DP2025-02

Reserve Bank of New Zealand
Te Pūtea Matua

Discussion Paper

Menu costs, inflation targets, and cross-sector price flexibility: evidence from a trimmed mean decomposition of New Zealand inflation, 1984-2022

Andrew Coleman

December 2025

Reserve Bank of New Zealand *Discussion Paper*
ISSN 1177-7567
The Reserve Bank of New Zealand's discussion paper series is externally refereed.
The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Reserve Bank of New Zealand
The Bank's *Discussion Paper* series is available here.

Reserve Bank of New Zealand
PO Box 2498
Wellington
NEW ZEALAND
www.rbnz.govt.nz

Abstract

This paper uses the distribution of sectoral inflation rates to decompose New Zealand's inflation rate history into a trimmed mean core inflation rate plus upper and lower trimmings (the highest and lowest price movements each quarter). This decomposition (i) shows why the trimmed mean inflation rate is biased, and how this bias can be corrected; (ii) identifies some fast-reversing transitory components of the inflation rate and (iii) helps explain how the variance and skewness of sectoral inflation rates depend on the inflation rate. The paper shows the distribution of sectoral inflation rates varies with inflation in a manner consistent with menu cost models: upward relative price flexibility steadily increases as inflation increases from zero, but downwards price flexibility first decreases and then increases. The evidence suggests an inflation target of 2 - 3 percent minimises downwards price flexibility.

JEL codes: E31

Keywords: core inflation, cross-sector price variability; price flexibility; trimmed mean estimators.

1. Introduction

The speed and size of the price increases that took place between 2021 and 2023 surprised central banks worldwide. Since then, economists have re-examined their assumptions and models to better understand why prices increased so quickly – and why central banks were surprised by the surge in inflation. At the macroeconomic level, economists are examining the circumstances whereby supply shocks and increased fiscal deficits trigger inflation. At the microeconomic level, they are examining the times when firms change prices to understand why macroeconomic events occasionally trigger an avalanche of price changes, but at other times do not (Cavallo, Lippi, and Miyahara 2024; Nirei and Scheinkman 2024).

Research on firm price setting behaviour has undergone a resurgence in the last two decades as new firm-level data sources have become available, particularly the firm-level data used to construct the consumer price indices. These data have allowed researchers to distinguish between the size and frequency of price changes, and to uncover new facts about the relative importance of time-dependent and state-dependent price changing decisions. In conjunction with advances in theoretical models about price changing frictions, especially menu-cost models, this research has helped us understand the circumstances when the frequency of price changes accelerates or slows down, and the circumstances where price changes are persistent rather than temporary (Nakamura and Steinsson 2008; Gagnon 2009; Berger and Vavra 2018; Alvarez et al 2019; Sheremirov 2020; Luo and Villar 2021).

One strand of the recent literature examines the way that inflation affects firm price-setting behaviour. Following Ball and Mankiw (1994, 1995) and Golosov and Lucas (2007), various authors have developed theoretical models that analyse how inflation changes the frequency and size of firm-level price changes. These models typically assume that it is costly for a firm to change prices, so it is not worthwhile making small changes. When inflation increases, the number of firms finding it worthwhile to increase prices will increase, while the number wanting to decrease prices will fall. In these circumstances prices are increasingly upwardly flexible, and economic shocks can quickly turn inflationary. Moreover, as inflation increases firms increasingly respond to the common macroeconomic inflation rate rather than idiosyncratic firm-level shocks. This description of pricing behaviour is broadly supported by firm-level datasets. In particular, there is evidence that the pattern of firm-level price changes is different in low and high inflation environments, and that large shocks can trigger rapid adjustments to the inflation rate (Gagnon 2009; Alvarez et al 2019).

While the firm-level data provide great insight into price-setting behaviour and inflation dynamics, they are rarely available to central banks in a timely fashion, and not available at all in many countries including New Zealand. Since both theory and evidence suggest the distribution of price changes is closely tied to the inflation rate, central banks may find it useful to use other measures of price dispersion to provide additional information about inflation dynamics when firm-level price data are not available. One obvious source is the sector-specific price indices generally released by statistical agencies when they publish the aggregate consumer price index. These indices indicate the average price level of the goods and services sold in each sector rather than the prices of specific products sold by particular firms. Nonetheless, these data may convey useful information about the sensitivity and persistence of the aggregate inflation rate to shocks when firms within a sector have similar cost and demand pressures and face similar price frictions. These data are already used by many central banks to calculate a variety of core inflation rate estimates to help them distinguish price changes that are likely to have temporary rather than persistent effects on the inflation rate.

There are many different estimates of the core inflation rate that are based on sector-specific price indices.[1] One popular class of core measures is calculated by excluding the prices of volatile sectors such as energy and food from the headline Consumer Price Index (CPI). A second class uses statistical methods to find common latent factors in sectoral price data. A third, the trimmed mean, calculates the core inflation rate from the goods and services that remain when the goods and services with the largest or smallest sector-specific price changes each quarter are excluded. In each case the distribution of sector-specific price changes is used to estimate a different aspect of the underlying inflation rate, such as the extent that different components are persistent rather than transitory. However, it is not clear how the insights provided by these sector-specific price indices are related to the insights provided by firm-level data.

This paper analyses New Zealand's inflation history since 1984 through the lens of a trimmed mean decomposition of the inflation rate that is based on sector-specific price indices. The decomposition calculates four composites: the trimmed mean inflation rate, which is an estimate of the core inflation rate; the upper and lower trim inflation rates (the contributions to the headline inflation rate of the goods and services with the highest or lowest price changes each quarter); and the headline-trimmed mean inflation gap (the difference between the headline and trimmed mean inflation rates). The upper trim, lower trim, and trimmed mean inflation rates are constructed using the price changes measured by the approximately 109 sectoral-specific subindices that underlie the headline CPI.

The analysis has four main parts. First, it examines how each of these composites and the variance and skewness of the cross-sector distribution of price changes vary with the inflation rate. By analysing how the relationship between the headline and trimmed mean inflation rates varies with the inflation rate, it is possible to better interpret core measures of the inflation rate such as the trimmed mean or the median should prices rapidly increase again. Secondly, it investigates the extent that the trimmed mean estimates of the core inflation rate, combined with estimates of the upper and lower trims, can be used to identify transitory and persistent components of inflation. Thirdly, it examines the extent that the headline-trimmed mean inflation gap can be used to improve forecasts of the headline inflation rate. Lastly, it investigates how the whole cross-sector distribution of price changes varies as the average inflation rate increases. The fourth part provides a cross-sector perspective on inflation dynamics that complements the perspectives obtained from firm-level data.

[1] The core inflation rate loosely refers to the "underlying" inflation rate, but there is no single definition. Central banks use a range of core inflation estimators because each core estimator emphasises a different aspect of the whole distribution of price changes.

The first issue considered in the paper is the way the upper and lower trims, and the headline-trimmed mean inflation gap vary with the inflation rate, which is examined using regression analysis. The estimates show that as the trimmed mean inflation rate increases by 1 percentage point, the upper trim inflation rate increases by approximately 2 percentage points while the lower trim inflation rate increases by just under 1 percentage point. In turn this means the headline-trimmed mean inflation gap increases by approximately 0.14 percentage points for every 1 percentage point increase in the trimmed mean. These results indicate that increasingly large price increases in a small number of sectors contribute an increasingly large fraction of the headline inflation rate when the inflation rate rises. These very large price changes further mean the variance and skewness of the cross-sector price change distribution increase with the inflation rate. These findings are consistent with other international evidence based on sector-specific price change data that (i) trimmed mean estimators of the inflation rate are downwardly biased estimators of the headline inflation rate and (ii) the cross-sector distribution of price changes is increasingly skewed to the right as the average inflation rate increases. They mean the downwards bias in trimmed estimates of the inflation rate gets larger as the inflation rate gets higher, making it a less reliable indicator of underlying inflation precisely when inflation increases.

The second issue addressed in the paper is the extent that large positive or negative sector-specific price changes are likely to be transitory rather than persistent. A second set of regressions indicates that whenever the upper or lower trims deviate from their expected values (conditional on the trimmed mean) or, equivalently, whenever the headline-trimmed mean inflation gap deviates from its mean value, the unusually large deviations are transitory and are rapidly reversed out of the headline inflation rate. These results are consistent with several papers based on United States data that have demonstrated that the headline-core inflation gap reflects largely transitory effects on the headline inflation rate (Mehra and Reilly 2009; Gamber, Smith, and Eftimoiu 2015; Atkinson et al 2021).

These two sets of regression estimates suggest univariate inflation forecasts based on the headline inflation rate could be improved by generating forecasts that utilise the headline-trimmed mean inflation gap, or, alternately, the trimmed mean inflation rate and the upper and lower trims. This is the third major issue considered in the paper. To test this contention, the forecast accuracy of a variety of time series models was evaluated by estimating rolling regressions over the ten-year period from 2012 to 2022. The results suggest that recursive forecasts from a vector autoregression that includes the headline-trimmed mean inflation gap or the trimmed mean inflation rate outperform univariate time series forecasts that only use headline inflation rate data. There is also a forecast improvement if the median cross-sector price change is used instead of the trimmed mean. The improvement arises because any difference between the trimmed mean and the headline inflation rates is reversed from the headline rate within a quarter, and is most noticeable within the first four quarters.

The fourth issue examined in the paper is related to the finding that when the trimmed mean inflation rate increases by 1 percentage point the average upper trim inflation rate increases by 2 percentage points but the lower trim inflation rate increases by just under 1 percentage point. To further investigate this finding, non-linear regression and non-parametric methods were used to examine how the whole cross-sector distribution of prices changes as the inflation rate changes. There are two sets of results. First, when the inflation rate is less than approximately 5% per annum the distribution is bimodal, with one peak near zero and another slightly above the median cross-sector price change. At higher inflation rates the spike near zero disappears and the distribution is single peaked. Secondly, the "width" of the upper half of the distribution, measured by the distance between different deciles and the median cross-sector price change, is increasing in the median cross-sector price change. In contrast, the "width" of the lower half of the distribution is first decreasing but then increasing in the median cross-sector price change.

The contraction and then expansion of the lower half of the distribution as the inflation rate increases reflects the position of the spike near zero and is consistent with a menu cost explanation of firm-level price changes. When the inflation rate is near zero, firms wishing to cut relative prices do so; but when inflation is a small positive number, many firms keep prices fixed and wait for inflation to reduce relative prices to their desired level. At the firm level, this should result in a 'spike' in the probability distribution near zero. At the sector-specific level, the situation is unclear, because the distribution measures average price changes in different sectors, and there is no reason to expect the distribution of average price changes across sectors to be the same as the distribution of price changes across firms unless sector-specific shocks comprise a dominant share of idiosyncratic firm-level shocks. Nonetheless, the data clearly suggest there is a 'spike' of sector-specific price changes near zero, and a contraction of the width of the lower half of the distribution as annual inflation rises from zero to approximately 2%. It is at least possible that these patterns occur because shocks to firms in the same sector are dominated by common sector-specific factors at a quarterly frequency. When inflation is sufficiently high, as it was in the 1980s, the spike near zero shrinks and then disappears, presumably because many firms that wish to reduce their relative prices still find it necessary to increase their actual prices, and the additional inflation "widens" the lower half of the distribution.

Even though the paper uses cross-sector rather than firm-level data, the distributional results are consistent with many of the insights about the price changes that arise in the firm-level pricing literatures. Two results stand out. First, when the median sectoral inflation rate is 2% or more, the cross-sector distribution of price changes has less downwards price flexibility (or, equivalently, greater downward price stickiness) and greater upwards price flexibility than when the inflation rate is 1% or less. This suggests that an economy with a 2% inflation target could experience deeper output declines in response to negative shocks than an economy with a 1% inflation target, because there is more downwards price stickiness. Secondly, a central bank may find it difficult to raise the inflation rate to 2% whenever the inflation rate declines to 1% or less. When the inflation rate is 1% or less, firms wanting to cut relative prices make larger price reductions relative to the median sectoral inflation rate than when the annual inflation rate is 2% or more; at the same time firms wanting to raise relative prices increase their prices by a smaller amount relative to the median. If prices are more flexible downwards and less flexible upwards when the median sectoral inflation rate is less than 1% than when it is greater than 2%, central banks will likely find it more difficult to raise the average inflation rate from 1% to 2% than from 2% to 3%. In the parlance of Caballero and Engel (2007), the price change hazard curve may be particularly low when the inflation rate is low, making it possible for inflation to be embedded at low rates.

The paper demonstrates that the distribution of sector-specific price changes becomes increasingly skewed to the right when the inflation rate increases, consistent with international findings that the distribution of firm-level price changes is increasingly skewed to the right as inflation increases. This means that the size of the changes of prices in the upper trim increases faster than the rest of the distribution when inflation increases. Many but not all of these changes are temporary and occur in sectors that are very volatile and are often found in the lower as well as upper trims. This paper demonstrates that the headline-trimmed mean inflation gap can be used to modestly improve short term inflation forecasts relative to those made only using the headline rate. However, it remains for further research to establish whether changes in the upper trim, or the right tail of the price change distribution more generally, can help explain movements in inflation expectations or the price change hazard curve, or whether they are particularly sensitive to changes in monetary policy.

The remainder of this paper is organized as follows. Section 2 provides the framework describing how the trimmed mean and the upper and lower trims are calculated, followed by the basic statistical properties of these series. Section 3 estimates the relationships between the headline inflation rate, the trimmed mean inflation rate, and upper and lower trims to decompose the inflation rate into temporary and persistent components. It includes an evaluation of the forecast performance of time series models that incorporate the headline-trimmed mean inflation rate gap and the upper trim, relative to univariate models that only use the headline inflation rate. It also provides an analysis of the sectors that are most frequently found in the upper and lower trims. Section 4 uses a mixture of regression models and non-parametric techniques to estimate how the whole distribution of sectoral inflation rates varies with the inflation rate. Finally, conclusions are offered in Section 5.

Figure 1: Splitting the CPI into trimmed mean and the trimmings.

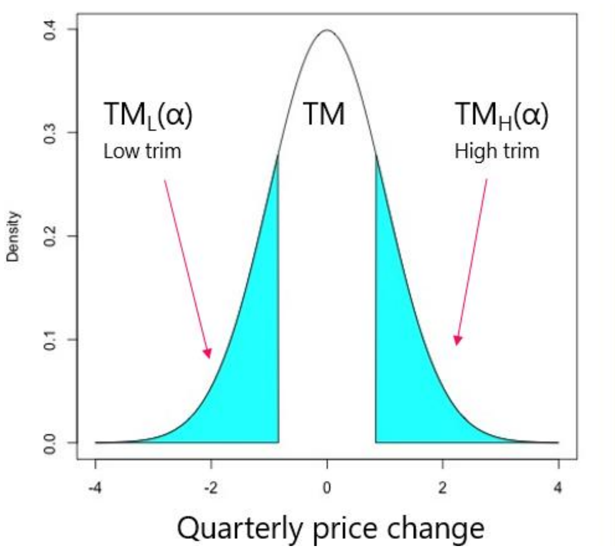

2. The trimmed mean estimator

2.1 Definition of the trimmed mean estimator

The trimmed mean decomposition of the inflation rate used in this paper is calculated by:
(1) ordering the quarterly price changes of each of the approximately 100 subgroups of the Consumer Price Index from smallest to largest;
(2) calculating the contribution to headline CPI inflation of the lowest α_L percent of the distribution ($TM_{L,t}$), the highest α_H percent of the distribution ($TM_{H,t}$) and the remaining $(1-\alpha_L-\alpha_H)$ percent of sectors in the middle of the distribution (TM_t); and
(3) calculating the average price increase each quarter in each of these components, $TM_{L,t}^*$, TM_t^*, and $TM_{H,t}^*$.

Formally, suppose there are M_t separate components of the CPI at time *t*, and that the *i*th component has a price change $\pi_t^i = P_t^i/P_{t-1}^i - 1$ and a weight w_t^i. Each period the series can be reordered by the size of the price change from $i_t^* = 1$ to $i_t^* = M_t$, with $\pi_t^{i^*} \leq \pi_t^{i^*+1}$. If there are $n_t^L(\alpha_L)$ sectors in the lower trim, where $\sum_{i*=1}^{n_L} w_t^{i*} \geq \alpha_L \; and \; \sum_{i*=1}^{n_L-1} w_t^{i*} < \alpha_L$ the contribution of the lower trim to the inflation rate is

$$TM_{L,t} = \sum_{i*=1}^{n_L-1} w_t^{i*}\pi_t^{i*} + (\alpha_L - \sum_{i*=1}^{n_L-1} w_t^{i*})\pi_t^{n_L*} \qquad (1)$$

and the average price increase of the sectors in the lower trim is

$$TM_{L,t}^* = TM_{L,t}/\alpha_L \qquad (1a)$$

The second term in equation 1 is an adjustment that occurs when the n_L^{th} sector is split between the lower trim and the trimmed mean inflation rate.

If there are $n_t^H(\alpha_H)$ sectors in the upper trim, where $\sum_{i*=n_H}^{M_t} w_t^{i*} \geq \alpha_H \; and \; \sum_{i*=n_{H-1}}^{M_t} w_t^{i*} < \alpha_H$, the contribution of the upper trim to the inflation rate is

$$TM_{H,t} = \sum_{i*=n_H+1}^{M_t} w_t^{i*}\pi_t^{i*} + \left(\alpha_H - \sum_{i*=n_H+1}^{M_t} w_t^{i*}\right)\pi_t^{n_H*} \qquad (2)$$

and the average price increase of the sectors in the upper trim is

$$TM_{H,t}^* = TM_{H,t}/\alpha_H \qquad (2a)$$

The contribution of the middle sections to the inflation rate is

$$TM_t(\alpha_L, \alpha_h) = \left[\sum_{i*=n_L+1}^{n_H-1} w_t^{i*}\pi_t^{i*} + (\sum_{i*=1}^{n_L} w_t^{i*} - \alpha_L)\pi_t^{n_L*} + \left(\sum_{i*=n_H}^{M_t} w_t^{i*} - \alpha_H\right)\pi_t^{n_H*}\right] \qquad (3)$$

The last two terms in this expression capture the contribution to the trimmed mean of the sectors that are partly split across the lower and upper trims. The trimmed mean estimate of the inflation rate is the average price change of the middle sections:

$$TM_t^*(\alpha_L, \alpha_h) = TM_t\;(\alpha_L, \alpha_h)/(1 - \alpha_L - \alpha_H) \qquad (3a)$$

This decomposition means the quarterly changes in the headline inflation rate π_t can be split into three components:

$$\pi_t = TM_{L,t} + TM_t \; + TM_{H,t}\,. \qquad (4)$$

In practice, there is a small difference between the headline inflation rate π_t calculated by equation 4 and the headline inflation rate produced by Statistics New Zealand, as the latter is calculated as the change in the weighted sum of sector-specific price levels (the consumer price index), whereas equation 4 is calculated as the weighted sum of sector-specific price changes. This difference is sufficiently small that it can be ignored.

The CPI-TM gap is the difference between the headline and trimmed mean measures: $CPITM_t = CPI_t - TM_t^*$.
This means there is a one-for-one relationship between the gap and the upper and lower trims:

$$CPITM_t = TM_{L,t} + TM_{H,t} - (\alpha_L + \alpha_H)TM_t^* \quad (5)$$

$$= \alpha_L(TM_{L,t}^* - TM_t^*) + \alpha_H(TM_{H,t}^* - TM_t^*) \quad (5a)$$

In the limit when $\alpha_L = \alpha_H = 0.5$, the trimmed mean estimate $TM_t^*(0.5,0.5)$ is equal to the median sectoral inflation rate.[2] Both the trimmed mean TM_t^* and the median sectoral inflation rate are frequently used by central banks to estimate the core inflation rate.[3] Trimmed mean estimators are typically better estimators of the underlying core inflation rate than the headline inflation rate when the distribution of price changes has many extremely large or extremely small changes, for the trimmed mean reduces the random effects that these large price changes can have on the headline rate (Bryan and Cecchetti 1994; Bryan, Cecchetti and Wiggins 1997). The optimal size of the trim depends on the underlying variability (particularly the kurtosis) of the distribution. This paper uses a symmetric 20% trim (i.e. $\alpha_L = \alpha_H = 20\%$) because it is the measure that most closely matches a measure of core inflation that is based on the average headline inflation rate in New Zealand over the period 1999 – 2022 (Greenaway-McGrevy and Jones 2022).[4] Formally, this estimate should be called $TM_t^*(20,20)$, but throughout the paper TM_t^* is used to refer to $TM_t^*(20,20)$ unless otherwise specified. For the 1999-2022 period the estimates are based on 109 subgroup classifications, and a similar number of classifications are used for the earlier period.[5]

Statistics New Zealand currently calculates several sector-based trimmed mean estimators based on 700 categories of price changes, dating back to 2003 (Goodchild 2003). The largest values for α_L and α_H used to calculate these estimators are $\alpha_L = \alpha_H = 0.15$. There are two reasons for calculating the new trimmed mean estimator used in this paper, rather than using the Statistics

[2] It follows from this definition that the median sectoral inflation rate (or, equivalently, the median cross-sector price change) is a weighted median equal to the quarterly price change of the $n_t^L(\alpha_L)$ sector when $\alpha_L = 0.5$. The weights are the CPI commodity weights, which vary over time.

[3] There is not a single definition for either the “core” or the “underlying” inflation rate any more than there is a single definition of a “smooth” series. Rather, there are a range of estimation techniques that attempt to distinguish price changes that are likely to have temporary rather than persistent effects on the inflation rate.

[4] Greenaway-McGrevy and Jones apply to New Zealand data the procedure used by the Federal Reserve Bank of Dallas to calculate the optimal trims and find the optimal upper and lower trims that match a rolling average of the headline CPI. They find the optimal lower and upper trims are 19% and 20%. These numbers are similar to the optimal trims found using U.S. data. See Dolmas (2005) for the Federal Reserve Bank of Dallas procedure.

[5] Because Statistics New Zealand periodically changes the data classifications it uses, it is not straight-forward to calculate the trimmed mean estimator back to 1984. The main data available from Statistics New Zealand for the period earlier than 1998 are available either in 70 categories or 300 categories, neither of which is easily comparable with the 109 categories used for the 1999-2022 period. To make the estimates, the disaggregated “300-category” data available for the 1984-1999 data were aggregated to approximately 100 categories. This was done separately for three sub-periods, 1983 – 1993, 1993-1996 and 1996 – 1999 as the product categories changed. Contemporaneous CPI weights were used each year. See Appendix 1 for details.

New Zealand estimators. First, the estimator in this paper can be backdated to the 1980s. The data Statistics New Zealand use to calculate their trimmed mean estimators are not available for the full period, but by using less disaggregated data a new estimator that covers the whole period could be calculated. Secondly, the evidence from Greenaway-McGrevy and Jones (2022) suggests that the optimal size of the trimmings is larger than the trimmed mean estimators calculated by Statistics New Zealand, because of the large number of extremely large price changes in the data.

2.2 Basic statistical properties of the price series

The graphs of the quarterly headline inflation rate and the TM_t^* estimator and the average inflation rate in the lower and upper trims ($TM_{L,t}^*$ and $TM_{H,t}^*$) are shown in figures 2 and 3. The graphs indicate that the average value of the headline inflation rate, the trimmed mean inflation rate, and the upper trim declined significantly after the Reserve Bank adopted an inflation targeting regime in 1990. Since it is reasonable to treat the adoption of inflation targeting as a regime change, the sample mean, standard deviation, and first order autocorrelation coefficient for the five series {π_t; TM_t^*; $TM_{L,t}^*$, $TM_{H,t}^*$; $CPITM_t$} and the corresponding correlation matrix are calculated separately for the 1984-1989 and 1990 – 2022 sub-periods. These statistics are shown in Table 1. [6]

The statistics in Table 1 confirm that there is a structural break in the data in 1990. Three aspects of the series are different in the two sub-periods. First, the mean of the (quarterly) headline and trimmed-mean inflation rates fell substantially after 1990, by 1.59 percentage points and 1.35 percentage points respectively. The mean price changes in the lower and upper trim declined by 1.20 percentage points and 2.70 percentage points. In each case the hypotheses that the mean of each series is identical in both sub-periods can be rejected at the 1% significance level. [7] It can be seen the decline in the upper trim is much larger than the decline in the other variables. Secondly, the variances of the (quarterly) headline inflation rate, trimmed mean inflation rate, and upper trim are all substantially smaller after 1990 than before 1990, and these differences are statistically significant at the 1% confidence level. In contrast, the variances of the quarterly changes in the lower trim and in the CPITM gap are almost the same in each period. Thirdly, there are some changes in the correlation matrix. The most notable is the decline in the correlation between the headline and trimmed mean inflation rates after 1990, and the increase in the correlation between the headline inflation rate and the CPITM gap. These changes reflect the decline in the variance of the trimmed mean inflation rate after 1990. As a result of these changes fluctuations in the upper and lower trims comprise a much larger fraction of the variation in headline inflation rate after 1990 than was the case before 1990.

[6] Note that the data used in Table 1 has been filtered to remove the effects of the three changes in the Goods and Services tax rate that occurred in 1986, 1989, and 2010. These changes are visible in figure 2, which is based on unfiltered data.

[7] A t-test is used to test the hypothesis that the means for the 1984-1989 and 1990-2022 periods are equal. Since the variances are unlikely to be the same across the two regimes, the t-statistic is calculated using the Welch test procedure.

Table 1: Descriptive statistics

					Correlation Matrix					First order correlation	
	Mean	Std	T-test[(a)]	F-test[(b)]	π	TM^*	TM^*_L	TM^*_H	CPITM	Coefficient	(se)
1984:1- 1989:4											
π	2.15%	1.14%			1	0.97	0.65	0.82	0.41	0.61	0.17
TM^*	1.86%	1.05%				1	0.54	0.76	0.16	0.72	0.14
TM^*_L	-1.30%	1.33%					1	0.22	0.55	0.25	0.21
TM^*_H	6.49%	2.14%						1	0.46	0.51	0.18
$CPITM^c$	0.28%	0.30%							1	-0.5	0.19
1990:1- 2022:4											
π	0.56%	0.53%	6.58^{**}	4.83^{**}	1	0.79	0.64	0.77	0.79	0.41	0.081
TM^*	0.51%	0.33%	6.12^{**}	10.20^{**}		1	0.26	0.60	0.24	0.67	0.066
TM^*_L	-2.49%	1.17%	4.03^{**}	1.34			1	0.19	0.75	0.1	0.087
TM^*_H	3.78%	1.19%	5.94^{**}	3.34^{**}				1	0.62	0.39	0.082
CPITM	0.05%	0.34%	3.42^{**}	0.84					1	0.12	0.087
1984:1 – 2022:4											
π	0.80%	0.87%			1	0.93	0.66	0.87	0.63	0.73	0.056
TM^*	0.72%	0.71%				1	0.46	0.78	0.30	0.86	0.042
TM^*_L	-2.31%	1.27%					1	0.34	0.73	0.22	0.078
TM^*_H	4.20%	1.69%						1	0.59	0.64	0.062
CPITM	0.08%	0.34%							1	0.1	0.08

Source: Author's calculations using Statistics New Zealand price data.
[a] "T-test" is the test that the means for the periods 1984-1989 and 1990-2022 are equal. Since the variances are unlikely to be the same across the two regimes, the t-statistic is calculated using the Welch test procedure. The degrees of freedom, which depends on the variances of each series, range from 23 – 33. A * indicates the hypothesis of equality can be rejected at the 5% significance level; ** means the hypothesis can be rejected at 1% significance level.
[b] "F-test" is the test that the variances for the periods 1984-1989 and 1990-2022 are equal. The test is the ratio of the variances and has degrees of freedom F(23, 131). A * indicates the hypothesis of equality can be rejected at the 5% significance level; ** means the hypothesis can be rejected at 1% significance level.
[c] $CPITM_t = \pi_t - TM^*_t$ = headline-trimmed mean inflation gap

__Figure 2: Quarterly changes in the headline CPI and trimmed mean inflation rates, 1984- 2022.__

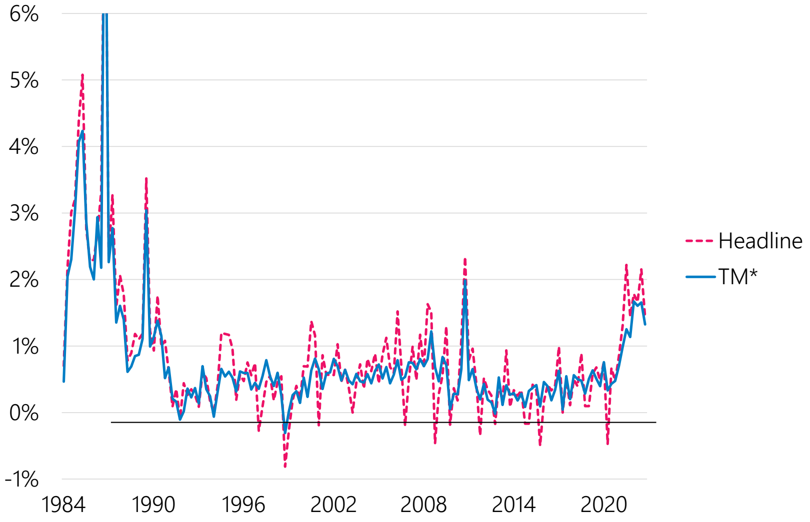

Source: Statistics New Zealand (CPI); author's calculations (TM).

__Figure 3: The upper and lower trim inflation rates, 1984-2022.__

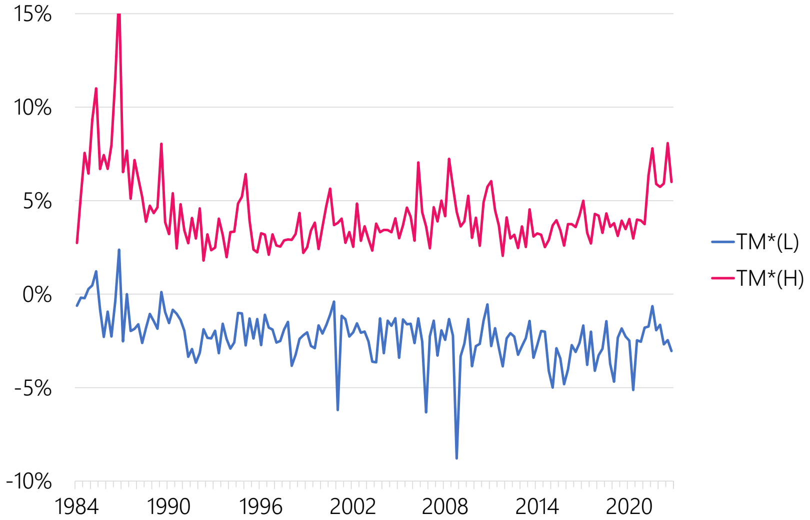

Source: Author's calculations using Statistics New Zealand data.

The final columns of Table 1 show the simple first order autocorrelation coefficients for each series (formal unit root tests are described below). Three results should be noted. First, the trimmed mean and headline inflation rates have considerable persistence (high first order autocorrelation coefficients), both for the full period (1984-2022) and for the post-1990 period. The first order correlation coefficients for the headline and trimmed mean inflation rates post-1990 were 0.41 and 0.67 respectively; for the periods 1984-1989 and 1984-2022 they all exceeded 0.60. The persistence of the trimmed mean is consistently higher than the headline inflation rate. Secondly, the first order autocorrelation coefficients of the CPITM gap are much closer to zero than the coefficients of either the headline inflation rate or the trimmed mean. This indicates that fluctuations in the difference between the headline inflation rate and the trimmed mean revert to the mean at a much more rapid pace than fluctuations in either the headline or trimmed mean inflation rates. In turn, this suggests the two series can be usefully described by an error correction model.[8] Thirdly, the lower trim has a much smaller first order autocorrelation coefficient than the upper trim and is much less correlated with the trimmed mean inflation rate or the headline inflation rate.

2.3 Unit root tests of the price series

Three sets of unit root tests were estimated for each of the five series. The first and most important set is for the post-1990 period. The preferred tests are the augmented Dickey Fuller tests applied to an autoregressive equation with either 1 or 4 lagged difference terms:

$$\Delta x_t = \alpha_0 + \alpha_1 x_{t-1} + \sum_{i=1}^{n} \theta_i \Delta x_{t-i} + e_t \qquad n = 1, 4 \tag{6}$$

The second set, which is included for completeness, covers the short period 1984 – 1989. The small number of observations and the high standard errors mean little is learnt from this exercise.

The third set covers the whole 1984-2022 period. The tests associated with this set are problematic because the evidence from Table 1 suggests that (i) the variance of the headline inflation rate, the trimmed mean inflation rate, and the upper trim series are significantly higher before 1990 than after 1990, and (ii) all of the series have different means before and after 1990. Since the variances are different in the two sub-periods, the augmented Dickey Fuller test will not have a standard Dickey Fuller distribution. A possible way to test the hypothesis $\boldsymbol{\alpha_1 = 0}$ is to conduct the "wild bootstrap" unit root test procedure developed by Cavaliere and Taylor (2009) that allows for heterogeneity in the variances of the innovations in the two subperiods. This procedure was attempted under the assumption that (i) the mean of the combined-period series was the same in both sub-periods, and (ii) the coefficients $\boldsymbol{\alpha_1}$ and $\boldsymbol{\theta_i}$ were the same in the two sub-periods under the alternative hypothesis that the series could not be described by a unit root process.[9] However, there is no reason to believe that either the means or the coefficients $\boldsymbol{\alpha_1}$ and $\boldsymbol{\theta_i}$ are the same in the separate subperiods as the monetary policy regimes before and after 1990 were so different.

[8] Data can be described by an error correction model even if they are stationary. Error correction models for stationary data are isomorphic with autoregressive distributed lag models, and are used as a means of describing the adjustment process when the difference of two series reverts to the mean difference more quickly than either series reverts to its mean (see Hendry, Pagan and Sargan (1984) for a theoretical perspective, and Alogoskoufis and Smith (1991) for a history of these models). Consider the following representations of two stationary series x and y:

$x_t = \alpha_1 x_{t-1} + \alpha_2 y_{t-1} + u_{1,t-1}$ and $y_t = \beta_1 x_{t-1} + \beta_2 y_{t-1} + u_{2,t-1}$

These can be transformed into

$\Delta x_t = \vartheta_1 (x_{t-1} - y_{t-1}) + \vartheta_2 x_{t-1} + v_{1,t-1}$ and $\Delta y_t = \theta_1 (x_{t-1} - y_{t-1}) + \theta_2 y_{t-1} + v_{2,t-1}$ where $\vartheta_1 = -\alpha_2$ and $\theta_1 = \beta_1$

The two pairs of equations are mutually consistent. However, the second pair of equations, which includes the error correction terms, often capture the underlying adjustment mechanisms better than the first pair of equations. This is the case when $\boldsymbol{\vartheta_2}$ and $\boldsymbol{\theta_2}$ are near zero. It is in this sense that I refer to either the third or fourth equations as an error correction model. This means a pair (or larger system) of stationary variables can be described by an error correction representation without the variables having unit roots or being cointegrated.

[9] The Cavaliere and Taylor test procedure was applied to the whole series (1984 – 2022) under the assumption that the data could be described by equation 6 with a common mean and common coefficients $\boldsymbol{\alpha_1}$ and $\boldsymbol{\theta_i}$. It was not possible to reject the hypothesis that the headline CPI, the trimmed mean, and the upper trim are unit root processes although the hypotheses that the CPITM gap and the lower trim are unit roots processes could be rejected. The results are available from the author on request. Note this test is of dubious value, since the distribution of the test statistic is unknown in the more likely case that the means and coefficients of the series are different before and after 1990.

This means the "wild-bootstrap" may not provide an accurate description of the distribution of the augmented Dickey Fuller test statistic as the assumed alternative hypothesis (that the mean and the coefficients α_1 and θ_i are the same in the two sub-periods) is unlikely to be appropriate. Moreover, if the variances, means, and autocorrelation coefficients are all different under the alternative hypothesis that the series are not described by a unit root process, there is little point combining the data from the separate sub-periods. [10] Consequently, while the results of the augmented Dickey Fuller test can be reported for the combined data, the distributions of the test statistics are unknown and inferences about their values should be drawn with caution.

The results of the Augmented Dickey-Fuller tests with one lag are shown in Table 2.[11] The most important result concerns the post-1990 sub-period when the Reserve Bank was targeting low inflation. In this sub-period it is possible to reject the five separate hypotheses that each series has a unit root. The estimated coefficients for α_1 range between -0.50 and -0.32 for the upper trim and the headline and trimmed mean inflation rates, and are less than -0.90 for the lower trim and the CPITM gap. All of these coefficients are at least 3.98 standard deviations away from zero. In the very short 1984-1989 period it is not possible to reject the unit root hypothesis for any of the variables except the CPITM gap, but the sample length is so short and the standard errors of the coefficients are so large that the failure to reject the null hypothesis provides little insight into the true value of the coefficients. The autoregressive coefficients for the full sample are all closer to zero than for the 1990-2022 sub-period, but are all estimated to be less than -0.20, and they are all more than 3 standard deviations from zero.

I have chosen to focus on the tests for the post-1990 period since the distributions of the augmented Dickey Fuller test statistics are not precisely known when the means, variances, and autoregressive coefficients of the variables are not the same before and after 1990. For the post-1990 period the evidence suggests that all five series are mean reverting rather than unit root processes. Even so, it is useful to use an error correction model to describe the headline and trimmed mean inflation rates as the CPITM gap has much lower serial correlation than either of the two individual series. In these circumstances error correction models provide a useful lens to analyse the dynamics of the price series as they reveal the speed at which divergences between the headline and the trimmed mean inflation rates are reversed (Hendry, Pagan and Sargan 1984). They also show the extent that the gap between the trimmed mean and the headline inflation rates is closed because of subsequent changes in the headline inflation rate rather than the trimmed mean inflation rate.

[10] See Perron (2006) for a description of the well-known difficulties of distinguishing between unit roots and stationary series with structural breaks.
[11] The results for the 4-lag tests are similar.

Table 2: Unit root tests for the main series

	Augmented Dickey Fuller test (1 lagged difference term)		
	α_1	se	t-test/ ADF test
	1984:1-1989:4		
Headline inflation	-0.36	0.096	-1.80
Trimmed Mean	-0.25	0.16	-1.60
CPITM	-2.09	0.36	-5.76**
Lower Trim	-0.69	0.28	-2.50
Upper trim	-0.44	0.22	-1.97
	1990:1-2022:4		
Headline inflation	-0.21	0.058	-3.62**
Trimmed Mean	-0.25	0.068	-3.70**
CPITM Gap	-0.88	0.11	-8.00**
Lower Trim	-0.65	0.097	-6.72**
Upper Trim	-0.25	0.069	-4.10**
	1984:1-2022:4		
Headline inflation	-0.50	0.097	-5.13**
Trimmed Mean	-0.32	0.080	-3.98**
CPITM	-0.90	0.118	-7.67**
Lower Trim	-0.94	0.118	-8.00**
Upper Trim	-0.43	0.098	-4.44**

Source: Author's calculations using Statistics New Zealand price data.
A '*' indicates the hypothesis of equality can be rejected at the 5% significance level; '**' means the hypothesis can be rejected at 1% significance level.

3. Key empirical features of the trimmed mean inflation decomposition, 1984 – 2022.

The trimmed mean estimator of the core inflation rate provides several insights into the distribution of sectoral price changes and the fluctuations in the headline inflation rate that are not obtained from other measures. This section explores four of these insights. In section 3.1, the relationship between the upper and lower trims and the trimmed mean inflation rate is estimated. These regressions indicate prices in the upper trim increase faster than the rest of the distribution as the trimmed mean inflation rate increases, which causes the headline inflation rate to increase faster than the trimmed mean inflation rate as the inflation rate increases. In section 3.2 error correction models are estimated to see whether a non-zero gap between the headline and trimmed mean inflation rates are subsequently reversed from the headline inflation rate or lead to changes in the trimmed mean inflation rate. Since the differences between the headline and trimmed mean inflation rate are typically fully reversed from the headline rate after a quarter, but have little effect on the subsequent trimmed mean inflation rate, this provides a means of identifying transitory and persistent components in the inflation rate. Section 3.3 evaluates the accuracy of the different inflation forecasting models that incorporate the trimmed mean inflation rate, or the median inflation rate, as well as the headline inflation rate. Lastly, Section 3.4 describes the sectors that are typically found in the upper and lower trims.

3.1 The relationship between the upper and lower trims and the trimmed mean inflation rate

Figures 4 and 5 show scatterplots linking the upper and lower trims (TM_H^* and TM_L^*) and the trimmed mean inflation rate (TM^*). The graphs distinguish the observations from 1984-1989 (in red squares) and those from 1990 – 2020 (in blue circles); the observations for 2021 and 2022 are highlighted separately. The figures show the size of the upper and lower trims increase in a linear fashion as the trimmed mean inflation rate increases.

Estimates of the relationships are shown in Tables 3a (the upper trim) and 3b (the lower trim). The regressions are all linear in the trimmed mean inflation rate, although quadratic specifications were tried and were rejected.[12] Three sets of regressions are estimated in each case, differing by lag structure and the sample period. The first two sets of regressions are variants of

$$TM_{Ht}^* = \alpha_{0H} + \alpha_{1H}\ TM_t^* + \alpha_{2H}\ TM_{Ht-1}^* + e_{Ht}$$

$$TM_{Lt}^* = \alpha_{0L} + \alpha_{1L}\ TM_t^* + \alpha_{2L}\ TM_{Lt-1}^* + e_{Lt} \qquad (7)$$

[12] Quadratic terms were added to various specifications of the regressions. The estimated coefficients were small, never statistically significant at the five percent significance level, and variable in sign depending on the specification and data period.

Table 3a: Regression coefficients, upper trim (TM^*_H) against trimmed mean inflation and lags

Upper Trim	1984-2022	1990-2022	1984-2022	1984-2022	1990-2022	1984-2022	1984-2022	1990-2022
Constant	0.029**	0.027**	0.027**	0.022**	0.022**	0.022**	0.020**	0.021**
	(0.0011)	(0.0015)	(0.0014)	(0.0027)	(0.0025)	(0.0025)	(0.0035)	(0.0031)
TM^*_t	1.87**	2.11**	2.12**	1.52**	1.88**	1.88**	1.55**	1.90**
	(0.14)	(0.30)	(0.30)	(0.13)	(0.30)	(0.25)	(0.14)	(0.29)
$TM^*_{H,\ t-1}$				0.22**	0.16*	0.16*	0.23*	0.16*
				(0.071)	(0.064)	(0.064)	(0.072)	(0.064)
$TM^*_{L,\ t-1}$							-0.061	-0.047
							(0.070)	(0.069)
Constant			0.0094*			0.035		
[1984-90]			(0.0047)			(0.082)		
TM^*_t			-0.58			-0.67		
[1984-90]			(0.36)			(0.34)		
TM^*_H (t-1)						0.11		
[1984-90]						(0.16)		
n	156	132	156	155	132	155	155	132
Durbin-h	2.49*	1.91	2.25*	-0.98	-0.04	-0.36	-1.08	-0.10
R^2	0.61	0.36	0.62	0.64	0.37	0.65	0.64	0.38
Chow-test[a]			2.08			1.20		

Source: Author's calculations using Statistics New Zealand price data.
[a] "Chow test" is the test that the coefficients in the periods 1984-1989 and 1990-2022 are identical.
A * indicates the hypothesis of equality can be rejected at the 5% significance level; ** means the hypothesis can be rejected at 1% significance level.

Table 3b: Regression coefficients, lower trim (TM^*_L) against trimmed mean inflation and lags

Lower Trim	1984-2022	1990-2022	1984-2022	1984-2022	1990-2022	1984-2022	1984-2022	1990-2022
Constant	-0.029**	-0.030**	-0.029**	-0.027**	-.028**	-0.028**	-0.020**	-0.022**
	(0.0011)	(0.0014)	(0.0014)	(0.0024)	(0.0025)	(0.0025)	(0.0034)	(0.0045)
TM^*_t	0.82**	0.90**	0.90**	0.79**	0.87**	0.87**	1.10**	1.10**
	(0.099)	(0.24)	(0.24)	(0.11)	(0.23)	(0.23)	(0.16)	(0.24)
$TM^*_{H,\ t-1}$							-0.21*	-0.16
							(0.084)	(0.10)
$TM^*_{L,\ t-1}$				0.056	0.066	0.066	0.089	0.087
				(0.073)	(0.077)	(0.077)	(0.069)	(0.080)
Constant			0.0037			-0.0005		
[1984-90]			(0.0046)			(0.0080		
TM^*_t			-0.21			-0.084		
[1984-90]			(0.30)			(0.033)		
TM^*_L (t-1)						-0.056		
[1984-90]						(0.23)		
n	156	132	156	155	132	155	155	132
Durbin-h	1.31	1.12	1.28	2.29*	2.21*	2.35*	-1.52	-0.45
R^2	0.21	0.07	0.21	0.22	0.07	0.22	0.26	0.09
Chow-test[a]			0.26			0.06		

Source: Author's calculations using Statistics New Zealand price data.
[a] "Chow test" is the test that the coefficients in the periods 1984-1989 and 1990-2022 are identical.
A * indicates the hypothesis of equality can be rejected at the 5% significance level; ** means the hypothesis can be rejected at 1% significance level.

Figure 4: Scatterplot of the upper trim (TM^*_H) versus the trimmed mean inflation rate, 1984-2022.

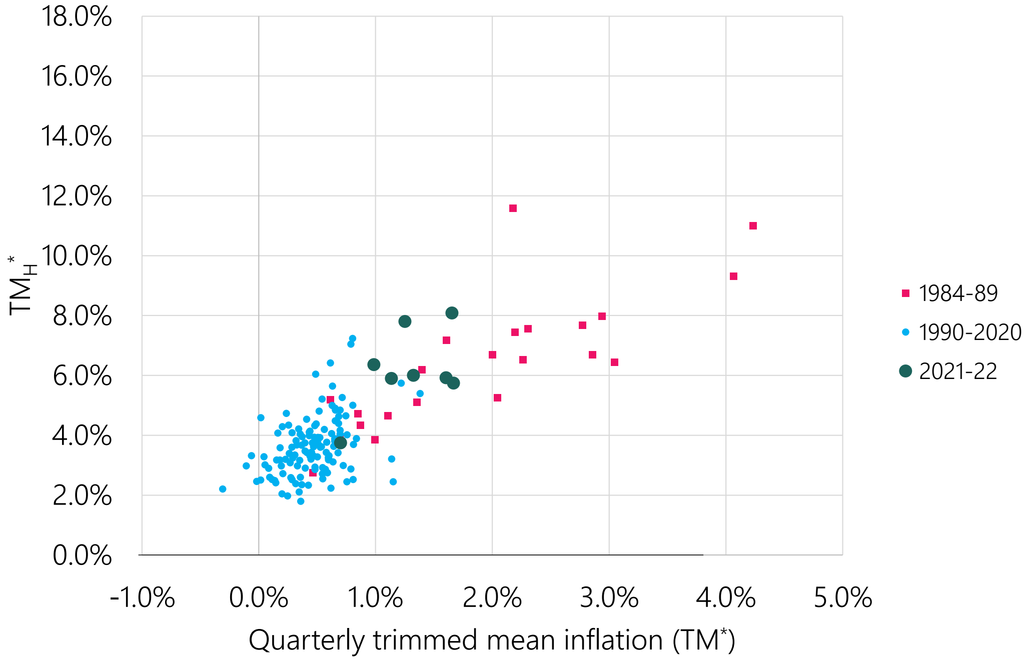

Source: Author's calculations using Statistics New Zealand data.

Figure 5 Scatterplot of the lower trim (TM^*_L) versus the trimmed mean inflation rate, 1984-2022.

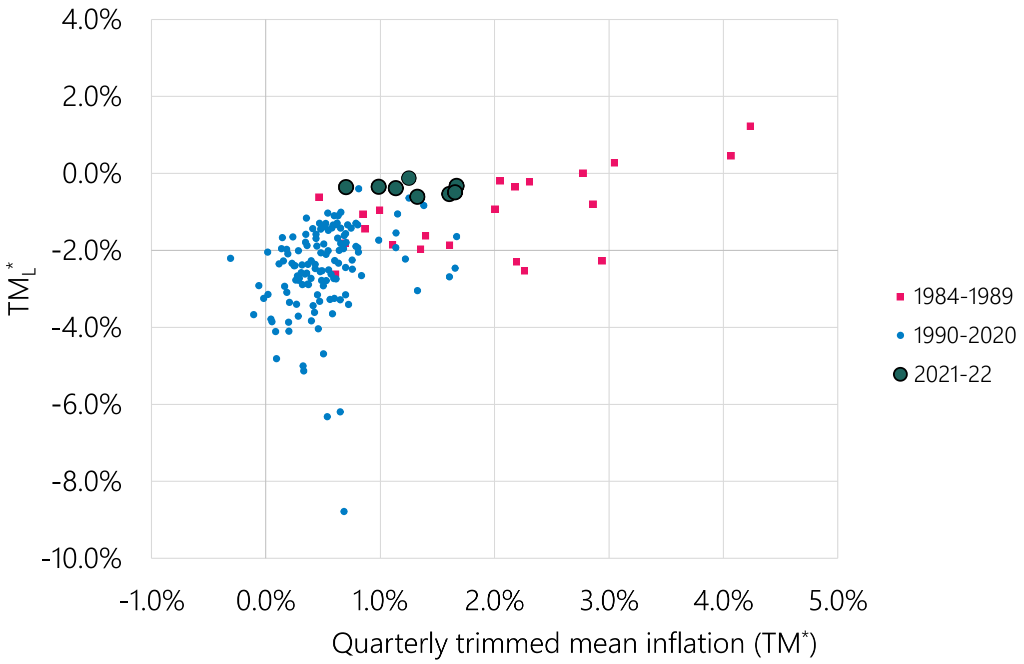

Source: Author's calculations using Statistics New Zealand data.

The first variant of each equation restricts the coefficient of the lagged dependent variable to equal zero (i.e. excludes the lagged dependent variable), while the second variant of each equation freely estimates all three coefficients. The four equations are estimated over the full period 1984 – 2022 and the low inflation period 1990 – 2022. A third version of each of the four equations is estimated over the full period but with separate coefficients for the two sub-periods. The reported standard errors are calculated using the Huber-White method to account for the potentially different residual variances in the two sub-periods.

For the whole period the preferred equation for the upper trim indicates that a 1 percentage point increase in the trimmed mean inflation rate is associated with a 1.52 percentage point contemporaneous increase in the upper trim inflation rate (with a 95% confidence interval of 1.25-1.75 percentage points),[13] and a total increase of 1.95 percentage points once lagged effects are incorporated.[14] This means that when the trimmed mean inflation rate increases by 1 percentage point the average price increase of the upper trim increases by 1.95 percentage points. For the post-1990 estimates, the estimates are somewhat larger, 1.88 for the contemporaneous effect and 2.25 for the total effect, but a Chow test does not reject the hypothesis that the coefficients for the two sub-periods are the same.

The increase in the lower trim inflation rate that accompanies a 1 percentage increase in the trimmed mean inflation rate is considerably smaller. For the whole period, the point estimate is 0.82 with a 95% confidence interval of 0.6 to 1.0 (column 1). Inspection of the data suggests that there are large numbers of near-zero price changes irrespective of the inflation rate. These near-zero price changes are consistent with menu cost models of price adjustment and seem to be the reason why average prices in the lower trim increase at a lesser rate as the trimmed mean inflation rate increases.[15] The point estimate is slightly larger in the post 1990 period, but with a Chow test it is not possible to reject the hypothesis that the coefficients in the two sub-periods are the same. It is noticeable that the explanatory power of the lower trim equation (measured by the R^2) is much smaller than that of the upper trim equation.

Together, these estimates suggest that the gap between the headline and trimmed mean inflation rates increases on average as the inflation rate increases, for the sum of the lower and upper trim coefficients multiplied by the weights $\boldsymbol{\alpha_L}$ and $\boldsymbol{\alpha_H}$ exceeds 0.4.[16] In Table 4 this contention is examined directly by estimating the equation

$$\boldsymbol{CPITM_t = \alpha_0 + \alpha_1 \; TM_t^* + e_t} \qquad (8)$$

[13] Column 4 of table 3a. It is preferred over column 1 as it has less serial correlation in the error term, and it is preferred over column 6 as the hypothesis that the coefficients in equation 6 that represent the two different periods are equal cannot be rejected.

[14] Calculated as 1.52/(1 – 0.22) = 1.95.

[15] This issue is explored at length in section 5.

[16] Recall that $CPITM_t = TM_{L,t} + TM_{H,t} - (\alpha_L + \alpha_H)TM_t^* = \alpha_L(TM_{L,t}^* - TM_t^*) + \alpha_H(TM_{H,t}^* - TM_t^*)$.

The estimates of α_1 are similar whether the equation is estimated by ordinary least squares or instrumental variables, using TM^*_{t-1} as an instrument for TM^*_t.[17] In both cases the estimates for α_1 are 0.14, with a 95% confidence interval from 0.06 to 0.23.

Since $\alpha_{1H} > \alpha_{1L}$ in Equation 7, the variance of sectoral inflation rates should also be increasing in the trimmed mean inflation rate. This result is confirmed directly in Table 4 (columns 3 – 5): all specifications indicate a statistically significant relationship, with the variance increasing by approximately 0.02 for each 1 percentage point increase in the inflation rate. The result that the variance of sectoral inflation rates is increasing in the inflation rate is consistent with a large number of studies that have shown that the mean and variance of sectoral inflation rates are positively correlated.[18] The result suggests that sector specific prices become more flexible as the average inflation rate increases, although as discussed further in Section 4 this result only applies when the inflation rate is less than 3%. If upward price flexibility increases with the inflation rate, high inflation may beget high inflation because an increasingly large number of firms are prepared to increase prices by large amounts when the inflation rate increases (Berger and Vavra 2018).

The more rapid increase of the upper trim than the lower trim as the trimmed mean inflation rate increases also suggests that the skewness of the price-change distribution should increase with the inflation rate, a result found in the U.S. (Rich, Verrugge and Zaman 2022). The regression results are consistent with this story, but the story is subtle. The estimates indicate that the skewness of the price-change distribution is highly and positively correlated with the headline inflation rate: a one percentage point increase in the inflation rate increases the skewness by 0.85 (Table 4: column 7). However, when the headline inflation rate is split into the trimmed mean inflation rate and the CPITM gap, the estimated coefficients are completely different, and an F-test of the hypothesis that the coefficients are equal is rejected at the 1 percent significance level (Table 4: column 8). Only the coefficient on the CPITM gap term is statistically significant, and it is large. In contrast, when the skewness measure is regressed against the trimmed mean inflation rate by itself, the coefficient is small (0.37) and only just statistically significant (Table 4: column 6). These results indicate that the skewness of the price-change distribution increases with the headline inflation rate precisely because there is a much faster increase in the upper trim than either the lower trim or the trimmed mean inflation rate as the inflation rate increases. The upper tail of the price change distribution becomes increasingly large as inflation increases, as firms in a small number of sectors increase prices by very large amounts.

[17] In the shorter 1990-2022 period the standard error of the IV estimate is significantly larger and while the coefficient is similar to the full period coefficient the 95% confidence interval includes zero.

[18] See Vining and Elsterwick (1976), Fischer (1981), and Bomberger and Makinen (1993) for evidence and an early review of the literature; Ball and Mankiw (1994, 1995) for theoretical models and empirical evidence that emphasise menu cost models; and Nakamura and Steinsson (2008) and Nakamura, Steinsson Sun and Villar (2018) for evidence based on firm-level rather than cross-sectional data.

Table 4: Regressions estimating the CPITM Gap, and the cross-sector variance and skewness, as a function of the trimmed mean and headline inflation rates, 1984:1 – 2022:4.

	Gap_t		$VAR(\pi_{it})$			$Skewness(\pi_{it})$		
	1.(OLS)	2.(IV)	3.	4.	5.	6.	7.	8.
constant	-0.0002	-0.0002	0.00054**	0.00056**	0.00054**	0.23	-0.18	0.33*
	(0.0003)	(0.0004)	(0.00008)	(0.00009)	(0.00008)	(0.21)	(0.21)	(0.17)
TM_t^*	0.14**	0.14**	0.021**		0.023**	37.0*		-23.0
	(0.038)	(0.043)	(0.0075)		(0.0075)	(17.7)		(14.2)
π_t				0.015			85.3**	
				(0.0081)			(19.1)	
Gap_t					-0.0092			415**
					(0.030)			(41)
$VAR(\pi_{it-1})$			0.22**	0.24**	0.22**			
			(0.083)	(0.085)	(0.083)			
N	155	155	155	155	155	155	155	155
R^2	0.09	0.09	0.142	0.125	0.144	0.015	0.12	0.41
Durbin-h	0.43	0.43	-1.90	-2.30*	-1.74	-0.36	0.31	0.09

Source: Author's calculations using Statistics New Zealand price data.
A * indicates the hypothesis of equality can be rejected at the 5% significance level; ** means the hypothesis can be rejected at 1% significance level.

The third set of equations in Table 3a and Table 3b allows for the possibility that prices in the upper (lower) trim depend on the size of the price changes in the lower (upper) trim the previous quarter:

$$TM_{Jt}^* = \alpha_{0J} + \alpha_{1J}\ TM_t^* + \alpha_{2J}\ TM_{Ht-1}^* + \alpha_{3J}\ TM_{Lt-1}^* + e_{JHt} \qquad J = Upper, Lower \qquad (9)$$

The results of equation 9 are reported in columns 7 and 8. They provide no support for the contention that large price falls are systematically followed by large price increases in the subsequent quarter (Table 3a), and only weak evidence that large price increases are systematically followed by large price falls (Table 3b).

3.2 Error correction models linking price changes to the CPI-TM gap.

The regression results presented in Table 2 indicate that the CPITM gap has smaller serial correlation than either the headline or trimmed mean inflation rates. This means it is useful to rearrange the dynamic equations for the quarterly change in the headline and trimmed mean inflation rates into error correction equations of the type[19]:

$$\Delta\pi_t = \delta_0 + \delta_1\ CPITM_{t-1} + \delta_2\pi_{t-1} + \delta_3\Delta\pi_{t-1} + \delta_4\Delta TM_{t-1}^* + e_{1t} \qquad (10a)$$
$$\Delta TM_t^* = \gamma_0 + \gamma_1 CPITM_{t-1} + \gamma_2 TM_{t-1}^* + \gamma_3\Delta\pi_{t-1} + \gamma_4\Delta TM_{t-1}^* + e_{2t} \qquad (10b)$$

[19] See the discussion in section 2.3, especially footnote 7, which follows the theoretical approach derived in Hendry, Pagan and Sargan (1984).

The first four columns of tables 5a and 5b present the estimates of four variations of equations 10a and 10b, differing as to whether they are univariate, include the lagged CPITM gap variable, or include the lagged gap variable and additional lagged difference terms. There are two additional variations. In column 5, the term $\alpha_H(TM^*_{H,t-1} - TM^*_{t-1})$ is included along with the CPITM gap term. This makes it possible to test whether the headline or trimmed mean inflation rates have different responses to unusually high price changes in the upper trim, or unusually low price changes in the lower trim.[20] In column 6, the CPITM gap is replaced by the difference between the headline inflation rate and median sectoral inflation rate (the 'median gap'), to see whether similar results are achieved when a different measure of the core inflation rate is used. These additional regressions are:

$$\Delta\pi_t = \delta_0 + \delta_1\, Median\ Gap_{t-1} + \delta_2\Delta\pi_{t-1} + \delta_3\Delta Median_{t-1} + e_{1t} \qquad (11a)$$

$$\Delta Median_t = \gamma_0 + \gamma_1 Median\ Gap_{t-1} + \gamma_2\Delta\pi_{t-1} + \gamma_3\Delta Median_{t-1} + e_{2t} \qquad (11b)$$

where 'median' is the median inflation rate and 'median gap' is the difference between the headline and the median inflation rates. Equations 10 and 11 are estimated over the low inflation period 1990:1 – 2022:4. The results found when the equations are estimated over the whole period from 1984:1 – 2022:4 are qualitatively similar but have much higher standard errors.[21]

The results across the different models strongly suggest that the gap between the headline and trimmed mean inflation rates measures transitory fluctuations in the inflation rate, for the gap is fully reversed from the headline inflation rate within a quarter and has no effect on the trimmed mean inflation rate. In all cases where the change in the headline inflation was regressed against the lagged CPITM gap, the coefficient was close to and insignificantly different from -1. In contrast, in all cases where change in the trimmed mean inflation rate was regressed against the lagged CPITM gap the coefficient was close to and insignificantly different from 0. These results also held when the CPITM gap was replaced by the median gap in the regressions. When the terms π_{t-1} and TM^*_{t-1} are entered separately into the headline inflation equation (column 2 of Table 5a), they are not significantly different from -1 and +1 respectively, and an F-test of the hypothesis that the coefficients are equal and opposite cannot be rejected at the five percent significance level. This suggests an error correction model is an appropriate representation of the data.

These results are similar with those found using U.S. data (Gamber et al 2015). The results using the median sectoral inflation rate rather than the trimmed mean inflation rate (column 6) are similar, reinforcing the contention that variations in the gap are caused by transitory variations in the headline inflation rate rather than movements in core inflation. When the term $\alpha_H(TM^*_{H,t-1} - TM^*_{t-1})$ is added to the regressions, the coefficients on this term were near to and insignificantly different from zero in both the headline and the trimmed mean inflation equations. This suggests that the responses of both the headline and trimmed mean inflation rates to the gap were the same whether it was caused by an unusually large value of the upper or lower trims. A result of this type has not previously been found.

[20] $\pi_t - TM^*_t = \alpha_L(TM^*_{L,t} - TM^*_t) + \alpha_H(TM^*_{H,t} - TM^*_t)$.

[21] The estimates for the sub-period 1984:1 – 1989:4 have much higher residual errors, although they are largely consistent with those estimated for the period 1990:1 – 2022:4.

Table 5a: Error correction models for changes in headline inflation, 1990 - 2022.

Dependent Variable	$1. \Delta\pi_t$	$2. \Delta\pi_t$	$3. \Delta\pi_t$	$4. \Delta\pi_t$	$5. \Delta\pi_t$	$6. \Delta\pi_t$
Constant	0.003**	0.002*	0.0005	0.0005	-0.0004	0.0005
	(0.0006)	(0.0007)	(0.0004)	(0.0004)	(0.002)	(0.0004)
π_{t-1}	-0.47**	-0.94**				
	(0.12)	(0.11)				
TM^*_{t-1}		0.72**				
		(0.20)				
$\Delta\pi_{t-1}$	-0.18*			0.009	-0.12	-0.04
	(0.088)			(0.14)	(0.08)	(0.12)
ΔTM^*_{t-1}				-0.32		
				(0.26)		
$CPITM_{t-1}$			-0.99**	-0.96**	-0.92**	
			(0.11)	(0.17)	(0.17)	
$\alpha_h(TM^*_{H,t-1} - TM^*_{t-1})$					0.13	
					(0.26)	
$MedianGAP_{t-1}$						-0.80**
						(0.15)
$\Delta Median_{t-1}$						-0.26
						(0.26)
N	132	132	132	132	132	132
Durbin-h	-0.95	-1.41	-2.0*	0.27	-1.44	-0.69
$\sum e^2$	0.00294	0.00275	0.00281	0.00273	0.00277	0.00279
R^2	0.31	0.36	0.34	0.36	0.35	0.35

Table 5b: Error correction models for changes in trimmed mean or median sectoral inflation rates, 1990 - 2022.

Dependent Variable	$1. \Delta TM^*_t$	$2. \Delta TM^*_t$	$3\Delta TM^*_t$	$4\Delta TM^*_t$	$5\Delta TM^*_t$	$6. \Delta Med_t$
Constant	0.001**	0.002*	0.0001	0.0001	-0.0002	0.0005
	(0.0004)	(0.0004)	(0.0002)	(0.0002)	(0.001)	(0.0004)
π_{t-1}		-0.03				
		(0.07)				
TM^*_{t-1}	-0.22*	-0.28*				
	(0.087)	(0.11)				
$\Delta\pi_{t-1}$				-0.02		-0.02
				(0.07)		(0.05)
ΔTM^*_{t-1}	-0.24*			-0.34**	-0.36**	
	(0.095)			(0.12)	(0.08)	
$CPITM_{t-1}$			-0.11	-0.04	-0.08	
			(0.07)	(0.09)	(0.10)	
$\alpha_h(TM^*_{H,t-1} - TM^*_{t-1})$					0.05	
					(0.17)	
$MedianGAP_{t-1}$						-0.03
						(0.05)
$\Delta Median_{t-1}$						-0.33**
						(0.11)
N	132	132	132	132	132	132
Durbin-h	-0.69	-2.81**	-4.45**	-1.88	-1.71	-2.49*
$\sum e^2$	0.00076	0.00082	0.00093	0.00081	0.0081	0.00068
R^2	0.20	0.15	0.02	0.14	0.14	0.12

Source: Author's calculations using Statistics New Zealand price data.
A * indicates the hypothesis of equality can be rejected at the 5% significance level; ** means the hypothesis can be rejected at 1% significance level.

Since these regressions indicate that the gap between the headline and trimmed mean inflation rates is transitory, the variations of the upper and lower trim around their expected values are also transitory. This provides a method of splitting the inflation rate into transitory and persistent components that does not rely on historical data, in contrast to the contention of Cukierman (1982, 2019). Consider the case when the coefficient on the gap variable is equal to $-1.$ The system can be written as

$$\pi_t = \alpha_0 + (1+\alpha_1)TM_t^* + e_{1t} \quad (12a)$$

$$TM_t^* = \beta_0 + (1-\beta_1)TM_{t-1}^* + e_{2t} \quad (12b)$$

$$\Delta\pi_t = \gamma_0 - (\pi_{t-1} - TM_{t-1}^*) + \gamma_1 TM_{t-1}^* + e_{3t} \quad (13)$$

where $\gamma_1 = (1+\alpha_1)\,(1-\beta_1) - 1$ and $e_{3t} = (1+\alpha_1)e_{2t} + e_{1t}$

In this case the persistent component of the headline rate is $\alpha_0 + (1+\alpha_1)TM_t^*$ (equation 12a). The fluctuations of the gap around the level $\alpha_0 + \alpha_1 TM_t^*$ are fully reversed from the headline inflation rate within a quarter, and thus are transitory. Note that the trimmed mean inflation rate is a biased estimator of the headline inflation rate, by a factor of $1/(1+\alpha_1)$; since the estimate of $\alpha_1 = 0.14$ (from Table 2, column 1), the trimmed mean is biased downwards by about 12%. This downward bias reflects that the trimmed mean inflation rate does not capture the extent that the upper trim inflation rate increases as inflation increases, for, as was shown in section 3.1, the headline rate is increasingly dominated by the price increases in the subset of sectors that are in the upper trim.[22] [23]

The results in Table 5b indicate that changes in the trimmed mean inflation rate can best be modelled using an autoregressive model with one or two lags and that excludes the headline inflation rate. Across the different models the coefficient on the first lag of TM_t^* is consistently although imprecisely estimated to be near -0.35, and the 95% confidence intervals of each estimate lie between -0.10 and -0.58. These results are consistent with those discussed in section 2 and suggest the trimmed mean has been a persistent but mean reverting series since 1990. Since the results in Table 1 indicate that the trimmed mean inflation rate has higher serial correlation and lower variance than the headline inflation rate, it is likely that the trimmed mean inflation rate can be used to improve forecasts of the next quarter's headline inflation rate, even though it is a biased estimator of this rate. This topic is investigated further in the next section.

[22] As firm level data are not available, it is impossible to tell whether the large price increases in the upper trim reflect very large increases made by firms once during the quarter or a series of small changes made sequentially in response to persistent cost shocks. U.S. studies indicate that when inflation increases the absolute size of price changes only changes by little, suggesting firms might be making several smaller increases in a row (see Nakamura et al 2018). Note that the bias in the median will depend on the frequency of the CPI observations. As the length of the observation period becomes smaller, the fraction of the firms not changing will increase, and the median price change will tend towards zero. In this case a larger fraction of the mean inflation rate will be due to the price changes of the few firms changing prices within the period, and many of these changes will be located in the upper and lower quantiles of the distribution. Were monthly data available, as it is in most countries, we could test how the bias changes with the frequency of the observations.

[21] Evidence from Argentina and Mexico suggests that when the inflation rate gets very high, most firms increase their prices by a similar amount because the common macroeconomic inflationary shock dominates the idiosyncratic shocks they receive. New Zealand inflation has not been high enough for this effect to be observed in the sectoral data. See Gagnon (2009) and Alvarez et al (2019).

3.3 Headline inflation forecasts

Since the quarterly change in the inflation rate is systematically related to the lagged CPITM gap, and the lagged median inflation gap, these gaps should be able to improve headline inflation forecasts. In this section this conjecture is evaluated by generating recursive time series forecasts based on vector autoregression and vector error correction models that incorporate combinations of the headline inflation rate, the trimmed mean inflation rate, the CPITM gap, the median gap, and the upper trim. The forecasts made from these systems are compared with recursive univariate forecasts that only utilise the headline inflation rate.

The basic procedure is as follows. Each of the eight models described below was first estimated over the period 1990:1 to 2012:4. The models are based on the models in section 3.2, differing only by the exclusion of statistically insignificant lagged variables. The models were selected by (i) estimating the model with up to n lags of the dependent variables, and then, if necessary, by (ii) re-estimating a smaller model that excluded lags whose coefficients were not significant at a 10% statistical significance level. Each model was used to generate recursive forecasts of the quarterly headline inflation rate for 12 quarters, and these forecasts were used to calculate a sequence of quarterly forecast errors. These errors were recorded.[24] The regression sample was then extended by one quarter, new coefficients were estimated, and these new coefficients were used to generate the forecasts for the next quarter.[25] The process was repeated until a set of 40 quarterly sets of forecasts were generated. The last forecast was made using data for the period 1990:1 to 2022:3.[26]

Eight models were evaluated. The first three forecast models are "benchmarks" and only include lags of the headline inflation rate. The equations were estimated in a differenced form, and non-significant coefficients were eliminated.

Model 1 is a random walk:

$$\Delta\pi_t = \theta_0 + e_{1t} \qquad (14.1)$$

Model 2 includes the first two lags of headline inflation:

$$\Delta\pi_t = \theta_0 + \theta_1\pi_{t-1} + \theta_2\Delta\pi_{t-1} + e_{1t} \qquad (14.2)$$

Model 3 includes up to five lags of the headline inflation rate; the final form was:

$$\Delta\pi_t = \theta_0 + \theta_1\pi_{t-1} + \theta_2\Delta\pi_{t-1} + \theta_3\Delta\pi_{t-4} + e_{1t} \qquad (14.3)$$

The next three models include the trimmed mean inflation rate or the CPITM gap as the second variable in a vector autoregression. As non-significant coefficients were suppressed, model 4 is:

$$\Delta\pi_t = \delta_0 + \delta_1(CPITM_{t-1}) + \delta_2 TM^*_{t-1} + e_{1t} \qquad (14.4a)$$

$$\Delta TM^*_t = \gamma_0 + \gamma_1 TM^*_{t-1} + e_{2t} \qquad (14.4b)$$

Model 5 is:

$$\Delta\pi_t = \delta_0 + \delta_1(CPITM_{t-1}) + e_{1t} \qquad (14.5a)$$

[24] If the forecasting model is a vector autoregression, forecasts of the other variable are also generated, but although they are used to generate the longer horizon forecasts of the headline inflation rate the accuracy of the forecasts of the additional variables has not been evaluated.

[25] The specification used for the initial period was used in all subsequent estimations, even if the coefficients were no longer significant at the 10% level.

[26] Note that the RMSE statistics for the 1 quarter ahead forecasts are calculated using 40 quarterly forecasts whereas the RMSE statistics for the 12 quarter ahead forecast errors are calculated using only 29 forecasts.

$$CPITM_t = \gamma_0 + \gamma_1 CPITM_{t-1} + e_{2t} \quad (14.5b)$$

Model 6 uses a vector autoregression to forecast the trimmed mean inflation rate and the CPITM gap, and constructs the headline inflation rate forecasts from these variables:

$$\Delta TM_t^* = \delta_0 + \delta_1(CPITM_{t-1}) + \delta_2 \Delta TM_{t-1}^* + e_{1t} \quad (14.6a)$$

$$CPITM_t = \gamma_0 + \gamma_1 CPITM_{t-1} + e_{2t} \quad (14.6b)$$

$$\pi_t = TM_t^* + CPITM_t \quad (14.6c)$$

Model 7 includes the upper trim as well as the CPITM gap, and thus has three equations plus identities:

$$\Delta\pi_t = \delta_0 + \delta_1 CPITM_{t-1} + \delta_2 \Delta\pi_{t-1} + \delta_3(TM_{H,t-1}^+) + e_{1t} \quad (14.7a)$$

$$CPITM_t = \pi_0 + \pi_1 CPITM_{t-1} + \pi_2 TM_{H,t-1}^+ + e_{2t} \quad (14.7b)$$

$$TM_{H,t}^+ = \rho_0 + \rho_1 CPITM_{t-1} + \rho_2 TM_{H,t-1}^+ + e_{3t} \quad (14.7c)$$

$$TM_{H,t}^+ = 0.2(TM_{H,t-1}^* - TM_t^*) \quad (14.7d)$$

Model 8 uses the median sectoral inflation rate to calculate the gap measure, rather than the trimmed mean inflation rate. The headline inflation forecast is calculated from the forecasts of the median sectoral inflation rate and the median gap. It is similar to model 6:

$$\Delta Median_t = \delta_0 + \delta_1(Median\ gap_{\ t-1}) + \delta_2 \Delta Median_{t-1} + e_{1t} \quad (14.8a)$$

$$Median\ gap_t = \gamma_0 + \gamma_1 Median\ gap_{t-1} + e_{2t} \quad (14.8b)$$

$$\pi_t = Median_t + Median\ Gap_t \quad (14.8c)$$

Tables 6a and 6b show the summary statistics for the eight forecasting models described above. Table 6a shows the root mean square errors (RMSE) for the forecasts of the quarterly inflation rate π_{t+n} made at time t = 2012:4 to 2022:3 for each horizon $n = 1,2,\dots,12$ for the eight models. Table 6b shows the root mean square errors for the forecasts of the accumulated price level change between t and t+n for the eight models.[27] The results are qualitatively similar to the results for the period 2012:4 to 2019:4 that excludes the Covid period.

Tables 7a and 7b shows the results of the Diebold-Mariano (DM) tests for selected pairs of forecasts (Diebold and Mariano 1995). The tables compare the best two vector autoregression/ vector error correction models (equations 14.5 and 14.6) and the median vector error correction model (equation 14.8) with the headline inflation random walk model (model 14.1) and the headline inflation autoregressive model (equation 14.3). The forecasts of the second model of each pair are superior to those from the first if the Diebold-Mariano statistic is large and positive. The tables show the Diebold-Mariano statistics for each forecast horizon and indicate those statistics that are statistically significant at the 5% and 1% significance levels. Table 7a shows the results for the n-quarter ahead quarterly inflation forecasts, and Table 7b shows the results for the accumulated price change over n quarters.

The results can be usefully split into three horizons: 1 quarter ahead; 2- 4 quarters ahead; and 5-12 quarters ahead.

[27] For example, for the 8 quarters horizon, Table 7a measures the average error made when forecasting the quarterly inflation rate π_{t+8} 8 quarters ahead. Table 7b measures the average forecast error of the accumulated change in the price level between t and t+8 i.e. the error in $\sum_{s=1}^{8} \pi_{t+s}$.

At the one quarter ahead horizon the best benchmark model was the autoregressive equation 14.3. The best multivariable models are slightly better than the benchmark univariate models, but the improvement is modest. A model using the CPITM gap, equation 14.5, was the best of all the models, but it only reduced the RMSEs by 9 percent relative to the best benchmark model (equation 14.3).

The improvement is not statistically significant, however. The forecast improvement is not particularly impressive given the in-sample fit of the vector error correction models estimated in section 3.2 showed that fluctuations in the CPITM gap were fully reversed from the headline inflation rate within a single quarter.[28]

Over the 2 – 4 quarter ahead horizon, the best multivariate models were equations 14.5 and 14.6, both of which use the CPITM gap. In terms of the quarterly inflation forecasts, neither of these models outperformed the best benchmark model, which at this horizon was the random walk model (equation 14.1), in a statistically significant way. However, in terms of the accumulated price difference, the vector error correction model 14.5 reduced forecast errors by 10 – 20% relative to the random walk, and this improvement was statistically significant. The improvement in the accumulated price change forecasts suggests the gap between the headline and trimmed mean inflation rate is not fully reversed out of the headline inflation rate within a quarter, but that it is more substantially reversed within a year.

At horizons between 5 and 12 quarters ahead, the multivariate model incorporating the median gap, equation 14.8, not the CPITM gap, was the best performing model.[29] The forecast improvement for individual quarters is modest at these horizons, but across the various horizons model 14.8 reduces the accumulated price change forecast errors by 10% - 21% relative to the best benchmark model (equation 14.3). In most cases these improvements are statistically significant at a 5% confidence level. The forecasts based on the median inflation gap also outperformed those based on the CPITM gap by a statistically significant amount. This is not surprising, since the reverse was true for accumulated price change forecasts over the first four quarters of the longer periods. A potential explanation is that the equations forecasting future values of both gap variables (for example equations 14.5b or 14.8b) had very little explanatory power, although they differed a little in terms of their estimated means. The accumulated effects of these differences in the forecast means over the longer forecast horizons appears to explain the outperformance of the forecasts based on the median gap.

[28] Note that Diebold (2015) argues that when models are used to generate forecasts, as they are in this case, the in-sample fit is a better measure of the performance of alternative models than the Diebold-Mariano statistic applied to pseudo out-of-sample forecasts, as the in-sample fit is based on the information from the whole period, not just a sample of it. Applying this rationale suggests that the full-sample regressions estimated in section 3.2 offer a better comparison of the different models than the out-of-sample forecasts discussed in this section.

[29] Note that the median cross-sector price change is the limit trimmed mean estimator when the trims are both 50%.

Table 6a: Root mean square forecast errors, quarterly headline inflation rate, 2012- 2022.

Single quarter errors, 2012- 2022

	Headline inflation rate only			Headline inflation, TM, and CPITM			Upper trim or median	
Quarter	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
1	0.087	0.083	0.080	0.079	0.073*	0.087	0.087	0.087
2	0.084	0.095	0.092	0.092	0.073*	0.075	0.078	0.077
3	0.101	0.104	0.100	0.102	0.091	0.088*	0.089	0.089
4	0.097	0.108	0.105	0.107	0.091*	0.091*	0.093	0.092
5	0.122	0.111	0.113	0.111	0.103	0.100	0.100	0.098*
6	0.113	0.115	0.115	0.114	0.102	0.103	0.104	0.099*
7	0.134	0.118	0.118	0.118	0.117	0.115	0.114	0.109*
8	0.133	0.122	0.121	0.121	0.121	0.119	0.118	0.114*
9	0.148	0.124	0.123	0.124	0.125	0.121	0.120	0.114*
10	0.143	0.125	0.125	0.125	0.121	0.123	0.122	0.117*
11	0.146	0.130	0.129	0.130	0.141	0.135	0.130	0.128*
12	0.132	0.134	0.134	0.134	0.134	0.133	0.131	0.127*

Source: Authors' calculations using Statistics New Zealand data.
A "*" indicates the column with the minimum RMSE for a particular horizon.
Model 1: Headline inflation random walk model.
Model 2. Headline inflation AR2 model.
Model 3. Headline inflation AR5 model.
Model 4: Vector Autoregression with headline inflation and trimmed mean inflation rate as dependent variables.
Model 5: Vector Autoregression with headline inflation and the CPITM gap as dependent variables.
Model 6: Vector Autoregression with the trimmed mean inflation rate and the CPITM gap as dependent variables.
Model 7: Vector autoregression with the headline inflation rate, CPITM gap and upper trim as dependent variables
Model 8: Vector autoregression with the headline inflation rate and the CPI-median gap as dependent variables.

Table 6b: Root mean square forecast errors, accumulated n-quarter headline inflation rate, 2012- 2022.

Accumulated n-quarter inflation forecast errors, 2012- 2022

	Headline inflation rate only			Headline inflation, TM, and CPITM			Upper trim or median	
Quarter	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
1	0.087	0.083	0.08	0.079	0.073*	0.087	0.087	0.087
2	0.103	0.114	0.112	0.108	0.082	0.080*	0.082	0.081
3	0.142	0.147	0.142	0.143	0.120	0.113*	0.115	0.114
4	0.155	0.170	0.166	0.167	0.134	0.129*	0.131	0.129*
5	0.214	0.192	0.195	0.190	0.167	0.163	0.162	0.160*
6	0.213	0.205	0.205	0.202	0.172	0.172	0.174	0.162*
7	0.268	0.219	0.219	0.217	0.210	0.207	0.205	0.192*
8	0.278	0.226	0.227	0.224	0.220	0.214	0.210	0.199*
9	0.320	0.233	0.233	0.231	0.230	0.221	0.218	0.199*
10	0.319	0.233	0.232	0.230	0.209	0.213	0.210	0.189*
11	0.310	0.241	0.240	0.240	0.262	0.242	0.229	0.216*
12	0.273	0.246	0.247	0.244	0.236	0.225	0.219	0.199*

Source: Authors' calculations using Statistics New Zealand data.
A "*" indicates the column with the minimum RMSE for a particular horizon.
Model 1: Headline inflation random walk model.
Model 2. Headline inflation AR2 model.
Model 3. Headline inflation AR5 model.
Model 4: Vector Autoregression with headline inflation and trimmed mean inflation rate as dependent variables.
Model 5: Vector Autoregression with headline inflation and the CPITM gap as dependent variables.
Model 6: Vector Autoregression with the trimmed mean inflation rate and the CPITM gap as dependent variables.
Model 7: Vector autoregression with the headline inflation rate, CPITM gap and upper trim as dependent variables
Model 8: Vector autoregression with the headline inflation rate and the CPI-median gap as dependent variables.

Table 7a: Diebold-Mariano test of significance in difference of quarterly headline inflation rate forecast errors, 2012- 2022.

Single quarter errors, 2012- 2022.

	Comparison with model 1			Comparison with model 3			Comparison with model 8	
Quarter	Models (1,5)	Models (1,6)	Models (1,8)	Models (3,5)	Models (3,6)	Models (3,8)	Models (5,8)	Models (6,8)
1	2.64*	1.22	2.32*	0.86	1.24	1.34	0.38	1.15
2	1.90	-0.23	1.05	1.83	1.77	1.82	-1.44	1.60
3	1.45	0.10	1.56	0.94	0.51	1.45	0.95	1.45
4	0.65	-0.55	0.47	1.71	-0.16	1.93	-0.34	1.88
5	1.61	0.91	1.73	1.11	1.60	1.81	1.61	1.72
6	1.25	-0.06	1.45	2.39*	1.75	2.85*	1.37	2.60*
7	1.80	1.57	2.12*	0.20	-0.26	2.06*	2.07	2.04*
8	1.30	1.14	1.73	0.01	-0.54	1.48	1.77	1.53
9	1.86	1.89	2.19*	-0.39	-0.56	1.78	2.55*	1.78
10	2.08*	1.61	2.02*	0.88	-0.64	1.54	0.89	1.54
11	0.78	1.91	2.73*	-2.41*	-1.12	0.13	3.42*	0.24
12	-0.21	-0.26	0.56	0.01	-0.78	1.41	1.50	1.55

Source: Authors' calculations using Statistics New Zealand data.
The DM statistic is positive if the forecasts of the second model has lower errors than the forecasts of the first model. The DM statistic has a t-distribution. A "*" means the DM statistic is significant and the 5% level; "**" means the statistic is significant at the 1% level.
Model 1: Headline inflation random walk model.
Model 3. Headline inflation AR5 model.
Model 5: Vector Autoregression with headline inflation and the CPITM gap as dependent variables.
Model 6: Vector Autoregression with the Trimmed mean inflation rate and the CPITM gap as dependent variables.
Model 8: Vector autoregression with the headline inflation rate and the CPI-median gap as dependent variables.

Table 7b: Diebold-Mariano test of significance in difference of accumulated n-quarter headline inflation rate forecast errors, 2012- 2022.

Accumulated n-quarter inflation forecast errors, 2012- 2022

	Comparison with model 1			Comparison with model 3			Comparison with model 8	
Quarter	Models (1,5)	Models (1,6)	Models (1,8)	Models (3,5)	Models (3,6)	Models (3,8)	Models (5,8)	Models (6,8)
1	2.645*	1.218	2.324*	0.860	1.236	1.340	0.377	1.155
2	2.299*	0.307	1.583	1.553	2.183*	1.708	-0.833	1.389
3	2.105*	0.303	1.690	1.297	2.125*	1.618	-0.107	1.425
4	1.672	0.179	1.401	1.410	3.362**	1.762	-0.053	1.513
5	1.940	0.725	1.835	1.287	3.037**	1.833	0.624	1.581
6	1.850	0.933	1.952	1.421	3.176**	2.101*	1.711	1.818
7	2.100*	1.752	2.308*	1.153	2.797**	2.369*	2.413*	2.068*
8	2.102*	1.983	2.271*	0.607	2.377*	2.695*	2.267*	2.401*
9	2.064*	1.972	2.242*	0.545	2.458*	2.724*	2.324*	2.409*
10	2.161*	1.999	2.241*	1.014	2.654*	2.650*	2.087*	2.326*
11	2.913**	2.282*	4.220**	0.624	1.989	2.560*	2.646*	2.263*
12	2.745*	2.745*	4.235**	0.192	1.838	2.654*	2.586*	2.555*

Source: Authors' calculations using Statistics New Zealand data.
The DM statistic is positive if the forecasts of the second model has lower errors than the forecasts of the first model. The DM statistic has a t-distribution. A "*" means the DM statistic is significant and the 5% level; "**" means the statistic is significant at the 1% level.
Model 1: Headline inflation random walk model.
Model 3. Headline inflation AR5 model.
Model 5: Vector Autoregression with headline inflation and the CPITM gap as dependent variables.
Model 6: Vector Autoregression with the Trimmed mean inflation rate and the CPITM gap as dependent variables.
Model 8: Vector autoregression with the headline inflation rate and the CPI-median gap as dependent variables.

Equation 14.7, which incorporated the upper trim as well as the CPITM gap, made only a small difference to the forecasts that only included the CPITM gap. There was a very small improvement relative to models 5 and 6 at horizons from 8 – 12 quarters, but this improvement was not statistically significant, and the longer horizon forecasts were not as good as those based on the median inflation gap. This model was never the best forecasting model over any horizon. Consistent with the results in section 3.2, splitting the gap into components associated with large sectoral price increases rather than large sectoral price decreases does not significantly improve the model fit even though, as demonstrated in the next section, there are clear differences in price dynamics at the upper and lower tails of the sectoral inflation rate distributions.

These results are consistent with the focus most central banks place on measures of core inflation rather than headline inflation to understand and forecast inflation. While there are many different measures of core inflation, since both the trimmed mean inflation rate and the median inflation rate are particularly easy to understand and communicate it is somewhat pleasing that inflation forecasts can be improved, even modestly, by such simple multivariate time series models.

3.4 Sectors frequently found in the upper and the lower trims

Figure 6 shows the distribution of the number of times different sectors are found in the upper and lower trims over the period 1999:2- 2022:4, weighted by each sector's average CPI weight.[30] For example, the graph shows that 25% (by weight) of the sectors appeared in the lower trim in 6 or fewer quarters (out of 95), while 25% appeared in the lower trim 34 or more times. The figure also includes a line ('Binomial') indicating the distribution that would be observed if each sector had an equal 20 percent probability of appearing in the upper or lower trim each quarter. The sectors with the smallest and largest number of appearances in the upper or lower trims are listed in Table 8.

Three observations can be made about the data. First, the probabilities of appearing in the upper and lower trims are not equally distributed across sectors. If the probabilities were equally distributed, there would be essentially zero chance of observing any sectors with fewer than 10 or more than 35 occurrences in either the upper or lower trim. Yet 20% of sectors by weight were either in the upper or lower trims (or both) more than 35 times, indicating that some sectors are particularly prone to experience very large price movements. In addition, 42% of sectors by weight were in the lower trim 10 or fewer times, and 25% were in the upper trim 10 or fewer times. These data indicate that a few sectors are responsible for a lot of the volatility in the headline inflation rate and the headline-trimmed mean inflation gap.

[30] The lists were compiled by counting the number of times each of the 109 sectors was in the upper trim or lower trim between 1999:Q2 and 2022:Q4 The analysis is only conducted for this period because the number of sectors and the definitions of the sectors were different prior to 1999 (see Appendix A). The weights are the average weights for each sector over the period.

Table 8: Sectors frequently appearing in the upper and lower trims.

	Least frequently found in the lower trim				Most frequently found in the lower trim		
	% Lower	% Upper	Weight		% Lower	% Upper	Weight
Restaurant meals	0%	2%	1.6	Purchase of bicycles	46%	32%	0.1
Ready-to-eat food	0%	5%	2.6	Men's clothing	47%	12%	0.9
Dental services	0%	19%	0.8	Major house appliances	48%	16%	0.9
Hairdressing services	0%	13%	0.7	Small electric appliances	48%	22%	0.1
Purchase of housing	1%	42%	6.5	Games, toys and hobbies	49%	17%	0.5
Vehicle service /repairs	1%	12%	1.1	Other personal effects	51%	15%	0.2
Veterinary services	1%	19%	0.3	Men's footwear	52%	21%	0.2
Property maintenance	2%	15%	1.8	Recording media	52%	20%	0.3
Hospital services	3%	20%	0.6	Computing equipment	78%	4%	0.5
Newspapers /magazines	3%	25%	0.7	Telecom. equipment	79%	12%	0.3
Vocational services	3%	13%	0.2	Audio-visual equipment	83%	14%	0.7

	Least frequently found in the upper trim				Most frequently found in the upper trim		
	% Lower	% Upper	Weight		% Lower	% Upper	Weight
Actual rentals/ housing	6%	0%	7.9	Accommodation services	34%	40%	0.9
Restaurant meals	0%	2%	1.6	Books	40%	41%	0.3
Telecom. services	44%	2%	2.7	Purchase of housing	1%	42%	6.5
Computing equipment	78%	4%	0.5	Real estate services	12%	44%	0.9
Ready-to-eat food	0%	5%	2.6	Fruit	42%	45%	1.0
Therapeutic appliances	6%	5%	0.3	Domestic air transport	39%	45%	0.7
Personal care appliances	27%	5%	1.4	International air transport	42%	45%	1.9
Other property services	18%	6%	0.0	Dwelling insurance	7%	45%	0.4
Overseas hotels etc	14%	6%	0.2	Vegetables	44%	47%	1.4
Wine	20%	7%	1.5	Vehicle fuels/lubricants	36%	52%	0.4
Women's footwear	32%	7%	0.4	Petrol	36%	53%	4.3

Source: Author's calculations using Statistics New Zealand price data.

Figure 6: Distribution of the number of quarters that sectors are found in the upper and lower trims, weighted by sector.

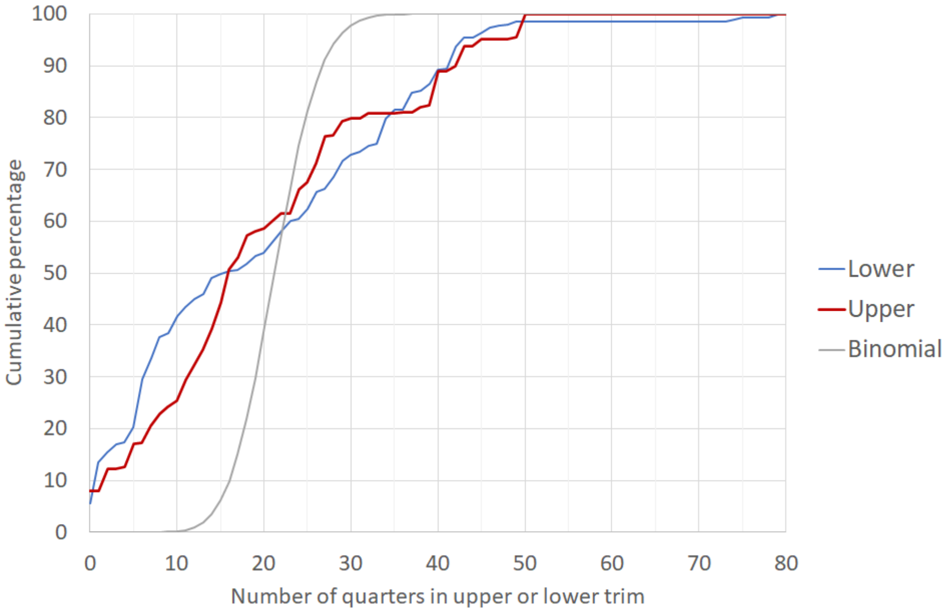

Source: Author's calculations using Statistics New Zealand data.
The 'Upper' line shows that 25% of sectors (by weight) appeared in the upper trim fewer than 10 times out of 95 quarters, and 10% of sectors appeared at least 40 times. The 'Binomial' line shows what the distribution of sectors would like if every sector had the same 20% probability of appearing in the upper and lower trims each quarter.

Secondly, the interquartile range of the number of times each sector was in the lower trim between 1999 and 2022 is greater than the interquartile range for the upper trim. The interquartile range of the lower trim was 6 – 34 occurrences; for the upper trim it was 10 – 27 occurrences. This means that the sectors that are frequently in the lower trim appear in it more often than the sectors that are frequently in the upper trim; correspondingly, there are more sectors that are rarely in the lower trim than rarely in the upper trim.[31] Sectors that are frequently found in the upper trim are primarily services, particularly services that are related to the construction, rental, sale and insurance of real estate. In contrast, the items that are most frequently found in the lower trim are manufactured goods, especially electronic goods.

Thirdly, there is a set of sectors that have volatile prices, and which frequently appear in the upper trim and the lower trim. The largest of these sectors are petrol and associated fuels, domestic and international airfares, and fruit and vegetables, which comprise 10% of the CPI basket in total. They also include two small sectors, books and bicycles, which may indicate sampling volatility measurement issues. The volatility of fuel and food prices is well known, which is why prices in these sectors are often excluded from measures of core inflation. The frequent oscillation of prices in these sectors is one of the reasons why the movements in the gap between the headline and trimmed mean inflation rates is quickly reversed out of the headline inflation rate by the subsequent quarter.

4. Price flexibility over the inflation cycle

It was shown in section 3.1 (Tables 3a, 3b and 4) that the upper trim, lower trim, and the variance of sectoral inflation rates are all increasing in the trimmed mean inflation rate. It was also shown that the regression coefficient between the upper trim and the trimmed mean inflation rate is over twice as large as the regression coefficient between the lower trim and the trimmed mean inflation rate. This result indicates that prices become more flexible upwards than downwards as the trimmed mean inflation rate increases, in the sense that the size of the largest upward sectoral price movements relative to the median price change increases much more rapidly than the size of the largest downwards sectoral price movements relative to the median price change. This section explores this result further.

The analysis is conducted in three stages. In the first stage the percentiles of each quarter's sectoral inflation rate distribution are estimated using the Harrell-Davis (1982) quantile estimator, with 2% intervals. The distributions are weighted using the contemporaneous CPI weights. This generates 156 separate estimates of the sectoral inflation rate distributions, each with 49 percentile

[31] It is also the case that the Gini coefficient of the number of times a sector in in the lower trim is higher than the Gini coefficient of the number of times a sector is in the upper trim.

Figure 7a: Distribution of sector-specific price changes by median sectoral inflation rate.

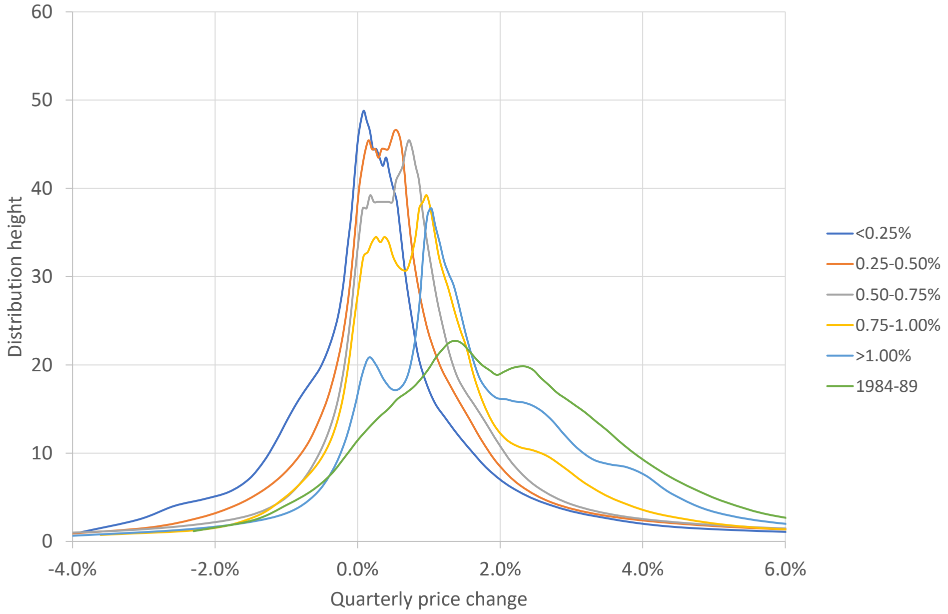

Source: Author's calculations using a Harrell-Davis estimator applied to Statistics New Zealand data.

Figure 7b: Distribution of sector-specific price changes relative to median, by median sectoral inflation rate.

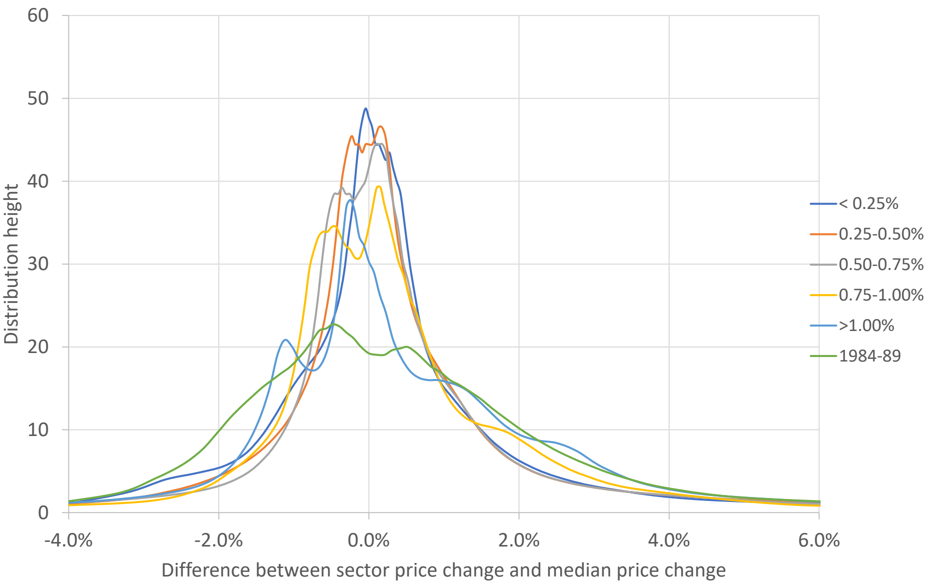

Source: Author's calculations using a Harrell-Davis estimator applied to Statistics New Zealand data.

Figure 8a: Average value of the 10-30 percentiles of the deviations of the sectoral inflation rate distribution from the median, sorted by the median sectoral inflation rate, 1984-2022.

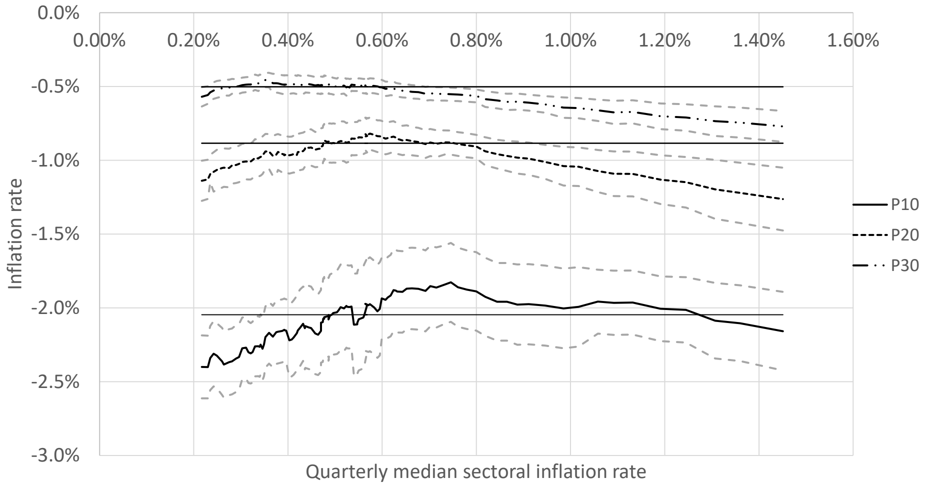

Source: Author's calculations using Statistics New Zealand data. The observations are smoothed using a 40-observation moving average applied to the data sorted by the median cross-sector price change. The dotted lines are two standard deviation bands.

Figure 8b: Average value of the 80-90 percentiles of the deviations of the sectoral inflation rate distribution from the median, sorted by the median sectoral inflation rate, 1984-2022.

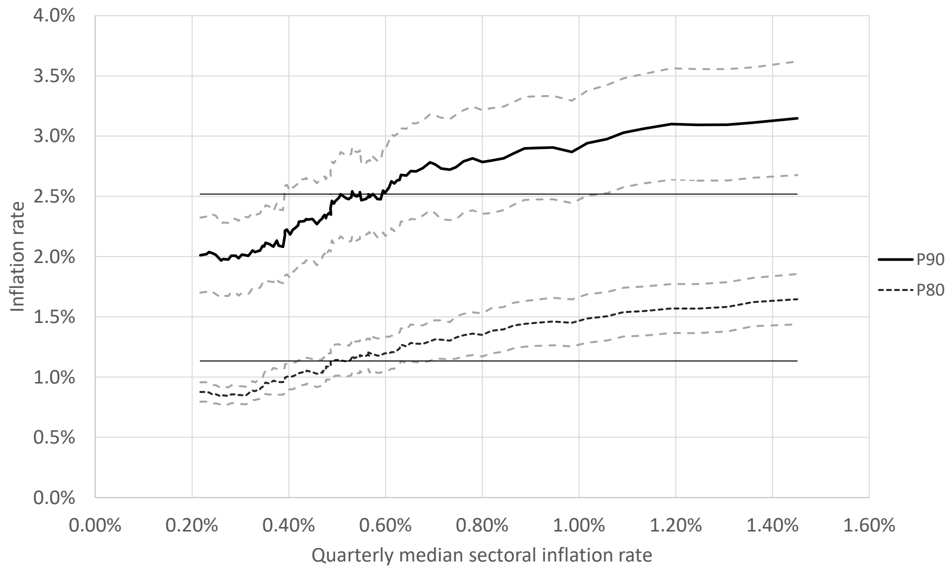

Source: Author's calculations using Statistics New Zealand data. The observations are smoothed using a 40-observation moving average applied to the data sorted by the median cross-sector price change. The dotted lines are two standard deviation bands.

observations (from 2% – 98%). In the second stage, these 156 distributions are sorted by the median sectoral inflation rate of each distribution and then divided into six groups. The average distribution for each of these six group is calculated. These are shown in Figure 7a. As discussed below, these data provide striking evidence that the distribution of price changes has a "spike" at zero that is consistent with standard menu costs price adjustment models. In the third stage the differences between the deciles and the median of each of the 156 quarterly distributions are calculated. These are measures of the "width" of the distribution. A series of regressions are estimated to determine how the "width" depends on the median sectoral inflation rate or the trimmed mean inflation rate. The regression results are shown in table 9 and illustrated for selected deciles in figures 8a and 8b.

Figure 7a shows the estimated sectoral inflation rate distributions grouped by the median sectoral inflation rate. The lines correspond to six inflation categories. The first five categories sort the data from 1990 – 2022 according to whether the quarterly median sectoral inflation rate was less than 0.25%, 0.26-0.50%, 0.51-0.75%, 0.76-1.00%, or more than 1%. The sixth category averages the distributions from the high inflation period 1984-1989, when the median quarterly inflation rate was 2.2%. Consider first the 1990 – 2022 categories. The distribution for the quarters where the median sectoral inflation rate was less than 0.25% has a single mode, at 0.1%. The remaining four distributions are all bimodal, with one peak at zero and another at the median sectoral inflation rate. The size of the peak at zero decreases as inflation increases. The last category, for the high inflation period 1984-1989, does not have a discernible spike at zero, but has a mode close to the median cross-sector price change. The shapes of these lines are strongly consistent with menu costs models such as Ball and Mankiw (1994, 1995). In these models there is a spike in the distribution at zero because when the inflation rate is positive many firms will choose not to reduce their relative prices, but simply wait for the average price level to increase to reduce their relative price to the desired level. These models further suggest that the spike at zero in the distribution should be decreasing in the inflation rate, for when the inflation rate is sufficiently high many firms wishing to reduce their relative prices will have to increase their actual price. At some point the inflation rate is so high that the spike at zero disappears.

In Figure 7b the deviations between the price change distribution and the median sectoral inflation rate are drawn, to standardise the different distributions. In the figure, this standardisation moves the "spike" at zero in Figure 7a to the left. Inspection of the graphs suggests that as the median sectoral inflation rate increases, the left-hand side of the distribution first 'contracts' inwards and then 'expands' outwards as the spike becomes more negative but smaller, and then disappears. This 'contraction' and 'expansion' is formally demonstrated below but can perhaps most easily be seen in figure 8a which shows how the difference between the median and the 10%, 20%, and 30% deciles of the quarterly sectoral inflation distributions varies with the median sectoral inflation rate. These graphs all indicate an inverted 'U-shaped' relationship as the median sectoral inflation rate increases. This pattern is not seen in the upper part of the sectoral inflation rate distributions, however, which are shown in figure 8b. As the median cross-sector price change increases the position of the difference between the median and the 70%, 80%, and 90% deciles steadily increases.

Table 9a shows the results of regressions that directly estimate how the position of each of the deciles depends on the median sectoral inflation rate.[32] For each decile other than the median the following regressions are estimated:

$$Decile_t^i - Median_t = \alpha_0^i + \alpha_1^i Median_t + e_t^i \quad (15a)$$

$$Decile_t^i - Median_t = \alpha_{01}^i + \alpha_{11}^i Median_t^1 + \alpha_{02}^i + \alpha_{12}^i Median_t^2 + u_t^i \quad (15b)$$

$$Median_t^1 = \begin{cases} Median_t & if Median_t \leq Median^* \\ 0 & otherwise \end{cases}$$

$$Median_t^2 = \begin{cases} 0 & if Median_t < Median^* \\ Median_t & otherwise \end{cases}$$

$Median^*$ is the median value of all of the quarterly median sectoral inflation rates, $Median^* = \mathbf{0.00504}$. $\boldsymbol{Median_t^1}$ and $\boldsymbol{Median_t^2}$ are variables that split the median sectoral inflation rate into two groups, according to their magnitude. If the median sectoral inflation rate of a particular quarter is less than the median value of all the quarterly medians, it is allocated into $\boldsymbol{Median_t^1}$; otherwise it is allocated into $\boldsymbol{Median_t^2}$. A standard Chow test was used to test the hypothesis that the coefficient estimates for the two halves of the data were the same. The regressions were estimated over the period 1984 – 2022.[33]

Two features of the results stand out. First, the coefficients α_1^i for the 60, 70, 80 and 90 deciles are positive and statistically significant at the 1% significant level; in addition, it is not possible to reject the hypothesis that the coefficients $\boldsymbol{\alpha_{11}^i}$ and $\boldsymbol{\alpha_{11}^2}$ are the same when the data are split according to the size of the median sectoral inflation rate (the Chow test). This means the data can be described by a linear relationship between the median decile position and the median sectoral inflation rate. The size of the coefficients is strictly increasing in the decile rank i.e., $\boldsymbol{\alpha_1^{60} < \alpha_1^{70} < \alpha_1^{80} < \alpha_1^{90}}$. This indicates that the width of the upper distribution steadily expands as the trimmed mean inflation rate increases, as the distance between the position of successive deciles gets wider and wider. Figure 8b shows the relationship between the 80 and 90 deciles and the median sectoral inflation rate.[34]

[32] The results are similar if the median sectoral inflation rate is replaced by the trimmed mean inflation rate. The results for these regressions are shown in table 9b.

[33] The observations for 1986:4 and 2001:1 are omitted from the decile 10 regression. In 1986 NZ introduced a VAT tax at a 10 percent rate, but housing was not included. This meant the observation for the 10 percent decile has an unusually large negative value that distorts the regression line. In 2001 the government significantly reduced state rentals, also inducing an extremely large outlier for the decile 10 line for this observation.

[34] The observations are smoothed using a 40-observation moving average applied to the data sorted by the median cross-sector price change. The dotted lines are two standard deviation bands.

Table 9a: Quarterly sectoral inflation distribution decile regressions: median, 1984- 2022.

	Const (all obs)	Median (all obs)	Const[1] (0-50)	Median[1] (0-50)	Const[2] (50-100)	Median[2] (50-100)	N	R^2	Durbin-h	Chow Test[a]
10	-0.020**	-0.26*								
percent	(0.0009)	(0.090)					153[b]	0.048	0.90	
			-0.025**	0.96	-0.017**	-0.41**				
			(0.0020)	(0.57)	(0.0018)	(0.11)	153[b]	0.093	0.24	3.68*
20	-0.008**	-0.36**								
percent	(0.0005)	(0.072)					155	0.214	0.70	
			-0.012**		-					
			(0.0009)	0.73**	0.0065**	-0.46**				
				(0.25)	(0.0008)	(0.09)	155	0.283	-0.37	7.17**
30	-0.004**	-0.26**								
percent	(0.0003)	(0.041)					155	0.332	-0.04	
			-0.006**	0.36	-0.004**	-0.29**				
			(0.0007)	(0.19)	(0.0004)	(0.05)	155	0.375	-0.72	5.18**
40	-0.002**	-0.14**								
percent	(0.0001)	(0.024)					155	0.302	-0.73	
			-0.002**	-0.03	-0.002**	-0.15**				
			(0.0002)	(0.11)	(0.0002)	(0.026)	155	0.306	-0.75	0.42
60	0.002**	0.16**								
percent	(0.0001)	(0.02)					155	0.332	0.85	
			0.002**	0.015	0.002**	0.17**				
			(0.0002)	(0.08)	(0.0002)	(0.021)	155	0.338	0.78	0.68
70	0.004**	0.32**								
percent	(0.0003)	(0.039)					155	0.363	1.83	
			0.005**	0.23	0.004**	0.32**				
			(0.0005)	(0.17)	(0.0005)	(0.044)	155	0.364	1.82	0.082
80	0.009**	0.48**								
percent	(0.0004)	(0.050)					155	0.335	0.02	
			0.009**	0.45	0.009**	0.46**				
			(0.0009)	(0.28)	(0.0008)	(0.054)	155	0.335	0.02	0.03
90	0.02**	0.68**								
percent	(0.0009)	(0.094)					155	0.218	-0.17	
			0.018**	1.43*	0.021**	0.64**				
			(0.002)	(0.67)	(0.002)	(0.07)	155	0.223	-0.29	0.59

Source: Author's calculations using Statistics New Zealand price data. $Median_t^1$ and $Median_t^2$ are variables that split the median sectoral inflation rate into two groups, according to their magnitude. If the median sectoral inflation rate of a particular quarter is less than the median value of all the quarterly medians, it is allocated into $Median_t^1$; otherwise it is allocated into $Median_t^2$.

[a] "Chow test" is the F(2,151) test that the coefficients in the periods 1984-1989 and 1990-2022 are identical.

[b] The observations for 1986:4 and 2001:1 are omitted from the regression. In 1986 NZ introduced a VAT tax at a 10 percent rate, but housing was not included. This meant the observation for the 10 percent decile has an unusually large negative value that distorts the regression line. In 2001 the government significantly reduced state rentals, also inducing an extremely large outlier for this observation.

A * indicates the hypothesis of equality can be rejected at the 5% significance level; ** means the hypothesis can be rejected at 1% significance level.

Table 9b: Quarterly sectoral inflation distribution decile regressions, trimmed mean inflation rate, 1984- 2022.

	Const (all obs)	TM^{*} (all obs)	Const[1] (0-50)	TM^{*1} (0-50)	Const[2] (50-100)	TM^{*2} (50-100)	N	R2	Durbin-h	Chow Test[a]
10	-0.020**	-0.19*								
percent	(0.0009)	(0.090)					153[b]	0.028	1.58	
			-0.025**	0.72	-0.016**	-0.43**				
			(0.0017)	(0.52)	(0.0017)	(0.11)	153[b]	0.107	0.51	6.59**
20	-0.009**	-0.30**								
percent	(0.0005)	(0.067)					155	0.160	1.13	
			-0.013**	0.82**	-0.006**	-0.46**				
			(0.0006)	(0.17)	(0.0009)	(0.09)	155	0.280	-0.86	12.58**
30	-0.004**	-0.23**								
percent	(0.0003)	(0.039)					155	0.264	0.32	
			-0.007**	0.44**	-0.003**	-0.29**				
			(0.0004)	(0.13)	(0.0005)	(0.05)	155	0.350	-1.07	9.97**
40	-0.002**	-0.13**								
percent	(0.0001)	(0.022)					155	0.253	-1.41	
			-0.002**	0.003	-0.002**	-0.14**				
			(0.0003)	(0.09)	(0.0003)	(0.03)	155	0.264	-1.52	1.16
60	0.002**	0.16**								
percent	(0.0001)	(0.02)					155	0.364	-1.40	
			0.002**	0.024	0.002**	0.15**				
			(0.0002)	(0.05)	(0.0002)	(0.017)	155	0.377	-1.66	0.68
70	0.004**	0.34**								
percent	(0.0003)	(0.035)					155	0.429	-1.32	
			0.004**	0.18	0.005**	0.30**				
			(0.0003)	(0.11)	(0.0005)	(0.037)	155	0.444	-1.66	2.00
80	0.009**	0.51**								
percent	(0.0004)	(0.050)					155	0.404	-0.25	
			0.009**	0.33	0.010**	0.45**				
			(0.0007)	(0.31)	(0.0009)	(0.054)	155	0.419	-0.56	1.96
90	0.02**	0.72**								
percent	(0.0009)	(0.094)					155	0.256	-1.10	
			0.019**	0.95*	0.022**	0.60**				
			(0.002)	(0.46)	(0.002)	(0.10)	155	0.268	-1.33	1.25

Source: Author's calculations using Statistics New Zealand price data. TM_t^1 and TM_t^2 are variables that split the trimmed mean inflation rate into two groups, according to their magnitude. If the trimmed mean inflation rate of a particular quarter is less than the median value of all the quarterly trimmed mean inflation rates, it is allocated into TM_t^1; otherwise it is allocated into TM_t^2.

[a] "Chow test" is the F(2,151) test that the coefficients in the periods 1984-1989 and 1990-2022 are identical.

[b] The observations for 1986:4 and 2001:1 are omitted from the regression. In 1986 NZ introduced a VAT tax at a 10 percent rate, but housing was not included. This meant the observation for the 10 percent decile has an unusually large negative value that distorts the regression line. In 2001 the government significantly reduced state rentals, also inducing an extremely large outlier for this observation. A * indicates the hypothesis of equality can be rejected at the 5% significance level; ** means the hypothesis can be rejected at 1% significance level.

The relationship between the position of the 10, 20, 30 and 40 deciles and the median sectoral inflation rate is quite different. It is possible to reject the hypothesis that the coefficients α_{11}^{i} and α_{12}^{i} are the same for the 10, 20 and 30 deciles. Rather, in each case the coefficient α_{11}^{i} (when the median sectoral inflation rate is small) is positive and the coefficient α_{12}^{i} (when the median sectoral inflation rate is large) is negative. This means that the left tail of the sectoral inflation rate distribution contracts as the median sectoral inflation rate increases from very low levels to its median level $Median^{*} = \mathbf{0.005}$, but it expands as the median sectoral inflation rate increases further.[35] This contraction and expansion reflect the spike of the sectoral inflation rate distribution at zero: there are more large downwards price changes when the inflation is very low than when it is at moderate 0.5% – 0.75% levels, because price-setters no longer have the option of waiting for generalised inflation to induce a drop in relative prices. The kinks in these curves are shown in Figure 8a. The behaviour of the 40 decile is different again, for in this case it is not possible to reject the hypothesis that α_{11}^{i} and α_{12}^{i} are the same, and the combined coefficient is negative. This indicates that this portion of the sectoral inflation rate distribution curves expands away from the median as the median sectoral inflation rate increases.

These findings have two interesting implications. First, the sectoral inflation rate distribution has less downwards price flexibility but greater upwards price flexibility when the annual median sectoral inflation rate is 2% or more than when the annual median sectoral inflation rate is less than 2%. This suggests that an economy with a 2% inflation target could experience deeper output declines in response to negative shocks than an economy with a 1% inflation target, because there is less downwards sector-specific price stickiness. Secondly, a central bank may find it difficult to raise the annual inflation rate to 2% whenever the annual median sectoral inflation rate is less than 2%. When the annual inflation rate is 2% or less, firms wanting to cut relative prices make larger price reductions relative to the median cross-sector price change than when the annual inflation rate is more than 2%; at the same time firms wanting to raise relative prices increase their prices by a smaller amount relative to the median. If prices are more flexible downwards and less flexible upwards when the median sectoral inflation rate is less than 2% than when it is greater than 2%, central banks will likely find it more difficult to raise the average inflation rate from 1% to 2% than from 2% to 3%.

This interpretation is consistent with the focus on price change hazard rates suggested by Callabero and Engel (2007). When the inflation rate is low, firms may be less willing to increase prices and more willing to cut prices than otherwise, making it possible for inflation to get embedded at low rates. Put differently, the whole price change hazard rate may shift downwards when the inflation rate is low. It is not possible to test their theories directly using the price information available in New Zealand, as there is a lack of firm-level price data and thus no means to estimate hazard rates directly. Nonetheless, the data demonstrate that the shape of the cross-sector price distribution curve changes as the median cross-sector price change changes in a manner that affects the headline rate by changing the frequency of very large and very small price changes.

[35] To be clear, the 'median' here is the median value over time of the distribution of quarterly trimmed mean values TM_t^{*}.

5. Discussion and conclusions

Why did consumer prices increase so quickly in 2021 and 2022 in New Zealand and around the world? While the answers to this question are far from obvious, it seems there was a sudden "avalanche" of large price increases at this time. Not only did prices in most sectors increase, but the magnitude of the largest price increases was much higher in 2021 and 2022 than at any time since the 1980s, when average inflation was also high. In the terminology of Caballero and Engle (2007), there was an upwards shift in the price change hazard curve, measured at the sectoral level.

This paper has provided new evidence on the extent that the price change hazard curve measured at the sectoral level shifts up or down when inflation increases. While this does not explain the cause of the shift in 2021, documenting the size of these shifts is useful, as there is evidence that unexpected shifts in the price change hazard curve are responsible for many of the forecasting errors made by central banks and private sector forecasters (Petrella, Santoro, and Simonsen 2019).

One of the most distinctive findings of this paper is the evidence that upwards price flexibility is systematically related to the median inflation rate. While it has long been known that the variance and skewness of sectoral inflation rates are both increasing in the inflation rate, this paper has shown that the right-hand side of the sectoral inflation rate distribution gets "wider" as the median inflation rate increases, in the sense that the difference between the median and the upper deciles of the distribution is an increasing function of the median sectoral inflation rate. This is direct evidence that the right-hand side of the sectoral inflation rate hazard curve increases as the median inflation rate increases. The effect is sizeable: as the median or trimmed mean inflation rate increases by 1 percentage point, the average inflation rate of the highest 20% of price increases (the upper trim) increases by 2 percentage points. In conjunction with the curious behaviour of the left-hand side of the sectoral inflation rate distribution, this means the average headline inflation rate increases by 1.14 percentage points for every 1 percentage point increase in the median or trimmed mean inflation rate. When the headline inflation rate is 6% or 7%, as it was in 2022, this means the median and trimmed mean inflation rates understate headline inflation by nearly 1 percentage point. The increase in the inflation rates of the sectors with the largest price increases was of a similar magnitude in 2022 as it was when inflation exceeded 6% or 7%, in the 1980s.

In contrast, the left-hand side of the sectoral price distribution closely resembles the distribution of firm level price changes in standard menu cost models. When there is positive inflation there is a spike in the distribution at zero, whose height is decreasing in the inflation rate. This means the distribution tends to have two modal values at low to moderate inflation rates, one at zero and one at the median inflation rate, although at high inflation rates the spike at zero disappears. In firm-level menu cost models, the spike at zero occurs because firms wanting to reduce their relative prices have an incentive for average prices to rise as other firms raise their prices rather than to cut their own prices. It is less obvious why the spike at zero is observed in sectoral-level data, as the average change in the price of the various goods and services produced by a large number of firms in a sector need not be very close to zero even if many of the firms in these sectors face menu costs. One explanation is that firms within a sector face highly correlated demand and cost shocks, so firms largely respond with similar price increases. Unfortunately, it is not possible to test this hypothesis in New Zealand as the appropriate firm-level data are not available. The spike in the distribution of sectoral inflation rates at zero means that prices have less downwards flexibility at the sectoral level than upwards flexibility. Moreover, since the lower trim inflation rate is less responsive than the upper trim inflation rate to changes in the median or trimmed mean inflation rates, the trimmed mean and median inflation rates become increasingly biased estimates of the headline inflation rate as the median inflation rate increases.

As the median and trimmed mean inflation rates increase, the differences between the left-hand and right-hand sides of the sectoral price distribution generate an asymmetric pattern of price adjustment. When the annual inflation rate increases from 0% to 2%, and then to 5%, upward price flexibility steadily increases but downwards price flexibility first reduces and then increases. This helps explain why a low inflation equilibrium can be self-reinforcing, for when the trimmed mean inflation rate is between 0% and 2% there is less upward price flexibility and more downwards price flexibility than when the inflation rate exceeds 2%. It also suggests that downwards price flexibility may be minimised when the inflation rate is around 2 percentage points, the centre of New Zealand's inflation targeting band. If this is the case, the current inflation target may have the surprising implication that it maximises output losses in the event of a contractionary shock to the economy. It would be interesting to find out whether this result is a peculiarity of the New Zealand data.

As a complement to the analysis of the distribution of sectoral inflation rates, many of the results have been presented in the context of a trimmed mean decomposition of the inflation rate, as the trimmed mean measure of core inflation is a popular measure used by central banks around the world including the Reserve Bank of New Zealand. Three conclusions stand out from this work.

First, a trimmed mean estimator of the core inflation rate, along with the upper and lower trims, provides a useful lens with which to better understand New Zealand's inflation dynamics. It helps identify some of the transitory components of inflation by identifying the extent that the upper and lower trims are unusually large, conditional on the trimmed mean. These transitory components are easily identified as they are simply the difference between the headline and trimmed mean inflation rates. This gap is transitory as it is fully and rapidly reversed out of the headline inflation rate within a quarter. Moreover, it is relatively easy to identify individual sectoral components that make up the gap, something that is not always immediately obvious with other measures of transitory or persistent inflation. The transitory components regularly include petrol, fruit, vegetables, domestic hotel accommodation and airfares, components that are often excluded from the headline inflation rate to calculate other core measures of inflation. But many other items also regularly appear in the upper and lower trims, so this methodology has the advantage of being systematic and comprehensive in its identification of volatile sectors.

Secondly, the trimmed mean estimator helps improve short term forecasts of the headline inflation rate, even though it is a biased estimator of the headline rate. This is not particularly surprising, as the purpose of a core inflation estimator is to filter out the most transitory and volatile components of the headline rate. Nonetheless, it proves that a simple vector error correction system can be used to reduce the error in short run forecasts by about 20 percent relative to univariate forecasts based only on the headline inflation rate. The key improvement occurs because the trimmed mean has higher persistence and lower volatility than the headline rate, as the most volatile components of the headline rate are identified and excluded. The method has the additional advantage of providing an explanation for the inflation forecast, as the transitory components are identified.

Thirdly, the paper has demonstrated that the size of the upper and lower trims increase at different rates as the inflation rate increases. This not only means the trimmed mean is a biased estimate of the headline inflation rate, which is well known, but that, within the range of inflation rates that has occurred in New Zealand over the last four decades, this bias has increased when the headline inflation rate has increased.[36] This does not negate its use, for, as noted above, variations in the difference between the headline and the trimmed mean inflation rates (the CPITM gap) are rapidly reversed out of the headline inflation rate and improve headline inflation forecasts. But the bias means that the trimmed mean and median inflation rates understate headline inflation as the inflation rate increases, because an increasingly large fraction of the overall inflation occurs due to very large price increases in a small number of sectors. As these very large increases appear to be a regular feature of the inflation rate when inflation increases, they should not be discounted or ignored simply because they appear to be "one-off" events hitting a small number of sectors. Rather, the varying size of this bias needs to be taken into account when interpreting the trimmed mean and median inflation rates, particularly in an environment where inflation is increasing and the trimmed mean inflation rate understates the headline inflation rate by an increasingly large amount.

[36] At very high levels of inflation, such as those sometimes seen in Argentia or Mexico, the trimmed mean and headline inflation rates may converge because in these circumstances most firms adjust prices at a similar frequency and by a similar amount in response to the dominant common inflation shock. See Alvarez et al (2019).

References

Alogoskoufis, G., & Smith, R. (1991). On error correction models: specification, interpretation, estimation. *Journal of Economic Surveys*, *5*(1), 97-128.

Alvarez, F., Beraja, M., Gonzalez-Rozada, M., & Neumeyer, P. A. (2019). From hyperinflation to stable prices: Argentina's evidence on menu cost models. *The Quarterly Journal of Economics*, *134*(1), 451-505.

Atkinson, T., Dolmas, J., Giannoni, M., & Koenig, E. F. (2021). *What the Trimmed Mean Says About Future Inflation: Broadening Price Pressures Ahead.* Federal Reserve Bank of Dallas: Commentary 1 July 2021

Ball, L., & Mankiw, N. G. (1994). Asymmetric price adjustment and economic fluctuations. *The Economic Journal*, *104*(423), 247-261.

Ball, L., & Mankiw, N. G. (1995). Relative-price changes as aggregate supply shocks. *The Quarterly Journal of Economics*, *110*(1), 161-193.

Berger, D., & Vavra, J. (2018). Dynamics of the US price distribution. *European Economic Review*, *103*, 60-82.

Bomberger, W. A., & Makinen, G. E. (1993). Inflation and relative price variability: Parks' study reexamined. *Journal of Money, Credit and Banking*, *25*(4), 854-861.

Bryan, M.F. and Cecchetti, S.G., 1994. Measuring core inflation. *Monetary Policy* (pp. 195-219). The University of Chicago Press.

Bryan, M. F., Cecchetti, S. G., & Wiggins, R. L. (1997). Efficient inflation estimation. NBER WP 6183.

Caballero, R. J., & Engel, E. M. (2007). Price stickiness in Ss models: New interpretations of old results. *Journal of Monetary Economics*, *54*, 100-121.

Cavaliere, G., & Taylor, A. R. (2009). Heteroskedastic time series with a unit root. *Econometric Theory*, 1228-1276.

Cavallo, A., Lippi, F., & Miyahara, K. (2024). Large shocks travel fast. *American Economic Review: Insights*, *6*(4), 558-574.

Coleman, A.M.G. 2007. "Tradables and non-tradables inflation in Australia and New Zealand," *Reserve Bank of New Zealand Bulletin* 70(1), 2007, 45-53.

Cukierman, A. (1982). Relative price variability, inflation, and the allocative efficiency of the price system. *Journal of Monetary Economics* 9 131-162.

Cukierman, A. (2019). Implications of the permanent-transitory confusion for New-Keynesian modeling, inflation forecasts and the post-crisis era. CEPR Discussion Paper series DP 13727

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3) 253-263

Diebold, F. X. (2015). Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of Diebold–Mariano tests. *Journal of Business & Economic Statistics*, *33*(1), 1-9.

Dolmas, J. (2005). Trimmed mean PCE inflation. *Federal Reserve Bank of Dallas Working Paper*, 506.

Fischer, S. (1981). Relative shocks, relative price variability, and inflation. *Brookings Papers on Economic Activity*, *1981*(2), 381-441.

Gagnon, E. (2009). Price setting during low and high inflation: evidence from Mexico. *Quarterly Journal of Economics* 124 (3) 1221-1263

Gamber, E. N., Smith, J.K., and Eftimoiu, R. (2015) The dynamic relationship between core and headline inflation. *Journal of Economics and Business* 81 38-53

Goodchild, J. (2003) *Trend Measures of Price Level Change: Trimmed Means and Weighted Medians* (Wellington: Statistics New Zealand)

Golosov, M., & Lucas Jr, R. E. (2007). Menu costs and Phillips curves. *Journal of Political Economy*, *115*(2), 171-199.

Greenway-McGrevy, R and Jones, J.A. (2022) Evaluating Limited influence estimators of core inflation in New Zealand. Auckland: Economic Policy Centre WP 10.

Harrell, F. E., & Davis, C. E. (1982). A new distribution-free quantile estimator. *Biometrika*, *69*(3), 635-640.

Hendry, D. F., Pagan, A. R., & Sargan, J. D. (1984). Dynamic specification. *Handbook of econometrics*, *2*, 1023-1100.

Luo, S., & Villar, D. (2021). The skewness of the price change distribution: A new touchstone for sticky price models. *Journal of Money, Credit and Banking*, *53*(1), 41-72.

Mehra, Y. P., & Reilly, D. (2009). Short-term headline-core inflation dynamics. *FRB Richmond Economic Quarterly*, *95*(3), 289-313.

Nakamura, E., & Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, *123*(4), 1415-1464.

Nakamura, E., Steinsson, J., Sun, P., & Villar, D. (2018). The elusive costs of inflation: Price dispersion during the US great inflation. *The Quarterly Journal of Economics*, *133*(4), 1933-1980.

Nirei, M., & Scheinkman, J. A. (2024). Repricing avalanches. *Journal of Political Economy*, *132*(4), 1327-1388.

Perron, P. (2006). Dealing with structural breaks. *Palgrave Handbook of Econometrics*, 1(2), 278-352.

Petrella, I., Santoro, E. and Simonsen, L. (2019) Time-varying price flexibility and inflation dynamics. CEPR Discussion Paper No. DP13027.

Rich, R. W., Verbrugge, R. J., & Zaman, S. (2022). Adjusting Median and Trimmed-Mean Inflation Rates for Bias Based on Skewness. *Economic Commentary*, (2022-05).

Sheremirov, V. (2020). Price dispersion and inflation: New facts and theoretical implications. *Journal of Monetary Economics*, *114*, 59-70.

Vining, D. R., & Elwertowski, T. C. (1976). The relationship between relative prices and the general price level. *The American Economic Review*, *66*(4), 699-708.

Appendix 1: Data used to calculate the trimmed mean inflation estimator 1984-2022

1. 1999-2022

The trimmed mean estimator is based on level 3 classification CPI data obtained from Statistics New Zealand, reference INFOS CPI013AA. There are 109 series, most of which begin 1999:2 and finish 2022:4. When observations were not available the inflation rate and the weights were set equal to zero. Contemporaneous weights were used: the weights were changed in 1999:2; 2002:2; 2006:2; 2008:2; 2011:2; 2014:2; 2017:3; 2020:2; 2021:2; and 2022:2.

2. 1983-1999

During this period Statistics New Zealand used a different classification. There were 66 categories and approximately 300 class 4 goods and services. These were re-aggregated into 107 categories so that the average size of a category was similar to the data used 1999-2022. Statistics New Zealand no longer provides access to the disaggregated data, so historical class 4 price data from the RBNZ (initially collected from contemporary Statistics New Zealand sources) were used instead. The 107 categories were based on the categories used by Coleman (2007) and are similar but not identical to the categories used from 1999-2022. (For example, there are fewer categories of education expenses, but more categories of clothing, in line with the changing importance of these items in the CPI basket.) Contemporary weights were used: the weights were changed in 1983:4; 1988:4; 1993:4; 1996:2. As the weights of the available class 4 goods and services did not add to exactly 100 percent, the weights were normalised before they were used to calculate the trimmed mean estimator.