NBER WORKING PAPER SERIES

INFLATION AND THE DISTRIBUTION OF PRICE CHANGES

Michael F. Bryan
Stephen G. Cecchetti

Working Paper 5793

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 1996

We gratefully acknowledge comments from Ben Craig, Margaret Mary McConnell, George Pennacchi, and seminar participants at the Reserve Bank of Australia and research assistance from Rod Wiggins. We thank N. Gregory Mankiw for supplying his data. The views stated herein are those of the authors and not necessarily those of the Federal Reserve Bank of Cleveland, the Board of Governors of the Federal Reserve System, or the National Bureau of Economic Research. This paper is part of NBER's research program in Monetary Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper 5793
October 1996

INFLATION AND THE DISTRIBUTION OF PRICE CHANGES

ABSTRACT

This paper reconsiders the empirical evidence connecting inflation to its higher order moments, and in particular, its third moment–the skewness of the price change distribution. Evidence on correlations between inflation and its moments goes back over thirty years, and was first used to reject the independence of relative price changes and inflation that is assumed in neoclassical models. More recently, New Keynesian macroeconomists have shown that the strong positive correlation between inflation and the skewness of the price change distribution is consistent with menu-cost models of price setting behavior. This is a fairly controversial result, prompting other researchers to demonstrate that the same correlation can be found in a multi-sector, flexible-price (real business cycle) model.

We examine the small-sample properties of the main empirical finding on which this work is based: the positive correlation between the sample mean and sample skewness of price change distributions. Our results show that this particular statistic suffers from a large positive small-sample bias, and demonstrate that the entirety of the observed correlation can be explained by this bias. To the extent that we find any relationship at all, it is that the correlation is negative. In other words, we establish that one of the most accepted stylized facts in the literature on aggregate price behavior, that inflation and the skewness of the price change distribution are positively linked, need not be a fact at all.

Michael F. Bryan
Research Department
Federal Reserve Bank of Cleveland
Cleveland, OH 44101
Michael.F.Bryan@clev.frb.org

Stephen G. Cecchetti
Department of Economics
Ohio State University
Columbus, OH 43210
and NBER
cecchetti.1@osu.edu

1 Introduction

Views of the sources, costs and consequences of inflation are often tied to evidence on the relationship among the cross-sectional moments of inflation. For example, one often hears the argument that higher inflation is more variable inflation (cross-sectionally), and this is costly. Evidence on the statistical relationship between mean inflation and its higher moments (variance and skewness) has been around for some time. Vining and Elwertowski (1976) is the standard reference for the United States. They find a large positive correlation between mean inflation and its higher-order moments.

More recently, Ball and Mankiw (1995) have revived interest in the relationship between the mean and skewness of inflation. It is their view that the positive correlation is evidence for a particular sticky-price model of inflation. Balke and Wynne (1996a, 1996b) propose an alternative theory for the correlation in which the same positive correlation between mean and skewness can arise in a multi-sectoral flexible-price model.

The purpose of this paper is to examine the small-sample properties of the sample statistics on which this work is based. We are interested in whether the observed sample (time-series) correlation between the mean and skewness of the cross-sectional distribution of price changes is significantly different from zero. The intuition for the existence of the small-sample bias is straightforward. To see it, consider a sample drawn from a zero-mean symmetric distribution. Take a case in which the initial observations have a sample mean of zero, and consider making one additional draw from the extreme positive tail of the distribution. The sample mean and the sample skewness are now both positive. And the further in the tail the added observation, the more positive both the sample mean and the sample skewness — there is a small-sample bias in correlation between the estimates of the first and third moments of a distribution. We show that the size of this bias depends on the fatness of the tails of the distribution from which the sample is drawn, that is, the higher the kurtosis of the distribution, the more positive the bias. Using Monte Carlo methods, we also demonstrate that as one increases the number of components in the cross-section used to calculate the sample moments, this bias dies out very slowly. Samples of even 1000 are still small.

Using consumer and producer price data at both monthly and annual frequencies, we estimate a large, positive small-sample bias in this statistic that accounts for virtually all of the observed sample correlation. This leads us to conclude that there is no empirical justification for either the Ball and Mankiw or the Balke and Wynne model.

Finally, insofar as there is any statistically significant correlation remaining after bias adjustment, it is negative. We interpret this to suggest the possibility that there is a nominal rigidity at zero. The rationale for this is simple. If price-setters are reluctant to lower their nominal prices, then as aggregate inflation falls, it should lead to a larger observed mass of prices changes near zero. This bunching at zero increases skewness in the price-change distribution at the same time it lowers the mean.

The remainder of the paper is divided into five sections. Section 2 describes the models in the literature. We follow with a preliminary examination of the data in Section 3. In Section 4, we examine the statistical theory of the small-sample bias in the mean-skewness correlation, and present results of Monte Carlo experiments. Section 5 reports the application of these techniques to the consumer and producer price data. The final section concludes.

2 Recent Models of Asymmetric Price Change

We begin with an overview of the theoretical models used to interpret the correlation between the mean and cross-sectional skewness of inflation. There are two such models in the literature. The first is a simple sticky-price model developed by Ball and Mankiw (1995), and the second is a flexible-price model due to Balke and Wynne (1996a, 1996b). We consider each of these in turn.

2.1 The Ball and Mankiw Model

Ball and Mankiw (1995) begin by assuming that there is a single period, at the beginning of which each firm in the economy fixes its price. Following this initial adjustment, each firm is subjected to a mean-zero (real) shock and can pay a menu cost to change its price a second time. Only some firms will experience shocks that are large enough to make the second adjustment worthwhile. As a result, the observed change in the aggregate price-level will depend on the shape of the distribution of idiosyncratic shocks.

In particular, Ball and Mankiw point out that if the distribution of real shocks is skewed, this will move the aggregate price level in the same direction. A positively skewed distribution, for example, will result in an aggregate price level increase. In general, asymmetrical distributions of relative price shocks will cause transitory fluctuations in the mean of prices, and so there will be a positive correlation between the mean and skewness moments of the price-change distribution.

In Bryan and Cecchetti (1994), we noted how we could remove these transitory fluctuations in aggregate prices by using limited-influence estimators, such as trimmed means and the median.[1]

The Ball and Mankiw model has additional implications they do not explore. For example, De Abreu Lourenco and Gruen (1995) note an implication for the relationship between the level and cross-sectional variance of inflation and the level of expected inflation. But in addition, the sticky-price model implies that the mean-skewness correlation should die out in the long run. In the Ball and Mankiw model, the correlation arises purely from short-run considerations. Once all the price-setters in the economy have had the opportunity to adjust to their relative-price shocks, the correlation should disappear. The implication is that as the length of time over which prices changes are measured increases from months to quarters to years, the correlation should weaken and eventually disappear.

2.2 The Balke and Wynne Model

Balke and Wynne (1996a, 1996b) build a multi-sector dynamic general equilibrium model with money and flexible prices. In this approach, a productivity (or supply) shock arising in one sector has two implications. First, it changes aggregate output. With the quantity of money fixed and a cash-in-advance constraint, this induces a change in the aggregate price level. Beyond this, the aggregate supply shock feeds through the input-output structure of the economy to create differential impacts on different sectors. Prices in the various sectors respond depending on the extent to which the supply shock affects their average level of productivity. Since the input-output matrix is not symmetric, the

[1]More recently, Cecchetti (1996) has extended this earlier analysis.

resulting price-change distribution is skewed.

Balke and Wynne show that their model implies a positive correlation between the mean and skewness of inflation. The reason is straightforward. The larger the initial supply shock, the larger the change in the aggregate price level. But since the coefficients of the input-output matrix are fixed, the larger the initial sectoral productivity disturbance, the more skewed its impact will be on the distribution of price changes.

While both the sticky- and flexible-price models imply a positive correlation between the mean and skewness of price changes, they differ in their predictions about the effects at short and long horizons. In the Balke and Wynne model, the effects need not die out in the long run. In fact, if the effect is predominantly a real one, it might become more pronounced over time.

It is worth making one more observation before moving on to the statistical issues involved in measuring the correlation between mean and skewness. A large and growing literature exists on the importance of inflation in facilitating relative price and wage changes.[2] Research has focused on whether firms and workers are reluctant to lower nominal wages and prices. If such a nominal rigidity existed at zero inflation, this would have implications for the relationship of the mean to the skewness of price-change distributions. Specifically, as inflation falls toward zero, one would expect the cross-sectional distribution of price changes to have increased mass near zero, and consequently increased skewness. The result is a *negative* correlation between the mean and skewness of inflation.

3 Preliminary Analysis of the Data

Our analysis of data begins with the simple calculation of sample moments. We study two data sets in detail. The first is composed of 36 components of the Consumer Price Index, available monthly from January 1967 to April 1996. The second utilizes 27 components of the Producer Price Index, monthly from January 1947 to December 1995. In both cases, the number of components chosen is the maximum that enables us to maintain a balanced panel over the entire sample period.

[2] See Akerlof, Dickens and Perry (1996) for one of the most recent contributions.

We begin by computing the inflation in each component over horizon k:

$$\pi_{it}^{k} = \frac{1}{k} \ln(p_{it}/p_{it-k}) \ , \tag{1}$$

where p_{it} is the index level for component i at time t. From this, we define the mean inflation in each time period as

$$\pi_t = \sum_i w_i p_{it} \ , \tag{2}$$

where the w_i's are weights. The higher-order central moments are then just

$$m_r(t) = \sum_i w_i (\pi_{it} - \Pi_t)^r \ . \tag{3}$$

Skewness are kurtosis are the scaled third and fourth moments:

$$\mathcal{S}_t = \frac{m_3(t)}{[m_2(t)]^{(3/2)}} \tag{4}$$

and

$$\mathcal{K}_t = \frac{m_4(t)}{[m_2(t)]^2} \ .$$

We compute two sets of moments for each data set that differ only in the weights (w_i's) used in the calculation. The first is labeled *unweighted*, and for these we assume that the weights are all equal. The second, called *weighted*, uses the actual 1985 expenditure weights to aggregate the 36 components of the CPI and the 1982 output value weights to aggregate the 27 components of the PPI.

Table 1 reports the results for weighted and unweighted moments of both indices computed over horizons of one to 24 months. The most important thing to notice is the kurtosis of these distributions. On average, monthly CPI inflation has a (weighted) kurtosis of more than nine. The kurtosis declines substantially with horizon, falling gradually to 4.7 for 24-month averages of monthly inflation. The PPI data show the same general pattern; kurtosis is high for one-month changes, falling gradually as the horizon increases.

Table 1: Sample Moments of Price-Change Distributions

	36 Components of the CPI Sample 1967:01 to 1996:04			
	$k=1$	$k=3$	$k=12$	$k=24$
	Mean			
Unweighted	5.015	5.021	5.068	5.137
Weighted	5.371	5.383	5.445	5.519
	Standard Deviation			
Unweighted	9.718	6.718	4.007	3.219
Weighted	7.953	5.631	3.455	2.807
	Skewness			
Unweighted	0.390	0.383	0.400	0.465
Weighted	0.322	0.233	0.176	0.164
	Kurtosis			
Unweighted	12.703	10.620	6.564	5.205
Weighted	9.169	8.069	5.460	4.709
	27 Components of the PPI Sample 1947:01 to 1995:12			
	$k=1$	$k=3$	$k=12$	$k=24$
	Mean			
Unweighted	3.385	3.388	3.336	3.328
Weighted	3.223	3.227	3.185	3.168
	Standard Deviation			
Unweighted	15.718	10.164	5.205	3.451
Weighted	13.989	9.212	4.635	3.192
	Skewness			
Unweighted	0.083	0.133	0.017	0.021
Weighted	0.118	0.153	0.051	0.097
	Kurtosis			
Unweighted	11.007	9.243	7.561	5.939
Weighted	8.554	7.958	7.067	6.051

Moments of price-change distributions are at annual rates.
k is the horizon, so $k = 12$ refers to 12-month changes.

4 Small-Sample Bias: Intuition

All evidence linking inflation and skewness comes from sample estimates. For example, the correlation between the weighted mean and weighted skewness in the CPI, at a monthly frequency, is 0.36. A standard test for whether this is statistically different from zero yields an asymptotic t-ratio of 6.735, which implies rejection at any meaningful level of statistical significance. But is such a conclusion warranted? What are the small-sample properties of this correlation, and how might they affect such an inference?

Some simple intuition suggests why asymptotics might be very misleading in this particular case. To see this, we begin with some definitions. First, take x_i to be a random draw from a symmetric distribution with mean zero and population moments μ_r. Then define the higher-order central sample moments as m_r, so that

$$m_r = \frac{1}{n}\sum_{i=1}^{n}(x_i - m_1)^r \tag{5}$$

for $r > 1$, and

$$m_1 = \frac{1}{n}\sum_{i=1}^{n} x_i \ . \tag{6}$$

The important insight is that draws that deviate from the population mean of the distribution, i.e., $x_i \neq 0$, can move the sample first and third moments systematically away from their population values, leading to a correlation. To see this, consider the simple case depicted in Figure 1. The x's represent draws from the symmetric distribution in the figure. Note that when one of the draws comes from the extreme positive tail of the distribution, it produces a positive sample mean and sample skewness: $m_1 > \mu_1 (= 0)$ and $m_3 > \mu_3 (= 0)$. Draws from the extreme negative tail have the opposite effect. In both cases, however, the result is a positive correlation that arises from observing only a small number of draws.

It is straightforward to compute the correlation between the first and third sample moments analytically. Consider a set of n draws from a mean-zero symmetric distribution. The relationship between the sample first moment is given by (6) and the sample third moment,

$$m_3 = \frac{1}{n}\sum_{i}(x_i - m_1)^3 \ , \tag{7}$$

FIGURE 1:
Sample Draw from a Symmetric Distribution

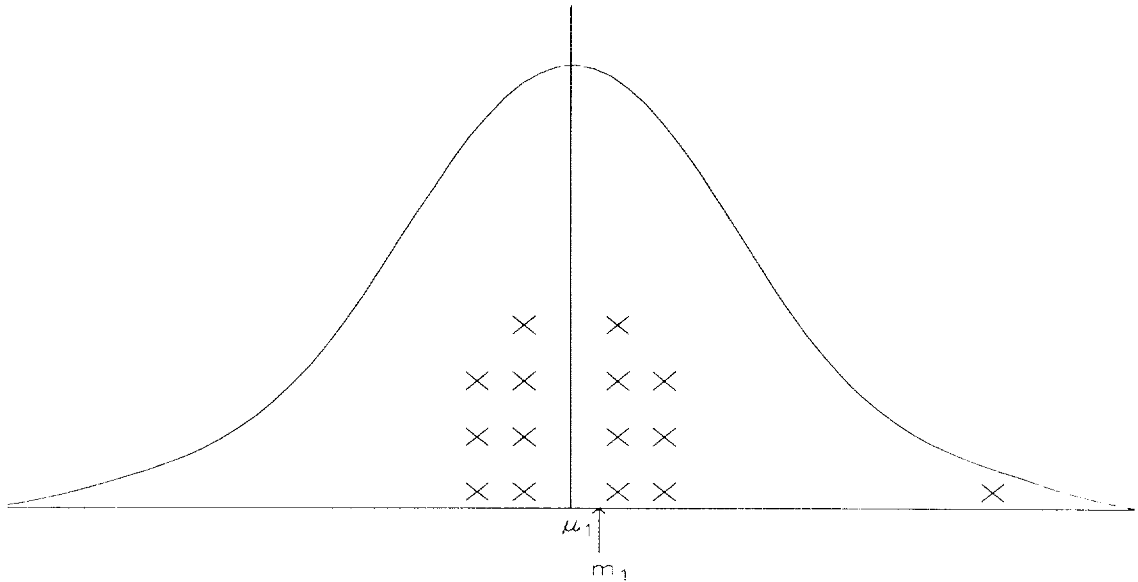

can be inferred by expanding (7). This yields

$$m_3 = \frac{\sum_i x_i^3 - 3m_1x_i^2 + 3x_im_1^2 - m_1^3}{n} \tag{8}$$

$$= \frac{\sum_i x_i^3}{n} - 3m_1\left(\frac{\sum_i x_i^2}{n}\right) + 2m_1^3 \tag{9}$$

$$= \frac{\sum_i x_i^3}{n} - 3m_1m_2 - m_1^3$$

Now consider the consequences of increasing the value of a particular draw x_1 to (x_1+1), for an initial sample with mean zero. Clearly this increases the sample mean, m_1, by $\frac{1}{n}$; that is, $\frac{dm_i}{dx_1} = \frac{1}{n}$. The impact on the sample third moment is

$$\left.\frac{dm_3}{dx_i}\right|_{m_1=0} = \frac{3}{n}(x_1^2 - m_2)\ . \tag{10}$$

This is positive if the absolute value of x_1 exceeds the sample standard deviation, $\sqrt{m_2}$. In other words, the sample correlation between the sample mean and the sample third moment is made more positive whenever there is a draw sufficiently far in the tail of the distribution.

What determines the likeliness it is that a sample will include extreme draws? The answer is that the fatter the tails of the probability distribution from which the draws are taken, the more likely that some of the x_i's will be outside one standard deviation. In other words, the kurtosis of the distribution is crucial for the size of the bias in the small-sample correlation of m_1 and m_3.

We can gain further insight into this issue by considering the case of the normal distribution. This case is of particular interest since, for a set of independent draws, the sample mean, m_1, and any function of deviations from the sample mean, $g(x_i - m_1, ..., x_n - m_1)$, are independent. Clearly, the sample third moment is a function of this form, implying that when the underlying probability distribution has a kurtosis of three, the sample correlation between m_1 and m_3 will be zero.

A simple Monte Carlo experiment helps to illustrate how large the small-sample bias is in the correlation between the sample mean and sample skewness. (The empirical work has focused on the sample skewness, not m_3.) All of our experiments are based on mean zero symmetric distributions, and so there is no population correlation between mean and skewness — both are zero. The discussion above suggests that the bias in the sample correlation should depend on the kurtosis of the density from which the sample observations are drawn; thus we should examine cases that differ depending on how fat-tailed the distributions are.

We consider the following mixture of a standard normal and a uniform distribution:

$$s * N(0,1) + (1-s)U(-A, A) \ , \tag{11}$$

where

$$Pr(s=1) = p \ .$$

That is, first we draw from a binomial to determine which distribution to draw from, then we draw from the distribution itself. The kurtosis of this distribution is

$$\mathcal{K} = \frac{3p + (1-p)\frac{A^4}{5}}{[p + (1-p)\frac{A^2}{3}]^2} \ . \tag{12}$$

FIGURE 2

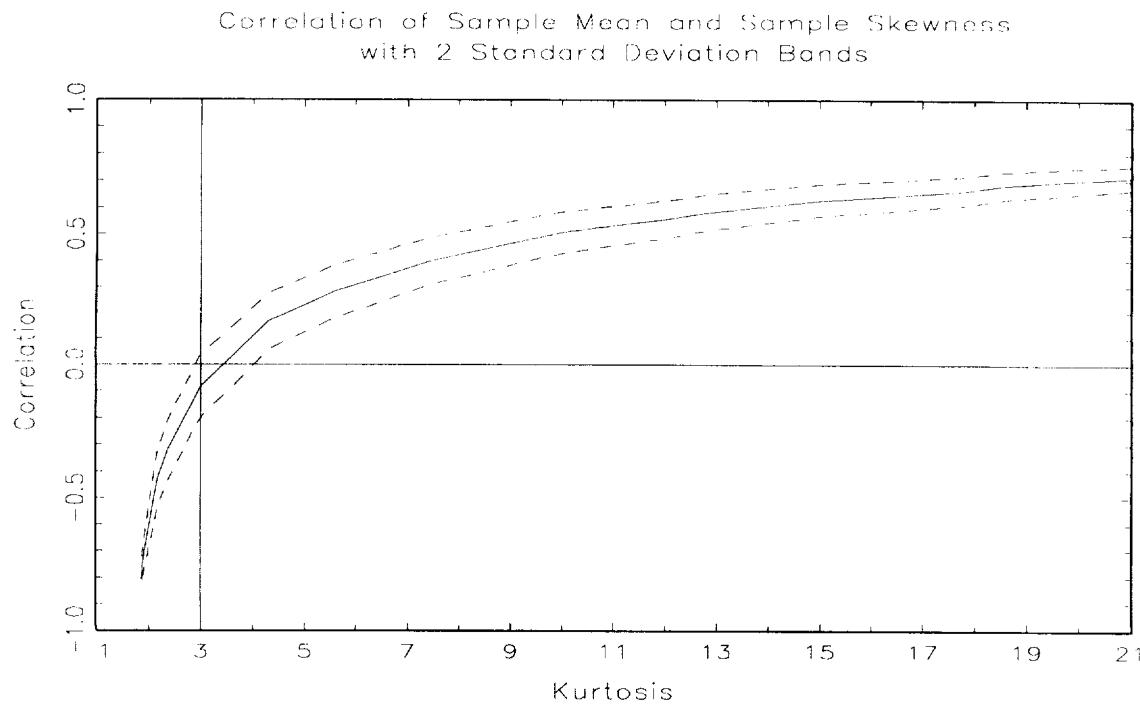

For a given p, this increases with A, with a minimum at 1.8.

The first set of experiments involves drawing 250 samples of 50 each for a given A. Throughout, p is set to 0.1. For each of the 250 samples, we calculate the sample mean and sample skewness. Then, using these 250 estimates of mean and skewness, we compute the sample correlation. This is repeated 2500 times for each value of A.

Figure 2 plots the mean and two-standard-deviation bands of the empirical distribution for the correlation between sample mean and sample skewness, as a function of the kurtosis of the distribution from which the simulated data are drawn. The results are striking in that they imply bias that is significantly positive whenever the kurtosis exceeds four! For kurtosis of 8.07, the mean of the sample value for the quarterly CPI data, the bias has a mean of 0.42 with a standard deviation of 0.05, implying a 95-percent confidence interval of approximately (0.32, 0.52).

But these results are for samples with only 50 cross-sectional observations and 250 time-series observations. If we increase either of these, what happens to the bias? To answer this question, we examine one case in more detail. Table 2 reports the results of a series of experiments, all of which are based on a distribution with a kurtosis of

Table 2: Small-Sample Bias in the Correlation of the Mean and Skewness: An Example
(Mean and standard deviation of small-sample distribution)

Time Periods	Number of Components 50	250	500	1000
25	0.4209	0.4112	0.4061	0.4054
	0.1697	0.1703	0.1726	0.1776
50	0.4222	0.4162	0.4107	0.4071
	0.1192	0.1181	0.1190	0.1242
100	0.4236	0.4170	0.4132	0.4102
	0.0836	0.0832	0.0832	0.0850
250	0.4247	0.4192	0.4153	0.4113
	0.0514	0.0518	0.0519	0.0525
500	0.4251	0.4190	0.4158	0.4130
	0.0364	0.0365	0.0371	0.0366

Mean and standard deviation of the small-sample distribution of the correlation of the mean and skewness of a distribution with kurtosis of 8.07. This matches the kurtosis in the three-month changes of the CPI and PPI data. With the exception of the last column, which has 2500 replications, all experiments have 5000 replications, where the number of components are equally weighted.

8.07. This matches the weighted kurtosis of the three-month changes in both the CPI and the PPI (see Table 1). The table reports results from increasing the cross-section to 1000 components with a time-series of up to 500 periods. The results are extremely interesting, as they show that increasing the size of the cross-section, within reasonable limits, reduces bias very slowly. The addition of time-series observations, on the other hand, reduces the spread of the empirical distribution of the bias, resulting in more precise estimates.

5 Small-Sample Bias: Price Statistics

The results of the previous example suggest that the estimates of the correlation between the sample mean and sample skewness of inflation may have large positive biases in them. We expect that these biases will become worse, as the tails of the price-change distributions become fatter. It is worth noting that, in general, heteroskedastic distributions will have fat tails. For example, a distribution constructed by mixing two normal densities will have leptokurtosis.

We follow Balke and Wynne (1996a, 1996b) and examine the properties of the simple correlation between the mean and the skewness of inflation. This avoids any need to take a position on whether certain real shocks *cause* inflation. By contrast, Ball and Mankiw (1995) focus on a regression with the mean on the left-hand side and skewness on the right-hand side.

5.1 The Experiments

To examine the actual price-change data, we use a series of Monte Carlo experiments to calculate the empirical distribution of the correlation between the sample mean and sample skewness. Each experiment has the following components: First, we specify a data-generating process from which to draw simulated inflation data. Each random draw has the same size data matrix as the actual panel of inflation data, $\{\pi_{it}^{j}\}$, where i is the component, t is the time period, and j is the number of the draw.

Using these draws, we compute the sample moments of interest. These are the time-series mean inflation rates, given by

$$\Pi_t^j = \sum_{i=1}^{36} w_i \pi_{it}^j \ , \tag{13}$$

and the higher-order moments needed to compute the skewness,

$$m_r^j(t) = \sum_{i=1}^{36} w_i (\pi_{it}^j - \Pi_t^j)^r \ . \tag{14}$$

The weights, w_i, are equal either to the weights used in the price index itself or to the inverse of the number of components. Using these, we compute the correlation, over time, of sample mean and sample skewness.

We conduct two sets of experiments that differ according to the weighting scheme used. In the first, the data are equally weighted; in the second we, use either the actual 1985 expenditure weights for the CPI data or the 1982 output-value weights for the PPI data. There are four cases for each of the two sets:

1. **I.I.D.**: All data are drawn from i.i.d. $N(0,1)$ distributions.

2. **I.N.I.D.**: All data are drawn from i.i.d. $N(0,\sigma_i^2)$ distributions, where σ_i^2 is the unconditional distribution of inflation in component i over the entire sample.

3. **I.N.I.D with serial correlation**: For each of the 36 components, sequences are built up in the following way. First, each is modeled as an AR(12) with innovations. These autoregressions are estimated from the actual data, yielding 36 parameter vectors and 36 time series of residuals. (Since inference is not an issue here, this is done using equation-by-equation OLS.) The 36 time-series of residuals are used to compute an estimated variance-covariance matrix for the system, which we call $\hat{\Sigma}$. A sequence of draws are then taken from a multivariate $N(0,\hat{\Sigma})$. These are used recursively to construct 36 data series, starting with the first 12 actual values of each series.

4. **Bootstrap**: This is the same as the previous experiment, with the exception that the draws are taken from the actual innovation series themselves, not from a multivariate normal distribution.

Finally, in addition to monthly changes, we consider overlapping 1-, 3-, 12-, and 24-month changes. When the data-generating process is assumed to be heteroskedastic, the actual p-values are constructed using robust methods. For simple i.n.i.d. cases, White's (1980) method is used. For the i.n.i.d. with serial correlation and bootstrap cases, Newey and West's (1987) procedure with bandwidth parameter set equal to $1.4k$ is used.

5.2 Results

We present results for three data sets: (1) 36 components of the CPI, monthly, from 1967 to 1996; (2) 27 components of the PPI, monthly, from 1947 to 1996; and (3) the Ball and Mankiw annual PPI data, which have a maximum of 301 components and a sample period from 1947 to 1989. For each of these, we present a table that includes the following results: the actual sample value of the correlation between the mean and the skewness of inflation; the p-value computed from the standard statistic (i.e. t-ratio) assuming that the asymptotic distribution is accurate; the bias-adjusted correlation value minus the median of the empirical distribution; and finally, an empirical p-value for the test that correlation observed in the data is equal to zero in the empirical distribution.

Table 3 includes the results obtained from the CPI data. Note that once we allow for heteroskedasticity, the bias is large. To see this, compare experiments 1 and 2 for $k = 1$. The sample correlation for the unweighted data is 0.548. If the data were truly generated by an i.i.d. normal distribution, then there would be no small-sample bias, and the null hypothesis that the correlation is zero would be rejected with a p-value of 0.00. However, once we introduce heteroskedasticity in experiment 2, the bias-adjusted correlation is estimated to be -0.017, with a p-value of 0.63!

Tables 3 and 4 each contain 24 experiments that allow for heteroskedasticity in the data-generating process for inflation (all but experiments 1 and 5). For the CPI data, 15 of the bias-adjusted correlation estimates are positive, but all of the p-values exceed 0.10. In fact, for the short horizons $k = 1$ and $k = 3$, the estimates are generally negative.

Moving to the PPI data, there is some evidence of a *negative* correlation between mean and skewness. All but two of the bias-adjusted correlation estimates are less than zero, and a number of the p-values are less than 0.10. This is especially true when we

Table 3: Bias-Adjusted Correlation Between Mean and Skewness
Consumer Prices, 36 Components, Monthly, 1967 to 1996

Experiment	Statistic	$k=1$	$k=3$	$k=12$	$k=24$
	Unweighted				
Sample values	Correlation	0.548	0.535	0.581	0.646
	p-value	0.00	0.00	0.00	0.00
1. i.i.d	Bias-Adjusted Correlation	0.547	0.533	0.581	0.641
	Empirical p-value	0.00	0.00	0.00	0.00
2. i.n.i.d.	Bias-Adjusted Correlation	−0.017	−0.026	0.014	0.071
	Empirical p-value	0.63	0.60	0.86	0.58
3. i.n.i.d. w/ Serial Corr.	Bias-Adjusted Correlation	−0.068	−0.087	0.051	0.158
	Empirical p-value	0.08	0.14	0.64	0.33
4. Bootstrap	Bias-Adjusted Correlation	−0.087	−0.098	0.049	0.160
	Empirical p-value	0.03	0.11	0.67	0.32
	Weighted				
Sample values	Correlation	0.360	0.376	0.428	0.465
	p-value	0.00	0.00	0.00	0.00
5. i.i.d.	Bias-Adjusted Correlation	0.801	0.818	0.876	0.914
	Empirical p-value	0.00	0.00	0.00	0.00
6. i.n.i.d.	Bias-Adjusted Correlation	0.015	0.032	0.083	0.115
	Empirical p-value	0.72	0.58	0.49	0.50
7. i.n.i.d. w/ Serial Corr.	Bias-Adjusted Correlation	−0.018	0.014	0.203	0.347
	Empirical p-value	0.76	0.87	0.19	0.13
8. Bootstrap	Bias-Adjusted Correlation	−0.057	−0.007	0.181	0.318
	Empirical p-value	0.33	0.91	0.23	0.16

Table 4: Bias-Adjusted Correlation Between Mean and Skewness
Producer Prices, 27 Components, Monthly 1947 to 1995

Experiment	Statistic	$k = 1$	$k = 3$	$k = 12$	$k = 24$
	Unweighted				
Sample values	Correlation	0.522	0.399	0.329	0.294
	p-value	0.00	0.00	0.00	0.00
1. i.i.d.	Bias-Adjusted Correlation	0.522	0.398	0.329	0.297
	Empirical p-value	0.00	0.00	0.00	0.02
2. i.n.i.d.	Bias-Adjusted Correlation	0.007	−0.117	−0.191	−0.231
	Empirical p-value	0.81	0.01	0.02	0.08
3. i.n.i.d. w/ Serial Corr.	Bias-Adjusted Correlation	−0.081	−0.090	−0.001	0.023
	Empirical p-value	0.00	0.06	0.99	0.87
4. Bootstrap	Bias-Adjusted Correlation	−0.077	−0.112	−0.0337	−0.0119
	Empirical p-value	0.01	0.02	0.74	0.94
	Weighted				
Sample values	Correlation	0.428	0.372	0.283	0.299
	p-value	0.00	0.00	0.00	0.00
5. i.i.d.	Bias-Adjusted Correlation	0.523	0.467	0.378	0.391
	Empirical p-value	0.00	0.00	0.00	0.00
6. i.n.i.d.	Bias-Adjusted Correlation	0.046	−0.011	−0.104	−0.096
	Empirical p-value	0.17	0.82	0.26	0.46
7. i.n.i.d. w/ Serial Corr.	Bias-Adjusted Correlation	−0.008	−0.045	−0.030	−0.033
	Empirical p-value	0.80	0.39	0.78	0.81
8. Bootstrap	Bias-Adjusted Correlation	−0.0518	−0.0671	−0.0548	−0.0505
	Empirical p-value	0.13	0.17	0.59	0.73

Table 5: Bias-Adjusted Correlation Between Mean and Skewness
Producer Prices, 301 Components, Annual 1947 to 1995

Unweighted		
Sample values	Correlation p-value	0.4371 0.00
1. i.i.d.	Bias-Adjusted Correlation Empirical p-value	0.4357 0.00
2. i.n.i.d.	Bias-Adjusted Correlation Empirical p-value	-0.0934 0.45
3. i.n.i.d. w/ Serial Corr.	Bias-Adjusted Correlation Empirical p-value	-0.1369 0.00
Weighted		
Sample values	Correlation p-value	0.4308 0.00
4. i.i.d.	Bias-Adjusted Correlation Empirical p-value	0.5091 0.00
5. i.n.i.d.	Bias-Adjusted Correlation Empirical p-value	0.0076 0.96
6. i.n.i.d. w/ Serial Corr.	Bias-Adjusted Correlation Empirical p-value	-0.2808 0.00

Results are from experiments with 2500 replications. All data are from Ball and Mankiw (1995), restricted to include only those series with non-zero weight in 1982, and with at least six years of data. Experiments 3 and 6 presume that each inflation series is a first-order autoregression.

---

use the unweighted data.

As is suggested by the experiments at the end of Section 4, increasing the number of components from 27 to 301 is of virtually no help in eliminating the small-sample bias in the correlation between mean and skewness. Once again, only negative estimates of the bias-adjusted correlation are significantly different from zero.[3]

Collectively, these results allow the following conclusions. First, the empirical implications of the Ball and Mankiw sticky-price model are not borne out in the data: No

[3]We use a restricted set of the components used by Ball and Mankiw. First, throughout the analysis, we use only those components with nonzero weights in 1982. Second, we employ only series with six or more years of data. Finally, we do not report results for bootstrap samples, as the gaps in some of the data series make their construction conceptually difficult.

evidence of a positive correlation between mean and skewness is more pronounced at short horizons. The failure to find a correlation that is both positive and significantly different from zero is also evidence against the mechanisms suggested by Balke and Wynne.

Our finding of a significant negative correlation in the monthly changes of the unweighted data, as well as in the annual changes in Ball and Mankiw's PPI data, is intriguing. It suggests the possibility that price-setters may face a rigidity at zero.

The rationale for this might be something like the following: If price-setters are reluctant to lower their nominal prices, then as aggregate inflation falls, it should lead to larger observed mass near zero in the distribution of price changes. This *bunching* at zero increases skewness in the price-change distribution at the same time it lowers the mean.

We caution, however, that our results are not robust across the data sets we examine. We observe no statistically significant negative correlations in the CPI data. And our strongest results are for the unweighted PPI data. Insofar as there is any pattern in the bias-corrected correlation estimates, it is that they become more positive as the horizon lengthens. This is exactly the opposite of what is implied by the menu-cost approach, which predicts that they should diminish over time.

6 Conclusions

Based on the evidence presented here, we believe that the recent focus on the correlation between the mean and skewness of the cross-sectional distribution of price changes is unwarranted. The positive small-sample bias explains the entire correlation observed in the data. The evidence is overwhelmingly against both the Ball and Mankiw model of inflation and the Balke and Wynne model of real supply shocks.

To the extent that we are able to document any statistically meaningful relationship between mean and skewness, it is a weak negative one. We view these results as suggestive and worthy of further study.

References

Akerlof, George A., William T. Dickens and George L. Perry, 'The Macroeconomics of Low Inflation,' *Brookings Papers on Economic Activity* (1996:1) 1-59.

Ball, Laurence and N. Gregory Mankiw, 'Relative-Price Changes as Aggregate Supply Shocks,' *Quarterly Journal of Economics* 110 (February 1995) 161-193.

Balke, Nathan S. and Mark A. Wynne, 'An Equilibrium Analysis of Relative Price Changes and Aggregate Inflation,' Research Department Working Paper 96-09, Federal Reserve Bank of Dallas, 1996 a.

______, and ______, 'Supply Shocks and the Distribution of Price Changes,' *Economic Review* of the Federal Reserve Bank of Dallas (First Quarter 1996b) 10-18.

Bryan, Michael F. and Stephen G. Cecchetti, 'Measuring Core Inflation,' in *Monetary Policy*, N. Gregory Mankiw, ed., Chicago: University of Chicago Press for NBER, 1994, 195–215.

Cecchetti, Stephen G., 'Measuring Short-Term Inflation for Central Bankers,' mimeo., Department of Economics, The Ohio State University, September 1996.

De Abreu Lourenco, Richard and David Gruen, 'Price Stickiness and Inflation,' Research Discussion Paper #9502 Reserve Bank of Australia, March 1995.

Newey, Whitney K. and Kenneth D. West, 'A Simple, Positive Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,' *Econometrica* 55 (May 1987): 703-708.

Vining, Daniel R., Jr. and Thomas C. Elwertowski, 'The Relationship between Relative Prices and the General Price Level,' *American Economic Review*, 66 (September 1976) 699-708.

White, Halbert, 'A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity,' *Econometrica* 48 (May 1980) 817-38.