

w o r k i n g
p a p e r

9 7 0 7

Efficient Inflation Estimation

by Michael F. Bryan,
Stephen G. Cecchetti, and
Rodney L. Wiggins II



FEDERAL RESERVE BANK OF CLEVELAND

Working Paper 9707

EFFICIENT INFLATION ESTIMATION

by Michael F. Bryan, Stephen G. Cecchetti, and Rodney L. Wiggins II

Michael F. Bryan is assistant vice president and economist at the Federal Reserve Bank of Cleveland. Stephen G. Cecchetti is executive vice president and director of research at the Federal Reserve Bank of New York and a research associate at the National Bureau of Economic Research. Rodney L. Wiggins II is a member of the Research Department of the Federal Reserve Bank of Cleveland. The authors gratefully acknowledge the comments and assistance of Todd Clark, Ben Craig, and Scott Roger.

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment. The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or of the Board of Governors of the Federal Reserve System.

Federal Reserve Bank of Cleveland working papers are distributed for the purpose of promoting discussion of research in progress. These papers may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications.

Working papers are now available electronically through the Cleveland Fed's home page on the World Wide Web:
<http://www.clev.frb.org>.

August 1997

Abstract

This paper investigates the use of trimmed means as high-frequency estimators of inflation. The known characteristics of price change distributions, specifically the observation that they generally exhibit high levels of kurtosis, imply that simple averages of price data are unlikely to produce efficient estimates of inflation. Trimmed means produce superior estimates of 'core inflation,' which we define as a long-run centered moving average of CPI and PPI inflation. We find that trimming 9% from each tail of the CPI price-change distribution, or 45% from the tails of the PPI price-change distribution, yields an efficient estimator of core inflation for these two series, although lesser trims also produce substantial efficiency gains. Historically, the optimal trimmed estimators are found to be nearly 23% more efficient (in terms of root-mean-square error) than the standard mean CPI, and 45% more efficient than the mean PPI. Moreover, the efficient estimators are robust to sample period and to the definition of the presumed underlying long-run trend in inflation.

1 Introduction

How should we interpret month-to-month changes in the measured Consumer Price Index? Over the years, this question has led to the construction of several measures of what has come to be called ‘core’ inflation. Common measures of core inflation regularly remove certain components from the construction of the CPI. In the U.S., ‘volatile’ food and energy price movements, are often ignored, and core inflation is synonymous with the CPI that excludes food and energy.¹ But is it truly the case that food and energy price changes never contain information about trend inflation? Or, for that matter, is it only the volatile food and energy components that distort attempts to measure the underlying inflation trend? Surely not. This leads us to consider how we might develop a systematic, statistical methodology for reducing the transitory noise in measured inflation indices.

This paper follows our recent work, largely beginning with Bryan and Cecchetti (1994), where we investigate the estimation of aggregate consumer price inflation using trimmed means of the distribution of price changes. These are estimators that are robust to the distributional anomalies common to price statistics. They are order statistics that are computed by trimming a percentage from the tails of a histogram, and averaging what is left. For example, the sample mean trims zero percent, while the median trims fifty percent, from each tail of the distribution of price changes.

Every student in introductory statistics learns that, when data are drawn from a normal distribution, the sample mean is the minimum variance estimator of the first moment. But price changes are not normally distributed. In fact, as we discuss in Bryan and Cecchetti (1996), the cross-sectional distribution of inflation has very fat tails, with a sample kurtosis that is often substantially above ten. Underlying leptokurtotic distributions create inferential difficulties, as they routinely produce skewed samples. In our earlier papers, we discuss how these problems lead to transitory movements in the sample mean, causing it to have a high small-sample variance.

Given what we know about the distribution of price changes, what is the most efficient

¹The 1997 *Economic Report of the President* is a prime example. Chart 2-6 on page 76, and accompanying text, use the now commonplace designation of core inflation as the ‘Consumer Price Index excluding the volatile food and energy components.’

estimator of the first moment of the price change distribution? How can we produce a reduced-noise estimate of aggregate inflation at high frequencies? Our answer is to trim the price change distribution, not by removing food and energy prices every time, but by ignoring some percentage of the highest and lowest price changes each month.

We study monthly changes in both consumer and producer prices in the U.S. Data availability dictate that we examine 36 components of the CPI from 1967 to 1996 and 29 components of the PPI over the same period. Throughout, we take as our benchmark the thirty-six month centered moving average of actual inflation. We evaluate the ability of candidate estimators to track the movements in the benchmark. Our conclusions are that the most efficient estimate of inflation at the consumer level comes from trimming 9% from each tail, while efficient estimation of producer prices trims 45%. By trimming a cumulative 18% of the consumer price distribution we are able to reduce the root-mean-square-error (RMSE) of aggregate inflation by nearly one-quarter. For the PPI, the improvement is even more dramatic, as the RMSE declines by over 45 percent!

The remainder of the paper is composed of five sections. Section 2 reports descriptive statistics for the distribution of CPI and PPI price changes. Section 3 discusses the statistical problems we attempt to overcome. Section 4 follows with by a discussion of the Monte Carlo results that guide our choice of the optimal trimmed estimator. We provide various robustness checks in Section 5. These include examining changes in sample period, changes in the degree of disaggregation of CPI data, and changes in the benchmark. Section 6 concludes.

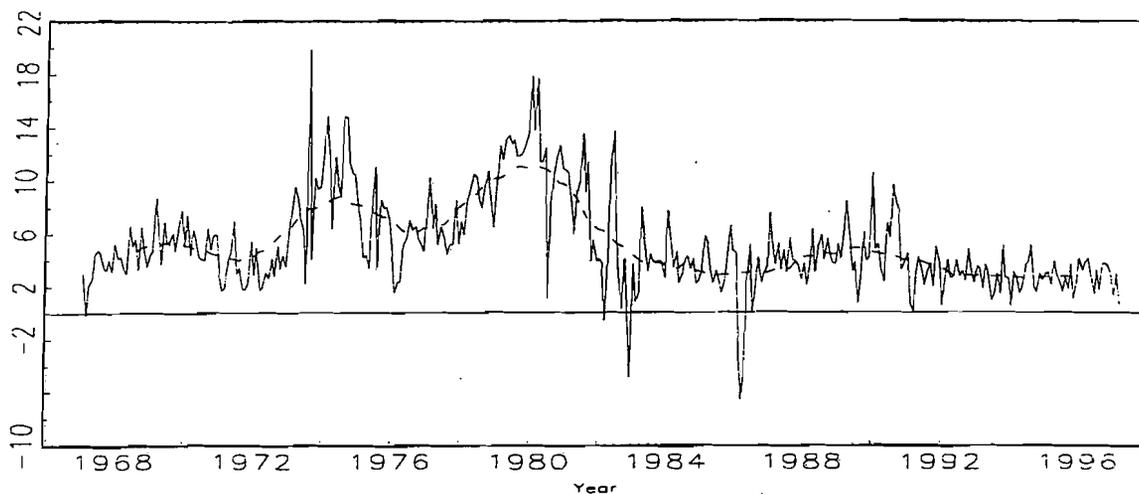
2 Characteristics of Price Change Distributions

By how much would the monthly measure of the consumer price index have to deviate from its recent trend for us to be relatively certain that the trend has changed? This is the question that is in most people's minds when the Bureau of Labor Statistics releases the price statistics each month.² Figure 1 plots the monthly changes in consumer and producer prices, at an annual rate, together with a three-year centered moving average,

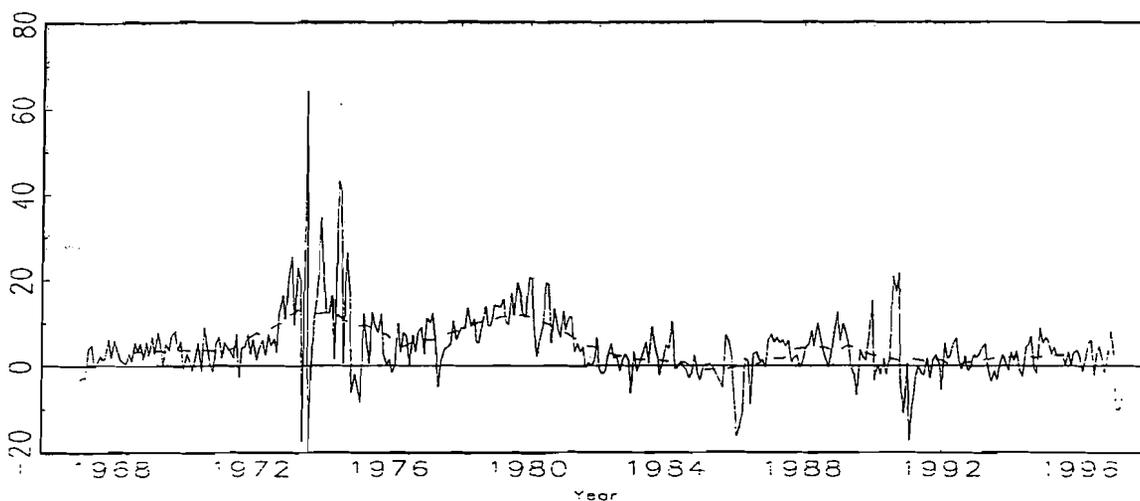
²Cecchetti (forthcoming) suggests a preliminary answer to exactly this question.

FIGURE 1

CPI Monthly with 36 Month Centered Moving Average



PPI Monthly with 36 Month Centered Moving Average



both for the period 1967:02 to 1997:04.³

As is evident from the figure, the monthly changes in both of these price indices contain substantial high-frequency noise. By this we mean that deviations of the monthly changes from the trend are quite large and often reversed. In fact, the standard deviation

³We use 36 components of the Consumer Price Index for Urban Consumers, seasonally adjusted by the BLS. These data are all available continuously, monthly, since 1967:01. The housing service component is based on the rental equivalence measure of owner occupied housing, and so prior to 1982, the series is essentially the experimental CPI-X1. The producer price is based on the PPI for commodities, and uses a set of between 29 and 31 components. All data are seasonally adjusted using the ARIMA X-11 procedures available with SAS. A detailed Appendix containing descriptions of the sources and construction of the data sets used is available from the authors upon request.

of the difference between the monthly and the moving average aggregate price change is 6.92 percentage points for the PPI and 2.50 percentage points for the CPI (both at annual rates). A look at the actual distributions shows that a 90% confidence interval for the CPI is from -3.92 to +3.76 percentage points, while for the PPI it is from -10.35 to +8.97 percentage points. In other words, since 1967, monthly changes in producer prices have been either more than 10 percentage points below or 8 percentage points above the thirty-six month moving average one in every ten months!

The common method of excluding food and energy simply does not help much. In fact, the standard deviation of the difference between the CPI ex food and energy and the thirty-six month average CPI is 2.31 percentage points, and the 90% confidence interval shrinks slightly to [-3.73,+3.76] percentage points. By contrast, for the PPI, excluding food and energy improves things, as the standard deviation of difference between the PPI excluding food and energy and the 36 month centered moving average of the actual PPI drops by about 40% to 4.14, and the 90% confidence interval shrinks by about the same amount to [-5.94,+4.76].

In an effort to better understand the nature of the transitory fluctuations in high-frequency inflation measurement, we begin by examining the characteristics of the price change distributions. It is useful to pause at this stage to introduce some notation. We define the inflation in an individual component price over an horizon k as

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

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

$$\Pi_t^k = \sum_i r_{it} \pi_{it}^k , \quad (2)$$

where the r_{it} 's are relative importances that are allowed to change each month to reflect the fact that the actual index is an arithmetic average.⁴

⁴It is straightforward to show that if the price level index utilizes fixed weights, call these w_i , then the percentage change in the aggregate index can be approximated by the weighted sum of the percentage

The higher-order central moments are then

$$m_{rt}^k = \sum_i r_{it} (\pi_{it}^k - \Pi_t^k)^r, \quad (3)$$

Skewness and kurtosis are the scaled third and fourth moments, respectively:

$$S_t^k = \frac{m_{3t}^k}{[m_{2t}^k]^{(3/2)}} \quad (4)$$

and

$$\mathcal{K}_t^k = \frac{m_{4t}^k}{[m_{2t}^k]^2}. \quad (5)$$

Table 1 reports numerous descriptive statistics for the cross-sectional distribution of monthly price changes at overlapping horizons of one to thirty-six months. Among the noteworthy characteristics is that the distributions are often skewed. The mean absolute value of the skewness, the mean of S_t^1 , in monthly CPI changes is 0.20 and in PPI changes it is 0.04, suggesting that the distributions are nearly symmetrical on average. There is little skewness in the distributions on average. But the standard deviation of S_t^1 is 2.35 for the CPI and 2.36 for the PPI, implying that distributions of one-month changes are often highly skewed. This standard deviation falls off as the horizon increases, implying that the distribution of longer-run changes are much less likely to exhibit skewness.⁵

The price change distributions also have very fat tails. The average kurtosis of monthly changes, the average value of \mathcal{K}_t^1 , is 11.24 for the CPI and 10.35 for the PPI. In fact, the weighted kurtosis of monthly price changes is in excess of 20 about ten percent of the time. See Figure 2.

These facts allow us to identify a potentially important source of high frequency noise in the measurement of inflation. In a given month, there is a high probability of observing some subset of prices change by a substantial amount — generating the skewness and kurtosis that we see. But, over time, these extreme changes are balanced out, reducing

changes in the components, where the weights change to reflect changes in relative prices. Defining the aggregate price level $P_t = \sum w_i p_{it}$, then $r_{it} = w_i (p_{it}/p_{t-1})$.

⁵For example the 5th and 95th percentiles of S_t^1 for the CPI are $[-3.52, 4.26]$. But the same percentiles for S_t^{36} are $[-2.39, 1.93]$.

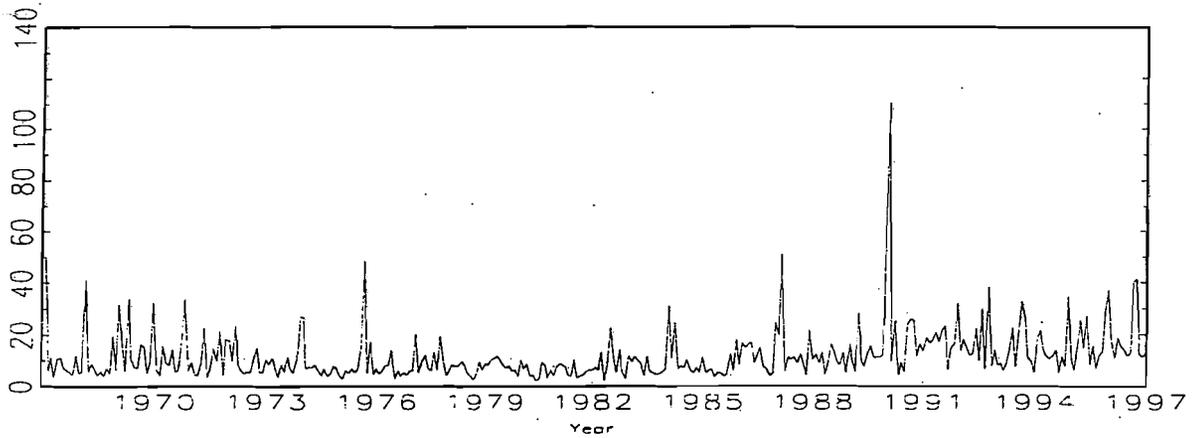
Table 1: Summary Statistics for Price Change Distributions

Deviations from 36 Month Moving Average					
Consumer Prices, 1967.01 to 1996.04					
36 Components					
	$k = 1$	$k = 3$	$k = 12$	$k = 24$	$k = 36$
Standard Deviation					
Average	9.18	6.64	4.06	3.36	3.14
Std. Dev.	190.45	79.80	25.49	11.81	8.83
Skewness					
Average	0.20	0.16	0.21	0.29	0.26
Std. Dev.	2.35	2.15	1.51	1.38	1.41
Kurtosis					
Average	11.24	9.56	5.72	4.52	4.23
Median	8.60	7.37	4.65	3.89	3.75
Std. Dev.	9.80	8.36	3.49	2.39	2.20
Producer Prices, 1967.02 to 1997.04					
29-32 Components					
	$k = 1$	$k = 3$	$k = 12$	$k = 24$	$k = 36$
Standard Deviation					
Average	15.59	10.60	6.24	4.82	4.44
Std. Dev.	955.66	266.64	86.85	35.08	22.82
Absolute Skewness					
Average	0.04	0.14	0.04	0.02	0.01
Std. Dev.	2.36	2.12	1.74	1.53	1.46
Kurtosis					
Average	10.35	8.80	7.26	5.47	4.03
Median	6.38	6.23	4.89	3.51	2.78
Std. Dev.	11.51	8.47	6.50	6.11	3.43

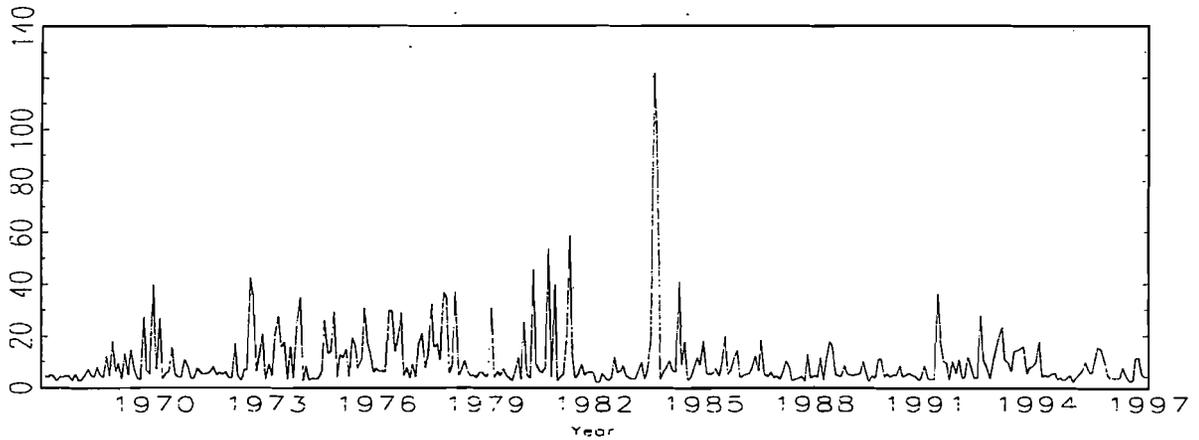
All data are at annual rates.

FIGURE 2

Weighted Kurtosis of Consumer Prices
Monthly Changes, 1967 to 1995



Weighted Kurtosis of Producer Prices
Monthly Changes, 1947 to 1995



the observed skewness.

One economic interpretation of these distributional characteristics is that if price change is costly, we will not observe the distribution of desired price changes each month. If the size and timing of price changes are based on two-sided state-dependent rules, as in Caballero and Engel (1991), or Caplin and Leahy (1991), what we observe will depend on the rule used by the price-setter and the history of the shocks to desired prices. As a result, we will rarely see prices that exactly equal the price that would be set in the absence of any price-adjustment costs. The amount of noise decreases over longer periods, when each price has changed numerous times. But for high frequencies of one quarter or one month, the problem can be a serious one.⁶

However, one need not necessarily attach oneself to a particular model of price-setting behavior in order to accept our conclusions. It is well known that a mixture of random draws from normal distributions with differing variances will produce a leptokurtic sample. As a statistical matter, then, we can show that the mean price-change statistic is unlikely to provide the efficient estimate of inflation, regardless of the price setting model that is assumed.

We can think of two possible approaches to handling the problem. One would be to actually model price-setting explicitly using the theory as it has been worked out. But this has substantial drawbacks, as it requires that we actually estimate the time-varying price change rules themselves. Alternatively, we can treat the complication presented by state-dependent price change rules as a statistical sampling problem. We view the monthly, skewed distributions as small-sample draws from the longer-horizon (roughly) symmetrical population distribution. The fact that the population has such high kurtosis leads us to consider a family of estimators that are robust to the presence of fat tails, a topic to which we now turn.

⁶An alternative interpretation is implied by Balke and Wynne (1996), who show that a multi-sector, dynamic general equilibrium model with money and flexible prices can produce similar characteristics in an environment of asymmetric supply shocks. A distinguishing feature of this model is that the 'noise' in the estimator need not significantly diminish at lower frequencies.

3 Robust Estimation

We begin by assuming that we have available a sequence of samples from a symmetric distribution with an unknown, and possibly changing, mean. At issue is the efficient estimation of the mean. We consider a set of estimators called trimmed means, that average centered portions of the sample. The method of averaging is to order the sample, trim the tails of the sample distribution, and average what remains.

To calculate the (weighted) α -trimmed mean, we begin by ordering the sample, $\{x_1, \dots, x_n\}$, and the associated weights, $\{w_1, \dots, w_n\}$. Next, we define W_i as the cumulative weight from 1 to i ; that is, $W_i \equiv \sum_{j=1}^i w_j$. From this we can determine the set of observations to be averaged for the calculation: the i 's such that $\frac{\alpha}{100} < W_i < (1 - \frac{\alpha}{100})$. We call this I_α . This allows us to compute the weighted α -trimmed mean as

$$\bar{x}_\alpha = \frac{1}{1 - 2\frac{\alpha}{100}} \sum_{i \in I_\alpha} w_i x_i. \quad (6)$$

There are two obvious special cases. The first is the sample mean, \bar{x}_0 , and the second is the sample median, \bar{x}_{50} .⁷

The efficient estimator of the mean, in the class of trimmed sample means, will depend on the characteristics of the data-generating process.⁸ If, for example, the data are drawn from a normal distribution, then we know that the sample mean is the most efficient estimator. That is, the sample mean is the estimator that has the smallest small-sample variance.

But when the data are drawn from leptokurtic distributions — distributions with much fatter tails than the normal — the sample mean will no longer be the most efficient estimator of the population mean, even in the class of trimmed sample means. It is relatively easy to see why this is so. With a fat-tailed distribution, one is more likely to obtain a draw from one of the tails of the distribution that is not balanced by an

⁷See Stuart and Ord (1987) pg. 50-51 and particularly Huber (1981) for general definitions of limited-influence estimators and their properties.

⁸For example, Yule and Kendall (1968) discuss the impact of changing kurtosis on the relative efficiency of the sample mean and the sample median. But we know of no general results concerning the relative efficiency of trimmed-mean estimators.

equally extreme observation in the opposite tail. That is to say, as the kurtosis of the data-generating process increases, samples have a higher probability of being skewed.⁹

The impact of kurtosis on the efficiency of trimmed-mean estimators is straightforward to demonstrate. To do so we construct a simple experiment in which we draw a series of samples from distributions with varying kurtosis and compute the efficiency of the entire class of trimmed-mean estimators, including the mean and the median.

In all of our experiments, the data-generating process is characterized by a two parameter distribution that is a mixture of two normals, one with unit variance, and one with changing variance. We consider a random variable z , such that

$$z = s * y_1 + (1 - s) * y_2 , \quad (7)$$

where

$$\begin{aligned} Pr(s = 1) &= p , \\ y_1 &\sim N(0, 1) , \text{ and} \\ y_2 &\sim N(0, A) . \end{aligned}$$

With probability p draws come from a standard normal and with probability $(1 - p)$ they come from a $N(0, A)$. The population mean, $E(z)$, is zero. The kurtosis of this distribution, $\frac{E(z^4)}{E(z^2)^2}$, varies with p and A :

$$\mathcal{K}(A, p) = \frac{3p + 3(1 - p)A^2}{[p + (1 - p)A]^2} . \quad (8)$$

We examine five cases, all with $p = 0.90$, and A set such that $\mathcal{K} = (3, 10, 15, 20, 30)$. In each of our experiments, we construct 10,000 replications of 250 draws each. We then compute the \bar{x}_α for $\alpha = \{0, 1, \dots, 49, 50\}$. This yields 10,000 estimates of all of the trimmed-mean estimators, which we label \bar{x}_α^j . From these we compute the root-mean-

⁹Bryan and Cecchetti (1996) demonstrate this point in another context. We can show that the standard deviation of the sample skewness increases with the kurtosis of the data-generating process.

square error (RMSE) and the mean absolute deviation (MAD). These are

$$\text{RMSE}_\alpha = \sqrt{\frac{1}{N} \sum_j^N (\bar{x}_\alpha^j)^2} \quad (9)$$

and (10)

$$\text{MAD}_\alpha = \frac{1}{N} \sum_j^N |(\bar{x}_\alpha^j)|.$$

Figure 3 plots the RMSE_α and the MAD_α for experiments based on distributions with varying kurtosis, $\mathcal{K}(A, p)$. To adjust for the fact that the variance of the distribution also changes with A and p , we have normalized RMSE_0 and MAD_0 to one for each case. The results clearly show that the efficient trim — the trim that minimizes either the RMSE or the MAD — increases with the kurtosis of the data generating process. As the kurtosis increases from 3 to 30, the efficient trim goes from 0 to 16%.

We caution that the results from these experiments are illustrative and apply only to the specific distributions we examine. We know of no general analytic result deriving the optimal trimmed mean estimator as a function of the moments of the underlying distribution and the size of the sample.

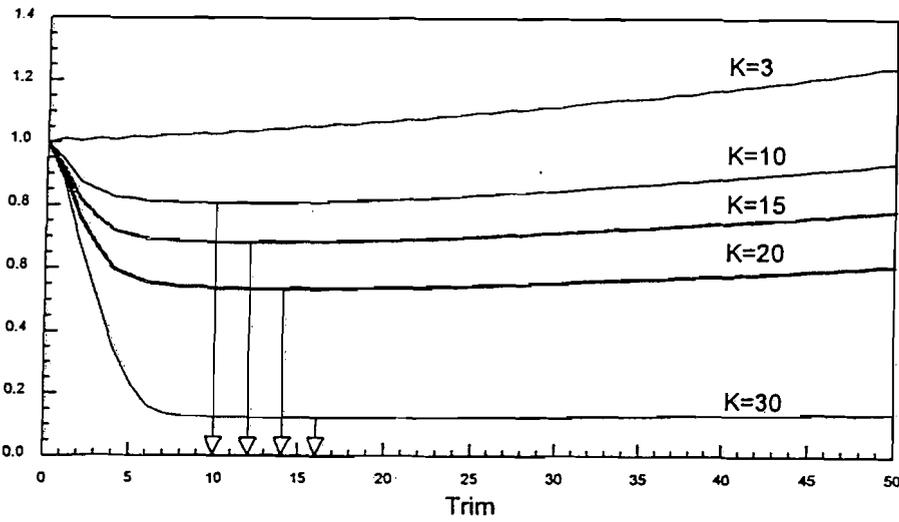
4 Efficient Estimation of Inflation: Preliminaries

We have now established one property of price data and a related statistical fact. First, the cross-sectional distribution of price changes, both in the CPI and the PPI, is fat-tailed. Second, trimmed-means are the efficient estimator of the mean of a leptokurtic distribution. We now combine these two insights and ask what is the most efficient estimator of inflation?

We begin with a preliminary examination of the data using a simple Monte Carlo experiment based on actual price data. In order to judge efficiency, we need to have a measure of the population mean we are trying to estimate. Following Cecchetti (forthcoming), we choose the thirty-six month centered moving average of actual inflation. This is an approximation of the long-term trend in inflation that is likely to be what

FIGURE 3

RMSE of Trimmed Estimators as Kurtosis Changes



MAD of Trimmed Estimators as Kurtosis Changes

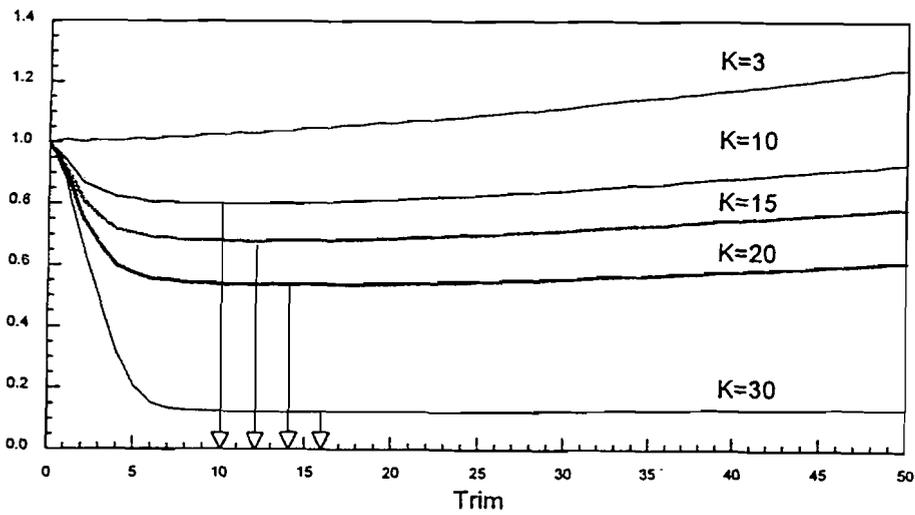
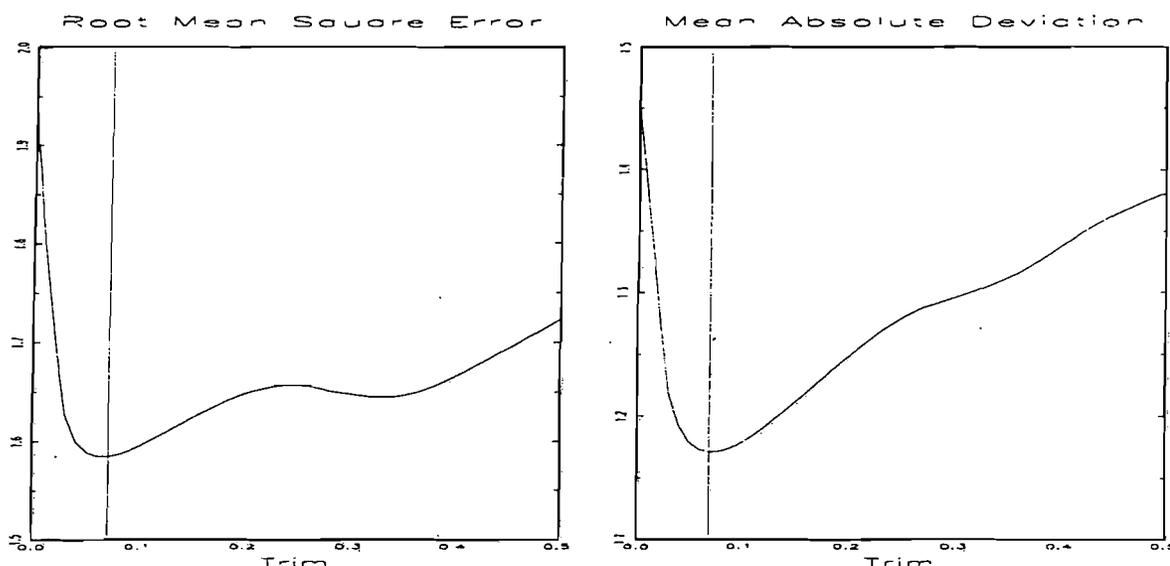


FIGURE 4: Consumer Prices
Efficiency of Trimmed Estimators, Monte Carlo Results



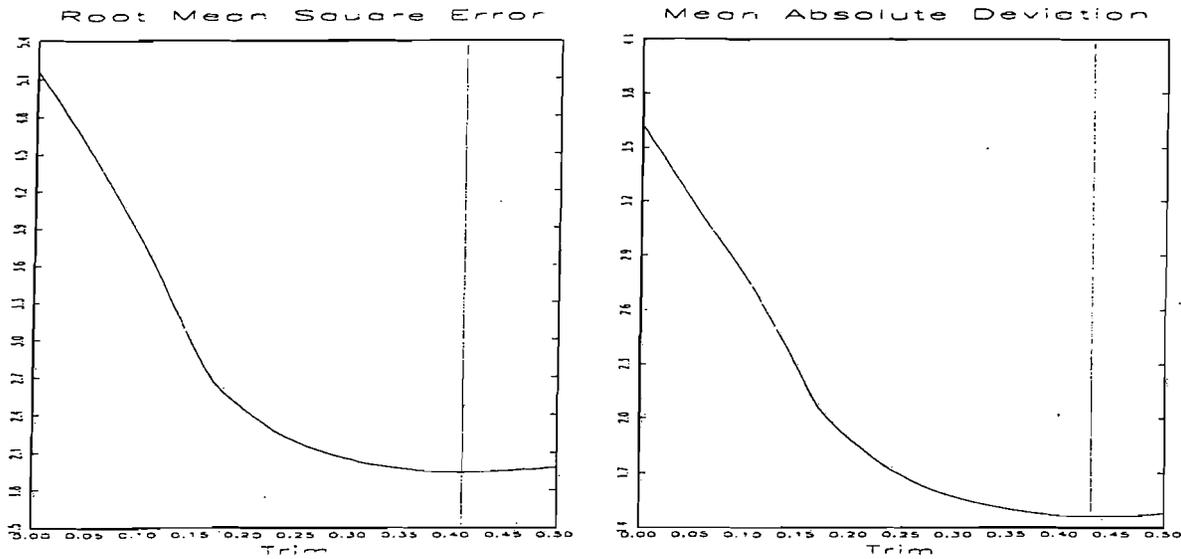
people have in mind when they attempt to construct measures they label *core* inflation.

To proceed, we take the deviation of monthly component price changes from this thirty-six month centered moving average of inflation. For the CPI, we use 36 components of the CPI-U over the period 1967.02 to 1997.04, with its 1985 weights. To simplify the experiments, we set the relative importances (r_{it}) equal to the 1985 weights (w_i), and leave them fixed throughout. For the PPI, we use a reduced set of 27 components also available over the 1967.02 to 1997.04 sample and their fixed 1982 weights. After subtracting each price change from the thirty-six month moving average change in the appropriate index, we have two matrices of relative price changes.

In each experiment, we randomly draw a series of samples by taking one observation for each of the component time-series — one draw from each column in the relative-price-change matrix. This is a bootstrap procedure from which we generate 10,000 samples, each with 36 relative price changes for CPI data, or 29 relative price changes for PPI data. We then compute the two measures of efficiency — the root-mean-squared error (RMSE) and the mean absolute deviation (MAD).

The results are reported in Figures 4 and 5. The weighted means are found to be the least efficient of all of the estimators. The efficiency of the inflation estimates greatly

FIGURE 5: Producer Prices
Efficiency of Trimmed Estimators, Monte Carlo Results



improves with even very small trims from the sample. For example, in the case of the CPI, trimming as little as 3% from each tail of cross-sectional distribution improves the efficiency of the estimator by over 15%. The most efficient estimator for monthly CPI data was the 7% trimmed mean where the efficiency gain is approximately 20%, although trims in the neighborhood of this estimator perform nearly as well.¹⁰

For the PPI, however, much larger trims of the sample distribution are necessary to achieve the efficient estimator. The optimal trim, which occurs in the range of 40%, has an RMSE that is only one-third that of the sample mean!

¹⁰The technique we suggest here is appropriate for cases in which the price-change distributions are symmetrical on average. We know of instances where this is not the case. For example, Roger's (1997) examination of New Zealand price data reveals a persistent, positive skewness in the price change distribution that produces a bias in the trimmed estimators of the mean. Roger constructs trimmed estimators centered on the mean percentile, or the percentile of the distribution corresponding to the mean of the distribution. That is, for New Zealand price data, Roger trims the tails of the distribution asymmetrically, centering on the 57th percentile. In this way, the trimmed estimator is an unbiased estimate of the CPI trend in New Zealand. Roger's insight implies a procedure in which the trim and centering parameter are chosen jointly to minimize either the RMSE or MAD criterion, subject to the estimator being unbiased in the sample.

5 Efficient Inflation Estimation: Historical Data

We now move to a more complete examination of the actual data. Here we will compare the relative efficiency of trimmed estimators using the historical time series, taking account of the changes in the relative importances [the r_{it} 's in equation (2)] over time. That is to say, we will compute the weighted distributions of inflation each month, where the weights vary based on changes in relative prices as well as the periodic rebasing done by the Bureau of Labor Statistics roughly once per decade.

In Section 5.1, we look for the optimal trimmed mean estimator using the entire 1967 to 1997 sample currently available. Are the results of the previous section robust to several obvious changes in methods? We examine three cases. In the first, reported in Section 5.2, we study more disaggregated CPI data over a shorter sample period. In Section 5.3, we look at the implications of changing the measurement benchmark from the thirty-six month centered moving average of actual inflation to moving averages of from twenty-four to sixty months. Finally, in Section 5.4, we study estimator stability by looking at optimal trims over varying sample periods. We conclude this section with a summary and comparison of the trimmed means with the inflation measures that arbitrarily exclude food and energy.

5.1 The Baseline Case

In this section we consider the time-series characteristics of the trimmed-mean estimators. We calculate the RMSE and the MAD for each trimmed estimator using monthly historical component price data. That is, we compute the trimmed-mean estimators of inflation month-by-month, and compare their deviations from the centered thirty-six month moving average. The results, reproduced in Figure 6 for the CPI, and Figure 7 for the PPI, are virtually identical to those in the Monte Carlo experiments shown in Figures 4 and 5.¹¹

It is easy to see how much inflation measures are stabilized by trimming. Figure 8

¹¹Throughout this section, the PPI data set uses a set of components that varies from 29 to 31 in number, depending on data availability

FIGURE 6: Consumer Prices
Efficiency of Trimmed Estimators, Historical Data

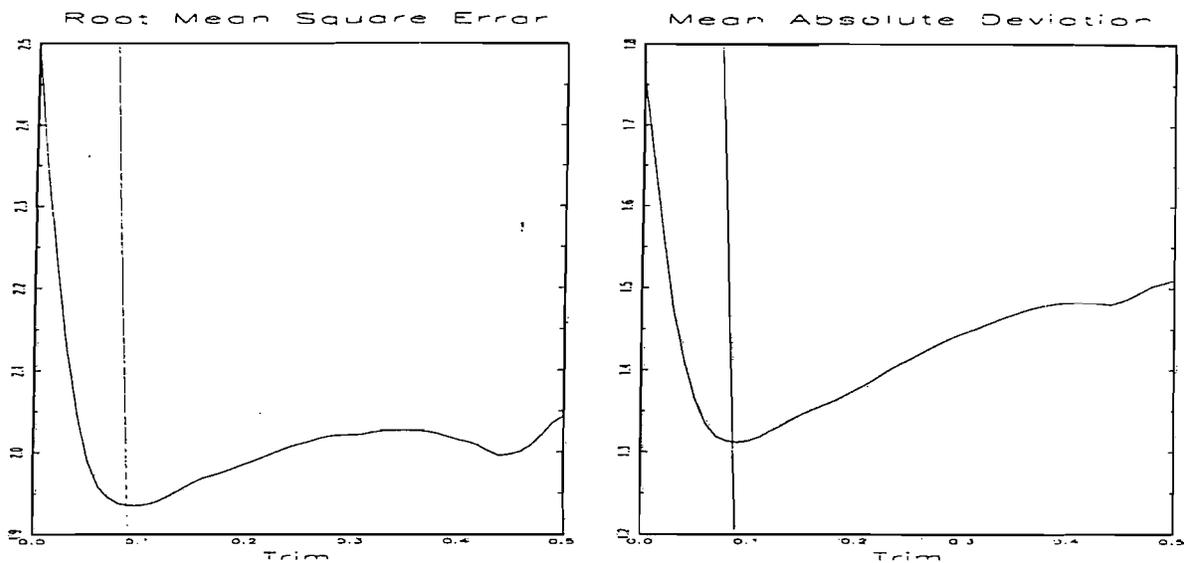


FIGURE 7: Producer Prices
Efficiency of Trimmed Estimators, Historical Data

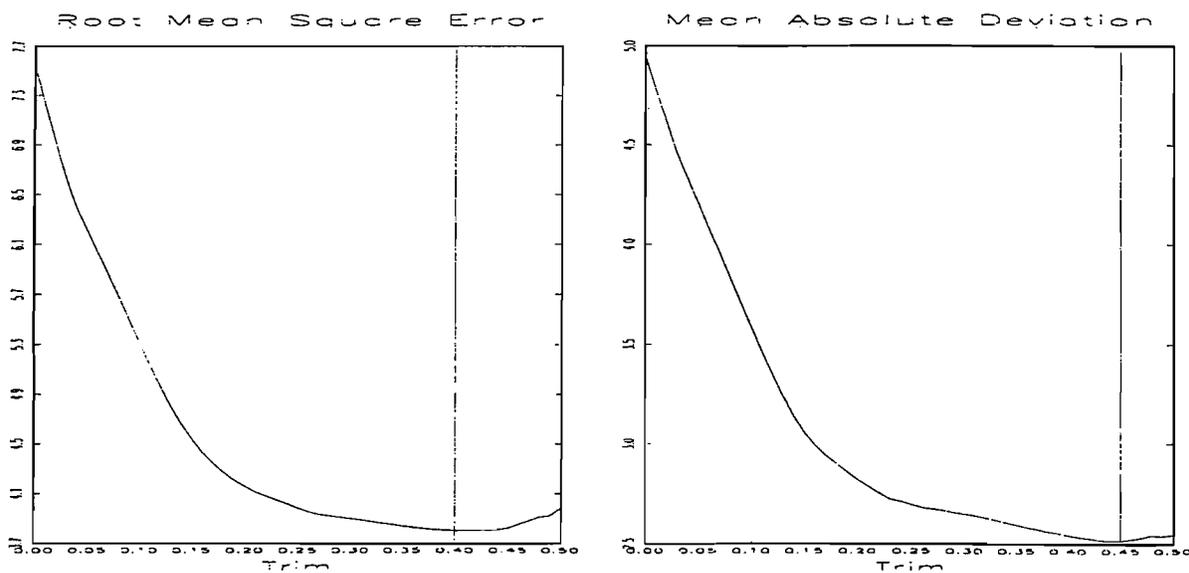
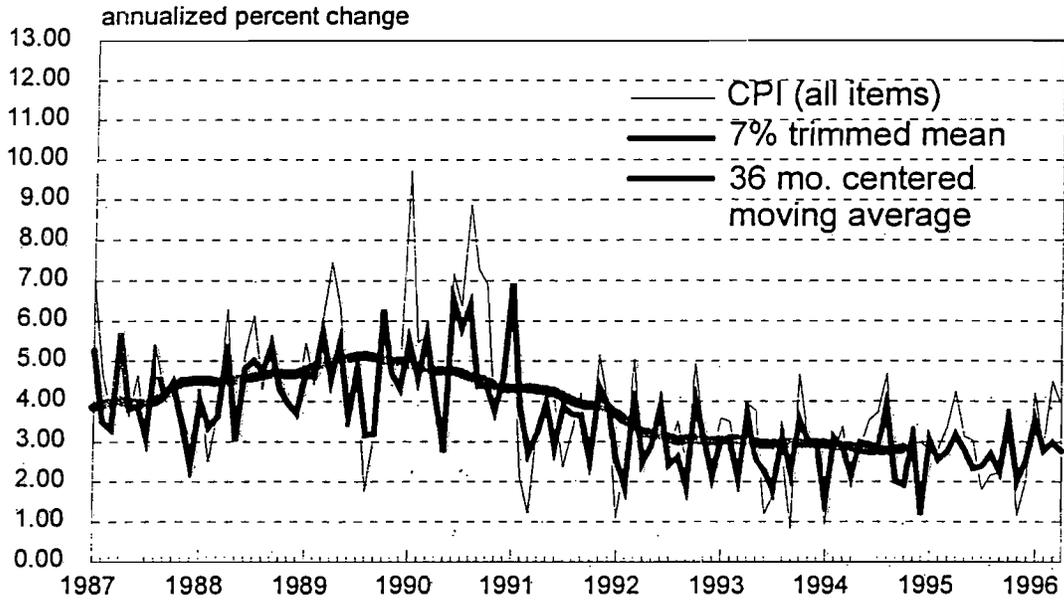


FIGURE 8
Monthly CPI Estimators



Monthly PPI Estimators

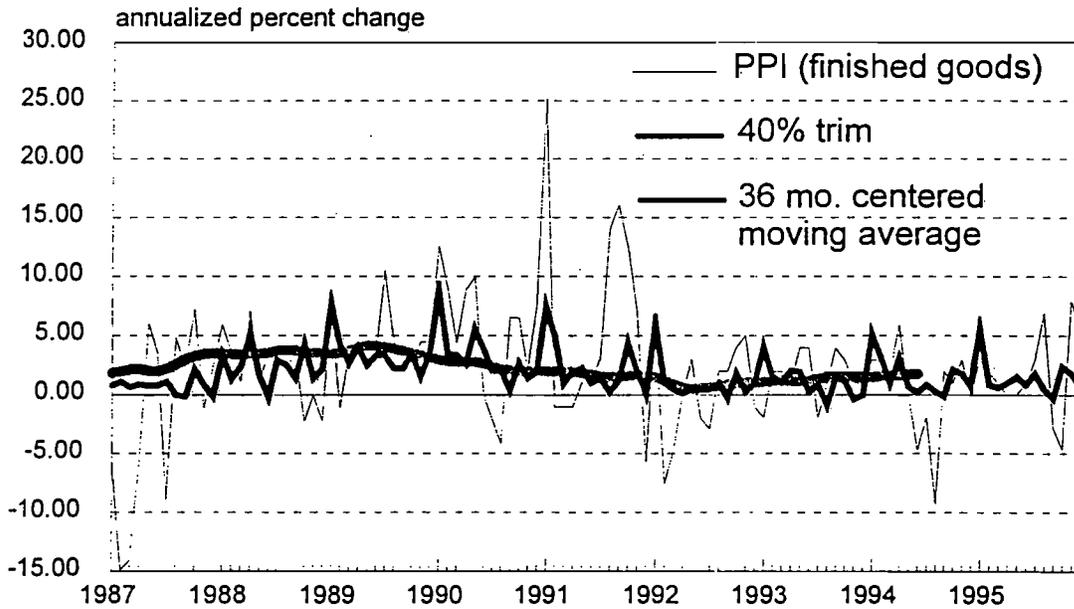


Table 2: Comparison of Inflation Estimators

	CPI		PPI	
	1967 to 1997		1967 to 1997	
	RMSE	MAD	RMSE	MAD
Mean (\bar{x}_0)	2.50	1.76	6.91	4.27
ex Food& Energy	2.31	1.62	4.14	2.55
Median (\bar{x}_{50})	2.04	1.51	3.98	2.55
Optimal Trim	1.93	1.31	3.80	2.52
Trim at Opt.	9%	9%	40%	45%

All values are computed from monthly changes as annual rates. Deviations are from the 36-month centered moving average. The optimal trim is the trim that minimizes either $RMSE_\alpha$ or MAD_α .

plots the mean, the thirty-six month centered moving average, and the efficient trimmed estimator for monthly CPI and PPI data for the January 1990 to December 1996 period.

Table 2 compares the properties of a number of commonly used estimators for consumer and producer price inflation. Focusing first on the CPI, we note that excluding food and energy produces little improvement in efficiency. The CPI excluding food and energy is only slightly more efficient than the CPI-U itself, reducing the RMSE from 2.50 to 2.31. But trimming clearly helps. Trimming 9% of the cross-sectional distribution of consumer prices reduces the RMSE by just under 23 percent.¹²

For producer prices, the improvements are even more dramatic. Using the long sample period, we find that trimming 40% of the distribution from each tail improves the RMSE by over 45 percent. Excluding food and energy from the PPI reduces the RMSE by less than 40 percent.¹³

¹²Bryden and Carlson (1994) also note that this trim produces the minimum time-series variance of any trimmed-mean estimator over the 1967 to 1994 period.

¹³A common technique for reducing the noise in the high frequency inflation estimates uses time-series averages. We have conducted experiments that combine trimming with time-averaging. We note that averaging the component price change data prior to trimming, or *pre-trim averaging*, decreases the amount of trimming necessary to produce a minimum RMSE estimator of the inflation trend. For example, using three-month average price changes of component CPI data, the minimum RMSE of the inflation trend is found by trimming 6% from the tails of the price change distribution, compared to the 9% trims required of monthly data. Similar results were found for *post-trim averages*, where we average the monthly trimmed means. That is, if we calculate the trimmed estimators, and compute a 3-month average of that result, the minimum RMSE estimate of the inflation trend is found by trimming 6% from

5.2 More Disaggregated Data

The price statistics are collected at a much more disaggregated level than what we have used thus far. Does the optimal trim change with the level of aggregation? The experiments in Section 3 suggest that the answer to this question will depend on what happens to the kurtosis of the cross-section distribution of price changes as we vary the level of aggregation.

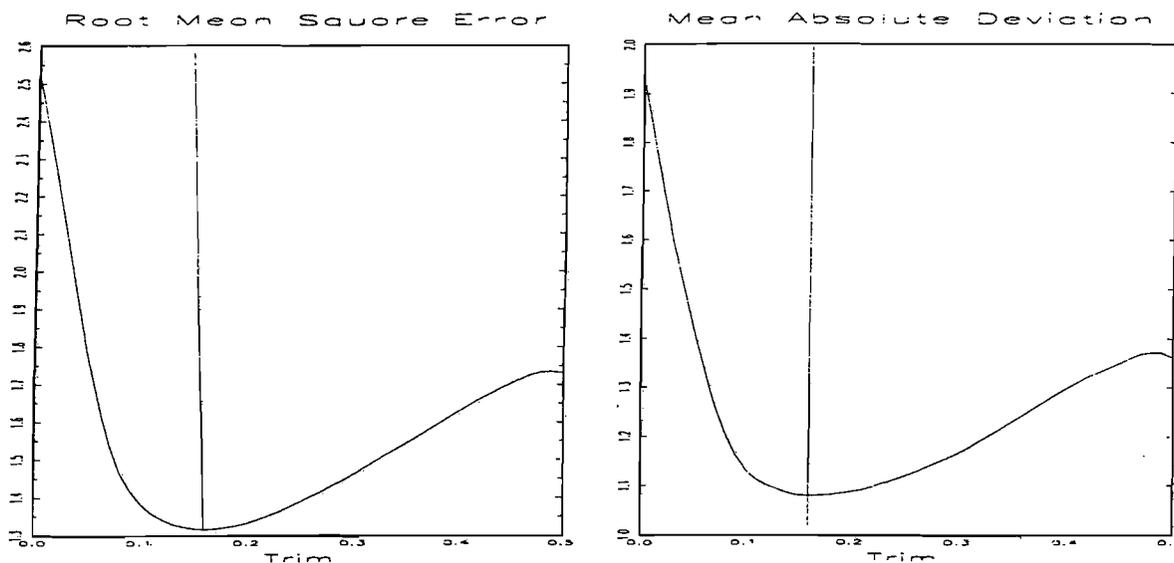
To examine this issue, we assembled a data set composed of between 142 and 175 components of the CPI-U from 1978 to 1996. The number of series (and the relative importance of each series) varies each month depending on availability. The weighted kurtosis of these data is much higher than that for the 36 component dataset examined in the previous section. For monthly changes, for example, Table 1 reports that inflation in the 36 components of the CPI-U has median kurtosis of 9.68. By contrast, the kurtosis in the more disaggregated data set has a median of 43.1!

As in Section 5.1, we construct, using historical data, the RMSE and MAD for each of the trimmed estimators, from $\alpha = 0$ to 50. These provide a gauge of the efficiency gains from trimming the outlying tails of the price-change distribution. The results in Figure 9 confirm that, in the case of consumer prices, the efficient estimation of inflation requires more trimming when more disaggregated data are used. In this experiment, the optimal trim is 16%, at which point the RMSE is cut nearly in half. But again, virtually any trimming helps. For example, trimming 9% from each tail — the optimal amount for the 36 component data set — reduces the RMSE by about 40%.

The practical implications of this exercise are fairly important. We have found that since the kurtosis of the price-change distributions depends on the level of disaggregation, so does the optimal trim. As a result, implementation of these techniques for the production of a core inflation index will depend critically on the exact dataset used.

each tail of the price change distribution. Even at relatively low frequencies, some amount of trimming of the price change distribution seems warranted. For example, using a 6-month component price change and a 6-month average of the trimmed estimators, the minimum RMSE estimator of the CPI trend is obtained by trimming 5% from each tail of the price change distribution. These alternative smoothing techniques address a somewhat different question from the one posed in this paper: How much new information does a monthly price report contain? We leave the investigation of this important area for future research.

FIGURE 9: Consumer Prices, 142 to 175 Components
Efficiency of Trimmed Estimators; Historical Data



5.3 Changes in the Benchmark

As we noted at the outset of the previous section, in order to assess efficiency, we must specify a goal: What is it we would ideally like to measure? Our second robustness check involves deviating from the thirty-six month centered moving average as the benchmark.

Table 3 reports optimal trims as a function of the length of the moving average specified for the benchmark, similar to those in Sections 4 and 5.1 for the optimal trim. Included are the optimal trims using the Monte Carlo methods, as well as those for the historical data. The table also reports an informal confidence interval constructed as the set of trims with RMSE or MAD within five percent of the minimum. For example, using the historical data in the case of the 36 components CPI data and the thirty-six month centered moving average benchmark, the minimum RMSE of 1.93 occurs at a trim of 9% (see Table 2). The fourth line in the first bottom panel of Table 3 reports that all of the trims between 5% and 48% have an RMSE below $1.93 \times 1.05 = 2.03$.¹⁴

Several patterns emerge from these results. First, the ‘point estimate’ of the optimal trim does not vary as we change the benchmark. But the approximate confidence intervals

¹⁴Note that there is no reason for the approximate confidence intervals to be either symmetrical or continuous. The ones reported in Table 3 all happen to be continuous.

Table 3: Optimal Trim for Changes in the Benchmark

Monte Carlo Results

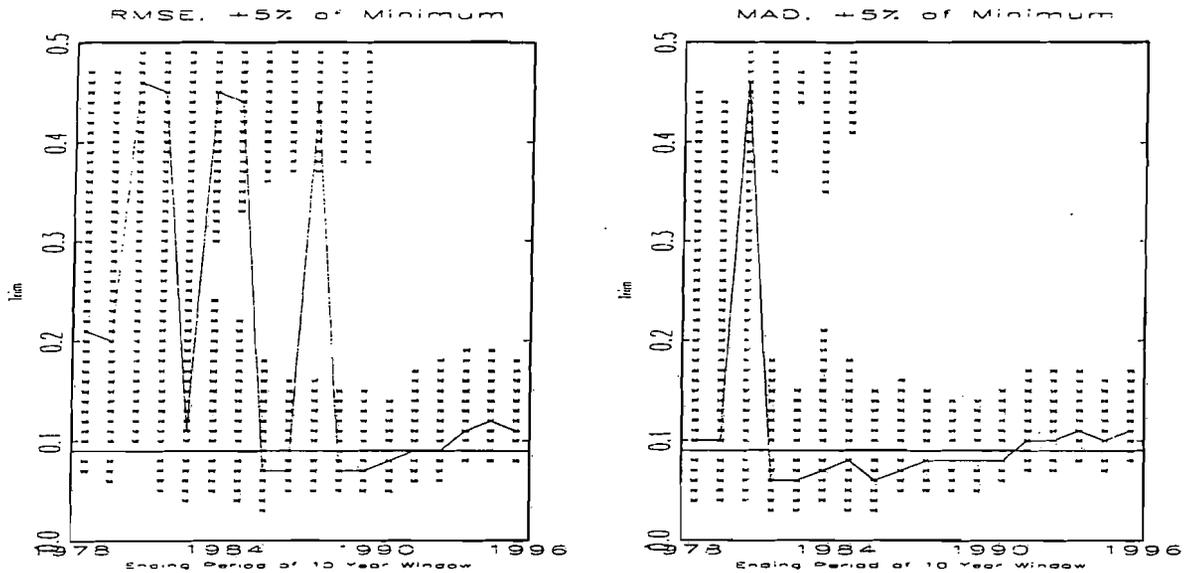
MA	CPI		PPI	
	RMSE	MAD	RMSE	MAD
24	0.07 (0.03,0.35)	0.07 (0.03,0.17)	0.43 (0.31,0.50)	0.45 (0.33,0.50)
36	0.07 (0.03,0.44)	0.07 (0.03,0.17)	0.41 (0.31,0.50)	0.43 (0.33,0.50)
48	0.06 (0.03,0.42)	0.07 (0.03,0.17)	0.43 (0.31,0.50)	0.46 (0.34,0.50)
60	0.06 (0.03,0.41)	0.07 (0.03,0.17)	0.42 (0.30,0.50)	0.45 (0.33,0.50)

Historical Data

MA	CPI 36 Components 1967 to 1997		PPI 29 to 31 Components 1967 to 1997		CPI 142 to 175 Components 1978 to 1996	
	RMSE	MAD	RMSE	MAD	RMSE	MAD
24	0.09 (0.05,0.25)	0.09 (0.05,0.17)	0.40 (0.25,0.49)	0.45 (0.30,0.50)	0.14 (0.08,0.23)	0.16 (0.09,0.24)
36	0.09 (0.05,0.48)	0.09 (0.05,0.19)	0.40 (0.25,0.50)	0.45 (0.31,0.50)	0.16 (0.10,0.24)	0.17 (0.11,0.26)
48	0.09 (0.05,0.50)	0.09 (0.05,0.21)	0.43 (0.25,0.50)	0.45 (0.29,0.50)	0.17 (0.11,0.25)	0.17 (0.12,0.25)
60	0.09 (0.05,0.50)	0.09 (0.05,0.23)	0.43 (0.25,0.50)	0.49 (0.27,0.50)	0.18 (0.12,0.26)	0.18 (0.12,0.28)

Numbers in parentheses are trims with RMSE or MAD within 5% of the value at the minimum. Monte Carlo experiments use 10,000 replications.

FIGURE 10: Consumer Prices, 36 Components
Efficiency of Trimmed Estimators, Changing Sample



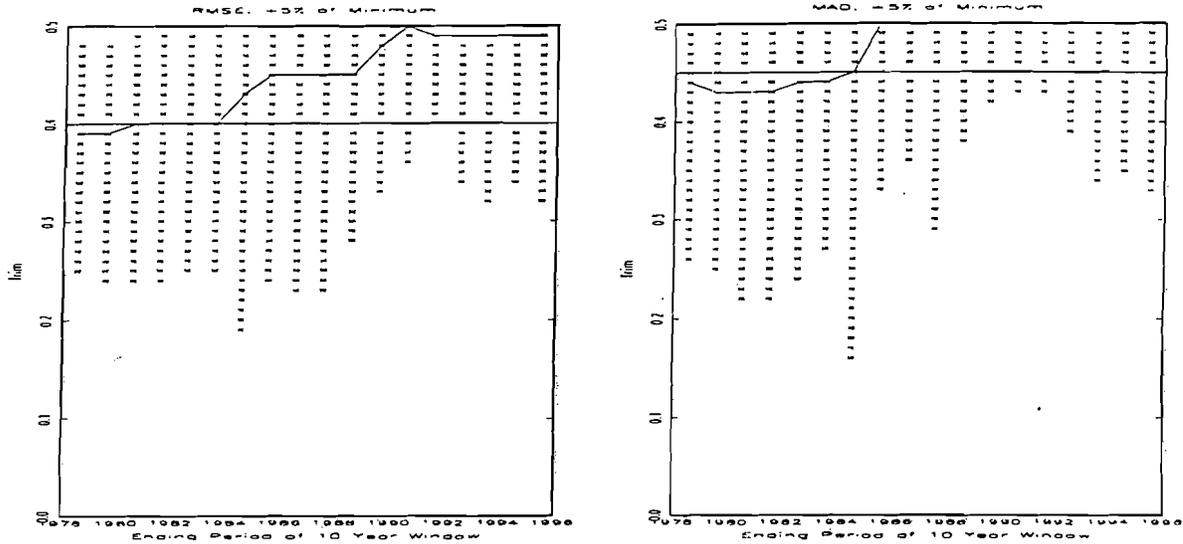
have a tendency to grow as the degree of the moving average increases. Second, for the PPI, there is little difference between the 'optimal trim' and the median. In all cases but one, the RMSE and MAD of the median are well within the 5% standard. Finally, for CPI at both levels of aggregation there is a large benefit to trimming a small amount.

5.4 Variations in the Sample Period

Next, we examine the sensitivity of the results to the sample period. This is analogous to asking whether the underlying distributional characteristics of the data are stable. To do this, we perform a series of Monte Carlo experiments comparable to those in Section 4, but instead of using the full sample from which to draw, we use rolling ten year samples. For example, in the case of the CPI we compute the optimal trim based on data from 1967 to 1976, then from 1968 to 1977, moving forward twelve months at a time.

Figures 10 and 11 report the results of these experiments. Each figure has a horizontal line at the optimal trim calculated using the full sample, together with a second line plotting the optimal trim based on each of the ten year samples. The horizontal axis shows the final date of the sample. To give some sense of precision, the X's in the figures represent the approximate confidence intervals constructed as all of the trims such that

FIGURE 11: Producer Prices
Efficiency of Trimmed Estimators, Changing Sample



the criterion, RMSE or MAD, is within 5 percent of the minimum.

The RMSE and MAD of the optimal full-sample trim are nearly always within 5 percent of the minimum value for the 10 year sub-samples. In fact, for the CPI, using the mean absolute deviation (MAD) criteria, the optimal trim is never outside of this rough confidence bound. For the PPI, there are thirty-six 10 year sub-periods. Using the RMSE criteria, the optimal full sample trim of 40 percent is within the confidence band in 33 of the 36 cases.

5.5 Summary and Comparisons

Given that the “CPI excluding food and energy” is the measure of core inflation in common use, it is useful to compare this measure of core inflation to ours. We do this in two ways. First, we ask which components we are trimming. And second, we look at a closer comparison of various candidate measures based on the RMSE criteria used above.

Table 4 examines which components we are trimming. For each month, we counted the frequency at which some portion of the weight of each component was trimmed using the optimal trim — 9% for the CPI and 40% for the PPI. We also note which components are systematically excluded by the ‘ex food and energy’ measures (highlighted in bold-faced

Table 4: Frequency That a Component is Trimmed: CPI 9% trim

CPI Component	Average Relative Importance	Percent of Sample period that a portion of the good is trimmed
Fruits and vegetables	2.26	69.61
Motor fuel	3.82	67.13
Fuel oil and other household fuel commodities	0.80	59.94
Used cars, etc.	2.27	58.84
Infants and toddlers apparel	0.11	54.97
Meats, poultry, fish and eggs	4.61	54.70
Womens and girls apparel	2.71	43.09
Public transportation	1.41	40.33
Other apparel commodities	0.58	37.85
Other private transportation commodities	0.67	37.85
Gas and electricity (energy services)	3.35	34.81
Tobacco and smoking products	1.63	33.43
Dairy products	1.92	28.73
Other private transportation services	3.35	24.59
Mens and boys apparel	1.90	23.48
Other utilities and public services	2.30	23.20
Personal and educational services	1.96	22.65
Toilet goods and personal care appliances	0.93	20.99
Medical care services	5.20	20.72
Other food at home	3.01	19.06
Footwear	1.02	19.06
Cereals and bakery products	1.86	17.96
School books and supplies	0.48	17.96
New vehicles	3.64	17.13
Housekeeping supplies	1.37	16.57
Housefurnishings	4.00	16.30
Entertainment services	1.88	15.47
Medical care commodities	1.01	14.92
Shelter	25.24	12.98
Housekeeping services	1.80	9.67
Entertainment commodities	2.37	7.46
Personal care services	0.94	7.18
Alcoholic beverages	1.73	6.91
Apparel services	0.92	5.25
Auto maintenance and repair	1.37	3.87
Food away from home	5.58	3.31
Mean of All Items		26.89
Mean of Food & Energy		39.93

Table 5: Frequency That a Component is Trimmed: PPI 40% trim

PPI Component	Average Relative Importance	Percent of Sample period that a portion of the good is trimmed
Farm products	7.47	98.90
Fats and oils	0.42	97.52
Meats, poultry, and fish	3.56	96.97
Prepared animal feeds	1.22	96.14
Fuels and related products and power	12.16	96.14
Metals and metal products	11.86	92.84
Hides, skins, leather, and related products	0.81	90.08
Lumber and wood products	2.40	88.98
Sugar and confectionery	1.04	87.88
Electronic computers and computer equipme	0.65	86.78
Transportation equipment	8.88	86.78
Chemicals and allied products	6.86	86.23
Processed fruits and vegetables	0.75	85.67
Dairy products	1.72	85.40
Cereal and bakery products	1.58	83.20
Miscellaneous processed foods	1.15	82.64
Miscellaneous Instruments	0.55	82.09
Beverages and beverage materials	1.90	81.54
Motor vehicles and equipment	7.01	81.54
Miscellaneous products	3.47	80.72
Electrical machinery and equipment	4.54	78.79
Construction machinery and equipment	0.74	78.51
Agricultural machinery and equipment	0.58	77.41
Textile products and apparel	5.33	77.13
Rubber and plastic products	2.56	77.13
Pulp, paper, and allied products	6.82	76.03
Nonmetallic mineral products	2.75	74.10
Miscellaneous machinery	1.73	72.45
Special industry machinery and equipment	1.19	71.35
Furniture and household durables	2.98	70.80
General purpose machinery and equipment	2.06	69.42
Metalworking machinery and equipment	1.24	66.39
Mean of All Items		83.05
Mean of Food & Energy		90.08

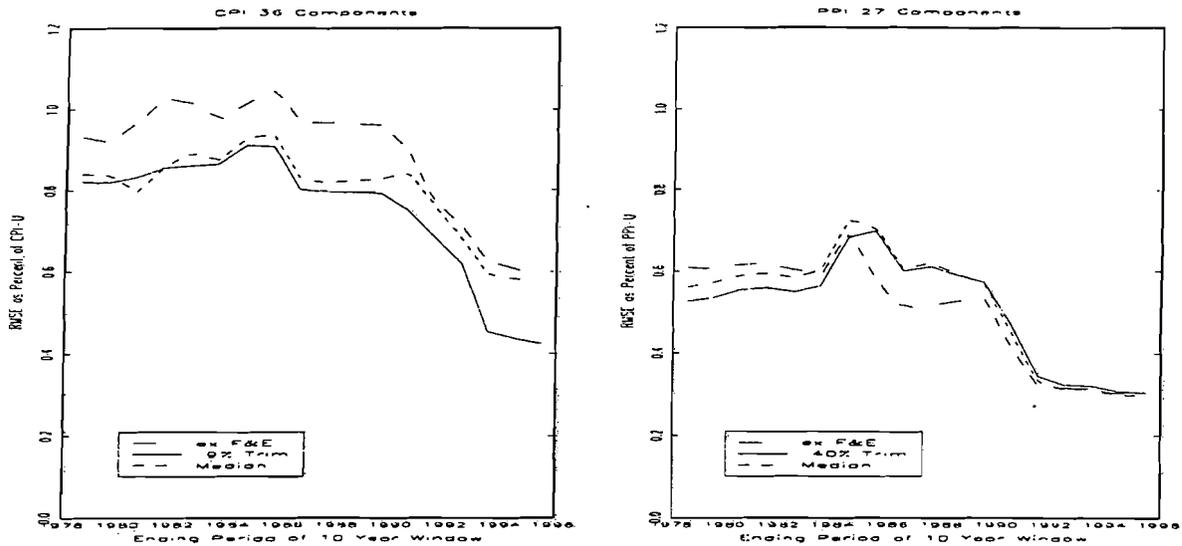
type). The results show that we often trim some of the food and energy prices. Indeed, for the CPI, food and energy components are trimmed from the efficient estimator nearly 40% of the time — nearly one and one-half times as frequently as the average component. Still, some food and energy goods, notably food away from home, appear to provide an efficient signal of core inflation as we define it here. In fact, of the 36 CPI components considered, food away from home was the least likely to be trimmed. Moreover, many non-food, non-energy goods appear to provide little information about the economy's inflation trend. Notable among these are used cars and infant and toddler apparel that are likely to be trimmed out of the efficient estimator nearly twice as frequently as the average good (the average component is trimmed out of the 9% trimmed mean in 27% of the months in the sample).

The components most likely to be included in the calculation of the efficient CPI estimator include a wide variety of services and the shelter component which, despite its high average relative importance of 25.24, is likely to be on one of the trimmed tails of the price change distribution only about 13% of the time.

Similarly for the PPI, food and energy goods tend to be trimmed from the efficient estimator a disproportionately large share of the time. But some food components, such as beverages and beverage materials and miscellaneous processed foods, are trimmed at the same frequency as the average component. The least frequently trimmed component, metalworking machinery and equipment, is still trimmed about two-thirds of the time. This is a relatively low proportion when one considers that, for any given month, 80% of the price change distribution is trimmed to produce an efficient estimator for PPI core inflation.

Finally, in Figure 12 we plot the ratio of the RMSE of various measures to the RMSE of the CPI-U and PPI themselves over different sample periods. For example, for the ten-year period ending July 1995, the RMSE for the CPI 'ex food and energy' was 57.8% than of the CPI-U itself — about the same as that of the median. But the RMSE of the 9% trim was 42.5% of the RMSE of the CPI-U. The main result is that, for the CPI, the 9% trim is always more efficient than the CPI excluding food and energy. But for the optimally trimmed PPI and the PPI 'ex food and energy' are very close.

FIGURE 12: Comparison of Various Estimators
Efficiency with Changing Sample



6 Conclusion

In this paper we challenge the conventional wisdom that core inflation can be measured by simply excluding food and energy from monthly price data. We show that price change distributions are highly leptokurtic, or 'fat-tailed,' and so commonly used measures, such as the sample-mean, are inefficient estimators of the population mean of interest. We demonstrate that trimmed-mean estimators significantly improve the efficiency of inflation estimates. Furthermore, we are able to show that as the kurtosis of the distribution increases, efficiency dictates trimming an increasing percentage of the sample.

We proceed to apply these insights to inflation data. For consumer prices beginning in 1967, we find that trimming 9% from each tail of the cross-sectional price-change distribution produces the minimum root-mean-square error and minimum mean-absolute deviation estimate of monthly inflation. This estimator provides efficiency improvements on the order of 23 percent relative to the mean. By contrast, the CPI excluding food and energy provides virtually no efficiency improvement at all.

More disaggregated data amplify the difficulties, as the kurtosis of the distributions

increases. Moving from a dataset composed of 36 components of the CPI to one with 185 components beginning in 1978, we show that the optimal trim nearly doubles to 16%. Here we find an efficiency gain of nearly 50 percent (although the sample period is substantially shorter). For producer prices beginning in 1947, where price-change distributions are more leptokurtic, trimming 40% to 50% from each tail produces the most efficient estimate of monthly aggregate price movements and improves efficiency by over 40 percent relative to the mean.

References

- 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.
- 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.
- _____ and _____, 'Inflation and the Distribution of Price Changes,' N.B.E.R. Working Paper No. 5793, October 1996 (revised June 1997).
- Caballero, Ricardo J. and Eduardo Engel, 'Dynamic (S,s) Economies,' *Econometrica* 59 (November 1991) p. 1659-1686.
- Caplin, Andrew and John Leahy, 'State-Dependent Pricing and the Dynamics of Money and Output,' *Quarterly Journal of Economics* 106 (August 1991) p. 683-708.
- Cecchetti, Stephen G., 'Measuring Short-Term Inflation for Central Bankers,' *Economic Review of the Federal Reserve Bank of St. Louis*, forthcoming. (NBER Working Paper No. 5786.)
- Carlson, John B. and Edward Bryden, 'On Disinflation since 1982: An Application of Change Point Tests,' *Economic Review of the Federal Reserve Bank of Cleveland* 30 (1994 Q1).
- Huber, Peter J. *Robust Statistics*. New York: John Wiley and Co., 1981.
- Roger, Scott, 'A Robust Measure of Core Inflation in New Zealand, 1949-96,' mimeo., Reserve Bank of New Zealand, March 1997.
- Stuart, Alan and J. Keith Ord. *Kendall's Advanced Theory of Statistics*. New York: Oxford University Press, 1987.
- Yule, G. Udny and M.G. Kendall. *An Introduction to the Theory of Statistics*. New York: Hafner Publishing Co., 1968.

Federal Reserve Bank
of Cleveland
Research Department
P.O. Box 6387
Cleveland, OH 44101

Address Correction Requested:

Please send corrected mailing label to the
Federal Reserve Bank of Cleveland,
Research Department,
P.O. Box 6387,
Cleveland, OH 44101

BULK RATE
U.S. Postage Paid
Cleveland, OH
Permit No. 385