RESERVE BANK OF NEW ZEALAND

DP2010/13

What drives core inflation? A dynamic factor model analysis of tradable and nontradable prices

Michael Kirker

December 2010

JEL classification: C11, C38, E31, E52

www.rbnz.govt.nz/research/discusspapers/

Discussion Paper Series

ISSN 1177-7567

DP2010/13

What drives core inflation? A dynamic factor model analysis of tradable and nontradable prices*

Michael Kirker†

Abstract

I develop a new estimate of core inflation for New Zealand and Australia based on a dynamic factor model. By using an over-identification restriction, the factors of the model are classified as tradable and nontradable factors. This innovation allows us to examine the relative contributions of tradable and nontradable prices towards core inflation. The results show that core inflation in both countries is primarily driven by the nontradable factor. The nontradable factor also explains significantly more of the variance in headline inflation relative to the tradable factor. Finally, both the tradable and nontradable factors show similar profiles across both countries suggesting common drivers.

* The Reserve Bank of New Zealand's discussion paper series is externally refereed. The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Reserve Bank of New Zealand. The author would like to thank Haroon Mumtaz, Özer Karagedikli, and an anonymous referee for their helpful comments and suggestions.

† Address: Economics Department, Reserve Bank of New Zealand, 2 The Terrace, PO Box 2498, Wellington, New Zealand. *email address:* Michael.Kirker@rbnz.govt.nz

ISSN 1177-7567

1 Introduction

The primary objective for most central banks is price stability. Usually price stability is expressed as a target for a measure of inflation such as the Consumer Price Index (CPI). However, as discussed by Bryan and Cecchetti (1993), the use of the CPI as a measure of inflation has two main drawbacks. First, the CPI is affected by temporary or short-lived sectorial shocks that the policymakers should 'look through' given the lags in the monetary policy transmission mechanism. And second, the CPI can be biased by its fixed expenditure weights and measurement issues.

To assist them in their analysis of the underlying inflation pressures, many central banks also look at various estimates of core inflation alongside the CPI measure when formulating policy decisions. Despite the widespread use of core inflation measures among central banks, there is no consensus on either the exact definition, or the most appropriate approach to measure core inflation (Wynne 2008). For example, Bryan and Cecchetti (1994) define core inflation as the inflation that causes money growth, and examine how a number of common core measures compare to this defintion; Quah and Vahey (1995) define core inflation in terms of inflation that has no long-run impact on output, and estimate it using a VAR system; and Cogley (2002) defines core in terms of responding to changes in mean inflation and estimates it using a constant gain update of mean inflation.

One approach to estimating core inflation that has been popular in the literature is the dynamic factor model (see Kapetanios 2004, Cristadoro et al 2005, and Giannone and Matheson 2007 as examples). This statistical approach combines cross-sectional as well as time-series information from a panel of individual prices to derive a measure of core inflation. In dynamic factor models, core inflation reflects the inflation component that is common to all price changes in the panel of data.[1] However, one of the major drawbacks of using a dynamic factor model to estimate core inflation is the lack of interpretation that can be placed upon the factors that drive core inflation.

In this paper I propose a new core inflation estimate based on a dynamic factor model that employs an over-identification restriction to interpret the factors as either tradable or nontradable factors. The advantage of this approach is it allows us to examine the relative importance of co-movement in tradable and nontradable prices in determining core inflation. This dynamic

[1] Reis and Watson (2010) extend the dynamic factor model approach to core inflation further to estimate the what they term "pure inflation" — the component of core inflation that causes and equal effect on all prices.

factor model is then applied to disaggregated CPI expenditure class data for both New Zealand and Australian, two typical small open economies.[2]

The main results show that the estimate of core inflation for both New Zealand and Australia is primarily driven by the nontradable factor. In addition, the nontradable factor explains significantly more of the variance in CPI inflation than the tradable factor. When comparing the tradable and nontradable factors across both countries, there are broad similarities in their time-series profiles. This is supportive of the idea of a common international driver of tradable prices and common business cycles in both countries.

In addition to these results, I find the estimate of core inflation for New Zealand is similar to other core inflation estimates published regularly by Reserve Bank of New Zealand, and for Australia, the estimate of core inflation is noticeably smoother than other core measures. The real-time estimate of core inflation produced by the dynamic factor model is broadly similar to the final full sample (ex-post) estimate for both countries.

The remainder of this paper is organised as follows. Section 2 outlines the New Zealand and Australian panels of data used in the model. Section 3 gives details on the estimation method and model. Section 4 discusses the results. And section 5 provides a brief conclusion.

2 Data

For both New Zealand and Australia, I focus on sample periods after the adoption of inflation targeting, when inflation was relatively low and stable. For New Zealand this corresponds to the period 1992Q1 to 2010Q3 period, and for Australia to the 1993Q2 to 2010Q3 period.[3]

Each panel of data consists of a larger number of individual, disaggregated

[2] An alternative approach to applying some interpretation to the factors would be to use a dynamic hierarchical factor model similar to Moench et al (2009). This approach would be well suited to a country that has a large number of sub-categories in the CPI basket. However, because I wish to focus on only the tradable nontradable distinction, and the number of sub-categories in the New Zealand and Australian CPI baskets are limited, I do not employ the more detailed structure of a dynamic hierarchical factor model in this paper.

[3] Although New Zealand adopted inflation targeting at the end of 1989, the period between 1989 and 1992 was characterised by a large recession and an inflation rate that was still declining from the high levels experienced in the 1980s. Therefore, I exclude this period from the sample range.

prices indices taken from the expenditure classes that make up the CPI basket. The indices of each series are converted to a quarterly percentage change, and then standardised to have a mean of zero and standard deviation of unity.

The series used in the New Zealand and Australian data panels are briefly discussed below.[4]

2.1 New Zealand data

The CPI basket in New Zealand comprises of 105 different expenditure classes of goods and services. From this panel of individual classes, I remove nine series that do not span the entire length of the sample. To the panel I add the headline CPI, tradable CPI, and nontradable CPI. Including headline CPI in the panel will enable the computation of core inflation. And including the tradable CPI and nontradable CPI series allows us to identify the sign of the latent factors (discussed in the section 3).

As a result of these adjustments, the New Zealand panel of data consists of 99 series (96 being individual CPI components). The panel of data is fairly evenly split between tradable and nontradable series. According to the official classifications by Statistics New Zealand, the panel contains 46 CPI expenditure classes that are classified as tradable price series, 39 that are classified as nontradable price series, and 11 that are classified as both tradable and nontradable price series.

2.2 Australian data

The Australian panel of data is constructed using the same approach as the New Zealand panel. The panel consists of 66 series: headline CPI, tradable CPI, nontradable CPI, and 63 individual CPI expenditure class components. The CPI expenditure class series include 37 tradable and 26 nontradable series as classified by the Australian Bureau of Statistics.

During 2000, the Australian Government introduced a Goods and Services Tax (GST) to replace a range of other taxes and duties. The majority of this tax change occurred during 2000Q3, and resulted in a one-off price rise in many of the expenditure class series in the panel. Without adjusting the panel of data for the introduction of GST, a dynamic factor model would

[4] A complete list of the series included and excluded from each panel can be found in appendix A.

interpret the GST effect as a spike in core inflation. This has the potential to bias the model's estimates of core inflation in other periods. Therefore, I use a panel of data that has been adjusted to remove the estimated effect of the GST introduction. In addition, the headline CPI series has been adjusted to exclude interest charges prior to 1998, making the composition of headline CPI more comparable over the whole sample.

Estimating the dynamic factor model on this adjusted panel of data has the advantage making the measure of core inflation from the dynamic factor model more comparable with the trimmed mean and weighted median measures of core inflation published by the Reserve Bank of Australia. Both of these other core measures are computed on the exact same adjusted data.

3 The Dynamic Factor Model

To construct a measure of core inflation, I assume that quarterly CPI inflation (π_t) can can be decomposed into two orthogonal components, a core inflation component (π_t^{core}) and a non-core or temporary component (π_t^{nc}):

$$\pi_t \equiv \pi_t^{core} + \pi_t^{nc}. \tag{1}$$

Such a decomposition of inflation lends itself naturally to the dynamic factor model framework. Below I discuss in more detail the dynamic factor model and how it relates to equation 1, the Bayesian estimation process, and how to construct an estimate of core inflation from the dynamic factor model results.

3.1 Methodology

Let the panel of inflation data be denoted as $\Pi_t = (\pi_{1,t}, \pi_{2,t}, \ldots \pi_{n,t})'$, where for convenience $\pi_{1,t}$ denotes headline CPI inflation, $\pi_{2,t}$ denotes tradable CPI inflation, $\pi_{3,t}$ denotes nontradable CPI inflation, and $\pi_{i,t}$ for $i = (4, \ldots, n)$ denotes the other inflation rates in the panel.

For our panel of data, the general form of a dynamic factor model can be denoted as:

$$\pi_{i,t} = \boldsymbol{\beta}_{\mathbf{i}}'\mathbf{F_t} + \nu_{i,t} \tag{2}$$

where $\mathbf{F_t}$ is a vector of (orthogonal) common factors that capture the co-movement across the entire panel, $\boldsymbol{\beta}_{\mathbf{i}}$ denotes the vector of factor loadings between series $\pi_{i,t}$ and the common factors, and $\nu_{i,t}$ denotes the idiosyncratic term associated with series $\pi_{i,t}$ (the part not explained by the common co-movement in the panel).

Equating equation 1 with equation 2 for CPI inflation ($i = 1$), we can see that core inflation can be expressed as the factor loadings for CPI inflation multiplied by the common factors ($\pi_t^{core} \equiv \boldsymbol{\beta_1}\mathbf{F_t}$), and non-core inflation can be expressed as the idiosyncratic term ($\pi_t^{nc} \equiv \nu_{1,t}$).

For the analysis in this paper, I choose to model two different types of dynamic factors within the model. The first type of factor is a tradable factor that captures co-movement in tradable prices, and the second type is a non-tradable factor that captures co-movement in nontradable prices.

Therefore, for the i^{th} inflation series in the panel of CPI data ($\pi_{i,t}$) the functional form for the dynamic factor model can be expressed as:

$$\pi_{i,t} = \boldsymbol{\beta_i^{tr}}\mathbf{F_t^{tr}} + \boldsymbol{\beta_i^{nt}}\mathbf{F_t^{nt}} + \nu_{i,t} \tag{3}$$

where $\mathbf{F_t^{tr}} = [F_{1,t}^{tr}, \ldots F_{n,t}^{tr}]'$ is a vector of n number of tradable factors, and $\mathbf{F_t^{nt}} = [F_{n+1,t}^{nt}, \ldots F_{n+m,t}^{nt}]'$ is a vector of m number of nontradable factors. $\boldsymbol{\beta_i^{tr}}$ and $\boldsymbol{\beta_i^{nt}}$ represent the vectors of tradable and nontradable factor loadings respectively. And $\nu_{i,t}$ is the idiosyncratic error term.

Each of the tradable and nontradable factors are assumed to follow an autoregressive process of order two:

$$F_{a,t}^k = \sum_{j=1}^{2} \rho_{a,j}^k F_{a,t-j}^k + \varepsilon_{a,t} \tag{4}$$

where k denotes either tradable or nontradable ($k = \{tr, nt\}$), a denotes an index number for each factor ($a = 1, \ldots, n+m$), and $\varepsilon_{a,t}$ is a white noise i.i.d. shock.

I also allow for the possibility of serial correlation in the idiosyncratic error term of the dynamic factor model (equation 3) by assuming $\nu_{i,t}$ follows an autoregressive process of order one:

$$\nu_{i,t} = \alpha_i \nu_{i,t-1} + \eta_{i,t} \tag{5}$$

where $\eta_{i,t}$ is a white noise i.i.d. shock.

3.2 Identification

To identify the latent factors in the dynamic factor model as a tradable and nontradable factors, I impose an over-identifying restriction on the factor loadings.[5]

In the over-identifying restriction, an inflation series $\pi_{i,t}$ is loaded onto the tradable (nontradable) factors, if and only if, it is classified as a tradable (nontradable) expenditure class by the national statistical agency.

This is achieved by imposing the condition $\boldsymbol{\beta_i^k} = \mathbf{0}$ if series i does not belong to category k (where $k = \{tr, nt\}$). Under such a condition, the tradable (nontradable) factors are independent of the common co-movement in the nontradable (tradable) CPI inflation series.

In the case of CPI inflation ($\pi_{1,t}$), the series is loaded onto both the tradable and nontradable factors. The tradable CPI series is only loaded onto the tradable factor ($\boldsymbol{\beta_2^{nt}} = \mathbf{0}$), and the nontradable CPI series is loaded only onto the nontradable factor ($\boldsymbol{\beta_3^{tr}} = \mathbf{0}$).

Imposing the over-identification restriction is not sufficient to uniquely identify the sign and size of the factor loadings and latent factors. This problem is formally known as the indeterminacy of factor rotation.[6] Thus we must impose two further identification restrictions for the scale and sign.

To identify the scale of the factor loadings and latent factors, I impose that the variance of each shock in equation 4 is unity (this restriction is also employed by Kose et al 2003 and many others in the literature). To identify the sign of the factor loadings, I impose that the sign of the factor loading for the tradable CPI series onto the tradable factors must be positive (i.e. $\boldsymbol{\beta_2^{tr}} > \mathbf{0}$), and the sign of the factor loading for the nontradables CPI series onto the nontradable factors must be positive (i.e. $\boldsymbol{\beta_3^{nt}} > \mathbf{0}$). These extra identification conditions are sufficient to uniquely identify the sign and size of the factor loadings and latent factors.

[5] This is an approach often used in the dynamic factor model literature to identify world and regional factors in cross country panels of data. See Kose et al (2003) as an example.

[6] The indeterminacy of factor rotation can be summarised as being unable to distinguish between the factor model $r_t = \alpha + \boldsymbol{\beta}\mathbf{F_t} + \varepsilon_t$, and the equivalent factor model $r_t = \alpha + \boldsymbol{\beta}^*\mathbf{F_t^*} + \varepsilon_t$ (where $\mathbf{F_t^*} = \mathbf{P}\mathbf{F_t}$ is a rotation of the original factor, and $\boldsymbol{\beta}^* = \mathbf{P}\boldsymbol{\beta}^*$ is a rotation of the original factor loadings), for any orthogonal matrix $\mathbf{P}$.

3.3 Number of tradable and nontradable factors

The number of dynamic tradable and nontradable factors to include in the model reflects a trade off between fit and parsimony. Including more factors would allow the model to explain more of the variation in the panel of data. On the other hand, having fewer factors will produce a more parsimonious model.

Within the dynamic factor model literature, the choice of the number of factors has received much attention. However, there seems to be no clear consensus on the approach to use when choosing the number of dynamic factors. For example, Bai and Ng (2007), Amengual and Watson (2007), and Hallin and Liška (2007) all propose different approaches to estimate the optimal number factors for a dynamic principal components model. All of these approaches would require modifications to be used with the over-identifying restrictions I am using to distinguish between tradable and nontradable factors.

Alternatively from Bayesian perspective, choosing the optimal number of factors to include can be seen as a problem of model selection. This approach has the advantage of easily being able to accommodate the distinction between tradable and nontradable factors that I make in the model. Bayesian model selection allows one to choose between two model (say one with a small number of factors and another with a larger number of factors) based on their fit and parsimony. Cogley and Sbordone (2006) use a generalisation of the Akaike Information Criterion know as the Bayesian Deviance Information Criterion (BDIC) to help choose between different Phillips curve specifications in their analysis of inflation.[7] I apply their general approach to the problem of identifying the number of each type of factor to include in a dynamic factor model.

Details of the BDIC and my results for New Zealand and Australia are discussed in appendix B. However, to summarise the findings, the BDIC recommends including three tradable and one nontradable factor for New Zealand, and four tradable and one nontradable factor for Australia.

In both the case of New Zealand and Australia, including the number of factors recommended by the BDIC produces a highly volatile core estimate that is statistically insignificant from headline inflation. To analyse why using

[7] Alternatively, one could used the Bayes factor to compare the models. However, as pointed out by Cogley and Sbordone (2006) and Del Negro and Schorfheide (2011) computing the marginal likelihood can be fraught with numerical issues.

the number of factors suggested by the BDIC does not provide a smoother, less volatile series than headline CPI, I examine the factor loadings on each factor. In the case of New Zealand, the third tradable factor has five series that have a factor loading much larger than the other series in the panel. In order of decreasing magnitude, the series with the largest factor loadings are: tradable CPI inflation, petrol, headline CPI inflation, fruit and vegetables, and other vehicle fuels and lubricants.

As discussed by Moench et al (2009), factor models tend to treat variance in groups of series as a common factor because the variance of the group contributes significantly to the total variance of the data. Given that petrol has a weighting of over five percent in the CPI basket, it is not surprising to find headline CPI inflation, tradable CPI inflation, and petrol prices all sharing some degree of co-movement. And therefore, this third tradable factor in the New Zealand data may reflect a group specific factor (say a petrol and oil factor) rather than a common tradable factor.

When I conduct a similar analysis for Australia, I also find evidence to suggest that the fourth tradable factor is a group related factor. Tradable CPI inflation, Automotive fuel, headline CPI inflation, motor vehicles, and food additives and sauces all have a much larger factor loading onto the fourth tradable factor than the other series in the panel.

As a result of this analysis, I do not believe it is appropriate to include the full number of factors suggested by the BDIC results. Therefore, in the spirit of parsimony I only use one tradable and one nontradable factor in my model specifications for New Zealand and Australia.[8]

3.4 Estimation

The dynamic factor model consisting of equations 3 to 5, is estimated using Bayesian methods. The posterior distribution is approximated by a Gibbs sampling algorithm. The priors on the factor loadings are centered on the factor loadings found using principal components. A detailed description of the Gibbs sampling algorithm can be found in appendix C.

I use 50,000 draws of the Gibbs sampling algorithm, and discard the first 45,000 draws as the burn-in period. Plotting the recursively calculated moments of the estimated parameters from the remaining 5,000 draws reveals

[8] Estimates of core inflation using various other numbers of factors are available upon request.

the recursive moments are relatively constant, suggesting the Gibbs sampling process has converged.[9]

3.5 Construction of the core inflation estimate

Once the dynamic factor model in equations 3 to 5 has been successfully estimated, we can compute the measure of core inflation by taking the tradable and nontradable factors and multiplying them by the factor loadings for CPI inflation ($\pi_{1,t}^{core} = \boldsymbol{\beta_1^{tr}} \mathbf{F_t^{tr}} + \boldsymbol{\beta_1^{nt}} \mathbf{F_t^{nt}}$). And because the data has been standardised in order to compute the dynamic factor model, we also need to multiply the core inflation estimate by the standard deviation of headline CPI inflation, and add back the sample mean of headline CPI inflation to scale it up to the original level of the inflation data.

This process is carried out for each iteration of the Gibbs sampling algorithm. The point estimate of core inflation is found as the median value of the distribution at each quarter, and the distribution of draws around the median provides an indication of the parameter and shock uncertainty inherent in the estimation process.

4 Results

In this section, I present the results of the core inflation estimate and the properties of the tradable and nontradable factors. Unless otherwise stated, the rescaled core inflation estimates are reported, and inflation rates are expressed in terms of an annual percent change.

4.1 Core inflation

The estimates of core inflation for New Zealand and Australia are plotted against their respective CPI inflation measures in figure 1.[10]

In the case of New Zealand (figure 1a), the core inflation estimate show an initial starting-point problem, whereby the estimate of core inflation is much

[9] These results are available upon request.

[10] For Australia, the headline CPI inflation adjusted for interest rate charges and the introduction of GST is plotted.

Figure 1: Core CPI inflation estimates

(a) New Zealand core inflation

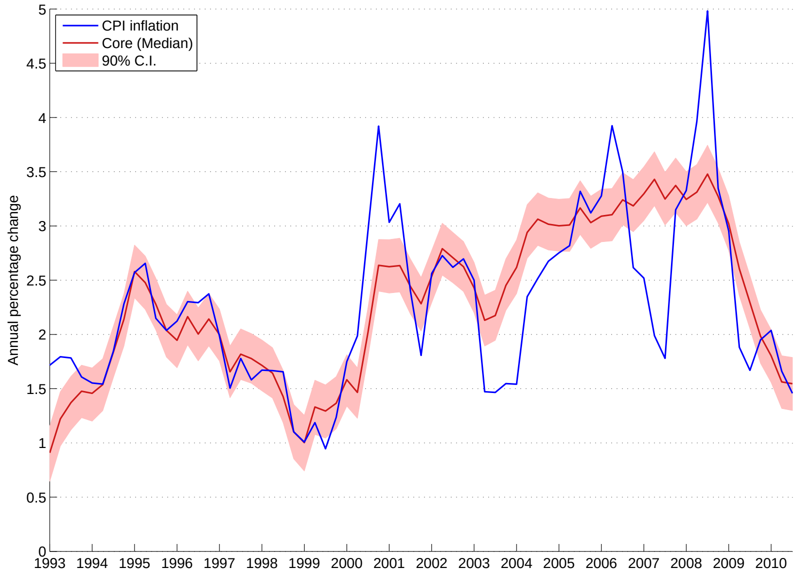

(b) Australian core inflation

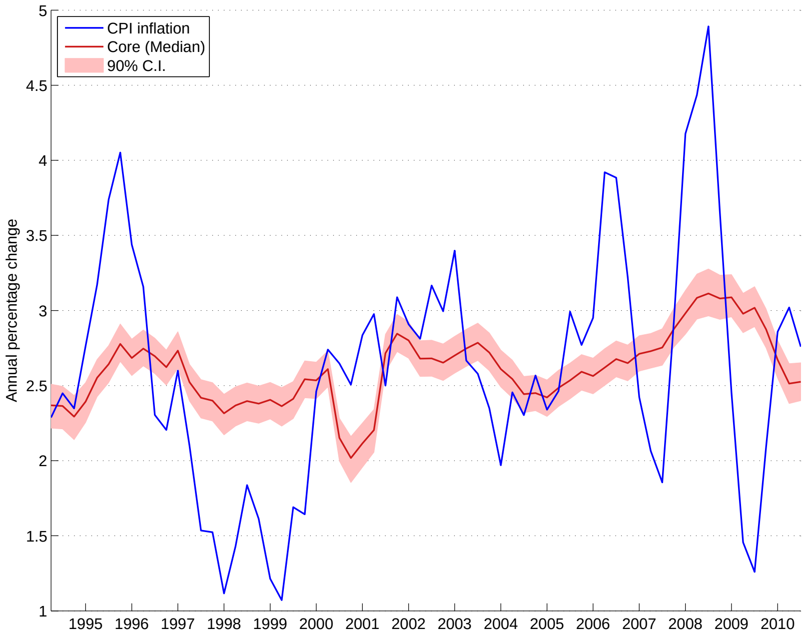

lower than CPI inflation for the first year.[11] However, after this initial period, the core inflation estimate produces a track that smooths out some of the larger spikes in CPI inflation, such as in 2000Q4 and 2009. Also of interest is that the estimate of core inflation shows a persistent increase from around 2004 up until 2008, and the onset of the global financial crisis.

For Australia, the estimate of core inflation (figure 1b) shows a relatively smooth profile that has remained inside the 2-3 percent inflation target band of the Reserve Bank of Australia for most of the sample. Like the core inflation estimate for New Zealand there is a gradual upwards trend in Australia's core inflation over the sample period. In 2000Q3, there is a noticeable one-off decline in the core inflation estimate. As will be apparent in the next section, this decline is caused by the nontradable factor. A similar effect is also observable in the Reserve Bank of Australia's weighted median measure of core inflation (see figure 2b). The one-off decline in the nontradable factor is discussed in more detail in the next section.

As a check of the plausibility of the dynamic factor model's estimate of core inflation, I plot it against a number of other common estimates of core inflation used by the Reserve Bank of New Zealand and Reserve Bank of Australia in figure 2. Despite the definition of core inflation differing between the various core estimates, one would still expect the dynamic factor model's estimate to be broadly consistent with the other estimates of core inflation.

In figure 2a, I compare the new estimate of core inflation for New Zealand with the estimate of core inflation found by Giannone and Matheson (2007) in their dynamic factor model, a weighted median, a trimmed mean, and an exponentially smoothed measure.[12] The estimate of core inflation is broadly similar to that of Giannone and Matheson (2007) over most of the sample period. This is despite differences in the number of factors employed in each model, and different estimation techniques.[13] In addition, the estimate is broadly similar to the other core inflation estimates over most of the sample. However, there is some slight deviation between 2003 and 2008 where the estimate of core inflation suggests slightly higher core inflation than the other core inflation estimates.

[11] This problem is exacerbated in part by graphing the estimate of core inflation in annual percentage change terms.

[12] The official weighted median and trimmed mean series are only available from 2006Q3 onwards.

[13] Another difference in methodology is that the dynamic factor model of Giannone and Matheson (2007) is designed to predict the two-year moving average of CPI inflation in real time. While the dynamic factor model in this paper is not.

In figure 2b, I compare the dynamic factor model's estimates of core inflation for Australia with the weighted median and trimmed mean measures published by the Reserve Bank of Australia (both estimates on the adjusted data). The estimate of core inflation from the dynamic factor model is significantly smoother than estimates from the other two measures. It does not fall by as much during 1997 to 2000 (the period following the Asian crisis), and it does not pick up the large increases in core inflation over 2008 like the other two measures do.

The difference in the dynamic factor model's estimate of core inflation does not appear to be the result of excluding series from the panel of data that do not span the entire sample length (these series are included in the computation of the trimmed mean and weighted median measures of core inflation). I computed a new trimmed mean and weighted median based on the same panel of data used by the dynamic factor model. Both of these showed a noticeable pickup in their measures of core inflation over 2008.

Instead, it appears that some of the difference in estimated core inflation can be explained by the use of different weighting schemes. The trimmed mean and weighted median measures of core inflation are based on the same fixed expenditure weights as the CPI basket, while the dynamic factor model weights are based on the relative co-movement in the panel data. By examining the distribution of price change for each expenditure class series (and their relative CPI weights) over this period, I find that those series with the largest weights in the CPI basket generally have increases in prices (automotive fuel being an example), while those with price declines tend to have a lower weight in the CPI basket. This skewness in distribution will tend to result in core inflation measure based on CPI weights producing a higher core estimate relative to measures based on other weighting schemes (such as the dynamic factor model approach).

4.2 The dynamic factors

One of the advantages of the dynamic factor model used in this paper is the story that comes from the estimated tradable and nontradable factors. Figure 3 shows the (unscaled) profiles of the tradable and nontradable factors for

Figure 2: Various core inflation estimates

(a) New Zealand

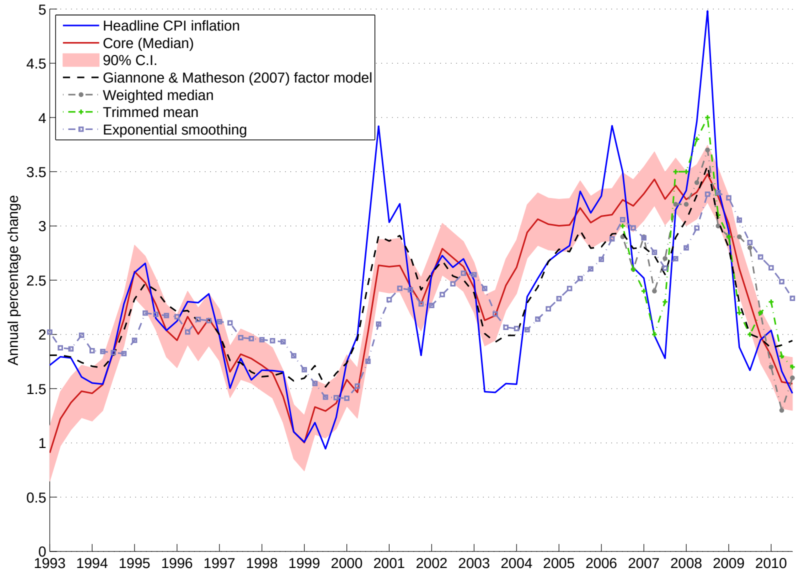

(b) Australia

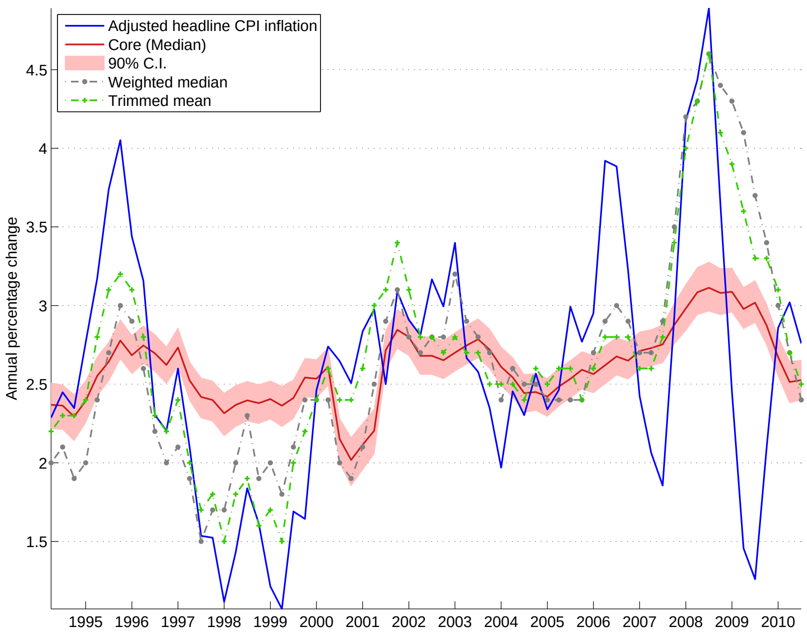

Figure 3: Dynamic factor estimates

(a) New Zealand: Tradable factor

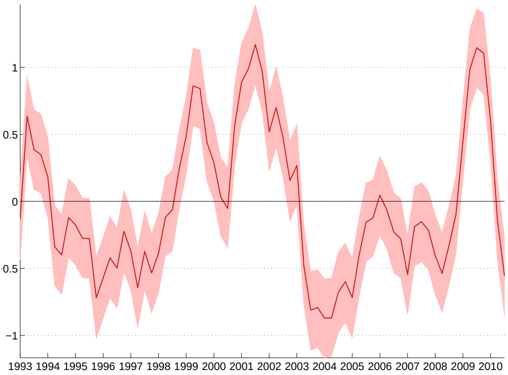

(b) New Zealand: Nontradable factor

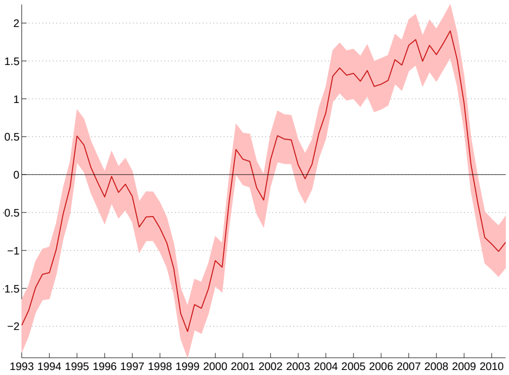

(c) Australia: Tradable factor

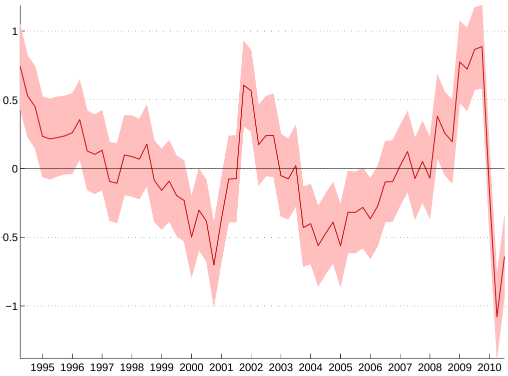

(d) Australia: Nontradable factor

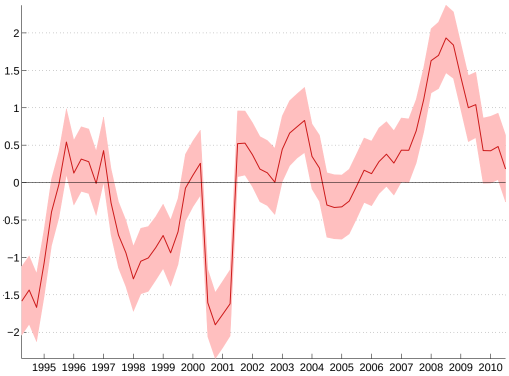

the final vintage (2010Q3) estimate for both New Zealand and Australia.[14]

A few key features stand out when analysing each of the factors in figure 3. First, the tradable factor for New Zealand (figure 3a) and Australia (figure 3c) show some broad similarities in their cycles with a decline early on in the sample, before spiking up to peaks around 2001, followed by a large decline before increasing again to a peak again in late 2009. Common cycles

[14] As a robustness check, I estimated a simple two factor Bayesian dynamic factor model without the sign restrictions to identify the factors as tradable and nontradable factors for New Zealand. The first factor from this model was visually similar to the nontradable factor shown in figure 3b, and the second was similar to the tradable factor shown in figure 3a. The estimate of core inflation from this model did not differ significantly from the identified model's estimate.

support the idea that tradable prices in both countries are influenced by a common external source (such as the exchange rate). Neely and Rapach (2008) also find evidence of a strong 'world factor' influence on New Zealand and Australian headline inflation.

Second, the nontradable factors in New Zealand (figure 3b) and Australia (figure 3d) also show similarities. There appears to be a general trend upwards in both factors, with both factors being above the historical mean over most of the second half of the sample. Because nontradable prices tend to be driven by domestic conditions, the results here suggest similar business cycles movements for New Zealand and Australia. The idea of similar business cycles between the two countries is a commonly held belief amongst economists and also supported by analysis carried out on key macroeconomic variables (see Hall et al 1998 as an example). The analysis from Neely and Rapach (2008) also shows that an 'Australasian factor' explains more of the variance in headline inflation in both New Zealand and Australia than the individual country factors, suggestion some common regional drive of inflation in both countrys.

Another similarity between the factors for both New Zealand and Australia is that the nontradable factor shows more persistence than the tradable factor (a characteristic also seen when comparing the tradable CPI and nontradable CPI inflation rates).

Figure 3b also suggests that the starting point problem for the estimate of New Zealand's core inflation occurs in the nontradable factor with the estimated factor being quite low at the start of the sample.

As discussed previously, the decline in Australian core inflation in 2000Q3 is a result of the nontradable factor. During 2000Q3, five of the 26 nontradable expenditure class series showed a decline in their quarterly percentage change larger than two standard deviations (on average, less than one series will show a decline this large each quarter). Aside from the introduction of GST (and removal of wholesale taxes), other factors impacted upon the CPI panel during 2000Q3. A number of other government tax and subsidies were also changed (such as the first home buyers grant and child care benefit scheme), new categories were added to the CPI basket, and the CPI basket was reweighted. As a result, it is very difficult to disentangle the drivers behind core inflation in 2000Q3.[15]

While analysis the time series of both the tradable and nontradable factors reveals some information about the underlying drivers of each factor, it does not tell us exactly how these factors relate to the estimates of core inflation see previously in figure 1. Even if a particular factor had large deviations over the sample, it would not correspond to large derivation in core inflation if the factor loading was relatively small. Therefore, to understand which of these factors has the largest contribution to deviations in core inflation from its mean, we must multiply each factor by the appropriate factor loading for CPI inflation ($\beta_1^k F_t^k$), and scale up the estimate by the standard deviation of CPI inflation. The results of this calculation are plotted for both countries in figure 4.

One of the key findings for both New Zealand (figure 4a) and Australia (figure 4b) is that the nontradable factor (shown by the blue bars) has generally a much large contribution to core inflation than the tradable factor (shown by the red bars). This result suggests that if central banks wish to target the persistent part of inflation, they should place more of their focus on nontradable prices.

4.3 Variance decomposition

Because the factors of the dynamic factor model are orthogonal, the variance of the i^{th} series in our panel ($\pi_{i,t}$) can be expressed as:

$$var(\pi_{i,t}) = \left(\boldsymbol{\beta_i^{tr}}\right)^2 var\left(\mathbf{F_t^{tr}}\right) + \left(\boldsymbol{\beta_i^{nt}}\right)^2 var\left(\mathbf{F_t^{nt}}\right) + var\left(\nu_{i,t}\right). \tag{6}$$

Using equation 6, we are able to decompose the variance of each inflation series in the panel into the share explained by each factor, and the share explained by the idiosyncratic component. The proportion of the variance

[15] As a robustness test, I tried two alternative approaches to remove the effects of the GST introduction: interpolating the 2000Q3 values for each series from the 2000Q2 and 2000Q4 values, and replacing the 2000Q3 value for each series with the sample median. Both of these approaches produced a more volatile core inflation series. However, implementing either of these adjustments in real time is problematic.
Another possible approach would be remove the 2000Q3 observations and treat them as missing data using an approach similar to that proposed by Jungbacker et al (2009). Future work will examine the feasibility of this approach, however, one problem will be that we could no longer use the principal component estimates of the factors to set the priors and starting values for the Gibbs sampling algorithm.

Figure 4: Factor contribution to core CPI inflation

(a) New Zealand

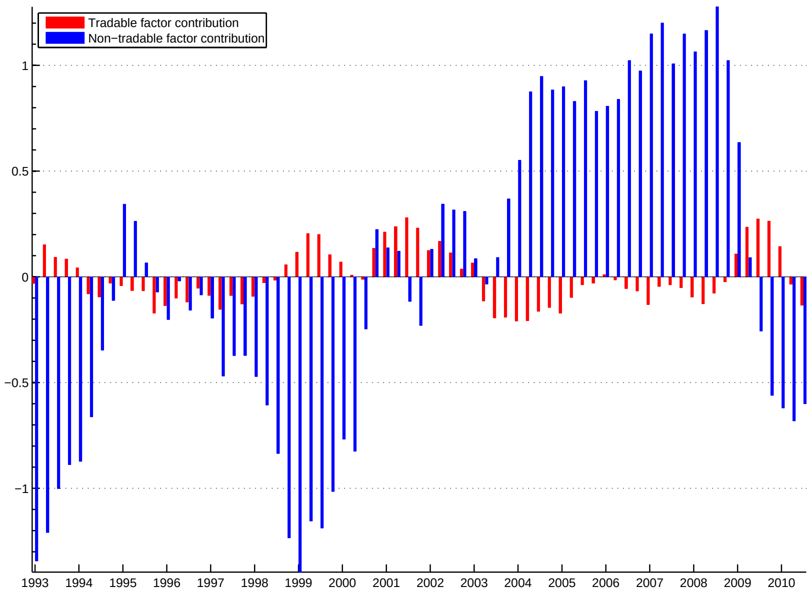

(b) Australia

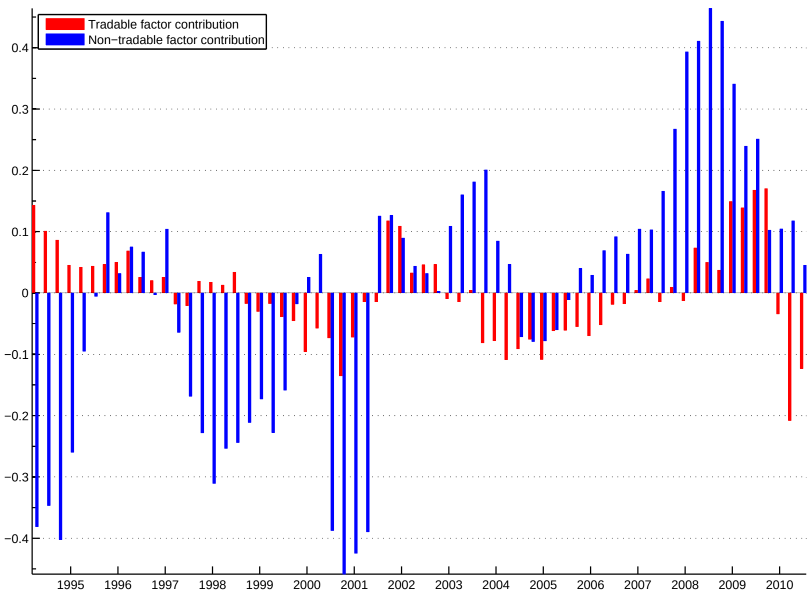

Table 1: Variance decomposition of annual inflation rates for New Zealand

	Variance explained by each factor		
	10th percentile	**Median**	**90th percentile**
	CPI inflation		
Tradable factor	1.46	2.47	3.84
Nontradable factor	57.12	63.31	68.91
Idiosyncratic	28.66	34.10	40.04
	Tradable inflation		
Tradable factor	9.01	12.47	16.74
Nontradable factor	0.00	0.00	0.00
Idiosyncratic	83.26	87.53	90.99
	Nontradable inflation		
Tradable factor	0.00	0.00	0.00
Nontradable factor	60.13	66.49	72.24
Idiosyncratic	27.76	33.51	39.87

explained by the tradable or nontradable factor is given by $\frac{(\beta_i^k)^2 var(F_t^k)}{var(\pi_{i,t})}$ (where $k = \{tr, nt\}$), and the proportion of variance explained by the idiosyncratic component is given by $\frac{var(\nu_{i,t})}{var(\pi_{i,t})}$.

Despite the dynamic factor model being orthogonal by design, sampling error in the computation of the Gibbs sampling algorithm results in the factors being weakly correlated. Therefore, to ensure the add up of equation 6, I follow the approach of Kose et al (2003) and orthogonalise the sampled factors at each iteration before computing the variance decomposition.[16]

The results of decomposing the variance of annual headline CPI inflation for New Zealand are present in table 1, and the results for Australia are presented in table 2. In addition to the variance decomposition of CPI inflation, I also show the variance decomposition of tradable CPI inflation and nontradable CPI inflation in each table.

In the case of New Zealand (table 1), nearly two thirds of the annual variance in CPI inflation can be explained by the nontradable factor (median variance explained = 63.31 percent), while the tradable factor explains very little of

[16] This extra step to orthogonalise the factors is carried out only for the computation of the variance explained by each factor. The rest of the results in the paper are carried out using the original factor draws.

Table 2: Variance decomposition of annual inflation rates for Australia

	Variance explained by each factor		
	10th percentile	**Median**	**90th percentile**
	CPI inflation		
Tradable factor	0.71	1.28	2.10
Nontradable factor	5.52	8.70	12.78
Idiosyncratic	85.83	89.93	93.19
	Tradable inflation		
Tradable factor	3.09	4.54	6.48
Nontradable factor	0.00	0.00	0.00
Idiosyncratic	93.52	95.46	96.91
	Nontradable inflation		
Tradable factor	0.00	0.00	0.00
Nontradable factor	35.20	43.81	52.64
Idiosyncratic	47.36	56.19	64.80

the annual variance in CPI inflation (2.47 percent). For Australia (table 2), the nontradable factor explains only around nine percent of the variance in annual inflation, while the tradable factor explains less (just over one percent). The differences between New Zealand and Australia highlight the smoothness of the Australia core measure.

The tradable factor explains around 12 percent of the variance of annual tradable CPI inflation in New Zealand, and around five percent in Australia. As discussed in section 3.3, headline CPI inflation and tradable CPI inflation is more strongly influenced by a number of small number of tradable series (such as petrol prices) than the common co-movement across the whole panel.[17]

For annual nontradable CPI inflation in New Zealand and Australia, around half the variance can be explained using only one factor.

[17] The tradable factor does explain a high amount of variance in some of the other series in the panel of data.

4.4 Real-time core inflation

The estimation of core inflation from the dynamic factor model employs the Kalman smoother which uses all the data available at any point in time to estimate the historical path of core inflation. As a result, the estimate of core inflation is subject to historical revisions. If these historical revisions are relatively small, and the interpretation of core inflation does not change significantly, then the dynamic factor model's estimate of core inflation would be suitable for use by a central bank for its policy analysis.

I examine the real-time performance of the Bayesian model over the sample range 2000Q1 to 2010Q3.[18] At each point in time, the estimate of core inflation for the current quarter is saved, and the results are plotted in figure 5 (in black) against the final ex-post core inflation estimate (in red) for New Zealand and Australia.

The real-time estimate of core inflation for New Zealand (shown in figure 5a) is broadly similar to the final ex-post estimate. However, there are two notable deviations. Between 2001-2002 and 2007-2008, the real-time core estimate is noticeably lower than the ex-post estimate (outside the 90 percent level of confidence). The deviation or the real time estimate 2001-2002 may indicate that there is still a short-sample problem that exists during this period.

For Australia (figure 5b), the real-time estimate is very similar to the final ex-post estimate for most of the sample. There is a larger deviation in 2001-2002 which may be the result of a short sample problem, or may be due to the adjustment for GST made in 2000Q3. There is also a deviation in 2008 where the real time measure indicated higher core inflation than the ex-post measure.

One approach to quantifying the size of the historical revisions to core inflation over time is to compare the mean revisions and/or mean absolute revisions between the historical vintage estimates of core inflation and the final full sample (ex-post) estimate. Formally, for each quarter (t) between 2000Q1 and 2010Q3, I take the core inflation estimate using data up to t, and for periods $t, t-1, t-2, t-3, \ldots, 1$ compute the difference between the estimated level of core inflation from that vintage and the ex-post estimate of core inflation. These differences are then averaged across all the results from the vintages 2000Q1 to 2010Q3. The results for the contemporaneous

[18] The real-time performance of the dynamic factor model prior to 2000Q1 is likely to be biased by the fairly short data sample. Thus I exclude it from this analysis.

Figure 5: Real-time vs. ex-post core inflation estimates

(a) New Zealand

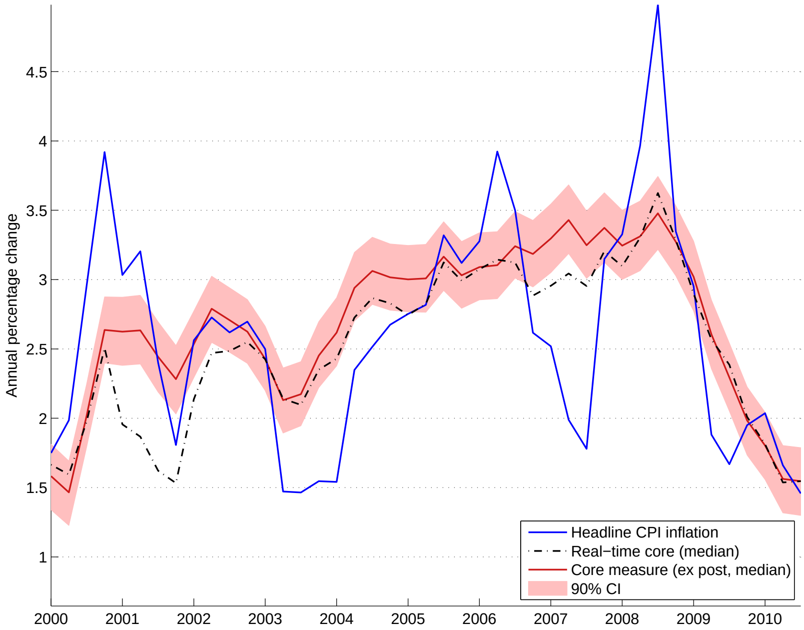

(b) Australia

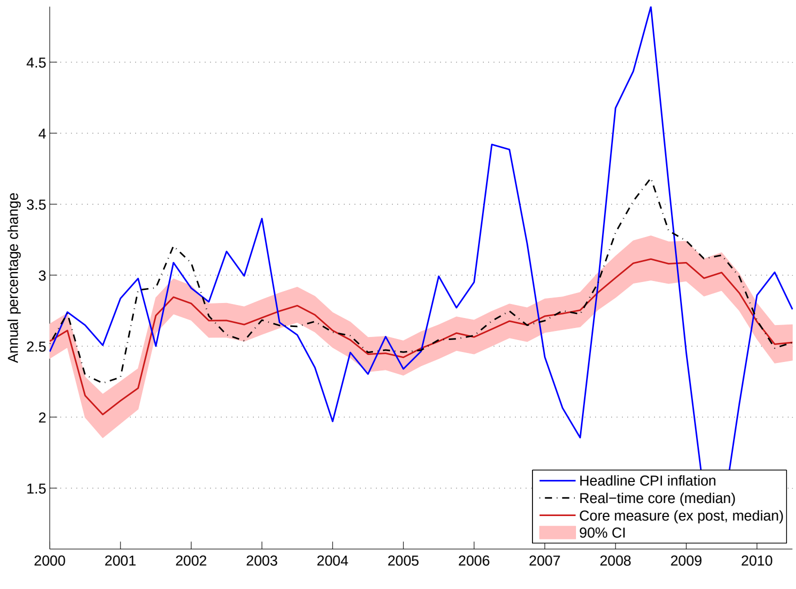

Table 3: Revision to the real-time annual core estimate (2000Q1-2010Q3) for New Zealand

	t	**t-1**	**t-2**	**t-3**	**t-4**
Mean revision	0.165	0.151	0.142	0.133	0.119
Mean absolute revision	0.191	0.191	0.194	0.187	0.181

Table 4: Revision to the real-time annual core estimate (2000Q1-2010Q3) for Australia

	t	**t-1**	**t-2**	**t-3**	**t-4**
Mean revision	-0.091	-0.087	-0.067	-0.038	-0.008
Mean absolute revision	0.168	0.163	0.168	0.182	0.185

quarter and first four lags are presented in tables 3 and 4.[19]

Table 3 reveals that on average, the estimated historical level of New Zealand's core inflation made at any point between 2000Q1 to 2010Q3 is around 0.165 percentage points lower than the full-sample estimate of the historical core level. As would be expected, the size of these revisions decline for longer lags.

Part of the bias in the revisions could be due to the rising mean inflation rate over our sample range. The dynamic factor model assumes a constant mean over the entire sample, thus a rising sample mean will bias the real-time estimate downwards. One potential driver for the rising mean of inflation over our sample could be changes in the inflation target of the Reserve Bank of New Zealand, which has twice been revised upwards over our sample range.[20]

For Australia (table 4), the results show the some small downwards revisions over time (row one of the table). At any point in time, the real time estimate is around 0.09 percentage points higher than the final ex-post estimate. This upwards bias in the real time estimate is driven primarily by the two large deviations in the real time estimate as seen in figure 5b.

However, one should be careful about reading too much into the exact number

[19] A positive number in the 'Mean revision' line indicates that the real-time estimate is revised upwards over time.

[20] As an extension, I estimated the model detrending each expenditure class inflation series with the mid-point of the inflation target before normalising the data in the panel. Doing so had some impact on the real time performance of the model (especially during the early part of the subsample). However, it is unlikely that all the bias can be accredited solely to the chances in inflation target.

presented in tables 3 and 4. The results will be dependent upon the size of the sample period used. In general, a much smaller real-time sample period will produce smaller revisions between the first and ex-post estimates.

5 Conclusion

In this paper I introduces a new approach to estimating core inflation. Combining an over-identifying restriction in a Bayesian framework with a dynamic factor model, I construct an estimate of core inflation that is drive by the common co-movements in tradable and nontradable prices. The major advantage of this approach is it allows for an interpretation of core inflation in terms of tradable and nontradable influences.

Applying this model to disaggregated CPI data New Zealand and Australian reveals that in both countries, the estimate of core inflation is primarily driven by the nontradable factor. In addition, the nontradable factor explains significantly more of the variance in headline CPI inflation than the tradable factor. And when compared across countries, both the tradable and nontradable factors of New Zealand and Australia show similar profiles, suggesting the possibility of a common drivers on the prices in both countries.

References

Amengual, D and M W Watson (2007), "Consistent estimation of the number of dynamic factors in a large N and T panel," *Journal of Business & Economic Statistics*, 25(1), 52–60.

Bai, J and S Ng (2007), "Determining the number of primitive shocks in factor models," *Journal of Business & Economic Statistics*, 25(1), 52–60.

Bryan, M F and S G Cecchetti (1993), "The consumer price index as a measure of inflation," *National Bureau of Economic Research, Inc, NBER Working Papers*, 4505.

Bryan, M F and S G Cecchetti (1994), "Measuring core inflation," in *Monetary Policy*, april-june, 195–219, National Bureau of Economic Reserach, Inc.

Carter, C and P Kohn (1994), "On Gibbs sampling for state space models," *Biometrika*, 81, 541–553.

Cogley, T (2002), "A simple adaptive measure of core inflation," *Journal of Money, Credit and Banking*, 34(1), 94–113.

Cogley, T and A M Sbordone (2006), "Trend inflation and inflation persistence in the new keynesian phillips curve," *Federal Reserve Bank of New York, Staff Reports*, 270.

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).

Del Negro, M and F Schorfheide (2011), "Bayesian macroeconometrics," in *Handbook of Bayesian Econometrics*, eds J Geweke, G Koop, and H van Dijk.

Giannone, D and T Matheson (2007), "A new core inflation indicator for New Zealand," *International Journal of Central Banking*, 3(4), 145–180.

Hall, V, K Kim, and R A Buckle (1998), "Pacific rim business cycle analysis: synchronisation and volatility," *New Zealand Economic Papers*, 32, 129–59.

Hallin, M and R Liška (2007), "Determining the number of factors in the general dynamic factor model," *Journal of the American Statistical Association*, 102(478), 603–617.

Jungbacker, B, S Koopman, and M van der Wel (2009), "Dynamic factor analysis in the presence of missing data," *Tinbergen Institute, Tinbergen Institute Discussion Papers*, 09-010/4.

Kapetanios, G (2004), "A note on modelling core inflation for the UK using a new dynamic factor estimation method and a large disaggregated price index dataset," *Economic Letters*, 85, 63–69.

Kim, C-J and C R Nelson (1999), *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, The MIT Press.

Kose, M A, C Otrok, and C H Whiteman (2003), "International business cycles: World, region, and country-specific factors," *American Economic Review*, 93(4), 1216–1239.

Moench, E, S Ng, and S Potter (2009), "Dynamic hierarchical factor models," *Federal Reserve Bank of New York*, *Staff Reports*, 412.

Neely, C J and D E Rapach (2008), "International comovements in inflation rates and country characteristics," *Federal Reserve Bank of St. Louis*, *Working Paper*, 25.

Quah, D and S P Vahey (1995), "Measuring core inflation?" *Economic Journal*, 105(432), 1130–44.

Reis, R and M W Watson (2010), "Relative goods' prices, pure inflation, and the phillips correlation," *American Economic Journal: Macroeconomics*, 2(3), 128–57.

Spiegelhalter, D J, N G Best, B P Carlin, and A van der Linde (2002), "Bayesian measures of model complexity and fit," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4), 583.

Wynne, M A (2008), "Core inflation: A review of some conceptual issues," *Federal Reserve Bank of St. Louis Review*, 90(3, Part 2), 205–28.

Appendices

A Inflation panel data

The following CPI expenditure class series were included in the New Zealand panel of inflation data:

Table 5: New Zealand panel

Series	TR	NT	Series	TE	NT
Consumer price index	1	1	Medical services	0	1
All groups - tradables component	1	0	Dental services	0	1
All groups - non-tradables component	0	1	Paramedical services	0	1
Fruit	1	0	Hospital services	0	1
Vegetables	1	0	Purchase of new motor cars	1	0
Meat and poultry	1	1	Purchase of second-hand motor cars	1	0
Fish and other seafood	1	0	Purchase of motorcycles	1	0
Bread and cereals	1	1	Purchase of bicycles	1	0
Milk, cheese and eggs	1	1	Vehicle parts and accessories	1	0
Oils and fats	1	0	Petrol	1	0
Food additives and condiments	1	0	Other vehicle fuels and lubricants	1	0
Confectionery, nuts and snacks	1	0	Vehicle servicing and repairs	0	1
Other grocery food	1	0	Other private transport services	0	1
Coffee, tea and other hot drinks	1	0	Domestic air transport	0	1
Soft drinks, waters and juices	1	1	International air transport	1	0
Restaurant meals	0	1	Postal services	0	1
Ready-to-eat food	0	1	Telecommunication services	0	1
Beer	0	1	Audio-visual equipment	1	0
Wine	1	1	Computing equipment	1	0
Spirits and liqueurs	1	1	Recording media	1	0
Cigarettes and tobacco	0	1	Major recreational and cultural equipment	1	0
Mens clothing	1	0	Games, toys and hobbies	1	0
Womens clothing	1	0	Equipment for sport, camping and outdoor recreation	1	0
Childrens and infants clothing	1	0	Plants, flowers and gardening supplies	1	0
Knitting and sewing supplies	1	0	Pet-related products	1	0
Clothing services	0	1	Recreational and sporting services	0	1
Mens footwear	1	0	Cultural services	1	1
Womens footwear	1	0	Veterinary services	0	1
Childrens and infants footwear	1	0	Books	1	0
Actual rentals for housing	0	1	Newspapers and magazines	1	1
Purchase of housing	0	1	Stationery and drawing materials	1	0
Property maintenance materials	1	0	Accommodation services	0	1
Property maintenance services	0	1	Early childhood education	0	1

continued on next page

Table 5: *continued*

Series	TR	NT	Series	TE	NT
Local authority rates and payments	0	1	Primary and secondary education	0	1
Electricity	0	1	Tertiary and other post school education	0	1
Gas	0	1	Hairdressing and personal grooming services	0	1
Solid fuels	1	0	Electrical appliances for personal care	1	0
Furniture and furnishings	1	0	Jewellery and watches	1	1
Carpets and other floor coverings	1	0	Other personal effects	1	0
Household textiles	1	0	Life insurance	0	1
Major household appliances	1	0	Dwelling insurance	0	1
Small electrical household appliances	1	0	Contents insurance	0	1
Repair and hire of household appliances	0	1	Health insurance	0	1
Glassware, tableware and household utensils	1	0	Vehicle insurance	0	1
Major tools and equipment for the house and garden	1	1	Direct credit service charges	0	1
Small tools and accessories for the house and garden	1	0	Vocational services	0	1
Cleaning products and other household supplies	1	0	Professional services	0	1
Other household services	0	1	Real estate services	0	1
Pharmaceutical products	1	1	Other miscellaneous services nec	0	1
Other medical products	1	0			

The following CPI expenditure class series were removed from the New Zealand panel as they did not span the entire sample range: Water supply; Refuse disposal and recycling; Therapeutic appliances and equipment; Rail passenger transport; Road passenger transport; Sea passenger transport; Telecommunication equipment; Package holidays; and Electrical appliances for personal care.

The following CPI expenditure class series were included in the Australian panel of inflation data:

Table 6: Australian panel

Series	TR	NT	Series	TE	NT
Consumer Price Index	1	1	Electricity	0	1
Tradables	1	0	Furniture	1	0
Nontradables	0	1	Floor coverings	1	0
Milk and Cream	0	1	Towels, lien and curtains	1	0
Cheese	1	0	Tools	1	0
Other dairy products	1	0	Household cleaning agents	1	0
Break	0	1	Veterinary services	0	1
Cakes and biscuits	0	1	Pet foods	1	0
Breakfast cereals	0	1	Postal services	0	1

continued on next page

Table 6: *continued*

Series	**TR**	**NT**	**Series**	**TE**	**NT**
Other cereal products	1	0	Telephone services	0	1
Beef and veal	1	0	Motor vehicles	1	0
Lamb and mutton	1	0	Automotive fuel	1	0
Pork	1	0	Motoring charges	0	1
Poultry	0	1	Tyres and tubes	1	0
Bacon and ham	0	1	Vehicle servicing repairs and parts	0	1
Processed meat	1	0	Urban transport fates	0	1
Fish and other seafood	1	0	Beer	0	1
Restaurant Meals	0	1	Wine	1	0
Take away food	0	1	Spirits	1	0
Eggs	0	1	Cigarettes and tobacco	1	0
Jams, honey and sandwich	0	1	Hospital and medical services	0	1
Tea coffee and food	1	0	Optical services	0	1
Food additives and sauces	1	0	Dental services	0	1
Other food	1	0	Pharmaceutical	1	0
Men's outer clothing	1	0	Toiletries and personal products	1	0
Men's underwear, nightwear	1	0	Hairdressing services	0	1
Women's outer clothing	1	0	Video and sound equipment	1	0
Women's underwear, nightwear	1	0	Holiday travel and accommodation in Australia	0	1
Men's footwear	1	0	Holiday travel and accommodation overseas	1	0
Women's footwear	1	0	Childcare fees	0	1
Children's footwear	1	0	Fats and oils	1	0
Dry cleaning and shoe repairs	0	1	Other household supplies	1	0
House repairs and maintenance	0	1	AV and computer media and services	1	0

The following CPI expenditure class series were removed from the Australian panel as they did not span the entire sample range: Clothing accessories & jewellery; house purchases; property rates & charges; water & sewerage; sports & recreational equipment; toys games & hobbies; sports participation; other recreational activities; fruit; vegetables; softdrinks waters & juice; snacks & confectionary; children's and infants' clothing; rents; gas & other household fuels; major household appliances; small electric household appliances; books; newspapers & magazines; insurance; preschool & primary education; secondary education; tertiary education; deposit & loan facilities; and other financial services.

B BDIC Results

The BDIC (developed by Spiegelhalter et al 2002) evaluates models by rewarding goodness of fit, and penalising model complexity. The BDIC can be defined as:

Figure 6: Inverted BDIC values

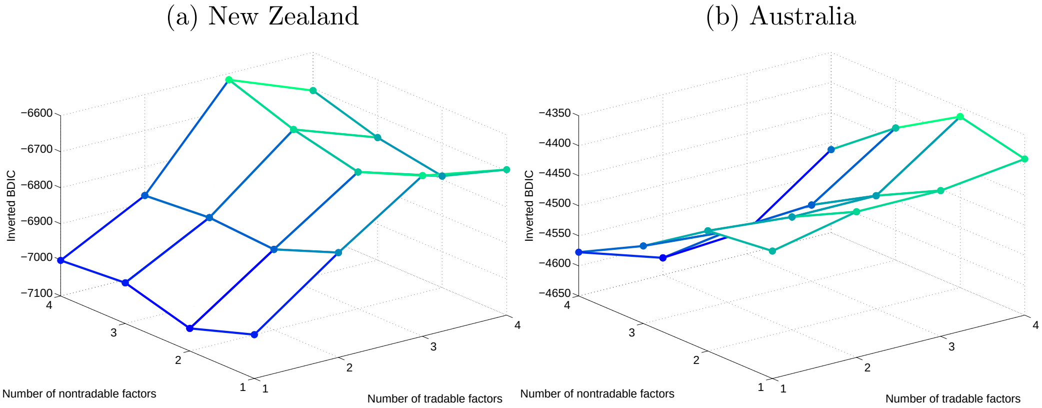

$$BDIC = \bar{D} + p_D \tag{7}$$

where

$$\begin{aligned} \bar{D} &= \frac{1}{N}\sum_{i=1}^{N}(-2lnL(\Xi_i)), \\ p_D &= \frac{1}{N}\sum_{i=1}^{N}(-2lnL(\Xi_i)) - \left(-2lnL(\frac{1}{N}\sum_{i=1}^{N}\Xi_i)\right). \end{aligned}$$

The first term ($\bar{D}$) is based on the log likelihood ($lnL(\Xi_i)$), and rewards fit (larger log likelihood) by producing a lower BDIC value. The second term (p_D) measures the model complexity by evaluating the average log likelihood of each iteration against the log likelihood at the posterior mean ($\frac{1}{N}\sum_{i=1}^{N}\Xi_i$). A more complex model (one with a higher number of effective parameters) will produce a higher p_D value. The relative sizes of these two terms ($\bar{D}$ and p_D) reflect the trade off between goodness of fit and parsimony. For this information criteria, a lower BDIC value refers to the better model.

For both the New Zealand and Australian panels of data, I estimate the dynamic factor model using between one and four tradable and non-tradable factors. For each specification I compute the BDIC of the model and compare them. The inverted BDICE results for both countries are shown in figure 6 (such that the highest point on the service refers to the lowest BDIC value and the best model). The actual BDIC numbers are reported in tables 7 and 8.

Table 7: BDIC results for New Zealand

		Number of tradable factors			
		1	**2**	**3**	**4**
Number of nontradable factors	**1**	6977.93	6809.25	6654.38	6696.23
	2	7037.36	6876.75	6720.71	6790.42
	3	6987.26	6865.04	6679.09	6760.06
	4	7001.48	6879.98	6617.47	6706.07

Table 8: BDIC results for Australia

		Number of tradable factors			
		1	**2**	**3**	**4**
Number of nontradable factors	**1**	4394.98	4407.55	4407.53	4390.03
	2	4450.01	4462.13	4461.82	4365.32
	3	4521.07	4532.1	4523.08	4430.06
	4	4577.14	4621.69	4611.85	4511.91

For the New Zealand panel, the lowest BDIC value (highest peak in figure 6a) is produced by the model specified with three tradable factors, and one nontradable factor. Figure 6a reveals that for any number of tradable factors in the model, we should only including one nontradable factor produces the lowest BDIC value possible (high peak in the graph). The graph also shows that independent of the number of nontradable factors, including three tradable factors produces a noticeably higher peak (lower BDIC value).

For the Australia panel, the lowest BDIC value is produced by four tradable factors, and two nontradable factors. However, the shape of the surface in figure 6b reveals that including more than one nontradable factor will generally worsen the BDIC score (the gain in fit is not enough to overcome the added model complexity.).

C Gibbs sampling algorithm

In this appendix, I outline in more detail the Gibbs sampling algorithm used in this paper.

C.1 The model

The entire dynamic factor model is given by:

$$\pi_{i,t} = \boldsymbol{\beta_i^{tr}} \mathbf{F_t^{tr}} + \boldsymbol{\beta_i^{nt}} \mathbf{F_t^{nt}} + \nu_{i,t} \tag{8}$$

where $\mathbf{F_t^{tr}}$ is a vector of n number of tradable factors, $\mathbf{F_t^{nt}}$ is a vector of m number of nontradable factors, $\boldsymbol{\beta_i^{tr}}$ and $\boldsymbol{\beta_i^{nt}}$ are vectors of factor loadings, and $\nu_{i,t}$ is the idiosyncratic error term.

Each factor is assumed to follow an AR(2) process:

$$F_{a,t}^k = \sum_{j=1}^{2} \rho_{a,j}^k F_{a,t-j}^k + \varepsilon_{a,t}, \qquad var\left(\varepsilon_{a,t}\right) = Q_a \tag{9}$$

where $k = \{tr, nt\}$ and a is an index number of factors ($a = 1, \ldots, K$).

And the idiosyncratic error term follows an AR(1) process

$$\nu_{i,t} = \alpha_i \nu_{i,t-1} + \eta_{i,t}, \qquad var(\eta_{i,t}) = R_i \tag{10}$$

In state-space form, the measurement equation for the model can be represented as:

$$\begin{pmatrix} \pi_{1,t} \\ \vdots \\ \pi_{N,t} \end{pmatrix} = \begin{pmatrix} \boldsymbol{\beta_1^{tr}} & \boldsymbol{\beta_1^{nt}} & -\alpha_1 \boldsymbol{\beta_1^{tr}} & -\alpha_1 \boldsymbol{\beta_1^{nt}} \\ \vdots & \vdots & \vdots & \vdots \\ \boldsymbol{\beta_N^{tr}} & \boldsymbol{\beta_N^{nt}} & -\alpha_N \boldsymbol{\beta_N^{tr}} & -\alpha_N \boldsymbol{\beta_N^{nt}} \end{pmatrix} \begin{pmatrix} \mathbf{F_t^{tr}} & \mathbf{F_t^{nt}} \\ \mathbf{F_{t-1}^{tr}} & \mathbf{F_{t-1}^{nt}} \end{pmatrix} + \begin{pmatrix} \eta_{1,t} \\ \vdots \\ \eta_{N,t} \end{pmatrix} \tag{11}$$

And the transition equation can be represented as:

$$\begin{pmatrix} \mathbf{F_t^{tr}} & \mathbf{F_t^{nt}} \\ \mathbf{F_{t-1}^{tr}} & \mathbf{F_{t-1}^{nt}} \end{pmatrix} = \begin{pmatrix} \mathbf{H1} & \mathbf{H2} \\ \mathbf{I_{n+m}} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{F_{t-1}^{tr'}} & \mathbf{F_{t-1}^{nt'}} \\ \mathbf{F_{t-2}^{tr'}} & \mathbf{F_{t-2}^{nt'}} \end{pmatrix} + \begin{pmatrix} \varepsilon_{1,t} \\ \vdots \\ \varepsilon_{N,t} \end{pmatrix} \tag{12}$$

where

$$\mathbf{H1} = \begin{pmatrix} \rho_{1,1} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \rho_{n+m,1} \end{pmatrix}, \qquad \mathbf{H2} = \begin{pmatrix} \rho_{1,2} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \rho_{n+m,2} \end{pmatrix},$$

and $\mathbf{I_{n+m}}$ is a $n+m$ by $n+m$ identity matrix.

C.2 Prior distributions

Priors ($\mathbf{B_{k,i}}$) for each factor loading ($\boldsymbol{\beta} = [\boldsymbol{\beta_1^{tr}} \ldots \boldsymbol{\beta_n^{tr}} \boldsymbol{\beta_1^{nt}} \boldsymbol{\beta_m^{nt}}]'$) are centered upon the factors found by via Principal Components. The priors are distributed normally:

$$\mathbf{B_{k,i}} \sim N\left(\mathbf{PC}, \theta \mathbf{I_{K \cdot N}}\right)$$

where $\mathbf{PC}$ is a vector of the factors found via principal components. I choose $\theta = 0.05$ to scale the variance of the factor loading priors.

To ensure that the tradable (nontradable) principal component factor is only capturing the co-movement in tradable (nontradable) prices, I exclude all series that load onto both factors when computing the principal component factors.

C.3 Simulating the posterior distributions

Step 0: Initial setup

Starting values for the factor loadings are also obtained from the principal component estimator. To assist the initialisation of the identification process, if the tradable CPI inflation series ($\pi_{2,t}$) is loaded negatively onto the tradable factors found by principal components, then I reverse the signs on both the factor and factor loadings for all tradable series. The same check is carried out for nontradable CPIinflation. This steps assists the model with imposing the positive factor loading sign restriction later in the algorithm.

The starting value of the AR coefficients (α_i) is set equal to the arbitrary starting value of zero. And the starting value for the variance of the idiosyncratic components (R_i) is set equal to unity.

The aim of the Gibbs sampling algorithm is to sample from the posterior distributions of the factor loadings ($\boldsymbol{\beta}$), serial correlation coefficients ($\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_N]'$), variance of the idiosyncratic components (R_i), factor transition coefficients ($\boldsymbol{\rho} = [\rho_{1,1}, \ldots, \rho_{K,1}, \rho_{1,2}, \ldots, \rho_{K,2}]'$), and factors ($\mathbf{F_t^{tr}}$ and $\mathbf{F_t^{nt}}$).

Step 1: Sample the factor loadings

Substituting equation 10 into equation 8, we are able to rewrite the system without any autocorrelation in the error term:

$$\underbrace{\pi_{i,t} - \alpha_i \pi_{i,t-1}}_{Y^*_{i,t}} = \boldsymbol{\beta_i^{tr}} \underbrace{(\mathbf{F^{tr}_t} - \alpha_i \mathbf{F^{tr}_{t-1}})}_{X^{*,tr}_t} + \boldsymbol{\beta_i^{nt}} \underbrace{(\mathbf{F^{nt}_t} - \alpha_i \mathbf{F^{nt}_{t-1}})}_{X^{*,nt}_t} + \eta_{i,t} \qquad (13)$$

For each series ($i = 1, \ldots, N$), the factor loadings can be sampled from the posterior distribution defined as:

$$\beta_i^k \sim N(\overline{\beta_i^k}, V_i)$$

where

$$\begin{aligned} \overline{\beta_i^k} &= \left(\theta^{-1} + \frac{1}{R_i}(X_t^{*,k})'(X_t^{*,k})\right)^{-1} \left(\theta^{-1} B_{k,i} + \frac{1}{R_i}(X_t^{*,k})'(Y^*_{i,t})\right) \\ V_i &= \left(\theta^{-1} + \frac{1}{R_i}(X_t^{*,k})'(X_t^{*,k})\right)^{-1} \end{aligned}$$

After each draw, the factor loadings a checked to ensure that tradable (non-tradable) CPI inflation is positively loaded onto the tradable (nontradable) factors. If it is not, the draw is discarded and a new random draw of the factor loadings is made. This imposes the sign identification onto the model.

The over-identifying restriction is also imposed at this step by fixing the factor loadings $\boldsymbol{\beta_i^k}$ (for $k = \{tr, nt\}$) to zero if series i is not classified as a tradable or nontradable series.

Step 2: Sample serial correlation coefficients

Using the updated factor loadings found in the previous step, we are able to compute the prediction error ($\nu_{i,t}$) of equation 8. And using equation 10 we are able to sample α_i from the following posterior distribution:

$$\alpha_i \sim N(\overline{\alpha_i}, V_i)$$

where

$$\overline{\alpha_i} = \left(\nu'_{i,t-1}\nu_{i,t-1}\right)\left(\nu'_{i,t-1}\nu_{i,t}\right)$$
$$V_i = \left(\nu'_{i,t-1}\nu_{i,t-1}\right)^{-1} R_i$$

When drawing the new estimates of α_i, only those draws with an absolute value less than unity are retained. This ensures the stationarity of the non-core series.

Step 3: Sample R_i

Using the new estimate of α_i for each series, we are able to compute the error term $\eta_{i,t}$ from equation 10 and drawn a new R_i estimates from Inverse Gamma posterior distribution:

$$R_i \sim N(\eta'_{i,t}\eta_{i,t}, T)$$

Where T is the sample length.

Step 4: Sample the coefficients of the factor transition equation

For each factor, the coefficients ($A = [\rho^k_{t-1} \quad \rho^k_{t-2}]'$) from equation 9 are sampled from the normal posterior distribution:

$$A \sim N(\overline{A}, Z)$$

where

$$\overline{A} = (X'X)^{-1}\left(X'F^k_t\right)$$
$$Z = (X'X)^{-1}$$

and where $X = [F^k_{t-1}F^k_{t-2}]$.

When drawing the new estimates of the coefficients, I only retain those draws that imply a stationary process for the factors.

We do not have to sample the variance of the shock to the transition equations ($var\left(\varepsilon_{a,t}\right) = Q_a$) as this is fixed to unity to ensure the scale of the factor and factor loadings can be uniquely identified.

Step 5: Sample factors using Kalman smoother

Using the state space representation of the model (equations 11 and 12), the factor estimates are updated, conditional on the other parameters found in steps one to four, using the Carter and Kohn (1994) algorithm, which employs the Kalman smoother. The reader is referred to chapter eight of Kim and Nelson (1999) for further details on the implementation of the Carter and Kohn (1994) algorithm.

Repeat steps 1-5

Steps one to five are repeated for each iteration of the Gibbs sampling algorithm. After an initial burn in period, the draws are saved and used in the computation of all the statistics and graphs in the paper.