A New Core Inflation Indicator for New Zealand*

Domenico Giannone[a] and Troy D. Matheson[b]
[a]European Central Bank, ECARES, and CEPR
[b]Reserve Bank of New Zealand

This paper introduces a new indicator of core inflation for New Zealand, estimated using a dynamic factor model and disaggregate consumer price data. Using disaggregate consumer price data, we can directly compare the predictive performance of our core indicator with a wide range of other "core inflation" measures estimated from disaggregate consumer prices, such as the weighted median and the trimmed mean. The medium-term inflation target of the Reserve Bank of New Zealand is used as a guide to define our target measure of core inflation—a centered two-year moving average of past and future inflation outcomes. We find that our indicator produces relatively good estimates of this characterization of core inflation when compared with estimates derived from a range of other models.

JEL Codes: C32, E31, E32, E52.

1. Introduction

Inflation is a key variable in the formulation of monetary policy. However, inflation data are often subject to transitory shocks, which do not require a policy response. The purpose of core inflation measures is thus to remove the influence of transitory shocks, revealing the unobserved, policy-relevant trend in inflation.

*Thanks to Andrew Coleman, Kirdan Lees, Michele Lenza, Lucrezia Reichlin, Christie Smith, and an anonymous referee for useful comments on earlier drafts. All errors and omissions are entirely ours, and the views expressed in this paper do not necessarily reflect those of the Reserve Bank of New Zealand or the European Central Bank. Corresponding author: Matheson: Economics Department, Reserve Bank of New Zealand, P.O. Box 2498, Wellington, New Zealand. E-mail: troy.matheson@rbnz.govt.nz.

The 2002 Policy Targets Agreement (PTA), signed by the Minister of Finance and the governor of the Reserve Bank of New Zealand (RBNZ), requires the Reserve Bank to keep annual inflation in the consumers price index (CPI) between 1 and 3 percent, on average, over the medium term. This suggests that the policy-relevant rate of inflation for New Zealand should remove short-term fluctuations from inflation by averaging over the medium term. But what is the "medium term"?

It is commonly thought that the lag between a change in monetary policy and its impact on inflation is one to three years, suggesting that the medium term is a period long enough to smooth out fluctuations in annual inflation lasting around two years in duration.[1] Core inflation could thus be defined as historical realizations of the inflation target, where the target looks something like an average of annual inflation over the past two years. The problem with this characterization of core inflation, however, is that it lags annual CPI inflation by almost a year.

Ideally, we would like our core inflation indicator to look like a moving average (MA) of annual inflation, but without the phase shift induced by taking an average of historical data. In other words, we want an indicator that looks like a *centered* moving average of annual inflation (or, equivalently, a forecast of a backward-looking average), but one that can be computed in real time. One way or another, this is precisely what all core inflation measures aim to do; they aim to produce a real-time estimate of the underlying, policy-relevant rate of inflation, something that is unobservable in real time.

We estimate core inflation using disaggregate consumer prices and a cross-sectional smoothing technique that was first used by Altissimo et al. (2001) to estimate a coincident indicator for the euro-area business cycle. Essentially, we aim to exploit the leading and lagging relationships among disaggregate prices to proxy for future CPI inflation, thus smoothing inflation and avoiding the phase shift induced by taking an average of historical data on aggregate prices.

Many central banks monitor a variety of measures of core inflation, typically constructed from disaggregate CPI data. The CPI

[1] Christiano, Eichenbaum, and Evans (2005), e.g., find that inflation responds in a hump-shaped fashion to a monetary shock, peaking at around about two years.

excluding food and energy, for example, is a very popular measure of core inflation in the United States. There is a multitude of other so-called exclusion measures used around the world, all with a common goal in mind—to eliminate the idiosyncratic noise from inflation by removing the most volatile components from the CPI. In addition to the exclusion measures, there is also a range of statistical measures of core inflation monitored by central banks. These measures aim to remove those components from inflation that are most volatile in a particular month, quarter, or year. Two popular statistical measures are the weighted median and the trimmed mean of inflation, proposed by Bryan and Cecchetti (1994). However, many of these measures of core inflation—both the exclusion and the statistical methods—tend to be too volatile and often fail to provide a reliable signal for underlying inflation (Cogley 2002).

Recently, there have been some advances in econometric theory that allow the decomposition of very large panels of data into a small number of common factors (Forni et al. 2000; Stock and Watson 2002; Forni et al. 2005). These methods are well suited to the problem of estimating core inflation, where the inflation signal of interest is both unobservable in real time and common to a large number of macroeconomic series.

Cristadoro et al. (2005) apply the dynamic factor model to extract the long-run component of inflation from almost 450 nominal, real, and financial indicators of inflation, producing a measure of core inflation for the euro area. Their core measure has the benefit of being smooth, without having the phase shift induced by taking an annual percentage change of inflation. Moreover, the Cristadoro et al. (2005) indicator is very competitive at forecasting annual inflation at horizons up to two years ahead. Similarly, Amstad and Fischer (2004) and Amstad and Potter (2007) use the dynamic factor model to compute core inflation measures for Switzerland and the United States, respectively. Unlike Cristadoro et al. (2005), Amstad and Fischer (2004) and Amstad and Potter (2007) allow for real-time estimates of core inflation as each new piece of data arrives (what the authors call "sequential information flow").

One of the key objectives of this paper is to compare dynamic factor model estimates of core inflation with a variety of methods of estimating core inflation from disaggregate consumer price data—including standard core inflation measures, estimates from pooling regressions, and estimates from bivariate forecasting models. Essentially, we want to ask, Given data on disaggregate consumer price movements, what is the best way to estimate core inflation? Thus, unlike Amstad and Fischer (2004), Cristadoro et al. (2005), and Amstad and Potter (2007) (who use broader data sets to estimate core inflation), we limit ourselves to using disaggregate consumer prices.[2]

We assess the performance of our core inflation indicator in real time by comparing how well it predicts a particular characterization of the RBNZ's inflation target—a centered two-year moving average of annual inflation. We compare the indicator with a variety of other estimates of core inflation based on disaggregate price data and find that the indicator performs well. Furthermore, when compared with a range of forecasting models used at the RBNZ, many of which use a wider range of information, the core inflation measure also compares favorably and is only bettered by models that incorporate judgment and utilize more up-to-date information.

2. Methodology

We have T time-series observations for N different inflation series from the CPI denoted π_{jt}, where $j = 1, \ldots, N$; $t = 1, \ldots, T$; and π_{jt} is the (log) seasonal change of the j-th price index. Further, let us add to this panel the (log) seasonal change in headline CPI, π_t. Headline inflation can then be represented as the sum of two unobserved components, a signal π_t^* and an error e_t:

$$\pi_t = \pi_t^* + e_t. \tag{1}$$

The objective is to estimate the signal π_t^* using all information in the panel of CPI price changes. We assume that each variable can be represented as two stationary, orthogonal, unobservable components—a common component χ_{jt} and an idiosyncratic component ε_{jt}:

$$\pi_{jt} = \chi_{jt} + \varepsilon_{jt}, \tag{2}$$

[2] In the case of New Zealand, the CPI data are not subject to revision and are more timely than most other key macroeconomic releases; the CPI is published more than one month before the producer price index and more than two months before the national accounts.

where the common component is driven by a small number of common factors (shocks).

We decompose the common component into a long-run component χ_{jt}^{L} and a short-run component χ_{jt}^{S} by removing high-frequency, short-run fluctuations up to a given critical period h (Cristadoro et al. 2005):

$$\pi_{jt} = \chi_{jt}^{L} + \chi_{jt}^{S} + \varepsilon_{jt}. \tag{3}$$

Specifically, the intertemporally smoothed (long-run) common component can be attained by summing waves of different periodicity between $[-\pi/h, \pi/h]$ using a spectral decomposition. The long-run common component is what we are after in estimating our measure of core inflation. This measure removes the idiosyncratic noise specific to each of the components of the CPI, as well as smoothing out the short-term fluctuations not requiring a monetary policy response.

Isolation of the unobserved common component can be achieved by assuming that the common components are driven by shocks that are pervasive in the cross-section of price movements, while the shocks driving the idiosyncratic terms are local and affect only a limited number of prices. This is called an approximate dynamic factor structure. The dynamic factor model is

$$\pi_{jt} = b_j(L) f_t + \varepsilon_{jt}, \tag{4}$$

where $f_t = (f_{1t}, \dots, f_{qt})'$ is a vector of q dynamic factors and $b_j(L)$ is of order s, for every dynamic factor $1, \dots, q$: the process for f_t is assumed to be stationary.[3] This model is said to have q common dynamic factors.

If we let $F_t = (f_t', f_{t-1}', \dots, f_{t-s}')'$, the dynamic factors have a static representation:

$$\pi_{jt} = \lambda_j F_t + \varepsilon_{jt}, \tag{5}$$

where $b_j(L) f_t = \lambda_j F_t$. Thus, a model with q dynamic factors has $r = q(s+1)$ static factors.

[3]Typically the data-generating process for the common factors is approximated by a finite-order vector autoregression (VAR). See, e.g., Forni et al. (2007).

3. Estimation

The dynamic factor model is estimated in the frequency domain using an eigenvalue decomposition of the spectrum smoothed over a range of frequencies. Estimation requires the specification of three parameters: two of the three parameters determining the number of static factors r, q, and s, and the size of the Bartlett lag window M.[4] Estimation of the static factor representation of the dynamic factor model (5), on the other hand, requires an eigenvector-eigenvalue decomposition of the variance-covariance matrix and requires r to be specified.[5] We use the dynamic factor representation to estimate our indicator of core inflation. More details on the estimation of the dynamic and static factor models are provided in appendix 1.

In order to get good estimates of the common factors, Cristadoro et al. (2005) recommend that the panel of data used to estimate the dynamic factor model should have series that lead and lag annual inflation. Their argument is summarized as follows.

Consider the case where there is only one common shock to inflation, f_t, which is loaded with different lags in a cross-section containing three variables:

$$\chi_{1t} = f_{t-1}, \quad \chi_{2t} = f_t, \quad \text{and} \quad \chi_{3t} = f_{t-2}. \tag{6}$$

In this case, it is clear that variable 2 leads variable 1, which in turn leads variable 3. If CPI inflation is χ_{1t}, then $\chi_{1t+1} = f_t$ and $\chi_{1t-1} = f_{t-2}$, so that the cross-sectional information hidden in the contemporaneous $\chi_{jt}s$ is exactly the time-series information contained in CPI inflation and its first lead and lag. CPI inflation can thus be smoothed in period t by using the leading variable as a proxy for future CPI inflation, which is unavailable at time t.

The methodology then consists of projecting the long-run common component of headline CPI inflation, χ_t^L, onto the (intertemporally smoothed) present and past common factors:

$$\hat{\chi}_t^L = \text{Proj}\big[\chi_t^L \big| f_{mt-k}, m = 1, \ldots, q; k = 1, \ldots, s\big]. \tag{7}$$

[4]For all the dynamic factor models estimated in this paper, we set $M = \sqrt{T}$, as in Forni et al. (2005).

[5]A VAR in F_t can then be used to estimate the dynamics of the model.

Thus, as seen in the example above, the method will produce good results if χ_t loads mainly with lags central to the interval $0, \ldots, s$. In other words, to get good estimates of the common factors, the data set must include variables that lead and lag χ_t. As noted by Cristadoro et al. (2005), leading variables are of particular importance, since lagging variables could be replaced by lags of existing variables, whereas the leading variables are irreplaceable.

As noted in the introduction, we characterize the target for our core inflation indicator as a centered two-year moving average of annual CPI inflation. Letting $MA(\pi_t, h)$ be a centered moving average with a window of $2h + 1$, then

$$MA(\pi_t, h) = \frac{1}{2h+1}(\pi_{t-h} + \pi_{t-h+1} + \ldots + \pi_t + \ldots + \pi_{t+h-1} + \pi_{t+h}). \tag{8}$$

Our objective is to construct an indicator of core inflation that matches the properties of this filter and does not require future information (information from period $t + 1$ to $t + h$) to compute. Estimation in the frequency domain allows us to remove short-term fluctuations from the common components to reveal a long-run common component of annual CPI inflation $\hat{\chi}_t^L$. Moreover, our target filter $MA(\pi_t, 4)$ sets the long-run frequencies that we can use to intertemporally smooth the common component $\hat{\chi}_t^L$. Specifically, we compute the intertemporally smoothed (long-run) common component of inflation by summing waves of different periodicity in the band $[0, 2\pi/9]$, removing all short-run cyclical fluctuations up to nine quarters in duration.[6] Henceforth, our characterization of the inflation target is denoted $\pi_t^{target} = MA(\pi_t, 4)$ and the dynamic factor model's estimate of core inflation is denoted $\pi_t^{core} = \hat{\chi}_t^L$.

4. Data and Dynamic Structure of the Data

We begin with quarterly data for all 264 subsections of the CPI for a period ranging from 1991:Q1 to 2006:Q2. To these series we also add headline CPI, tradable CPI, and nontradable CPI.[7]

[6] See Giannone and Matheson (2006) for further discussion.

[7] Statistics New Zealand publishes a split of the CPI regimen into tradable and nontradable price indexes. We include these indexes in our panel to enable

The data are filtered in five steps. First, we remove all series from panel that do not span the entire sample. Second, we take the natural logarithm and seasonally difference all series to achieve stationarity. Third, we remove outliers from each series by replacing observations more than six times the interquartile range with the median of the series. Fourth, we remove the series whose prices change less than once a year on average.[8] The filtered series are then standardized to have zero mean and a unit variance. After filtering, the panel contains headline inflation, tradable inflation, nontradable inflation, and inflation data for 228 subsections of the CPI.[9] Core inflation is obtained by reattributing the mean and the variance of each series; i.e., the estimate of π_t^{core} is rescaled by multiplying it by the standard deviation of π_t and by adding the mean of π_t.

The left side of figure 1 displays a histogram of the leads/lags at which the highest absolute correlation with inflation occurs in our panel. The figure shows that the structure of the data is well balanced and includes a similar number of leading and lagging variables. Because of the importance of the leading variables in estimation, we also examine the predictive power of the series in our panel by way of a Granger causality test—testing the null hypothesis that series j does not Granger-cause headline CPI inflation. The right side of figure 1 displays a histogram of the p-values from this test. Around 30 percent of the series in the panel Granger-cause headline inflation at the 10 percent level of significance. With a similar number

the RBNZ to more readily compute core measures of tradable and nontradable inflation. The inclusion of these aggregate variables should not pose a problem for estimating core inflation, because an approximate factor structure allows for local correlation among the idiosyncratic components.

[8]Correcting for outliers and removing series that change less than once a year on average are not crucial steps. The core indicator estimated with the raw, unadjusted data is very similar to the indicator estimated with the adjusted data: the root mean-squared forecast error of the core indicator estimated without adjustments is only marginally (0.0052 percent) higher than that presented in table 4 (in section 6). Nevertheless, we choose to make these corrections to insure against any coding or typographical errors (made by the RBNZ or Statistics New Zealand) that could adversely affect quarterly updates of core inflation in the future. The core inflation indicator has been published in the RBNZ's quarterly *Monetary Policy Statement* since September 2006.

[9]During the filtering process, a total of 9.21 percent of the CPI regimen is removed from the original panel and thirty-five observations are classified as outliers. See Giannone and Matheson (2006).

Figure 1. The Correlation Structure of the Panel

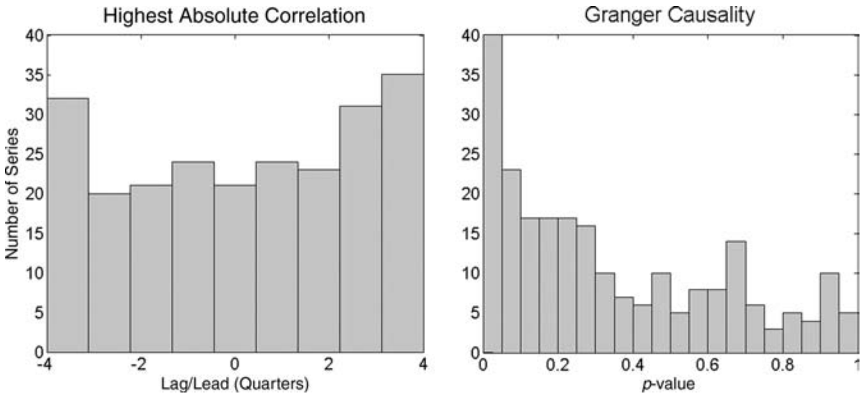

of leading and lagging variables and with a sizable proportion of the series having some predictive ability with respect to annual CPI inflation, this panel seems well suited to estimating a dynamic factor model. A more detailed description of our panel, including statistics relating to the dynamic structure of the data discussed in the next section, can be found in Giannone and Matheson (2006).

4.1 The Dynamic Structure of the Data

Determining the unobserved factors driving prices is difficult. As with many statistical problems, the choice of the number of factors requires a trade-off between parsimony and fit. The more factors that are added to the model, the more variation in the data set that will be explained by the factors. Fewer factors, on the other hand, will produce a smoother indicator of core inflation, but at a cost of poorer fit to the data. Moreover, the selection of the number of static factors r is further complicated since there is more than one configuration of the parameters q and s that has similar fit to the data.

Bai and Ng (2002) show that the number of static factors can be consistently estimated with an information criterion. Unfortunately, the Bai and Ng criterion does not appear to be suitable for our empirical application and chooses T static factors. Effectively, this means that, statistically, our panel is driven by very many factors—so many, in fact, that if one were to project inflation on the estimated common components, one would get precisely the original data back, rotated on the factors.

In an approximate factor structure with q common shocks and r common factors, there are q dominant eigenvalues from the spectral-density matrix (dynamic rank) and r dominant eigenvalues from the covariance matrix (static rank). Thus, some insight into a number of the common shocks can be found by examining the behavior of the dynamic eigenvalues, which represent the variance of the dynamic factors. Figure 2 plots the first ten dynamic eigenvalues from the spectral-density matrix of our data over frequencies $[0, \pi]$. The first four eigenvalues are considerably larger than the others, particularly at the long-run frequencies with which we are concerned. It thus seems that the long-run common comovement in these data can be adequately captured by the first four dynamic factors.

Another approach to determine the number of factors is to describe the comovements of the panel using the percentage of the

Figure 2. The First Ten Dynamic Eigenvalues from the Spectral-Density Matrix

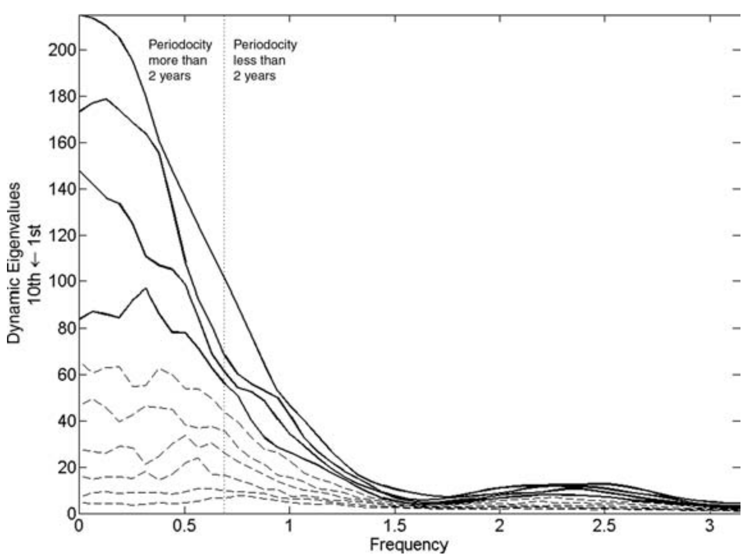

Table 1. Percentage of Total Variance Explained by the First Twelve Dynamic and Static Factors

	1	**2**	**3**	**4**	**5**	**6**	**8**	**10**	**12**
q	0.23	0.42	0.57	0.69	0.78	0.84	0.92	0.95	0.97
r	0.13	0.26	0.35	0.42	0.48	0.53	0.63	0.70	0.76

variance of the panel accounted for by the common factors, as suggested by Forni et al. (2000). Table 1 reports the percentage of the total variance in our panel explained by the first twelve dynamic factors q (estimated using dynamic principal components) and the same number of static factors r (estimated using principal components).[10] The table shows that a small number of factors explain a large amount of the variation in our panel. Moreover, there is a discrepancy between the variance explained by the dynamic and static factors, suggesting, as with the results in section 4, that there are some rich dynamics at play in our panel. Because the rank of the covariance matrix of the panel is always $r = q(s+1)$ and the rank of the spectral-density matrix is always q, the difference between the variance explained by r and q reflects the lagged factors s. Selecting the number of dynamic factors q so that the marginal contribution from adding one more factor is less than 10 percent, as suggested by Forni et al. (2000), produces $q = 4$. Selecting the number of static factors r to explain the same amount of variation as $q = 4$ produces $r \approx 10$, implying that s is somewhere between 1 and 2.

For our indicator of core inflation, we choose to set $q = 4$, $s = 2$, and $r = 12$.[11] As a robustness check, we estimate the indicator recursively with different configurations of q and s over a period from 2000:Q1 to 2005:Q2 and find that our chosen parameterization performs comparatively well. Indeed, with the exception of when $s = 0$, the models with $q = 4$ outperform all other models for a given s, justifying our assumption of four dynamic factors. The

[10]See appendix 1 for a description of principal-components estimators.

[11]Equation (5) implies that $r = q(s + 1)$. However, it is worth noting that this holds only in the case where the order of the MA in the common shocks is finite and that, in general, $r \geq q(s + 1)$; see Forni et al. (2007). In choosing our parameter configuration, we implicitly assume that the order of the MA is finite.

performance of the indicator with $q = 4$ and $s = 1$ is the same as our chosen parameterization.[12]

Giannone and Matheson (2006) show that the first four common factors explain more than 75 percent of the variability of headline CPI inflation. Moreover, at the periodicities with which we are concerned (longer than two years), the degree of commonality is even higher, with the common factors explaining over 80 percent of the variation.

5. The Core Inflation Indicator

The core inflation indicator is displayed in figure 3 alongside annual CPI inflation and our target variable, a centered two-year moving average of annual CPI inflation. We find that the core inflation indicator smooths much of the noise from annual CPI inflation, and it closely tracks the target variable.

Figure 3. CPI Inflation, the Inflation Target, and the Core Indicator

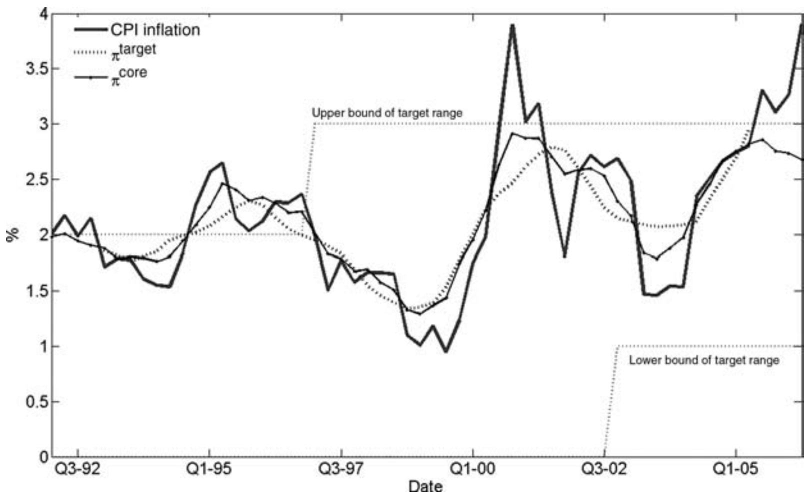

[12]See section 6 for more details on the real-time prediction experiment and appendix 4 for the results for different configurations of q and s.

So how well does our core inflation indicator do in predicting the target measure of inflation? Table 2 compares the core inflation indicator with weighted-median inflation; trimmed-mean inflation (with a 10 percent trim); the CPI excluding food, administration changes, and gasoline; and the exponentially smoothed measure of core inflation proposed by Cogley (2002). All of these measures of core inflation are described in greater detail in appendix 2 of the paper. For the moment, we only discuss our core inflation indicator estimated up to the end of the sample. The real-time indicator in the final row of the table, Core (Real Time), is discussed in the next section.

All of the core measures average around 2 percent, similar to headline CPI. However, there are some large differences in the variability of the measures. In particular, the measures of the weighted median, trimmed mean, and CPI excluding food, administration charges, and gasoline are much more volatile than our core indicator and the exponentially smoothed indicator. In fact, these core measures seem to do a bad job at smoothing inflation, having standard deviations higher than headline CPI inflation itself. Our core indicator has a standard deviation that closely matches the target variable, whereas the exponentially smoothed measure seems to smooth inflation too much and has a much lower standard deviation than the target.[13]

Looking at absolute correlations with the target measure, π_t^{target}, at leads and lags of up to one year, we find that our core inflation indicator correlates highly and is in phase with the target variable. The exponentially smoothed measure also correlates highly with the target, although with a lag of two quarters. Likewise, the weighted median, the trimmed mean, and the CPI excluding food, administration charges, and gasoline tend to lag the target by a couple of quarters.

It is interesting to look at the concordance of each core inflation measure with our target variable, π_t^{target}. Concordance measures the percentage of the sample where changes to the indicator and changes to the target variable have the same sign, i.e., the percentage of time that changes to the core indicator accurately reflect changes in the

[13] The standard deviation of π_t^{target} is 0.46.

Table 2. Descriptive Statistics for Various Measures of Core Inflation

						$\pi_{t+4} - \pi_t = \alpha + \beta(\pi_t^{core} - \pi_t) + \epsilon_{t+4}$		
	Mean	**Std.**	**Max.**	**Lag**	**Conc.**	α	β	R^2
CPI	2.07	0.72	0.77	0	0.70	—	—	—
Core	2.14	0.51	0.95	0	0.79	–0.06	2.58*	0.71
Weighted Median	2.22	0.79	0.82	3	0.55	0.13	0.03	0.00
Trimmed Mean	2.08	0.76	0.83	1	0.58	0.13	–0.27	0.01
CPI excl. Food, Admin., and Gasoline	1.89	0.88	0.83	2	0.64	0.05	–0.45	0.04
Exponentially Smoothed	2.04	0.29	0.94	2	0.70	0.11	1.17*	0.38
Core (Real Time)	2.03	0.32	0.92	1	0.82	0.13	1.31*	0.44

Notes: Mean of the series (Mean); Standard deviation (Std.); Maximum absolute correlation with π_t^{target} (Max.); Lag at which the maximum absolute correlation occurs (Lag); Concordance with π_t^{target} (Conc.); coefficient on the deviation from each core measure to predicting $\pi_t+4-\pi_t(\beta)$; * denotes significance of β at the 1 percent level (standard errors have been adjusted for heteroskedasticity and autocorrelation using the Newey and West 1987 estimator); R^2 is the coefficient of determination for the regression. All statistics are estimated for thirty-four observations from 1997:Q1 to 2005:Q2—the longest period for which all data are available.

unobserved target variable (Harding and Pagan 2002). For all of the indicators, concordance is above 50 percent—the proportion of time that a coin toss would accurately predict the correct change in the target variable. Concordance is highest for our core inflation indicator, at 0.79 percent. The exponentially smoothed measure also predicts changes in the target 70 percent of the time.

Following Cogley (2002), the predictive ability of core inflation can be evaluated using the following regression:

$$\pi_{t+4} - \pi_t = \alpha + \beta\left(\pi_t^{core} - \pi_t\right) + \epsilon_{t+4}, \tag{9}$$

where π_t is headline CPI inflation, π_{t+4} is headline CPI inflation one year into the future, π_t^{core} is the core inflation measure, and ϵ_{t+4} is idiosyncratic noise. The regression estimates whether the current gap between core inflation and headline inflation predicts future changes in headline inflation, where predictive power is indicated by $\beta > 0$ (α should also be equal to 0).

We find that our core indicator and the exponentially smoothed indicator have significant predictive power, and they explain nontrivial amounts of the variation of future changes in headline inflation, according to R^2. The other core measures, in contrast, perform very poorly by this metric.

5.1 The Real-Time Properties of the Core Indicator

Because our core indicator is estimated, unlike the other core inflation measures displayed in table 2, it is subject to revision as more data become available. However, this may not impact dramatically on the relative predictive performance of the indicator.

To examine the real-time properties of the indicator, we estimate it recursively for each quarter from 1997:Q1 to 2005:Q2 (the next section describes the real-time estimation procedure in more detail). An indicator of core inflation that is not subject to revision can then be created using the real-time estimates of core inflation in each quarter; i.e., the series $\pi_{t,real}^{core}$ can be compiled using π_t^{core} estimated each quarter from 1997:Q1 to 2005:Q2.

Comparing these real-time estimates of core inflation with core inflation estimated using all of the data (the last row of table 2), we find that the indicator retains many desirable features. It remains

Table 3. Revisions to the Core Indicator in Real Time

	t	$t-1$	$t-2$	$t-3$	$t-4$
Mean Revision	0.12	0.09	0.08	0.07	0.06
Mean Absolute Revision	0.26	0.17	0.10	0.08	0.07

highly correlated with the target, albeit now with a lag of one quarter. Interestingly, the concordance with the target increases with the real-time core indicator.

Revisions to the core inflation indicator over time can be analyzed by comparing the real-time estimates of core inflation with the core inflation indicator estimated using all of the data. To be concrete, each quarter t, the estimate of core inflation can change not only for period t but also for all ω historical periods prior to period t, i.e., periods $t-1, \ldots, t-\omega$, where ω denotes the first observation used to estimate the indicator each quarter. Table 3 displays the mean revisions and the mean absolute revisions to the core indicator at periods $t, t-1, \ldots, t-4$ (revisions are defined to be the core indicator estimated up to the end of the sample minus the real-time estimates).

Not surprisingly, the revisions are larger for the most recent estimates of core inflation. The final five estimates of core inflation have generally been revised up since 1997:Q1, with the revisions to the most recent estimate, period t, being just under twice as large as those from a year prior to the most recent estimate. The magnitude of each revision in period t is 0.26 percent, almost four times more than the estimate from period $t-4$.

Notwithstanding any changes to the dynamic relationships within our panel, with such a small sample to begin with (twenty-one observations), the core inflation will suffer from biases relating to changes to time-series properties of headline CPI inflation over the sample. This can be seen in figure 3, where we see that the mean and variance of the series are higher in the second half of the sample, perhaps as a result of the changes to the inflation target of the Reserve Bank of New Zealand. Recall that, after the dynamic factor model is estimated, the mean and the standard deviation of headline CPI inflation are reattributed to π_t^{core} to make it comparable to

Figure 4. Revisions to Real-Time Estimates of Core Inflation

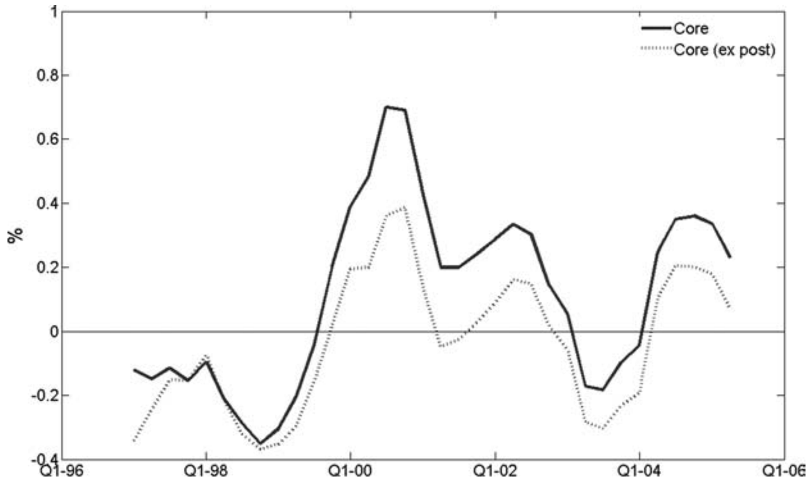

headline CPI inflation. Mean and variance shifts are thus a source of revision to our core indicator.

To further examine the impact of this type of revision, we compare the real-time estimates of the core inflation, Core, with the real-time estimates of core inflation scaled with the mean and standard deviation of headline inflation estimated over the entire sample, Core (ex post). In this way, we get a sense of how mean and variance shifts over the sample have influenced the size of the revisions. The revisions to the most current estimate of core inflation in real time (period t) from these two core indicators are displayed in figure 4.

Revisions to the standard core indicator were particularly large over 2000, when there was a substantial shift in headline inflation, although the impact of that shift did not influence Core (ex post) by as much. Indeed, generally speaking, Core was not revised as much as Core (ex post), suggesting that mean and variance shifts over our sample period were partly to blame for the revisions to the real-time estimates of core inflation. However, this is only one source of revision to our core indicator. There are many others not considered here, including changes to the dynamic structure of the panel over the sample period. Notwithstanding any changes of this type, it is reasonable to expect better estimates of the mean and standard deviation with which we scale the core inflation indicator as more data come to hand, which suggests that revisions to the indicator will likely be smaller in the future.

We have seen that, regardless of being subject to (sometimes substantial) revision, the core indicator compares favorably with a range of other, more standard, measures of core inflation in predicting an unobserved centered two-year moving average of headline CPI inflation—an approximation of the medium-term inflation target of the RBNZ. In the next section, we further examine the real-time properties of the inflation indicator by way of a real-time prediction experiment for the unobserved target variable.

6. A Real-Time Prediction Experiment

To simulate the predictive performance of our indicator, we estimate the dynamic factor model each quarter from $T_0 = 1999{:}Q4$ to $T_1 = 2006{:}Q2$, using exactly the information that was available in real time. In New Zealand, the CPI data are not subject to revision and do not require seasonal adjustment since they are expressed in log seasonal differences. The data are outlier-adjusted and standardized prior to estimation in each quarter.

The predictive ability of our indicator is compared with two broad categories of estimators of core inflation. The first category contains methods of estimating core inflation that utilize only CPI data: dynamic factor forecasts, time-series forecasts, conventional core inflation indicators, and static factor model forecasts. The other category contains indicators based on the real-time forecasts from a suite of models used in the policy process at the RBNZ, many of which use a much broader data set than is used to compute our core inflation indicator.

By definition, our core inflation indicator produces an estimate of the unobservable centered moving average of annual inflation. Some of the other methods of estimating core inflation we consider, however, are not tailored to this representation of inflation; forecast-based methods, e.g., do not yield estimates of inflation that are compatible with our target measure. To address this problem, we adopt a forecast-based measure of the target, which averages historical inflation and forecasts of inflation. Specifically, each quarter, inflation is forecast four periods ahead, and an estimate of the unobserved inflation target, π_t^{target}, is computed as a centered moving average of the historical and forecast data. Specifically, the forecast-based estimate of the inflation target based on model i is

$$\hat{\pi}_t^{target}(i) = \frac{1}{9}(\pi_{t-4} + \pi_{t-3} + \ldots + \pi_t + \ldots + \hat{\pi}_{t+3}(i) + \hat{\pi}_{t+4}(i)), \quad (10)$$

where all inflation observations after period t are forecasts from model i.

The estimate of the target is then compared to the ex post target, π_t^{target}; the number of prediction errors is $T_1 - T_0 - 4$. We compute each indicator's root mean-squared error (RMSE) and compare these statistics with the RMSEs from a range of other real-time estimates of the target variable. The RMSE of model i is defined as

$$RMSE(i) = \sqrt{\frac{1}{T_1 - T_0 - 4} \sum_{t=T_0}^{T_1-4} \left(\pi_t^{target} - \hat{\pi}_t^{target}(i)\right)^2}. \quad (11)$$

We use the Diebold and Mariano (1995) statistic to assess the quality of our results, testing whether the MSEs (the square of the RMSE) of the competing models are statically different from the dynamic factor model estimate of the target. The test statistic for model i is

$$STAT(i) = \frac{\bar{d}(i)}{\sqrt{\hat{V}(\bar{d}(i))}}, \quad (12)$$

where $\bar{d}(i)$ is the mean difference in squared errors $(e_t^2(i) - e_t^2)$ between model i and the dynamic factor model estimate of the target, $e_t(i)$ is the forecast error from model i, and e_t is the forecast error from the dynamic factor model.[14]

To test whether the forecasts make a useful contribution to an estimate of the inflation target estimated using the published forecasts of the RBNZ (discussed below), we also compute a variant of

[14]In practice, $STAT(i)$ is tested using a t-test with $T_1 - T_0 - 4$ degrees of freedom. The variance of the error differentials $V(\bar{d}(i))$ is adjusted for heteroskedasticity and autocorrelation using the Newey and West (1987) estimator, with a truncation lag of 3.

the Chong and Hendry (1986) encompassing test. This is based on the following forecast combination regression:

$$\pi_t^{target} = \lambda \hat{\pi}_t^{target}(i) + (1-\lambda)\hat{\pi}_t^{target}(RBNZ) + \epsilon_t, \qquad (13)$$

where $\hat{\pi}_t^{target}(i)$ is the prediction from model i, $\hat{\pi}_t^{target}(RBNZ)$ is the prediction from RBNZ's published forecasts, and ϵ_t is an idiosyncratic error term. If $\lambda = 0$, model i adds nothing to the RBNZ predictions, and if $\lambda = 1$, the RBNZ predictions add nothing to the predictions from model i.[15]

We define the following forecasting model for annual headline CPI inflation:

$$\pi_{t+h} = \beta_0 + \sum_{j=0}^{p} \beta_{1j}\pi_{t-j} + \sum_{j=0}^{k}\sum_{m=1}^{r} \beta_{2jm}x_{m,t-j} + \epsilon_{t+h}, \qquad (14)$$

where π_{t+h} is inflation h periods into the future $(h = 1, \ldots, 4)$, π_t is inflation in period t, and $x_{m,t}$ is a variable used to predict inflation.

6.1 The Models Based on Disaggregate CPI Series

6.1.1 Core Inflation Indicator

The dynamic factor model is estimated with $q = 4$ and $s = 2$, intertemporally smoothing the common component by summing waves of frequency between $[-\pi/9, \pi/9]$. The estimate of the target is simply $\hat{\chi}_t^L$.[16]

6.1.2 Dynamic Factor Model Forecasts

As with the core indicator, the dynamic factor model is estimated with $q = 4$ and $s = 2$. Here, however, we forecast the common component $\hat{\chi}_t$ itself (Forni et al. 2005 show how this done in practice—a brief summary of the procedure can be found in appendix 1).

[15] As with the Diebold and Mariano (1995) test, the standard errors of the regression are adjusted using the Newey and West (1987) estimator.

[16] Appendix 4 displays the RMSEs for different configurations of q and s.

These forecasts are concatenated with historical inflation data and averaged using equation (10) to produce $\hat{\pi}_t^{target}$.[17]

6.1.3 Standard Core Inflation Indicators

We compute two estimates of the inflation target using a range of standard core inflation indicators: trimmed-mean inflation; weighted-median inflation; the CPI excluding food, administration charges, and gasoline; median inflation; double-weighted inflation (Wynne 1997) ; and exponentially smoothed inflation (Cogley 2002). Appendix 2 describes these indicators in greater detail. The first estimate of the target is the “naive” prediction based on each indicator’s raw estimate of core inflation, $\hat{\pi}_t^{target} = \pi_t^{core}$. The second estimate uses equation (14) to forecast inflation four quarters into the future. Specifically, in (14), $m = 1$ and $x_{m,t} = \pi_t^{core}$. In each quarter, and at each forecasting horizon, the lag orders p and k are selected using the Schwartz-Bayesian information criteria (BIC), where $p = 0, \ldots, 3$ and $k = 0, \ldots, 3$ (so that the maximum number of lags of each variable is four).[18] These forecasts are then used to compute $\hat{\pi}_t^{target}$, based on equation (10). This is called the “scaled” standard core inflation indicator.

6.1.4 Time-Series Forecasts

Three time-series models are used to forecast headline inflation one year ahead: autoregressive, random walk, and random walk in mean. In each quarter, and at each forecasting horizon, the autoregressive forecast is made using equation (14), where $\beta_{2jm} = 0$, and the autoregressive order p is chosen using the BIC, with $p = 0, \ldots, 3$.

---

[17]A core inflation indicator similar to this produced the best forecasts of annual inflation in Cristadoro et al. (2005). Cristadoro et al. (2005), however, concatenate their dynamic factor model forecasts with the in-sample predicted values from the dynamic factor model. We chose to concatenate with historical data instead, because the resulting indicator produced slightly better predictive performance.

[18]The BIC for model i and horizon h is defined as $BIC^h(i) = T\ln(S/T) + 2k\ln(T)$, where T is the number of usable time-series observations, k is the number of estimated parameters, and S is the sum of squared errors of the regression, $S = \sum_{t=1}^{T}(\pi_{t+h} - \hat{\pi}_{t+h})^2$.

The random-walk forecast is made using equation (14) with $\beta_0 = 0$, $p = 0$, $\beta_{10} = 1$, and $\beta_{2jm} = 0$, and the random-walk-in-mean forecast takes the mean of the series as the forecast at all horizons (all terms in equation (14) are set to 0 except for β_0). Once the forecasts are made, they are concatenated with historical data to yield a prediction for the inflation target, based on equation (10).

6.1.5 Pooling Regressions

We make inflation forecasts from one to four quarters ahead using equation (14) and each of the i components of the CPI, where $m = 1$, $x_{m,t} = \pi_{i,t}$, and $i = 1, \ldots, n$. The lag orders p and k are selected using the BIC at each horizon h, where $p = 0, \ldots, 3$ and $k = 0, \ldots, 3$. The resulting 227 forecasts for $h = 1, \ldots, 4$ are then weighted using three methods: a simple average (equal weights), a BIC-weighted average, and an expenditure-weighted average. To be explicit, at each forecasting horizon, let the weighted-average forecast for headline CPI inflation be

$$\hat{\pi}_{t+h} = \sum_{i=1}^{n} \Omega^h(i)\hat{\pi}_{t+h}(i), \tag{15}$$

where $\Omega^h(i)$ is the weight attached to model i's forecast at horizon h and $\hat{\pi}_{t+h}(i)$ is the forecast from model i at horizon h. The average forecast sets the weights equal across the n forecasts so that $\Omega^h(i) = 1/n$ for all h. The BIC-weighted average weights each forecast according to the fit that the underlying estimated model had to the data, weighting models with better fit more highly. Here, the weight attached to each forecast i at each horizon h is calculated as

$$\Omega^h(i) = \frac{\exp(-0.5BIC^h(i))}{\sum_{j=1}^{n} \exp(-0.5BIC^h(j))}, \tag{16}$$

where $BIC^h(i)$ is the minimum BIC for model i at horizon h over models estimated with $p = 0, \ldots, 3$ and $k = 0, \ldots, 3$.[19]

[19]The BIC-weighted-average forecast is approximately the Bayesian-model average forecast arising from equal model priors and diffuse coefficient priors.

The expenditure-weighted average weights the forecast associated with CPI component i with that component's expenditure weight, w_i, across all horizons h, $\Omega^h(i) = w_i$.

The three sets of pooled forecasts are concatenated with historical data and averaged using equation (10) to yield predictions for the target measure of inflation.

6.1.6 Static Factor Model Forecasts

The static factor model forecasts for headline inflation are made using the first r static principal components estimated from the panel of CPI data, $f_{1,t}, \ldots, f_{r,t}$. For each quarter, and each forecasting horizon h, equation (14) is estimated with $x_{m,t} = f_{m,t}$, where the orders of p, k, and r are selected using the BIC, with $r = 1, \ldots, 4$; $p = 0, \ldots, 3$; and $k = 0, \ldots, 3$. These forecasts are then concatenated with historical data and averaged using equation (10).

6.2 Core Indicators Based on Some of the Forecasting Models Used at the RBNZ

A prediction of the inflation target at time t is computed using these forecasts in the same way as for the time-series forecasts from above: the real-time estimate of core inflation at time t is $\hat{\pi}_t^{target}$, constructed using equation (10), where all observations up to and including period t are actual data and all observations beyond period t are forecasts. Some of these models use a much broader information set than the estimates of the target rate of inflation described above (see appendix 3 for a detailed description of these forecasts). The forecasting models are as follows:

- the real-time published forecasts of the RBNZ
- an average of real-time private-sector forecasts
- a Bayesian VAR forecast (BVAR)
- a BIC-weighted VAR forecast
- a factor-model forecast
- an indicator forecast

The published forecasts and the average of private-sector forecasts require further discussion, as they can be characterized as

Figure 5. Calender of CPI Publication and Forecasts

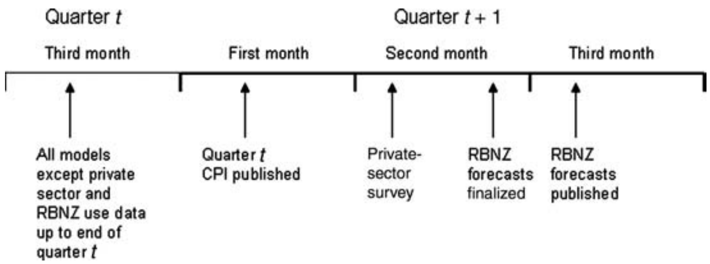

being judgment-based forecasts—unlike the other forecasts discussed so far.

With the exception of the core indicators based on these two forecasts, which can incorporate information dated in period $t+1$, all other core indicators discussed so far use information dated up to the end of period t—the same period for which the latest CPI data are available. This can be better illustrated in figure 5. The majority of the estimators of core inflation are made with information up to the end of quarter t; estimates of $\hat{\pi}_t^{target}$ can be constructed using most models with the arrival of period t's CPI data, published in the middle of the first month of period $t+1$. However, the RBNZ and private-sector forecasts incorporate information up to the second month of period $t+1$, using more up-to-date information than the other models.

7. Empirical Results

The results from the real-time prediction experiment are displayed in table 4.

Looking first at the indicators based on the disaggregate CPI data, the core indicator and the scaled exponentially smoothed indicator compare favorably with the other estimates of core inflation, producing the lowest RMSEs (around 0.26 percent) of the price-based indicators examined. The other indicators perform much worse than these two indicators, with several indicators having predictive performance significantly worse than the core indicator (at the 10 percent level). Aside from the core indicator and the exponentially smoothed indicator, the standard core inflation indicators perform

Table 4. Prediction Statistics

Indicators Based on Disaggregate Prices		RMSE		Encompassing (λ)	
Dynamic Factor Models	Core Indicator		0.257		0.270*
	Forecast		0.310		0.114*
Time-Series Models	AR		0.323*		0.087*
	Random Walk		0.342		0.051
	Random Walk (Mean)		0.342		0.108*
		Naive	*Scaled*	*Naive*	*Scaled*
Core Inflation Indicators	Weighted Median	0.630*	0.390	−0.052	−0.058
	Trimmed Mean	0.500*	0.791	−0.049	−0.033
	CPI exl. Food, Admin., and Gasoline	1.011*	0.358	−0.090	−0.004
	Median	0.640*	0.408*	−0.015	0.049
	Double Weighted	0.399	0.328	−0.033	0.060
	Exponentially Smoothed	0.310	0.258	0.146*	0.269*
Pooling Regressions	Average		0.325*		0.085*
	BIC Weighted		0.325*		0.085*
	Expenditure Weighted		0.332*		0.067
Static Factor Model	Static Factor Model		0.337		0.034
Indicators Based on RBNZ Forecasts					
RBNZ Models	Published		0.186		—
	External Average		0.123*		1.002*
	BVAR		0.291		0.065
	VAR		0.290		0.165*
	Factor Model		0.264		0.123
	Indicator Median		0.259		0.165*

Note: * denotes statistical significance at the 10 percent level. RMSE: A Diebold and Mariano (1995) test is used, testing whether the MSE of model i is statistically different from the core indicator. Encompassing (λ): A Chong and Hendry (1986) test is used, testing whether there is predictive content in model i over and above the predictive content of the indicator based on the published forecasts of the RBNZ, $\lambda \neq 0$.

particularly poorly, though (generally speaking) the performance of these indicators improves when scaled (used in a forecasting regression). The remainder of the models based on the disaggregate CPI data have comparable predictive ability, each with RMSEs of between 0.30 and 0.35 percent.

Overall, the indicators based on the published forecasts of the RBNZ and the private-sector forecasts are the best predictors of the target variable, followed by the core inflation indicator and the scaled exponentially smoothed indicator: the indicator based on the private-sector forecasts is particularly good and statistically better than the core inflation indicator. The relatively good performance of the RBNZ and private-sector indicators comes as no surprise, given that they are based on an information set that is both broader and more up-to-date than is used by the other indicators. Indeed, aside from the core inflation indicator and the scaled exponentially smoothed indicator, the indicators based on disaggregate CPI data are worse than the indicators based on RBNZ forecasts, perhaps reflecting the broader range of information that is incorporated in these indicators (e.g., real variables are incorporated into all RBNZ models).

Despite being altogether worse than the published forecasts of the RBNZ according to RMSE, some of the core indicators are able to improve the predictive performance of the RBNZ indicator. Looking at weights attached to each of the indicators relative to the RBNZ indicator from the encompassing regression λ, we find that the largest weights are attached to the external average, the core indicator, and the scaled exponentially smoothed indicator. Moreover, because the core inflation indicator and the scaled exponentially smoothed indicator can be computed as soon as the CPI data are available, these indicators are more timely than the indicator based on published RBNZ forecasts, which are finalized more than one month later.

8. Summary and Conclusion

This paper introduced a new indicator of core inflation for New Zealand, estimated using a dynamic factor model. Defining core inflation to be consistent with the medium-term inflation target of the RBNZ, we found that our indicator produced relatively accurate estimates of core inflation, compared with a range of other indicators of core inflation. Estimates of core inflation derived from the RBNZ's published forecasts and an average of private-sector forecasts produce more accurate measures of core inflation than our indicator. However, our core indicator is more timely and can be computed as soon as the CPI data are published—around a month before the RBNZ forecasts are finalized.

Appendix 1. Estimation of Static and Dynamic Factor Models

Static Factor Model

Assume there are T time-series observations for N cross-section units denoted x_{it}, where $i = 1, \ldots, N$ and $t = 1, \ldots, T$. We let X be the $T \times N$ matrix of observations, x_t is a row denoting all N observations at time t, and x_j is a column vector denoting all T observations for cross-section unit j.

The static factor model is estimated using an eigenvector-eigenvalue decomposition of the sample covariance matrix—principal components. Specifically, let V be the eigenvectors corresponding to the r largest eigenvalues of the $N \times N$ matrix $\hat{\Gamma} = \frac{1}{T}\sum_{t=1}^{T} x_t x_t'$. The static principal components estimator yields

$$\hat{F} = XV, \quad \hat{\Lambda} = V, \quad \text{and} \quad \hat{\chi} = \hat{F}\hat{\Lambda}', \tag{17}$$

where $\hat{F}$ is a $T \times r$ matrix of common factors, $\hat{\Lambda}$ is an $r \times N$ matrix of factor loadings, and $\hat{\chi}$ is a $T \times N$ matrix of common components.

Dynamic Factor Model

The dynamic factor model is estimated using an eigenvalue decomposition of the spectrum smoothed over a range of frequencies—dynamic principal components. Estimation proceeds as follows:

(i) Estimate the spectral-density matrix of X, using a Bartlett lag window of size M; i.e., compute the autocovariance matrices $\hat{\Gamma}_X(k) = \frac{1}{T}\sum_{t=k+1}^{T} x' x_{t-k}$, where $k = 1, \ldots, M$, multiply them by the weights $\omega_k = 1 - \frac{|k|}{M+1}$, and apply the discrete Fourier transform:

$$\hat{\Sigma}_X(\omega_m) = \frac{1}{2\pi} \sum_{k=-M}^{M} \omega_k \Gamma(k) e^{-i\omega_m k}. \tag{18}$$

(ii) For each frequency, $\omega_m = \frac{2\pi m}{2M+1}$, $m = -M, \ldots, M$, let $D_q(\omega_m)$ be the diagonal matrix with the q largest eigenvalues of $\hat{\Sigma}_X(\omega_m)$ on the diagonal, and let $U_q(\omega_m)$ be the associated matrix of eigenvectors. Use the discrete inverse Fourier transform on $\hat{\Sigma}_\chi(\omega_m) = U_q(\omega_m) D_q(\omega_m) U_q(\omega_m)'$ to obtain

$$\hat{\Gamma}_\chi(k) = \frac{2\pi}{2M+1} \sum_{m=-M}^{M} \hat{\Sigma}_\chi(\omega_m) e^{i\omega_m k}. \tag{19}$$

(iii) The covariance matrix of the idiosyncratic part is then estimated as a residual:

$$\hat{\Gamma}_\epsilon(k) = \hat{\Gamma}_X(k) - \hat{\Gamma}_\chi(k). \tag{20}$$

(iv) Let Z be the r generalized eigenvectors (with eigenvalues in descending order) of $\hat{\Gamma}_\chi(0)$ with respect to $\hat{\Gamma}_\epsilon(0)$ with the normalization that $Z_j \hat{\Gamma}_\epsilon(0) Z_i' = 1$ if $i = j$ and 0 otherwise. This is generalized principal components, which is essentially weighted principal components, where each principal component is weighted by the inverse of the size of its idiosyncratic noise.

(v) The estimated dynamic factors and common components are, respectively,

$$\hat{F} = XZ \quad \text{and} \quad \hat{\chi}_{t+h} = \hat{\Gamma}_\chi(h) Z (Z' \hat{\Gamma}_X(0) Z)^{-1} Z' x_t \tag{21}$$

for $h = 0, \ldots, M$.

The covariance matrix of the long-run common components, $\hat{\Gamma}_\chi^L(\omega_m)$, can be estimated by applying the inverse Fourier transform (step ii above) to the spectral-density matrix (18) over the frequency band of interest $[-\pi/\tau, \pi/\tau]$ (where τ denotes the periodicity of the shortest cycle allowed). The estimated long-run common components are then

$$\tilde{\chi}^L_{t+h} = \hat{\Gamma}^L_\chi(h) Z(Z'\hat{\Gamma}_X(0)Z)^{-1} Z' x_t. \tag{22}$$

Appendix 2. The Standard Core Inflation Measures

Let headline CPI inflation be the weighted average of the inflation rates of n disaggregate subgroups of the CPI:

$$\pi_t = \sum_{i=1}^{n} w_{it}\pi_{it}, \tag{23}$$

where π_t is headline CPI inflation, π_{it} is inflation in the i-th component of the CPI, and w_{it} is the expenditure weight of the i-th component.

Trimmed Mean

Each quarter, the trimmed mean is calculated as follows:

(i) Compute the annual percentage change in each of the i components of the CPI.

(ii) Sort the resulting series from smallest to largest, along with their associated weights w_i.

(iii) Compute the cumulative sum of the weights of the ordered prices.

(iv) Exclude the series with cumulative weights either less than 5 percent or with cumulative weights greater than 95 percent.

(v) Compute the trimmed-mean inflation rate as

$$[1/\sum_{i=first}^{last} w_i] \sum_{i=first}^{last} w_i \pi_{it}, \tag{24}$$

where *first* and *last* denote the first and last CPI components in the truncated list of ordered price changes.

Weighted Median

The weighted-median inflation rate is calculated using steps (i)–(iii) above. The weighted median is then the first percentage change in price where the cumulative weight is greater than or equal to 50 percent.

Median

Median inflation is simply the median rate of inflation from the n components of the CPI.

CPI Excluding Food, Administration Charges, and Gasoline

This measure is computed using

$$\sum_{i=\Omega}^{n} w_i \pi_{it}, \tag{25}$$

where Ω is the number of CPI components categorized as food, administration charges, and gasoline, and the weights w_i have been adjusted to sum to 1.

Double-Weighted Inflation

This measure of core inflation is computed by each of the components of the CPI—a weight inversely proportional to its variability. Double-weighted inflation is calculated as

$$\frac{\sum_{i=1}^{n} w_i v_i \pi_{it}}{\sum_{i=1}^{n} w_i v_i}, \quad \text{where} \quad v_i = \frac{1/\sigma_{it}}{\sum_{i=1}^{n} 1/\sigma_{it}}, \tag{26}$$

and where σ_{it} is the standard deviation of inflation in component i relative to headline CPI inflation, $(\pi_{it} - \pi_t)$.

Exponentially Smoothed Inflation

Exponentially smoothed inflation is

$$\phi \sum_{i=1}^{j} (1-\phi)^j \pi_{t-i}, \tag{27}$$

where ϕ, the expectations adjustment parameter, is set to 0.125, as in Cogley (2002).

Appendix 3. Some Forecasts Used at the RBNZ

The Published Forecasts

These are the real-time forecasts published in the Reserve Bank's quarterly *Monetary Policy Statement* (*MPS*). The forecasts are a combination of model-based forecasts and judgment. Broadly speaking, the near-term forecasts can be characterized as being judgment based and indicator based. The longer-term forecasts, on the other hand, are made with the help of a large-scale macroeconomic model, the Forecasting and Policy System (FPS).

External Average

The external-average forecast is the average forecast from a group of private- and public-sector institutions. These forecasts are attained by an informal survey conducted by the RBNZ in the middle of each quarter.

BVAR

The BVAR has a Minnesota prior, shrinking the VAR coefficients toward univariate unit roots (Doan, Litterman, and Sims 1984), and contains five endogenous variables n (GDP, CPI, ninety-day rates, the trade-weighted index [TWI], and the terms of trade [TOT]). The VAR can be written as

$$y_t = \Phi_0 + \Phi_1 y_{t-1} + \ldots + \Phi_p y_{p-t} + u_t, \tag{28}$$

where u_t is a vector of one-step-ahead forecast errors and y_t is $n \times 1$. We assume that the innovations in (28) have a multivariate normal distribution $N(0, \Sigma_u)$. The VAR can also be expressed in matrix form as

$$Y = X\Phi + U, \tag{29}$$

where Y is a $T \times n$ matrix with rows y_t'; X is a $T \times k$ matrix with rows $x_t' = [1, y_{t-1}', \ldots, y_{t-p}']$, where $k = 1 + np$; U is a $T \times n$ matrix with rows u_t'; and $\Phi = [\Phi_0, \dots, \Phi_p]'$. The Minnesota prior is implemented as in Del Negro and Schorfheide (2004):

$$\bar{\Phi} = (I \otimes (X'X) + \iota H^{-1})^{-1}(\text{vec}(X'Y) + \iota H^{-1}\Phi), \tag{30}$$

where the ι denotes the weight of the Minnesota prior, Φ is the prior mean, and H is the prior tightness. The values for Φ and H are the same as in Doan, Litterman, and Sims (1984). All variables, except ninety-day interest rates, enter the VAR in log-differences. Thus, to be consistent with the Minnesota prior, the prior for the mean of the first lag of all growth variables is 0; the prior for the mean of the first lag of the ninety-day interest rate is 1. The prior is augmented with a proper inverse Wishart prior for Σ_u. The overall tightness of the prior is determined by the hyperparameter ι. Following Del Negro and Schorfheide (2004), this parameter is chosen ex ante to maximize the marginal data density; i.e., each quarter,

$$\hat{\iota} = \arg\max_{\iota} p_{\iota}(Y). \tag{31}$$

VAR

The VAR forecast is a BIC-weighted average of VAR forecasts from over 500 different VAR specifications. The VAR specifications differ in three respects: (i) they have different variables in them, (ii) they embody a different number of lags (between one and three lags for each variable), and (iii) some of the VARs are in difference form while others use levels data. The base VAR uses GDP, CPI, the ninety-day bank-bill rate and the (real/nominal) TWI. This base model is then augmented with a world sector, commodity prices, migration and housing, and hours worked. The world sector is sometimes represented using U.S. GDP, U.S. CPI, and U.S. interest rates, and sometimes using the inputs that feed into FPS—world GDP (defined to be a twelve-country weighted average), short world interest rates, and world CPI. The commodity prices also take various forms: the ANZ SDR commodity price index, (import and) export prices denominated in New Zealand dollars, and a U.S.-dollar oil price. Because of limitations in the data, not all of these series can be incorporated in the model simultaneously, which is why the VAR forecast is obtained over an average of models.

Factor Model

The factor-model forecast is derived from a data set comprising almost 400 macroeconomic series. All series in the data set are seasonally adjusted using X12 (additive). The series are then transformed to account for stochastic and deterministic trends; the I(1) series are logged and then differenced, and the I(0) series are left as levels. See Matheson (2006) for a more detailed description of these data.

We estimate static factors from this data set using the same method as described above for the static factor models and use them in the following h-step-ahead regression:

$$\pi_{t+h} = \phi + \beta(L)f_t + \gamma(L)\pi_t + e_{t+h}, \tag{32}$$

where $\pi_{t+h} = \ln(p_{t+h}/p_t)$ is h-period inflation in CPI P_t and $\pi_t = \ln(P_t/P_{t-1})$, ϕ is a constant, $\beta(L)$ and $\gamma(L)$ are lag polynomials, f_t is a vector of factors (estimated using static principal components), and e_{t+h} is an error term.

The algorithm that is used to produce factor-model forecasts each quarter tailors the raw data set X to the task of forecasting inflation at different horizons: note that the factors from X are the same regardless of the horizon being forecast. Following the method described in Matheson (2006), the data set X is reduced by removing those series that do not have a high correlation with inflation at the horizon being forecast. Essentially, we regress the inflation rate that we are trying to forecast on each series in the data set. We then rank the resulting R-squareds (coefficients of determination) and remove those series that are least informative—keeping a proportion θ of the series at each horizon h. The factors are then extracted from this reduced data set X^*.

Due to uncertainty about the particular cut-off criterion to choose, we average the factor-model forecasts over a variety of different criteria: $\theta = (5, 10, 20, 50, 100)$.[20] Aside from the size of the data set, the five factor-model forecasts are constructed in the same way using (32). We estimate (32) using the factor with the largest

[20]Note that when $\theta = 100$ all of the data are used to extract factors, as in Stock and Watson (2002).

eigenvalue (the principal component) and do not allow lags of the factor ($\beta(L) = \beta$). The number of lags of inflation that are included at each horizon is chosen using the BIC, with lags varying from zero to four.

Indicator Forecast

The indicator median forecast uses the same data set and forecasting methodology as the factor-model forecast. Bivariate regressions are run for each series in the data set, where series x_i replaces f_t in (32). The number of lags of each indicator is allowed to vary from one to four and the number of lags of inflation is allowed to vary from zero to four, with all lags selected with the BIC. The BICs from these regressions are then ranked. The indicator median forecast is the median forecast from the top 10 percent of the ranked bivariate regressions.

Appendix 4. Core Indicator: RMSE for Different Configurations of *q* and *s*

	s	**0**	**1**	**2**	**3**	**4**
q						
1		0.465	0.493	0.461	0.471	0.478
2		0.432	0.370	0.364	0.362	0.368
3		0.340	0.320	0.327	0.319	0.320
4		0.241	0.257	***0.257***	0.274	0.290
5		0.235	0.265	0.279	0.293	0.294
6		0.254	0.264	0.290	0.300	0.302

References

Altissimo, F., A. Bassanetti, R. Cristadoro, M. Forni, M. Hallin, M. Lippi, L. Reichlin, and G. Veronese. 2001. “EuroCOIN: A Real Time Coincident Indicator of the Euro Area Business Cycle.” CEPR Discussion Paper No. 3108.

Amstad, M., and A. Fischer. 2004. “Sequential Information Flow and Real-Time Diagnosis of Swiss Inflation: Intra-Monthly DCF Estimates for a Low-Inflation Environment. CEPR Discussion Paper No. 4627.

Amstad, M., and S. Potter. 2007. “Real Time Underlying Inflation Gauges for Monetary Policy Makers.” Memo, Federal Reserve Bank of New York.

Bai, J., and S. Ng. 2002. “Determining the Number of Factors in Approximate Factor Models.” *Econometrica* 80 (1): 191–221.

Bryan, M. F., and S. G. Cecchetti. 1994. “Measuring Core Inflation.” In *Monetary Policy*, ed. N. G. Mankiw, 67–123. Chicago: University of Chicago Press.

Chong, Y. Y., and D. F. Hendry. 1986. “Econometric Evaluation of Linear Macro-Economic Models. *Review of Economic Studies* 53 (4): 671–90.

Christiano, L. J., M. Eichenbaum, and C. Evans. 2005. “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy.” *Journal of Political Economy* 113 (1): 1–45.

Cogley, T. 2002. “A Simple Adaptive Measure of Core Inflation.” *Journal of Money, Credit and Banking* 34 (1): 94–113.

Cristadoro, R., M. Forni, L. Reichlin, and G. Veronese. 2005. “A Core Inflation Indicator for the Euro Area.” *Journal of Money, Credit and Banking* 37 (3): 539–60.

Del Negro, M., and F. Schorfheide. 2004. “Priors from General Equilibrium Models for VARs.” *International Economic Review* 45 (2): 643–73.

Diebold, F. X., and R. S. Mariano. 1995. “Comparing Predictive Accuracy.” *Journal of Business and Economic Statistics* 13 (3): 253–63.

Doan, T., R. Litterman, and C. Sims. 1984. “Forecasting and Conditional Projections Using Realistic Prior Distributions.” *Econometric Reviews* 3 (1): 1–100.

Forni, M., D. Giannone, M. Lippi, and L. Reichlin. 2007. “Opening the Black Box: Structural Factor Models with Large Cross-Sections.” ECB Working Paper No. 712.

Forni, M., M. Hallin, M. Lippi, and L. Reichlin. 2000. “The Generalized Dynamic Factor Model: Identification and Estimation.” *Review of Economics and Statistics* 82 (4): 540–54.

———. 2005. "The Generalized Dynamic Factor Model: One-Sided Estimation and Forecasting." *Journal of the American Statistical Association* 100 (471): 830–40.

Giannone, D., and T. D. Matheson. 2006. "A New Core Inflation Indicator for New Zealand." Reserve Bank of New Zealand Discussion Paper No. 2006/10.

Harding, D., and A. Pagan 2002. "Dissecting the Cycle: A Methodological Investigation." *Journal of Monetary Economics* 49 (2): 365–81.

Matheson, T. D. 2006. "Factor Model Forecasts for New Zealand." *International Journal of Central Banking* 2 (2): 169–237.

Newey, W. K., and K. D. West. 1987. "A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–708.

Stock, J. H., and M. W. Watson. 2002. "Macroeconomic Forecasting Using Diffusion Indexes." *Journal of Business and Economic Statistics* 20 (2): 147–62.

Wynne, M. 1997. "Commentary for 'Measuring Short-Run Inflation for Central Bankers.'" *Review* (Federal Reserve Bank of St. Louis) 79 (3): 161–67.