. Our results suggest that Fischer's comment could equally well be applied now to the observed correlations between the first three moments of the distribution of price changes; the menu-cost approach does not by itself provide a satisfactory explanation for them.³

Our paper is organized as follows. In section 1 we sketch BM's static menu-cost model; in section 2 we extend it to allow for trend inflation, maintaining the static framework; in section 3 we develop a dynamic menu-

cost model with trend inflation and forward-looking behaviour by firms; in section 4 we present our simulation results; and in a final section we offer brief summary.

1. THE BALL-MANKIW MODEL

BM assume an economy consisting of a continuum of industries, each with a continuum of imperfectly competitive firms. Firms within the same industry have identical *desired* prices but face heterogeneous menu costs. For simplicity, initial desired and actual price levels (in logs) are assumed equal to zero. Desired price changes across industries are then subject to a mean-zero shock, θ , which is distributed across industries with density $f(\cdot)$. Since each price is in logarithms and has an initial equilibrium of zero, θ is the rate of change of the equilibrium price. A firm which re-sets its price incurs a menu cost of C ; a firm which does not incurs a loss of θ^2 .⁴

A firm will therefore adjust its price only if $|\theta| > \sqrt{C}$. BM assume \sqrt{C} is distributed across firms with distribution function $G(\sqrt{C}) = 1 - \exp(-a\sqrt{C})$ where a is parameter. For an industry confronted with shock θ , the proportion of firms that re-set will be $G(|\theta|)$. Given the static framework, re-setting firms select a price θ . BM derive π , the inflation rate, as the average of these price changes across industries,

$$\pi = \int_0^{\infty} \theta G(\theta) [f(\theta) - f(-\theta)] d\theta \quad (1)$$

If the density of $f(\theta)$ is symmetric around zero, $f(\theta) - f(-\theta)$ and π is zero. But if $f(\theta)$ is asymmetric (i.e. skewed), then although the mean *desired*

price change is zero, the *observed or actual* inflation rate, π , will not be. BM use numerical methods to demonstrate the impact of the asymmetry of $f(\theta)$ on π by assuming that $f(\cdot)$ is skew-normal and a takes the value 7.⁵ They derive the mean, standard deviation and skewness of *actual* price changes under varying assumptions about the standard deviation and skewness of the shocks to *desired* prices. They find that increased skewness in the underlying shocks raises the inflation rate, an effect amplified as the standard deviation of the shocks rises. Similarly, a rise in the standard deviation raises π , an effect which is amplified by higher skewness of the shocks.⁶

2. THE BALL-MANKIWI MODEL WITH TREND INFLATION

BM use (1) to show that if θ is *symmetrically* distributed around zero then, regardless of whether there are menu costs, inflation will not diverge from the mean value of θ , zero: the “range of inaction” caused by menu costs - the range of shocks to which firms will not respond by changing price - will prevent just as many price falls as equivalently-sized price rises. Our first point is that when the mean of θ (μ_θ) is positive, this is no longer the case: whilst the distribution of θ has been shifted to the right the range of inaction will still be centred around zero; and so there will be fewer firms who wish to lower their price but, because of menu costs, choose not to do so than there are firms wishing to raise them but choosing not to. Hence, in the presence of menu costs, if μ_θ is positive, inflation will, in general, be below μ_θ even when the distribution of θ is symmetric. Only in the limiting case of hyperinflation, when the distribution of θ has been shifted so far to the right that the loss of

profits from not resetting overwhelms all firms' menu costs, will this effect disappear and the higher moments of price changes have little or no effect on the mean.⁷ More generally, positive trend inflation can be expected to change the co-movements of mean, standard deviation and skewness. Figure 1 illustrates the point. For clarity, the upper and lower bounds that trigger re-setting behaviour are assumed to be the same for all firms and assumed densities, and, when they re-set, firms set price to the desired level. In panel (a) the desired price shocks are symmetric around a zero mean. The shaded area indicates the densities of the prices being re-set: firms hit by shocks greater than the upper threshold raise their prices; those hit by shocks below the lower threshold lower them. The net effect is zero and the mean of *actual* price increases is equal to the mean of the *desired* price increases: both are zero.

In panel (b) we assume a trend inflation of 0.03, so the mean of the symmetric distribution is now positive (and indicated by the long-dash vertical line). The unshaded area of the density captures the price changes that would have been made by firms who, because of menu costs, have not reset. In this case the unshaded area above zero (to the right of the short-dash vertical line) is greater than that below zero, so a greater proportion of non-resetting firms would, in the absence of menu costs, have re-set at a higher rather than lower price. Because these firms do not re-set, the actual inflation rate (the bold vertical line) will be below the mean of desired price changes. At higher trend inflation rates this effect will gradually disappear as the number of firms not re-setting falls.

In panel (c) of Figure 1 we repeat BM's case of positive skewness in the desired price distribution with zero trend inflation. The area of the unshaded density above zero is less than that below. Because these firms do not re-set, the mean of actual price changes (bold vertical line) is above the mean of desired price changes (long-dash vertical line). This is how BM explain the co-movement of the skewness of price changes and their mean. In panels (d) through (f) the densities are all positively skewed, but the mean is progressively raised from 3% to 20%. As the distribution moves to the right, the respective unshaded densities above and below zero change. In panel (e) for example the proportion of firms who would have raised their price in the absence of menu costs is far greater than the very small proportion who would have lowered them, and hence the mean of the actual price change (bold vertical line) is lower than the mean of desired price changes (long-dash vertical line). When the mean of desired price changes is 0.2, skewness, menu costs and their interaction become irrelevant: virtually all firms re-set and the means of actual and desired price changes become equal.

The graphs illustrate that the effect of menu costs on the moments of price changes is likely to be more complex when there is trend inflation. To demonstrate this complexity more formally, consider an economy where θ is symmetrically distributed. And assume, with BM, that all initial prices are equal to their desired values of zero and that firms, when they re-set, set a price of θ .⁸ What will be the effect of variations in the mean and standard deviation of θ on the mean and higher moments of actual prices? We numerically generate the distribution of *actual* price changes, $\theta G(|\theta|)$, and calculate the mean (μ_p), standard deviation (σ_p) and skewness (k_p) of this

distribution given different assumptions about the mean (μ_θ) and standard deviation (σ_θ) of θ .⁹ Note (from Figure 1b for example) that whereas there is no skewness in the distribution of *desired* price changes, the failure of some firms to re-set when confronted by sufficiently small shocks means that the distribution of *actual* price changes will be positively skewed if $\mu_\theta > 0$. Table 1 presents the results. The first three rows in the table are based on $\mu_\theta = 0$.¹⁰ The inflation rate (μ_p) and skewness (k_p) are both zero and independent of σ_θ and σ_p , confirming BM's results for this case. The remaining rows set out the moments of actual prices when $\mu_\theta > 0$. With 5% trend inflation, a rise in the standard deviation of θ raises σ_p and μ_p but lowers k_p . Looking across the 5% rows one would be tempted to expect skewness (k_p) and inflation (μ_p) to be *negatively* correlated. With the more extreme 25% trend inflation, a rise in the standard deviation of θ again raises σ_p and μ_p but now it also generally raises k_p .

The moments set out in Table 1 are based on the simplest possible framework: a static model in which firms are confronted by symmetric shocks to their desired prices. Yet even here the relationships between the moments of actual price changes are noticeably more complex than those BM derived for the zero trend inflation case. And the degree of complexity would be even greater had we allowed for skewness in the distribution of θ .

3. A DYNAMIC MENU-COST MODEL

An obvious weakness in BM's original model, and in the model allowing for trend inflation that we have just described, is that they are both static, or single-period, models. Such a framework precludes firms from taking into account the potential *future* benefits from re-setting. One likely consequence of this weakness is that the model will underestimate the proportion of firms that re-set: the menu cost of changing a price is incurred in a single period whereas the benefits could be felt over all future periods; firms which do not find the benefit of resetting in the current period worth the cost might find it profitable if they took into account future benefits.¹¹

A second likely consequence is that a forward-looking, re-setting firm will not set a price of θ . The firm will anticipate that, in the presence of trend inflation, its desired price will be higher in the future. Recognizing that it may not re-set its price for some time, a firm which is re-setting its price will therefore typically set a price *above* its currently desired level, accepting higher current losses in return for lower expected losses in the future. Furthermore, because it is more likely not to be re-setting its price for some time thereafter, a re-setting firm with a *high* menu cost will set a price higher than a re-setting firm with a low menu cost. Less obviously, firms faced with skewed shocks will not re-set at θ even with zero trend inflation. For example, a firm faced with a positively skewed distribution of shocks will expect frequent small negative shocks, each of which may not justify a price re-set, and infrequent, larger positive shocks which do. If it re-sets its current price below θ the small negative future shocks will, for some time at least,

nudge its desired price towards its actual price; and any positive shock is more likely to be large enough to justify re-setting even if the current price were set at θ rather than below it. So there are possible gains from setting the current price below θ and possibly little to lose. Symmetrically, a firm subject to shocks from a negatively skewed distribution will typically set a price above θ .

It follows from all this that, when a firm is forward-looking, the mean, standard deviation and skewness of the shocks will have effects on its pricing behaviour that will be quite different from those implied by a static model. To explore the implications of forward-looking behaviour by firms we develop a dynamic version of the menu-cost model which we then use to simulate the relationship between inflation and the higher moments of the distribution of price changes.

Our “benchmark” model retains many features of BM's static model. We assume a model economy consisting of a large number of industries, each with a large number of imperfectly competitive firms. Firms within each industry share the same desired price but they have different menu costs, with \sqrt{C} distributed across firms with distribution function $G(\sqrt{C}) = 1 - \exp(-a\sqrt{C})$.

Firm i in industry j is assumed to incur a loss in each period t of $(p_{i,j,t} - p_{j,t}^*)^2$

where $p_{i,j,t}$ is the actual (log) price and $p_{j,t}^*$ is the equilibrium price. The change in equilibrium price (θ in BM's notation) has a known distribution, but in our first departure from BM's model we assume that it has a positive mean for all industries. In a second departure we assume that firms are forward looking. This requires us to derive a set of price-setting rules which are

appropriate for a dynamic framework: it is no longer valid to assume that firms will reset when $(p_{i,j,t} - p_{j,t}^*)^2 > C$ and that a firm which does reset will reset at $p_{i,j,t} = p_{j,t}^*$.

3.1 Price-setting rules

We assume that the firm has an infinite planning horizon and is faced with i.i.d. shocks to its desired price. Conditional on its menu cost and the distribution of the desired price shocks, the optimal policy of such a firm can be characterized by three numbers: an upper threshold t_U , such that if $p - p^* > t_U$ the firm changes its price; a lower threshold t_L , such that if $p - p^* < t_L$ the firm also changes its price; and a reset value, r , that $p - p^*$ is set to whenever price is changed. Because we assume that each firm's menu cost is fixed over time, the firm's strategy is conveniently summarized in these three parameters. For firms with zero menu costs, $t_U = t_L = r = 0$. For positive menu-cost cases we obtain estimates of t_U, t_L and r by a numerical grid-search method. For computational tractability, we derive estimates of the three parameters for ten menu-cost “nodes”, corresponding to the first nine deciles and the 99.9 percentile of the assumed distribution of C , and use linear interpolation to derive estimates of the three parameters for any particular menu cost. For a firm i with menu-cost node C_j , the realised present value of the loss function is,

$$PVL_{i,t} = \sum_{k=0}^T \beta^k \left[D_{i,t+k} C_i + (p_{i,t+k} - p_{t+k}^*)^2 \right] \quad (2)$$

where $D_{i,t+k}$ is an indicator variable which takes the value 1 when the firm re-sets its price and zero at other times; and $\beta \left(\equiv (1-\delta)^{-1} \right)$ is a discount parameter, with δ assumed to be 2%. We approximate the *PVL* of infinite-horizon firms by setting $T = 200$. We draw $M (= 1,000)$ paths of p^* over T periods and grid-search¹² over values t_u , t_l and r that minimize,

$$\frac{1}{M} \sum_{i=0}^M PVL_{i,t}$$

3.2 Model Calibration

Our main aim is to investigate whether the dynamic menu-cost model can account for the observed correlations of mean, standard deviation and skewness found in US producer price indices (PPI) over the period 1948 to 2003.¹³ To do so we calibrate our benchmark model so that it matches as closely as possible key features of these data, namely the mean inflation rate, the standard deviation across commodities of their price changes, the standard deviation of this standard deviation over time, the skewness of price changes and the standard deviation of skewness over time. Once the model has been satisfactorily calibrated we then use it to predict the correlations between inflation, the standard deviation and skewness of price changes. The test of the menu-cost approach is therefore the success it has in predicting these correlations. Following BM, we use PPI largely because of its high

degree of disaggregation: price data were readily available at the four-digit level. The number of industries covered each year varies, ranging from 203 in 1947 to 368 in 1995.¹⁴ 151 industries appeared every year. We give details of the inflation rate and weighted¹⁵ standard deviations, skewness and the Q statistic of PPI inflation rates in the appendix. Q captures both variance and skewness in a single descriptive statistic. It is defined as,

$$Q = \int_{-\infty}^{\infty} |q| \cdot qh(q) dq \quad (3)$$

where q is the commodity's price change minus the inflation rate and $h(q)$ is the density of q . Q averages the product of each relative price change and its own absolute value. It takes the value zero for a symmetric distribution, moves positively with skewness and is amplified by a larger variance.

The mean “all-commodity” PPI inflation rate was 3% over the data period, peaking at over 17% in 1974. All three moments of price changes - mean, standard deviation and skewness - vary substantially over the period. This is particularly true of skewness, which frequently switched sign from one year to the next: for example (weighted) commodity price changes were negatively skewed in 2002 (-2.4) and equally positively skewed a year later (2.6). The mean weighted standard deviation was 6.1%, and there was also substantial annual variation, ranging from 2.6% in 1963 to over 14% in 1974. The simple correlations with inflation over the full period are 0.608 for weighted skewness, 0.411 for weighted standard deviation and 0.805 for weighted Q . It is these correlations that we seek to explain using a dynamic menu-cost model.

To capture the time variation in PPI skewness, we follow a suggestion of Bryan and Cecchetti (1999): shocks to desired price changes are drawn from

a distribution exhibiting excess kurtosis. In any single year, a relatively small sample of such shocks will display positive skewness if a number of shocks are drawn from the upper tail of the distribution. And similarly negative skewness will be present in a small sample containing a number of draws from the lower tail without any counter-balancing draws from the upper tail. In this way we can mimic the time variation of skewness exhibited in the US PPI data.

In fact Bryan and Cecchetti (1999) argued that such an underlying data-generating process can explain why skewness is correlated with the mean: a small-sample containing draws from the upper (lower) tail will raise (lower) both mean and skewness, with no underlying economic mechanism required. Ball and Mankiw (1999) respond with an appeal to the standard theory of price determination: the aggregate inflation rate is determined by monetary factors whilst real factors (such as demand and costs in the various industries) determine $N-1$ relative prices, where N is the number of industries. According to BM, Bryan and Cecchetti “merely offer a statistical version of the layman's misconception” - that sectoral shocks determine all N nominal prices and that the aggregate price level is simply determined by the average of these prices.¹⁶ To avoid this misconception, we assume that the change in the *aggregate* equilibrium price level is determined by an exogenous “monetary process” - characterized by a given mean and variance. The N desired *relative* prices are required to have a mean-zero distribution exhibiting excess kurtosis.¹⁷

We assume that the change in industry j 's equilibrium price is influenced by two factors: a “monetary” component ($\mu + \varepsilon_t$, common to all industries) and an industry-specific relative price shock ($v_{j,t}$), so we write,

$$p_{j,t}^* - p_{j,t-1}^* = \mu + \varepsilon_t + v_{j,t} \quad (4)$$

where the money component has a mean of μ and a variance of σ_ε^2 ; and the relative price shock $v_{j,t}$ has a variance σ_v^2 which is assumed common to all industries. The presence of the common shock will mean that desired and actual price changes will be correlated across industries. The mean of the “all-commodity” PPI over the period 1948-2003 was 0.03 and we adopt this in the benchmark model as our estimate of μ . To derive an estimate of σ_ε^2 we assume that the variance of all-commodity PPI inflation is partly attributable to variations in the moments of industry price changes (their skewness and standard deviation as the menu-cost model predicts) and partly due to variations in the common shock. σ_ε^2 is estimated as the residual variance in a regression of inflation on the skewness and the standard deviation of producer prices. Over the period 1948-2003 the standard deviation of the all-commodity inflation rate was 0.04 and the standard deviation of the regression residual was 0.03, which we take as our estimate of σ_ε .¹⁸

The relative price shocks ($v_{j,t}$) of N industries are assumed to be drawn from a mean-zero Student's t -distribution with 2.5 degrees of freedom, scaled to replicate the average standard deviation of observed price changes.¹⁹ As we later note, this degree of freedom generated small sample distributions with patterns of skewness and standard deviation in simulated actual industry

price changes similar to those observed in the data. The relative *desired* price shocks are assumed to have a standard deviation $\sigma_v = 0.08$, as this replicates the observed average standard deviation of actual price changes in the PPI data. Although the relative prices were drawn from an assumed mean-zero distribution, aggregation consistency is ensured by scaling the relative prices so they sum to zero.

Trend inflation since the early 1980s has been substantially lower than that over the full period. When applied to the PPI data, a Markov switching model²⁰ identifies two distinct inflation “regimes”: a low inflation regime covering the periods 1953-72 and 1983-2003 and a high inflation regime over the other years. The mean inflation rate in the low-trend period was 1.5% and over 7% in the 15 years of high-trend inflation. During the low-trend period the mean standard deviation of the all-commodity inflation rate was substantially lower - 2% compared with 4% over the full period. And the average standard deviation of commodity inflation rates was a little lower with a mean of 5.4% (compared with 6.1% over the full period). The simple correlations with inflation over the low-inflation sub-periods are (with the full period correlations in brackets) 0.718 (0.608) for weighted skewness, 0.025 (0.411) for weighted standard deviation and 0.749 (0.805) for weighted Q.²¹ In a further set of simulations, we therefore re-calibrated the model to capture these features of the low-inflation regime. Following the approach described above, the parameters for low trend inflation calibration are: $\mu = 0.015$, $\sigma_v = 0.07$ and $\sigma_\varepsilon = 0.015$.

3.3 Simulation procedure

Our simulation results presented below are based on 500 replications of the simulated behaviour over 66 years of 10,000 firms equally divided into 100 industries. For these simulations we first, from its assumed distribution, randomly assign to each firm a menu cost which is assumed constant over time. By comparing this cost with the relevant pair of adjacent menu-cost “nodes”, we derive the firm's unique t_U , t_I and r by linear interpolation. Following BM, in the benchmark model we set the menu-cost parameter $a=7$. This implies (at least in the static context) that the maximum deviation between actual and desired prices the average firm will tolerate is 15% which BM argue is consistent with the observed frequency of price adjustment. In our dynamic model this parameter value leads to a distribution of contract lengths which, as we explain below, is broadly in line with US experience.

We next set the initial equilibrium price for each commodity to zero but we assign an initial price to each firm within each industry as follows. First, for each menu-cost node we compute a numerical steady-state distribution of $p-p^*$.²² Each firm's initial price is then a random draw from the numerical distribution of the node which is the closest match to its own menu cost. From these initial prices we then generate 66 years of observations for the 10,000 firms. To further remove any influence of initial conditions, we drop the first 10 years observations: the remaining 56-year period then matches the 1948-2003 PPI sample.²³

The simulated commodity price is then the mean price across firms in the industry. Our simulations of skewness and standard deviation are based on the distribution of the implied series on *changes* in these prices.

4. SIMULATION RESULTS

The key simulation results of the benchmark model are set out in Table 2. In panel (a) we present the estimated values of t_u , t_l and r for the ten menu-cost nodes and we present the first nine graphically in Figure 2. A firm with the median menu cost will not reset price if $-0.076 < p - p^* < 0.099$. This range of inaction is not symmetric around zero and this asymmetry is more pronounced as menu costs rise. The reasons for this were first identified by Ball and Mankiw (1994): if the firm's current price is above its equilibrium, inflation (which implies a rise in p^*) will close the future gap; but if the current price is below equilibrium, inflation will further widen the gap; hence the firm will tolerate a larger positive gap between their current and desired price than a negative one. A re-setting firm with median menu cost re-sets at $p = p^* + 0.0163$: anticipating a rise the desired price in the future, the firm re-sets above the current equilibrium in an attempt to avoid incurring future menu costs. Figure 2 compares these parameters with those implied by the static model, where $-\sqrt{C} < p - p^* < +\sqrt{C}$ is the range over which a firm will not re-set and $r = 0$. The outer dashed lines are the upper and lower reset thresholds in the static model and these are symmetric functions of menu costs. The dynamic thresholds (the inner solid lines) are lower than their static equivalents in absolute terms: firms are more inclined to adjust their prices because they gain both in the present and the future.

According to Blinder (1994), 40% of the US prices he sampled were reset annually, 10% were reset more frequently and 50% less frequently. In

their more recent analysis of consumer goods prices, Bils and Klenow (2004) found that half of goods prices lasted less than 4.3 months. Blinder's survey mainly comprised firms producing intermediate goods (79% of their sales) whereas Bils and Klenow's coverage was exclusively consumer goods retail prices. As our calibrations are based on producer price data, the Blinder findings are, perhaps, more relevant. In our simulations the median "contract length" was one year. 53.7% of prices lasted only one year and 46.3% remained unchanged over two or more years. Noting that in our simulations price contracts of less than one year are not possible, these results are broadly in line with Blinder's findings.

In panel (b) of Table 2 we compare the model's simulations with the US data. The 3% trend inflation rate in the data is confirmed in the simulations. The assumption of excess kurtosis in the distribution of price shocks does lead to time-variation in the skewness and standard deviation of industry prices similar to those found in the data: the mean simulated standard deviation of skewness over time is 2.5 compared with 2.4 in the data and the mean simulated cross-time standard deviation of the standard deviation of prices is 0.029, compared with 0.027 in the data. The mean simulated standard deviation of actual price changes was 0.067 (compared with a historical mean of 0.061). The model is less successful at capturing the time variation in inflation: the mean simulated value of the standard deviation of inflation is 0.024 whereas in the data it is 0.043. As we shall see this arises because in the simulations, skewness and the standard deviation have no significant effect on inflation and hence the variability of inflation over time arises only because of the common shocks.²⁴

In panel (b) we also compare (with their weighted historical equivalents) the simulated correlations of inflation with skewness (SK), standard deviation (SD) and Q . In the benchmark model the simulated correlations are all substantially lower than those found in the data, particularly in the cases of SK and Q . Interpreting the 2.5th and 97.5th percentile range as the bootstrapped 95% confidence interval, the correlations of SK and Q observed in the data are significantly above those we would expect from the dynamic menu-cost model and the observed correlation with SD is close to the simulated upper bound. In this (and later) simulations we find that when the historical correlation of inflation with SK is outside the model's 95% confidence interval, the same is true of Q .

We repeated the simulation exercise for a model economy with zero and 6% trend inflation, all other model parameters set to their benchmark values. The results are shown in the top panels of Table 3. For the case of zero trend inflation (panel (a)), the simulated correlations of inflation with skewness and Q are stronger than those in the benchmark case though still substantially weaker than those observed. For example the observed correlation of inflation with SK is 0.608, whereas with zero trend inflation we would expect a correlation in the range -0.006-0.452. The correlation with the standard deviation is now substantially weaker than both the benchmark case and its actual value. With trend inflation at 6% (panel (b) in Table 3), the simulated correlation between the standard deviation of price changes and inflation is more in line with its actual value but the mean simulated correlation of skewness and inflation is now mildly negative - something we might have expected given the results of the static model with symmetric shocks reported

in Table 1. Looking across the 5% rows in Table 1, the inflation rate is positively related to the standard deviation in actual prices but negatively correlated with their skewness. So, BM's prediction of a correlation of skewness with inflation appears to be mainly a feature of an economy with zero trend inflation, and even in this case the correlation predicted by our menu-cost model with forward-looking firms is distinctly weaker than that observed.

In the lower panels of Table 3 we report two robustness checks. In panel (c) we raise menu costs by assuming $a = 4$, which implies that the maximum deviation between actual and desired prices the average firm will tolerate is approximately 25% (compared with 15% in the benchmark case).²⁵ The mean simulated correlations are only marginally affected by this change.

In panel (d) we report simulated correlations from a model where the standard deviation of desired price changes is twice that of the benchmark model, i.e. 0.16 compared with 0.08. The mean simulated correlation of inflation with *SK* (0.136) is higher than the benchmark and that with *SD* (0.068) substantially lower. In this high variance case the menu-cost model fails to explain either the correlation of inflation with *SK* or its correlation with *SD*.

The results of a model calibrated to capture the low trend inflation periods are set out in Table 4. In this case the assumed trend inflation is lower and monetary and relative price shocks less volatile than the benchmark case. Because of these combined changes, the thresholds t_u , t_l and r are very similar to those of the benchmark case: the zone of inaction is narrower but the differences are marginal. Whilst the mean simulated correlation of

inflation with SK (0.200) is higher than the benchmark (0.088), the historical correlation (0.718) is still well outside the model's 95% confidence interval (-0.115 to 0.493). So, in a dynamic setting with trend inflation as low as that experienced in the US since the early 1980s a menu-cost model provides a poor explanation of the observed correlations amongst the moments of commodity price changes, especially between the first and third moments.²⁶

The failure of menu-cost effects to match the data is potentially due to one or all of the following: (i) the effects are weakened when we move from a static model of the firm to dynamic one; (ii) trend inflation complicates the relationships among the moments of the distribution of price changes; (iii) allowing for common shocks introduces “noise” into the inflation process and this masks the expected correlations between the moments. We present evidence on the separate importance of these three influences in Table 5.

In panel (a) of this table we report correlations based on a static model of the firm under differing assumptions about the common shock's mean and standard deviation.²⁷ The simulated correlations of inflation with skewness are far stronger than those with the standard deviation but the former is weakened when the variance of the common shock is raised. In this static framework, when the mean of θ is zero the observed correlation of inflation with skewness can only be explained if no common shocks are assumed; and the observed correlation with the standard deviation of producer prices cannot be explained. When the mean value of θ is assumed to be 0.03 the correlation of inflation with skewness is reduced whilst the correlation of inflation with the standard deviation of product prices is raised. As in the zero trend case the

correlations with both SK and SD are substantially weakened when the variance of the common shocks is raised.

In panel (b) we report simulated correlations in a dynamic model with forward-looking firms. Moving from static to dynamic modelling of firm behaviour has only marginal effects on the correlations. The presence of common shocks and the assumption of a positive trend inflation thus both substantially weaken the predicted correlation of skewness and inflation. The assumption of forward-looking behaviour by firms has a relatively minor effect.

5. CONCLUSIONS

The relationship between first- and higher-order moments of the distribution of price changes has been a puzzle to economists for many years. To explain it Ball and Mankiw (1995) put forward a menu-cost model whose distinctive prediction is of a correlation between the skewness of price changes and inflation rate. We have shown that when appropriately modified for a world of non-zero trend inflation and forward-looking behaviour by firms, and when calibrated to capture the main features of US data over the period 1948-2003, the menu-cost model predicts a substantially weaker correlation between inflation and the skewness of price changes than that found in the US Producer Price data. Our conclusions are much the same when we calibrate using data from the low inflation sub periods, 1953-1972 and 1983-2003. Overall, our results suggest that when the menu-cost model is modified to allow for non-zero trend inflation and forward-looking behaviour by firms it does not by itself provide a satisfactory explanation for the observed correlations between the moments of the distribution of price changes.

¹ The relationship between individual prices and the aggregate price level was first noted in Mills's (1927) analysis of the US price system.

² An alternative explanation is provided by Balke and Wynne (1996, 2000) who introduce skewed technology shocks into a flexible-price multi-sector real business cycle model. In this model the skewness of sectoral price changes is correlated with the inflation rate, providing there are strong contemporaneous correlations of the sectoral technology shocks.

³ In both the Ball-Mankiw and Balke-Wynne models, the direction of contemporaneous causation runs from relative-price shocks to inflation. In a series of Granger-causality tests (using US and German data), Fischer (1982) attempted to shed some light on causality - but his results were inconclusive.

⁴ The profit loss from not re-setting is given by the square of distance between the actual and equilibrium price. BM argue this is “equivalent to taking a second-order approximation to a general profit function” (p169).

⁵ The skew-normal class of densities was given its name by Azzalini (1985) - see also Azzalini and Capitanio (1999). This distribution reduces to normal when skewness is zero and to half-normal when skewness approaches an

upper bound of $z = 0.5 \cdot (4 - \pi) \left(\frac{2}{\pi - 2} \right)^{1.5} \approx 0.995$ and a lower bound of $-z$.

⁶ These features apply symmetrically to negatively skewed distributions, but with the signs reversed.

⁷ In their review of modern hyper-inflations, Fischer *et al* (2002) find that the evidence “supports the notion that nominal rigidities are weakened as inflation reaches higher levels” (p877).

⁸ Even if the true mean of θ were positive, re-setting at θ would be an optimal strategy for a firm *expecting* its own θ to have a zero mean.

⁹ We follow BM in assuming $a = 7$.

¹⁰ The first three rows are identical to the first three rows of BM's Table 1, page 171.

¹¹ This would be true even in a world of zero trend inflation. Firms would be more likely to re-set if, by so doing, they reduce future as well as current losses.

¹² The grid search sweeps over the parameters in steps of 0.001 for t_u and t_l , and 0.0001 for r which requires finer precision, and we ensure that the solutions are not on any boundary. The initial p^* is set to zero and following

BM's suggestion (see page 172) the initial p is assumed to be distributed triangularly between $-\sqrt{C}$ and $+\sqrt{C}$. To further avoid undue influence of 'invalid' initial conditions, period t in equation (2) is year 51 in the grid-search exercise, at which stage starting prices can be viewed as being drawn numerically from the steady-state distribution of $p - p^*$.

¹³ We use similar data to BM, revised and updated to 2003. The data are available at this URL: <http://www.ecn.bris.ac.uk/www/ecdd/bm/>.

¹⁴ BM's own sample covered a range of industries from 213 in 1949 to 343 in 1989.

¹⁵ The weights are those used by the Bureau of Labor Statistics in computing the overall PPI (broadly measuring the "relative importance" of the industry in 1997).

¹⁶ Ball and Mankiw (1999) cite a Newsweek article by Milton Friedman (1975) on the effects of oil-price rises in the early 1970s. Friedman asks, "Why should the average level of all prices be affected significantly by changes in the prices of some things relative to others?"

¹⁷ The N desired relative prices were scaled by subtracting the mean of the initial simulated series from all simulated values to ensure that they did in fact sum to zero.

¹⁸ We later consider alternative assumptions about the variance of the common or monetary shock.

¹⁹ The standard deviation of the Student's t -distribution is $\sigma_t = \sqrt{\frac{df}{df-2}}$, where

df is the degree of freedom. Each draw of the distribution is scaled by $\frac{\sigma_v}{\sigma_t}$

where σ_v is selected to replicate the standard deviations of actual price changes in the data.

²⁰ We applied the two-state Markov switching model described in Hamilton (1990).

²¹ Very similar features apply to the 1983-2003 period on its own. In particular the trend was just over 1.5% and the correlations of inflation with skewness, standard deviation and Q are (respectively) 0.828, 0.051 and 0.839.

²² In deriving the numerical distribution of $p - p^*$ we assume that the initial price of a firm with menu cost node C_i is distributed triangularly between $-\sqrt{C_i}$ and $\sqrt{C_i}$. We then generate 6,000 'years' of shocks to the industry's desired price and apply the firm's price-setting rules to generate 6,000 values for $p - p^*$. We drop the first 1,000 values to purge any influence of the presumed initial distribution.

²³ For the low inflation period we carried out the same procedure but for only 51 years, and again by dropping the first 10 years observations: the remaining 41-year period then matches the low trend inflation sub-samples (1953-72 and 1983-2003).

²⁴ The mean standard deviation of simulated inflation is less than the assumed standard deviation of the common shocks. This is because some firms do not reset their price when hit by a common shock to their desired price. The mean skewness predicted by the model is greater than actual skewness but, as noted above, skewness is a very volatile variable and so this result is not surprising.

²⁵ With higher menu costs price contracts lengthen: the median contract length is now 2 years compared with 1 in the benchmark case.

²⁶ Our menu-cost model assumes that all firms within an industry share the same desired price for their product but set different prices because they have different menu costs. Firms also have different initial prices, drawn from numerical steady-state distributions appropriate for the firm's menu cost. We also explored a model suggested by recent contributions by Midrigan (2005, 2006) and Golosov and Lucas (2007) in which firms incur identical menu costs but face idiosyncratic price shocks in addition to 'monetary' and industry shocks shared by other firms in the same industry. Specifically we assumed that the firm's desired price is subject to three influences: an aggregate

'money' shock which affects all firms in all industries, an industry shock which affects firms in the same industry and an idiosyncratic shock which is peculiar to each firm. Firms within the same industry share the same menu cost, but menu costs vary across industries, the distribution of which is the same as that assumed earlier across firms. This captures the notion that price-setting technologies are more likely to vary across industry rather than across firms within the same industry. After appropriate calibration - details available from the authors - the model produced predicted correlations almost identical to those presented in Table 4.

²⁷ Industry desired price shocks are again drawn from a Student's t distribution with a standard deviation of 0.08 and with 2.5 degrees of freedom, as in the benchmark case.

Table 1 Actual Price Moments
[Symmetric Distributions of θ]

μ_θ		σ_θ				
		0.05	0.10	0.15	0.20	0.25
0	μ_p	0.0000	0.0000	0.0000	0.0000	0.0000
	σ_p	0.0217	0.0657	0.1171	0.1708	0.2248
	k_p	0.0000	0.0000	0.0000	0.0000	0.0000
0.05	μ_p	0.0229	0.0325	0.0386	0.0422	0.0445
	σ_p	0.0286	0.0699	0.1197	0.1723	0.2258
	k_p	1.2529	0.7344	0.3824	0.2114	0.1260
0.10	μ_p	0.0562	0.0688	0.0787	0.0853	0.0897
	σ_p	0.0403	0.0801	0.1263	0.1765	0.2285
	k_p	0.7602	0.8559	0.5863	0.3668	0.2317
0.15	μ_p	0.1007	0.1106	0.1215	0.1295	0.1353
	σ_p	0.0489	0.0915	0.1352	0.1826	0.2325
	k_p	0.3819	0.6608	0.5997	0.4370	0.2980
0.20	μ_p	0.1521	0.1582	0.1673	0.1757	0.1820
	σ_p	0.0537	0.1008	0.1442	0.1897	0.2376
	k_p	0.1756	0.4229	0.5025	0.4301	0.3241
0.25	μ_p	0.2071	0.2098	0.2163	0.2235	0.2298
	σ_p	0.0558	0.1068	0.1519	0.1964	0.2429
	k_p	0.0611	0.2282	0.3647	0.3740	0.3139

Notes: This table sets out the moments of *actual* price changes given different assumptions about the moments of *desired* price changes in BM's static model of the firm. μ_θ and σ_θ are the mean and standard deviation of desired price changes (the distribution is assumed symmetric). μ_p , σ_p and k_p are the mean, standard deviation and skewness of actual price changes.

Table 2 Benchmark Model
 $\mu = 0.03$; $\sigma_{\varepsilon} = 0.03$; $\sigma_{\nu} = 0.08$; $a = 7$

Menu Cost Percentile	(a) Threshold Parameters		
	t_u	t_l	r
10	0.015	-0.015	0.0002
20	0.033	-0.030	0.0016
30	0.053	-0.045	0.0052
40	0.071	-0.061	0.0111
50	0.099	-0.076	0.0163
60	0.124	-0.095	0.0277
70	0.150	-0.113	0.0295
80	0.179	-0.125	0.0445
90	0.214	-0.162	0.0643
99.9	0.507	-0.335	0.1587

(b) Simulated inflation and its moments

	Actual ^(c)	Mean Simulated
Mean Inflation ^(a)	0.030	0.030
Standard Deviation of Inflation ^(b)	0.043	0.024
Mean Standard Deviation ^(a)	0.061	0.067
SD of Standard Deviation ^(b)	0.027	0.029
Mean Skewness ^(a)	0.235	0.817
SD of Skewness ^(b)	2.431	2.498

Correlations with inflation

	SK	SD	Q
Actual (1948-2003) ^(c)	0.608	0.411	0.805
Simulations:			
Mean	0.088	0.200	0.103
2.5pc ^(d)	-0.195	-0.094	-0.191
97.5pc ^(d)	0.346	0.462	0.370

Notes: Desired price changes have a common component with a mean and standard deviation of 0.03 and a standard deviation across commodities of 0.08. *SK* is skewness and *SD* is standard deviation. *Q* is defined in equation (3). (a) Means over time of the three cross-commodity moments. (b) Standard deviations over time of the three cross-commodity moments. (c) The actual figures are based on US Producer Price data, 1948-2003, presented in the appendix below. (d) Percentiles drawn from 500 replications of model simulations, representing the 95% confidence interval.

Table 3 Alternative Models

	Correlations with inflation		
	<i>SK</i>	<i>SD</i>	<i>Q</i>
Actual (1948-2003) ^(a)	0.608	0.411	0.804
(a) Zero Trend: $\mu = 0.0$; $\sigma_v = 0.08$; $\sigma_\varepsilon = 0.03$; $a = 7$			
	<i>SK</i>	<i>SD</i>	<i>Q</i>
Mean	0.253	-0.003	0.190
2.5pc ^(b)	-0.006	-0.265	-0.101
97.5pc ^(b)	0.452	0.267	0.435
(b) High Trend: $\mu = 0.06$; $\sigma_v = 0.08$; $\sigma_\varepsilon = 0.03$; $a = 7$			
	<i>SK</i>	<i>SD</i>	<i>Q</i>
Mean	-0.070	0.231	-0.008
2.5pc ^(b)	-0.347	-0.060	-0.271
97.5pc ^(b)	0.200	0.499	0.257
(c) High Menu Costs: $\mu = 0.03$; $\sigma_v = 0.08$; $\sigma_\varepsilon = 0.03$; $a = 4$			
	<i>SK</i>	<i>SD</i>	<i>Q</i>
Mean	0.086	0.259	0.127
2.5pc ^(b)	-0.176	-0.056	-0.155
97.5pc ^(b)	0.329	0.544	0.406
(d) High Variance: $\mu = 0.03$; $\sigma_v = 0.16$; $\sigma_\varepsilon = 0.03$; $a = 7$			
	<i>SK</i>	<i>SD</i>	<i>Q</i>
Mean	0.136	0.068	0.121
2.5pc ^(b)	-0.123	-0.220	-0.171
97.5pc ^(b)	0.385	0.347	0.387

Notes: This table presents actual and simulated correlations for alternatives to the benchmark model. *SK* is skewness and *SD* is standard deviation. *Q* is defined in equation (3). (a) The actual figures are based on US Producer Price data, 1948-2003, presented in the appendix below. (b) Percentiles drawn from 500 replications of model simulations, representing the 95% confidence interval.

Table 4 Low Inflation Regime Calibration
 $\mu = 0.015$; $\sigma_{\varepsilon} = 0.015$; $\sigma_{\nu} = 0.07$; $a = 7$

Menu Cost Percentile	t_U	t_I	r
10	0.015	-0.015	0.0004
20	0.032	-0.030	0.0022
30	0.050	-0.043	0.0042
40	0.065	-0.055	0.0078
50	0.085	-0.069	0.0118
60	0.106	-0.085	0.0191
70	0.132	-0.102	0.0282
80	0.151	-0.122	0.0414
90	0.205	-0.155	0.0504
99.9	0.483	-0.294	0.0715

	Actual ^(c)	Mean Simulated
Mean Inflation ^(a)	0.015	0.015
Standard Deviation of Inflation ^(b)	0.020	0.011
Mean Standard Deviation ^(a)	0.054	0.056
SD of Standard Deviation ^(b)	0.025	0.026
Mean Skewness ^(a)	-0.007	0.807
SD of Skewness ^(b)	2.118	2.681

	Correlations with inflation		
	SK	SD	Q
Actual ^(c) [1953-1972, 1983-2003]	0.718	0.025	0.749
Simulations:			
Mean	0.200	0.134	0.194
2.5pc ^(d)	-0.115	-0.244	-0.170
97.5pc ^(d)	0.493	0.447	0.503

Notes: Desired price changes are assumed to have a common component with a mean and standard deviation of 0.015 and a standard deviation across commodities of 0.07. *SK* is skewness and *SD* is standard deviation. *Q* is defined in equation (3). (a) Means over time of the three cross-commodity moments. (b) Standard deviations over time of the three cross-commodity moments. (c) The actual figures are based on US Producer Price data, 1953-72 and 1983-2003, presented in the appendix below. (d) Percentiles drawn from 500 replications of model simulations, representing the 95% confidence interval.

Table 5 Common Shock Assumptions
 $\sigma_v = 0.08$; $a = 7$

	Correlations with inflation					
	Mean	SK		Mean	SD	
		2.5pc	97.5pc		2.5pc	97.5pc
(a) Static Model						
$\mu = 0.0$						
σ_ε						
0	0.771	0.660	0.865	-0.028	-0.653	0.634
0.01	0.394	0.162	0.593	0.004	-0.430	0.383
0.02	0.310	0.064	0.529	-0.005	-0.309	0.302
0.03	0.283	0.022	0.511	-0.004	-0.262	0.267
(b) Dynamic Model						
$\mu = 0.03$						
σ_ε						
0	0.370	0.094	0.597	0.339	-0.140	0.108
0.01	0.104	-0.158	0.390	0.176	-0.141	0.479
0.02	0.069	-0.188	0.331	0.198	-0.076	0.466
0.03	0.072	-0.176	0.344	0.222	-0.098	0.503
(b) Dynamic Model						
$\mu = 0.03$						
σ_ε						
0	0.357	0.104	0.581	0.302	-0.094	0.649
0.01	0.095	-0.198	0.363	0.152	-0.141	0.429
0.02	0.074	-0.238	0.359	0.170	-0.128	0.450
0.03	0.088	-0.195	0.346	0.200	-0.094	0.462


Notes: This table reports correlations of the moments of price changes with alternative assumptions about the mean (μ) and standard deviation (σ_ε) of the shocks common to all industries. The dynamic model assumes that firms are forward-looking in their behaviour and the static model assumes they are not. *SK* is skewness and *SD* is standard deviation.

Figure 1: Desired Price Densities

(a) Zero Mean, Symmetric



(b) Mean = 0.03, Symmetric



(c) Zero Mean, Skewed



(d) Mean = 0.03, Skewed




(e) Mean = 0.10, Skewed



(f) Mean = 0.20, Skewed



Figure 2: Static and Dynamic Threshold Parameters



Data Appendix

Year	Inflation	Weighted		
		<i>SD</i>	<i>SK</i>	<i>Q</i>
1948	0.0788	0.0734	1.3948	0.0037
1949	-0.0519	0.0871	-2.0974	-0.0037
1950	0.0373	0.0564	4.8784	0.0023
1951	0.1076	0.0596	2.4920	0.0033
1952	-0.0267	0.0586	-8.2840	-0.0020
1953	-0.0136	0.0644	-3.1541	-0.0025
1954	0.0034	0.0476	-2.7084	-0.0010
1955	0.0000	0.0529	-0.6548	-0.0001
1956	0.0336	0.0441	-0.4041	0.0002
1957	0.0293	0.0475	-1.4134	-0.0001
1958	0.0127	0.0497	1.5365	0.0007
1959	0.0032	0.0370	-1.2571	-0.0002
1960	0.0000	0.0361	-2.1393	-0.0003
1961	-0.0032	0.0318	-1.2753	-0.0002
1962	0.0032	0.0257	-1.4089	-0.0001
1963	-0.0032	0.0256	1.6891	0.0000
1964	0.0000	0.0282	1.5641	0.0002
1965	0.0219	0.0343	2.8668	0.0007
1966	0.0305	0.0299	2.4672	0.0006
1967	0.0030	0.0338	-3.7033	-0.0005
1968	0.0237	0.0312	1.7710	0.0003
1969	0.0401	0.0343	1.7285	0.0006
1970	0.0359	0.0401	2.2400	0.0007
1971	0.0320	0.0428	-0.0086	0.0003
1972	0.0437	0.0381	4.9556	0.0011
1973	0.1228	0.1026	3.6863	0.0092
1974	0.1730	0.1419	2.1772	0.0141
1975	0.0876	0.1060	-0.9776	0.0014
1976	0.0452	0.0714	1.6814	0.0019
1977	0.0603	0.0641	2.2571	0.0021
1978	0.0742	0.0530	1.6951	0.0017
1979	0.1186	0.0741	2.2952	0.0043
1980	0.1319	0.0975	2.3394	0.0066
1981	0.0874	0.0704	1.1850	0.0022
1982	0.0202	0.0715	-1.2809	-0.0011
1983	0.0129	0.0474	0.0881	0.0001
1984	0.0234	0.0384	0.7755	0.0003
1985	-0.0048	0.0526	-3.2655	-0.0016
1986	-0.0295	0.1147	-3.3712	-0.0098
1987	0.0256	0.0551	0.0144	0.0001
1988	0.0391	0.0689	1.1821	0.0015
1989	0.0484	0.0404	1.4904	0.0006
1990	0.0359	0.0623	1.6150	0.0014
1991	0.0017	0.0505	-1.2485	-0.0008
1992	0.0060	0.0443	-2.6791	-0.0005
1993	0.0144	0.0447	1.2301	0.0003
1994	0.0125	0.0519	1.9610	0.0002
1995	0.0351	0.0689	1.3817	0.0011
1996	0.0238	0.0883	-1.0220	0.0012
1997	-0.0008	0.0453	-2.1578	-0.0004
1998	-0.0254	0.0975	-2.3259	-0.0055
1999	0.0088	0.0740	1.7170	0.0021
2000	0.0558	0.1347	2.4449	0.0122
2001	0.0112	0.0709	-1.0232	-0.0002
2002	-0.0234	0.0770	-2.3818	-0.0034
2003	0.0520	0.1120	2.5989	0.0081

REFERENCES

- Azzalini, Adelchi. (1985) "A class of distributions which includes the normal ones." *Scandinavian Journal of Statistics*, 12, 171-178.
- Azzalini, Adelchi, and Antonella Capitanio. (1999) "Statistical applications of the multivariate skew-normal distribution." *Journal of the Royal Statistical Society, Series B*, 61, 579-602.
- Balke, Nathan, and Mark A. Wynne. (1996) "Supply shocks and the distribution of price changes." *Economic Review of the Federal Reserve Bank of Dallas*, First Quarter, 10-18.
- Balke, Nathan, and Mark A. Wynne. (2000) "An equilibrium analysis of relative price changes and aggregate inflation." *Journal of Monetary Economics*, 45, 269-292.
- Ball, Laurence, and N. Gregory Mankiw. (1994) "Asymmetric price adjustment and economic fluctuations." *Economic Journal*, 104, 247-261.
- Ball, Laurence, and N. Gregory Mankiw. (1995) "Relative-price changes as aggregate supply shocks." *Quarterly Journal of Economics*, 110, 161-93.
- Ball, Laurence, and N. Gregory Mankiw. (1999) "Interpreting the correlation between inflation and the skewness of relative prices: a comment on Bryan and Cecchetti." *Review of Economics and Statistics*, 81, 197-198.
- Bils, Mark, and Peter J. Klenow, (2004) "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy*, 112, 947-985.
- Blinder, Alan S. (1994) "On sticky prices: academic theories meet the real world." In *Monetary Policy*, edited by N. Gregory Mankiw, Chicago: University of Chicago Press.

- Bryan, Michael F., and Stephen G. Cecchetti. (1999) "Inflation and the distribution of price changes." *Review of Economics and Statistics*, 81, 188-196.
- Domberger, Simon. (1987) "Relative price variability and inflation: a disaggregated analysis." *Journal of Political Economy*, 95, 547-566.
- Fischer, Stanley. (1982) "Relative price variability in the United States and Germany." *European Economic Review*, 18, 171-196.
- Fischer, Stanley, Ratna Sahay and Carlos A. Végh. (2002) "Modern hyper- and high inflations." *Journal of Economic Literature*, 40, 837-880.
- Friedman, Milton. (1975) "Perspectives on Inflation.", *Newsweek*, June 24, 73.
- Golosov, Mikhail and Robert E. Lucas Jr (2007). "Menu costs and Phillips Curves." *Journal of Political Economy*, 122, 171-199.
- Hamilton, James D. (1990) "Analysis of Time Series Subject to Changes in Regime." *Journal of Econometrics*, 45, 39-70.
- Midrigan, Virgiliu. (2005) "Is firm pricing state or time-dependent? Evidence from US manufacturing.", Department of Economics, Ohio State University mimeo.
- Midrigan, Virgiliu. (2006) "Menu costs, multi-product firms and aggregate fluctuations.", Department of Economics, Ohio State University mimeo.
- Mills, Frederick. (1927) *The Behavior of Prices*, New York: NBER.
- Parks, Richard W. (1978) "Inflation and relative price variability." *Journal of Political Economy*, 86, 79-96.
- Vining, Daniel R. Jr. and Thomas C. Elwertowski. (1976) "The relationship between relative prices and the general price level." *American Economic Review* 66, 699-708.