BANK OF CANADA
BANQUE DU CANADA

Discussion Paper/Document d'analyse
2015-12

A Comprehensive Evaluation of Measures of Core Inflation for Canada

by Mikael Khan, Louis Morel and Patrick Sabourin

www.bank-banque-canada.ca

Bank of Canada Discussion Paper 2015-12

October 2015

A Comprehensive Evaluation of Measures of Core Inflation for Canada

by

Mikael Khan,[1] Louis Morel[2] and Patrick Sabourin[1]

[1]Canadian Economic Analysis Department
[2]International Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
mkhan@bankofcanada.ca
lmorel@bankofcanada.ca
psabourin@bankofcanada.ca

ISSN 1914-0568

Acknowledgements

The authors are grateful to Sharon Kozicki, Marc-André Gosselin, Stephen Murchison, Rhys Mendes, Eric Santor, Jean Boivin and Gregor Smith for comments and helpful suggestions. This paper also benefited from discussions with participants at two Bank of Canada research seminars. The authors would like to extend a special acknowledgement to Laurence Savoie-Chabot for her work on trimmed mean measures of inflation. David Schefter provided great research assistance.

Abstract

This paper evaluates the usefulness of various measures of core inflation for the conduct of monetary policy. Traditional exclusion-based measures of core inflation are found to perform relatively poorly across a range of evaluation criteria, in part due to their inability to filter unanticipated transitory shocks. In contrast, measures such as the trimmed mean and the common component of CPI perform favorably, since they better capture persistent price movements and tend to move with macroeconomic drivers. All measures of core inflation, however, have limitations – consequently, there is merit in monitoring a set of measures. Moreover, core inflation measures are best viewed as complements to, rather than substitutes for, the thorough analysis of inflation and capacity pressures that informs the monetary policy process.

JEL classification: E31, E52
Bank classification: Inflation and prices; Monetary policy framework

Résumé

Dans cette étude, les auteurs évaluent l'utilité de diverses mesures de l'inflation fondamentale pour la conduite de la politique monétaire. À la lumière de toute une série de critères d'évaluation, ils constatent que les mesures traditionnelles à exclusion de composantes font plutôt pâle figure, ce qui s'explique en partie par leur incapacité à faire abstraction des chocs temporaires inattendus. En revanche, les mesures telles que la moyenne tronquée et la composante commune de l'IPC (indice des prix à la consommation) se démarquent avantageusement, car elles saisissent mieux les variations durables des prix et suivent habituellement l'évolution des déterminants macroéconomiques. Toutes les mesures de l'inflation fondamentale présentent toutefois des lacunes. Il y a donc lieu de surveiller un ensemble de mesures. Par ailleurs, il vaut mieux reconnaître les mesures de l'inflation fondamentale non comme des substituts, mais comme des compléments à l'analyse approfondie de l'inflation et des pressions sur la capacité de production, analyse qui vient éclairer le processus d'élaboration de la politique monétaire.

Classification JEL : E31, E52
Classification de la Banque : Inflation et prix; Cadre de la politique monétaire

1. Introduction

Since 1991, monetary policy in Canada has been conducted within an inflation-targeting framework. The inflation target in Canada is expressed in terms of the annual rate of increase of the consumer price index (CPI), with the Bank of Canada aiming to keep inflation at the 2 per cent midpoint of a target range of 1 to 3 per cent. In the short run, movements in the total CPI tend to be noisy, so measures of core inflation can be used to help the Bank focus on the underlying trend in inflation.

CPIX, a measure of core inflation that excludes eight of the most volatile components of the CPI and the effect of indirect tax changes on the remaining components, has served as the Bank of Canada's main operational guide for monetary policy since 2001. The Bank also monitors several other measures of core inflation. Along with CPIX, these measures have been periodically evaluated to assess their usefulness as indicators of underlying inflation. The most recent evaluation (Armour 2006) concluded that, despite some important limitations, CPIX retained advantages over the alternatives.[1] It was also stressed that many of the measures of core inflation contain relevant information about underlying inflation, and that regular research would be required to ensure the Bank continues to use the most reliable measures.

The contribution of this paper is threefold. First, we extend the evaluation criteria previously used to assess measures of core inflation in Canada. Second, we evaluate a recently introduced measure of core inflation – the common component of CPI – which had yet to undergo a comprehensive evaluation alongside other competing measures. Finally, we take advantage of the fact that we now have over two decades of data spanning Canada's inflation-targeting regime to examine the properties of core measures of inflation.

Our main findings are:

- Traditional exclusion-based measures of core inflation, such as CPIX, have the advantage that they are easy to understand and explain. However, they perform relatively poorly across a range of evaluation criteria, in part due to their inability to filter unanticipated transitory shocks.
- Trimmed mean measures and the common component of CPI perform favorably, since they better capture persistent price movements and tend to move with macroeconomic drivers.
- Overall, no single measure of core inflation dominates across all the evaluation criteria, providing support for the practice of monitoring a set of measures to help assess underlying inflation.

The following section provides a brief discussion of the intuition behind different methods of measuring core inflation. We then lay out the criteria across which these measures are evaluated, present and discuss key results, and provide some concluding remarks.

[1] For instance, electricity prices and auto insurance premiums, which are not excluded from CPIX, had undergone periods of heightened volatility shortly after CPIX was first introduced.

2. Measures of Core Inflation for Canada

There are various ways to measure core inflation. This section describes six measures used at the Bank of Canada, which are plotted alongside total CPI inflation in Figure 1.

CPIX, CPIXFET

A common approach to filtering the noise in total CPI inflation is to exclude a fixed set of components from the CPI basket that have historically displayed particularly volatile price movements. The intuition for this approach is that price movements in such items are likely to be transitory, and therefore of little relevance to monetary policy actions that can take a considerable time to affect inflation. For instance, retail gasoline prices are sensitive to shocks affecting global crude oil prices that often reverse themselves relatively quickly. The Bank of Canada's main core inflation measure (CPIX) is partly based on this intuition, as was its predecessor (CPI excluding food, energy and indirect taxes).

Trimmed mean

Traditional exclusion-based measures omit a pre-specified list of items from the CPI basket, but in practice large transitory price movements can occur in other components not excluded from such measures. In addition, the excluded components may at times contain relevant information about underlying inflation. A more flexible approach is to exclude the impact of different components each month based on whether their price changes are extreme at *that specific point in time*. More specifically, trimmed mean measures of inflation exclude components whose rates of change in a given month are located in the tails of the distribution of price changes.

CPIW

Another approach to reducing the volatility in total CPI inflation is to simply attribute lower weights to volatile components rather than exclude them altogether. This entails assigning a weight to each CPI component that is inversely proportional to its historical volatility. The volatility-weighted measure used at the Bank of Canada is called CPIW.

Wmedian

A common theme of the aforementioned measures of core inflation is that they can be thought of as weighted averages of disaggregated inflation rates. However, if the distribution of price changes is skewed, then the median may be a better measure of central tendency than the mean. The weighted median (Wmedian) measure of core inflation simply corresponds to the price change located at the 50th percentile (in terms of CPI basket weights) of the distribution of price changes in a given month. Note that the weighted median represents an extreme case of a trimmed mean, since all but the midpoint of the distribution is trimmed.

Common component of CPI

A notable development in more recent years has been the growing application of factor models to the task of measuring core inflation. These methods remove the influence of idiosyncratic or sector-specific price movements by extracting the part of inflation that is common across the individual prices comprising the CPI. Common movements in prices are likely to be indicative of aggregate demand fluctuations in the economy. The *common component* of CPI is a factor model-based measure of underlying inflation constructed for Canada by Khan, Morel and Sabourin (2013).

The remainder of this paper focuses on the six measures of core inflation described above. It is important to note that measures of these types are not exclusive to the Bank – they are used by other

major central banks (Table 1). In fact, most central banks in major advanced economies tend to consider numerous measures of core inflation.[2] While the focal measure is often an exclusion-based measure, a number of central banks appear to give equal importance in their official publications to the trimmed mean and weighted median.[3] Factor model-based measures are less common, but such a measure is used prominently in New Zealand.

Table 1. Select central bank practices regarding core inflation[a]

	Exclusion-based	Trimmed mean	Weighted median	Volatility-weighted	Factor model
U.S. Federal Reserve System	✪	☑	☑		☑
European Central Bank	✪				
Bank of England	☑				
Bank of Japan	✪	☑			
Swiss National Bank	☑	☑			
Reserve Bank of Australia	☑	✪	☑		
Reserve Bank of New Zealand		☑	☑		✪
Sveriges Riksbank	✪	☑		☑	
Norges Bank	✪	☑	☑		

a. The Bank of England and the European Central Bank have recently published analyses using a broader set of measures than what is shown in the table. These were excluded since they are not yet a regular feature of these central banks' communications.

✪ = focal measure

[2] In undertaking this work, many other measures of core inflation were also evaluated. These included variants of the six measures, deflator-based measures (such as the personal consumption deflator excluding food and energy), and persistence-weighted measures (following Bilke and Stracca 2008). None of these altered the conclusions of what follows.

[3] Note that the versions of the trimmed mean and weighted median evaluated in this paper incorporate some technical improvements relative to the ones previously used at the Bank. These modifications are detailed in the appendix.

Figure 1. Measures of core inflation (year-over-year percentage change)

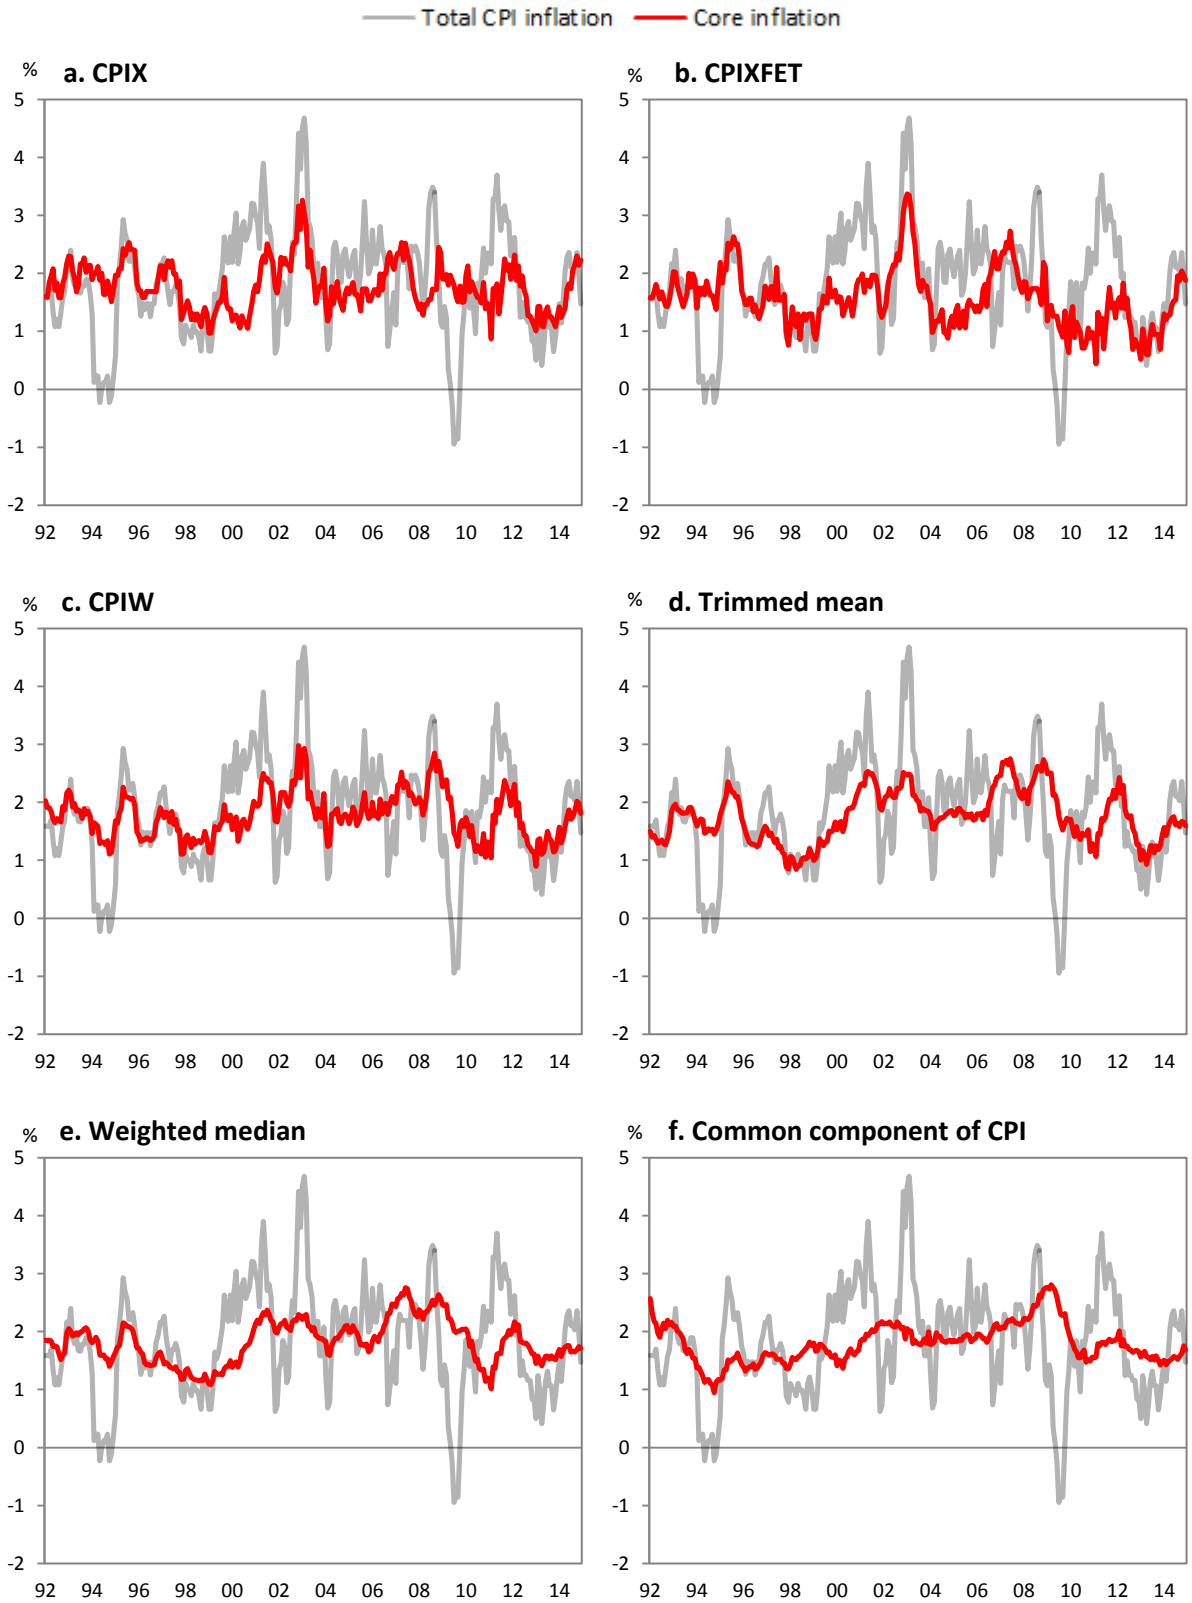

3. Evaluation of Core Inflation Measures

As discussed by Côté (2014), there are a number of criteria that an effective measure of core inflation must satisfy. We begin by examining whether the various measures of core inflation have desirable statistical properties. In particular, we investigate whether these measures track long-run movements in total inflation (i.e. whether they are unbiased), are less volatile than total inflation, and capture persistent movements in prices. This is followed by an empirical evaluation of the relationship between measures of core inflation and the output gap. We also test whether core inflation measures can reliably predict future total inflation. Finally, we discuss the extent to which different measures are easy to understand and explain to the public.

3.1 Statistical properties

The first step in evaluating the various measures of core inflation is to assess their statistical properties. Our framework for this evaluation is based on the following decomposition, which expresses total CPI inflation as the sum of core inflation (intended to capture the signal) and a residual (the noise):

$$\pi_t = \pi_t^C + e_t. \quad (1)$$

If a measure of core inflation is to serve as a useful operational guide for monetary policy, there are certain preferable statistical properties associated with this decomposition. In this section we focus on three statistical properties: bias, volatility and persistence.

Bias with respect to total CPI inflation

If core inflation is a useful proxy for the underlying trend in total CPI inflation, one would expect the two to share a similar long-term mean (i.e. e_t would on average be close to zero). This would support the credibility of core inflation as an operational guide in the conduct of monetary policy. In practice, the presence of bias in a core inflation measure can be accounted for by the central bank. However, differences in the average inflation rates of core and total inflation can complicate communication regarding the underlying rate of inflation (Mishkin 2007). Therefore, a desirable statistical property of core inflation is that it should be an unbiased measure of total inflation.

Table 2. Bias of core inflation measures[a]

CPIX	-0.01
CPI ex food, energy, taxes	-0.23
CPIW	-0.02
Trimmed mean	-0.03
Weighted median	0.03
Common component of CPI	0.00

a. Average difference in Y/Y growth rates of core and total CPI excluding tax (percentage points) from 1992-2014

To assess the potential bias in core inflation measures, their average inflation rates are compared to that of total CPI inflation. As shown in Table 2, there is little to choose between alternative core inflation measures on the basis of this criterion alone. With one exception, the bias of each measure is less than a tenth of a percentage point. The exception is CPI excluding food, energy and indirect taxes, which has averaged 0.2 percentage points lower than total CPI inflation since the inception of inflation targeting in Canada. This reflects the fact that food and energy prices have grown at a particularly fast pace in recent years, consistent with the large and persistent increases in commodity prices observed over the same period.

Volatility and persistence

It is common in the construction of core inflation measures to assume that minimizing volatility is the equivalent of reducing "noise." Therefore, core inflation measures also tend to be evaluated in part on the basis of their volatility relative to that of total CPI inflation. It is intuitive that core inflation should be less volatile than total inflation, but volatility and persistence are statistically distinct concepts. Thus, reducing the volatility in inflation does not necessarily amount to reducing the influence of transitory price changes.

To illustrate this point, Figure 2 displays two fictitious time series with the same volatility (as measured by their standard deviation), but with different degrees of persistence. Based solely on volatility, one would be indifferent between the two. However, it is difficult to argue that both warrant equal importance when considering policy actions that can take a considerable time to affect inflation. For this reason, we argue that the volatility of core inflation should not be seen independently from its persistence.

Figure 2. Two series with the same volatility but different persistence

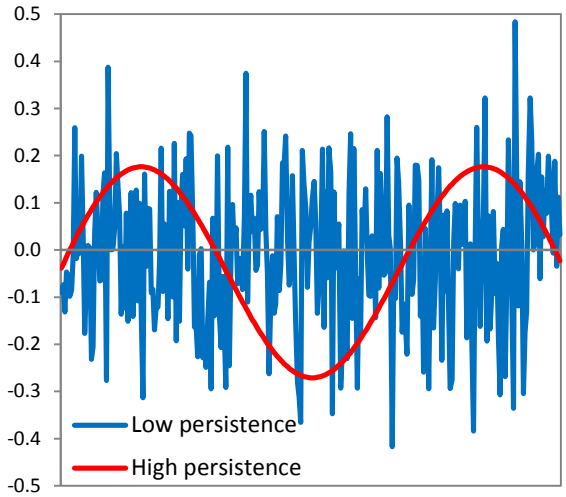

Table 3 shows both the volatility and estimated persistence of the core inflation measures and their respective residual components. Volatility is measured by the standard deviation of year-on-year inflation rates, while persistence is estimated as the sum of autoregressive coefficients from univariate regressions of quarter-on-quarter inflation rates.

We see that all the core inflation measures are less volatile than total inflation. The common component of CPI is found to be the least volatile of the measures. This reflects the fact that common price movements among components of the CPI basket account for only a small portion of the variability in total inflation. Indeed, Khan, Morel and Sabourin (2013) find that about 80 per cent of the variance of inflation is instead driven by sector-specific price movements. Not surprisingly, measures that are designed to exclude or down-weight volatile items (CPIX, CPIW) also rank among the least volatile.

In terms of persistence, the common component of CPI again ranks at the top. The trimmed mean, weighted median and CPIXFET are also found to exhibit a statistically significant degree of inflation persistence. Other measures display little to no persistence. This is also true of the residual components. However, a noteworthy result is that the portion of inflation excluded from CPIX is more persistent than CPIX itself. In other words, it appears that CPIX may have on average included a greater share of transitory price changes than it has excluded.

Table 3. Volatility and persistence of underlying inflation measures				
	Volatility[a]		Persistence[b]	
	π_t^c	e_t	π_t^c	e_t
CPIX	0.39	0.74	-0.04	0.10
CPI excl. food, energy, indirect taxes	0.48	0.66	0.40*	-0.12
CPIW	0.38	0.61	0.11	0.03
Trimmed mean	0.44	0.63	0.60*	-0.36
Weighted median	0.37	0.72	0.65*	-0.22
Common component of CPI	0.34	0.83	0.80*	-0.08
Total CPI (excl. indirect taxes)	0.80	---	-0.17	---

a. Standard deviation of Y/Y growth rate over 1992Q1-2014Q4
b. Sum of first five autoregressive coefficients on Q/Q inflation over 1992Q1-2014Q4
*Statistically significant at the 10 per cent level

3.2 Macroeconomic content

Examining the statistical properties of core inflation measures is an important first step, but desirable statistical properties are useful only to the extent that they are indicative of economically meaningful phenomena. We first evaluate the macroeconomic content of the core inflation measures through an empirical investigation of their relationship with macroeconomic drivers, focusing in particular on the output gap. With some caveats, we then explore whether the core inflation measures help predict future total inflation.

Relationship with macroeconomic drivers

In order to assess empirically the relationship between core inflation and the output gap, we begin by plotting simple correlations. Figure 3 shows the correlation between the level of the Canadian output gap[4] and the various measures of core inflation, expressed in 12-month rates of change, at different horizons. For instance, *t*+2 denotes the correlation of the output gap at time *t* with core inflation two quarters ahead.

Figure 3. Correlations with the output gap

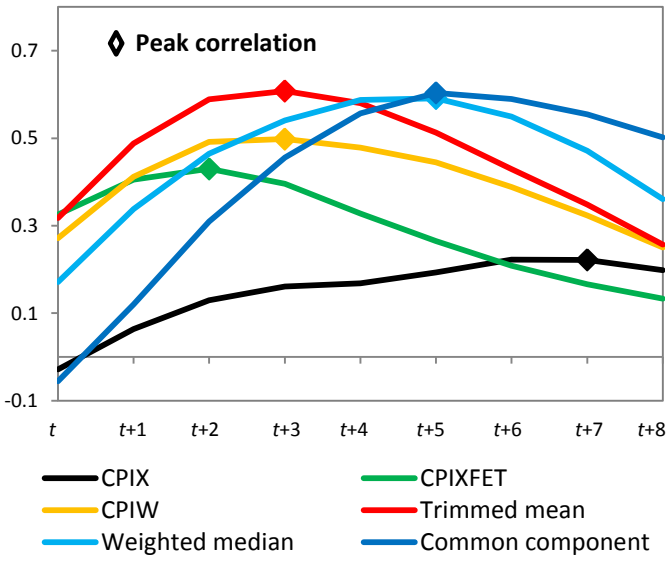

The measures that display the strongest correlation with the output gap are the common component of CPI, the trimmed mean and the weighted median. The peak correlations for these measures occur after 3 to 5 quarters, suggesting that they react with a lag to business cycle fluctuations. In contrast, CPIX inflation is only weakly correlated with the output gap at all horizons.

[4] All results reported in this section use the Bank's conventional estimate of the output gap but are found to be robust to using the structural estimate of the output gap instead.

Next, we report estimates from reduced-form Phillips curve models. We employ a standard specification in which quarter-on-quarter inflation is regressed on lags of the output gap and a set of control variables. An important element of this exercise that warrants further discussion is the role of mortgage interest cost (MIC) in some measures of core inflation. CPIX, by definition, excludes MIC on the theoretical basis that a decrease (increase) in the policy rate aimed at boosting (lowering) inflation could have the perverse impact of decreasing (increasing) inflation in the short run through its direct impact on MIC. However, no other measure of core inflation excludes MIC. This is important because MIC declined significantly during the recent recession as the Bank lowered its policy rate. Therefore, the procyclicality of measures that attribute a high weight to MIC may be overstated. To shed light on the importance of this phenomenon, Table 4 reports results using core inflation measures as they are defined as well as those from modified versions of the measures that exclude MIC.

Table 4. Phillips curve estimates (1992-2014)

	Baseline		Excluding MIC	
	Sum of coefs. on output gap	Adj. R^2	Sum of coefs. on output gap	Adj. R^2
CPIX	0.02	-0.01	0.02	-0.01
CPIXFET	**0.28***	0.13	0.07	0.09
CPIW	**0.12***	0.05	0.05	0.00
Trimmed mean	**0.28***	0.27	**0.11***	0.07
Weighted median	**0.23***	0.23	**0.11***	0.08
Common component	**0.19***	0.36	**0.17***	0.33

* Statistically significant at the 10 per cent level

Consistent with the insights from correlation analysis, it is difficult to detect a statistical relationship between CPIX inflation and the output gap. In contrast, the results for all other measures are generally favorable. However, for CPIXFET and CPIW the exclusion of MIC makes a marked difference, completely eliminating what otherwise appears to be a positive and statistically significant relationship with the output gap. The estimated coefficients on the trimmed mean and weighted median also decline once MIC is excluded, but importantly remain positive and statistically significant. The relationship between the common component of CPI and the output gap is the most robust to whether or not MIC is included. This is an intuitive result, since a measure of inflation intended to capture common price movements should be relatively insensitive to removing one item.

Ability to forecast total inflation

The ability of core inflation measures to help predict future total inflation is typically viewed as a particularly desirable characteristic.[5] However, with over two decades of inflation targeting in Canada the appropriateness of this criterion is now debatable. For instance, Rowe and Yetman (2002) argue that nothing but the inflation target should help forecast inflation under a successful inflation-targeting regime, since the predictive ability of variables should be fully exhausted by the central bank when making policy decisions aimed at keeping inflation at target. Short-term deviations from target can arise, but the ability to forecast these deviations using measures of core inflation would indicate that they have not been optimally incorporated into monetary policy decisions, thus leaving room for improvement. Similar arguments have been made by Clinton (2006) and Rich and Steindel (2007).

[5] See for instance Bryan and Cecchetti (1994), Blinder (1997) and Cogley (2002).

An intuitive way to assess the predictive ability of core inflation measures is to simply use lagged values of these measures as forecasts for future total inflation and compute root mean square errors (RMSEs) as follows:

$$RMSE_i = \sqrt{\frac{\sum_{i=1}^{N}\left(\pi_{i,t}^{k} - \pi_{CPI,t+k}^{k}\right)^2}{N}}, \tag{2}$$

where $\pi_{i,t}^{k}$ denotes core inflation measure *i* over *h* quarters to time *t*. Figure 4 displays results for horizons of up to 12 quarters. The RMSEs are plotted relative to those obtained from forecasting total inflation with only its own lags (i.e. univariate forecasts). Values below one indicate superior forecasting performance relative to univariate forecasts.

Figure 4. Relative RMSEs
(1 = RMSE obtained using total inflation only)

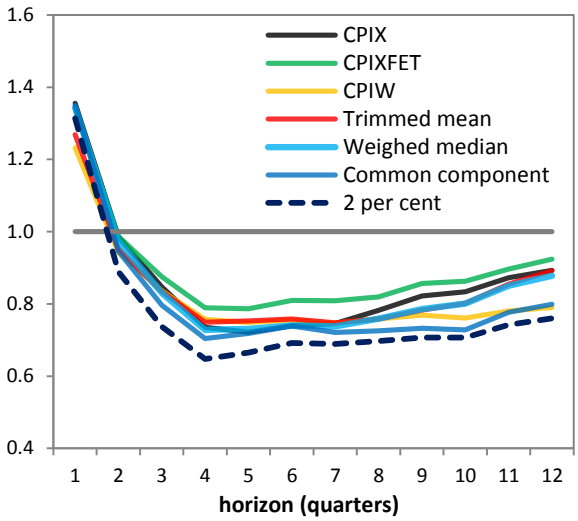

There are two main take-aways. First, we see that measures of core inflation are generally better predictors of total inflation than total inflation itself. Among the measures, the common component of CPI has the lowest RMSE at most horizons. Second, and most importantly, the best prediction of future total inflation is simply the inflation target: 2 per cent. This finding lends support to the arguments presented above, and suggests that forecasting performance is no longer a particularly useful criterion to discriminate between different measures of core inflation.[6]

4. Discussion of Empirical Findings

The results presented thus far have documented the empirical performance of core inflation measures across an extended list of criteria relative to previous Bank evaluations. In addition, this work has drawn upon almost 10 years of additional data since the last evaluation. Overall, these results appear to favor the trimmed mean, weighted median and common component of CPI. In contrast, there is relatively less evidence in support of CPIX, CPIXFET or CPIW.

One reasonable criticism of this finding is that CPIX may have been a victim of its own success. In other words, because CPIX has served as the Bank of Canada's main operational guide, the Bank's own policy actions may have neutralized the impact of various shocks on this measure so as to reduce its signal. We certainly cannot rule this out definitively, but we are of the view that there are more fundamental reasons why CPIX performs poorly in the empirical evaluation. Indeed, it is telling that the three

[6] Of course, there are other methods that can be used to evaluate this criterion. We have also conducted tests using a regression framework and have examined both in- and out-of-sample results. None of those exercises altered the conclusion presented in this section.

measures that perform poorly are the ones that rely on historical volatilities of CPI components in their construction.

The experience of the past several years has shown that many different components of the CPI can undergo periods of extreme volatility, even if their historical behavior to that point had been relatively stable. Take, for instance, the surge in automobile insurance premiums in the early 2000s.[7] In the span of 12 months, the inflation rate of this category rose from 2.0 per cent to 30.6 per cent. At its peak this one component of the CPI basket was adding almost a full percentage point to CPIX inflation, and also contributed significantly to CPIXFET and CPIW. As seen in Figure 1, these three measures were running at close to 3 per cent during this episode. On the other hand, there was no material impact on the three other measures. The trimmed mean excluded auto insurance throughout 2002, while the weighted median was little affected given that insurance premiums were a clear outlier in the overall distribution of price changes. The common component of CPI was the least affected, since by construction this measure reduces the influence of sector-specific shocks. Thus, we interpret the underperformance of CPIX and other measures built on similar intuition as a consequence of their inability to filter such unanticipated transitory shocks that have pushed Canadian inflation away from target in recent years.

5. Practical Considerations

A critical step in evaluating measures of core inflation is to assess whether they help to articulate the conduct of monetary policy in an easy and transparent way. This is admittedly more subjective than other aspects of our evaluation, but is nevertheless important given that core inflation measures are used in part to help communicate policy decisions.

Traditional exclusion-based measures of core inflation are perhaps the most easily understood and accepted. However, being able to effectively communicate the construction of a core inflation measure can be quite independent of whether that measure facilitates effective communication of policy decisions. For instance, consider the case in which an idiosyncratic shock to a component not excluded from an exclusion-based measure temporarily pushes inflation away from target. The central bank then needs to communicate that it is excluding an additional component from a measure that is built on the assumption that what to exclude can be known beforehand. This can make the use of a core measure seem somewhat ad hoc.

Trimmed mean measures of inflation are quite appealing in this regard, since their flexibility to filter unanticipated shocks makes them relatively immune to this shortcoming of traditional exclusion-based measures. At the same time, their method of construction is similar enough to that of traditional exclusion-based measures so as to not create further complications in communications. The weighted median does, however, require some basic statistical knowledge to be understood. In contrast, factor-model based measures of core inflation, such as the common component of CPI, present a unique challenge to central banks. On the one hand, these measures have both intuitive appeal and compelling empirical support. However, the fact that their computation requires knowledge of advanced statistical methods means they are not easily understood or replicable.

It is also important to consider that no measure of core inflation can account for every type of shock that the central bank may wish to look through in conducting monetary policy. For example, in a small

[7] There are numerous other examples of this, such as the pronounced deflation in auto prices in 2007 and the run-up in meat prices in 2014.

open economy like Canada, currency movements can have a profound impact on consumer goods prices, which comprise about half of the CPI basket. Exchange rate pass-through (ERPT) can have a permanent effect on the level of prices but is transitory for the rate of inflation, and so does not necessarily warrant policy action as long as inflation expectations remain anchored.[8] At the same time, most measures of core inflation will tend to be affected to varying degrees by ERPT, since none of them is likely to be insensitive to price movements impacting almost half of the CPI basket.

Thus, choosing one or more measures of core inflation to help guide monetary policy involves important trade-offs. Ultimately, there is no substitute for the thorough analysis of inflation and capacity pressures that must inform the monetary policy process.

6. Conclusion

In this paper we have evaluated measures of core inflation across several different dimensions. Among the measures, the trimmed mean, weighted median and common component of CPI stand out as the top performers according to our empirical criteria: they capture persistent price movements and tend to move with macroeconomic drivers. Results for the other measures are not as favorable. While traditional exclusion-based measures, including CPIX, are easy to explain and have played a useful role, they tend to perform relatively less well in our empirical assessment. This is in part due to the failure of such measures to effectively filter unanticipated transitory shocks. Nevertheless, all measures have strengths and limitations, thus providing merit to monitoring a set of measures.

We would emphasize, however, that regardless of the measure(s) used to help guide monetary policy, core inflation is just one of many inputs into this process. There are times when no measure of core inflation will provide a fulsome assessment of underlying inflation, which is why it is important to consider them together with a detailed analysis of the drivers of inflation and broader measures of capacity pressures. These measures include, but are not limited to, labour market indicators, output gap estimates and insights gleaned from surveys of businesses and consumers.

[8] For more details, see Savoie-Chabot and Khan (2015).

References

Armour, J. 2006. “Evaluation of Core Inflation Measures.” Bank of Canada Working Paper No. 2006-10.

Bilke, L. and L. Stracca. 2008. “A Persistence-Weighted Measure of Core Inflation in the Euro Area.” European Central Bank Working Paper No. 905.

Blinder, A. 1997. “Commentary.” Federal Reserve Bank of St. Louis *Review* (May-June): 157–60.

Bryan, M. and S. Cecchetti. 1994. “Measuring Core Inflation.” In *Monetary Policy*, 195–215. Edited by G. Mankiw. Chicago: University of Chicago Press.

Clark, T. 2001. “Comparing Measures of Core Inflation.” Federal Reserve Bank of Kansas City *Economic Review* (Second Quarter): 5–31.

Clinton, K. 2006. “Core Inflation at the Bank of Canada: A Critique.” Queen’s Economics Department Working Paper No. 1077.

Cogley, T. 2002. “A Simple Adaptive Measure of Core Inflation.” *Journal of Money, Credit, and Banking* 43 (1): 94–113.

Côté, A. 2014. “Inflation Targeting in the Post-Crisis Era.” Remarks to the Calgary Chartered Financial Analyst (CFA) Society, Calgary, Alberta, 18 November.

Khan, M., L. Morel and P. Sabourin. 2013. “The Common Component of CPI: An Alternative Measure of Underlying Inflation for Canada.” Bank of Canada Working Paper No. 2013-35.

Mishkin, F. S. 2007. “Headline versus Core Inflation in the Conduct of Monetary Policy.” Speech at the Business Cycles, International Transmission and Macroeconomic Policies Conference, HEC Montréal, Montréal, Quebec.

Rich, R. and C. Steindel. 2007. “A Comparison of Measures of Core Inflation.” Federal Reserve Bank of New York *Economic Policy Review* (December): 19–38.

Rowe, N. and J. Yetman. 2002. “Identifying a Policymaker’s Target: An Application to the Bank of Canada.” *Canadian Journal of Economics* 35 (2): 239–56.

Savoie-Chabot, L. and M. Khan. 2015. “Exchange Rate Pass-Through to Consumer Prices: Theory and Recent Evidence.” Bank of Canada Discussion Paper No. 2015-9.

Appendix. Revisiting trimmed mean measures of inflation in Canada

Recent research conducted at the Bank of Canada revisited the methodology used to compute trimmed mean measures of inflation. It was concluded that both Meanstd and Wmedian stand to benefit from methodological changes along two main dimensions:

i. Frequency of price changes

A key issue with the calculation of trimmed mean measures concerns the choice of the price change frequency. The Bank has previously been using the distribution of year-on-year (y/y) price changes to calculate both Meanstd and Wmedian, but the use of month-on-month (m/m) price changes is much more common. In the latter case, a year-on-year measure of core inflation can simply be derived by cumulating the monthly changes. There are reasons to suspect that this higher frequency of price changes might produce a more informative measure. For example, if a component undergoes a big one-time price level shock, a trimmed mean measure derived from year-on-year data may exclude it in the month in question but also in the following 11 months. In contrast, a measure derived from month-on-month data would exclude the impact of a one-off shock at its onset while allowing potentially informative price movements to re-enter the calculation thereafter.

One issue with using month-on-month changes is that they must be seasonally adjusted, and seasonally adjusted data are subject to revision. However, tests concluded that revisions are likely to be small, averaging about 0.1 percentage point. These are comparable to the historical revisions associated with the common component of CPI.[9]

ii. How much to trim?

The second issue is what percentage of the CPI basket to trim. There is no agreed-upon criterion for determining the optimal trim, so practices tend to vary by country. For example, the Federal Reserve Bank of Cleveland's measure trims 8 per cent on each side of the distribution, while the Reserve Bank of Australia trims 15 per cent from both ends.

Meanstd excludes prices whose rate of year-on-year change is over or under 1.5 standard deviations from total inflation. This means that Meanstd actually trims a different percentage of the CPI basket each month. On average, Meanstd corresponds roughly to a 5 per cent trimmed mean, though the total amount trimmed in any given month can range from 1 per cent to 21 per cent. In previous work, the justification for this approach was that it ensures that only values far from the mean are excluded in a given month. In the end, deciding on the appropriate trim level requires a trade-off. Smaller trims might not exclude enough of the CPI basket to sufficiently eliminate all idiosyncratic shocks, but bigger trims might do the opposite in excluding potentially informative price movements. The full spectrum of potential trims was evaluated to arrive at a sound judgment.

[9] Seasonal adjustment was conducted at the bank using the X12 procedure. In the future, seasonally adjusted data at the required level of detail will likely be provided by Statistics Canada.

Evaluation

To evaluate alternative trimmed mean measures of inflation, the same criteria used in this paper to evaluate different measures of core inflation were employed. Figures A1 to A4 display how alternative trimmed mean measures perform across these criteria. Trimmed mean measures derived from the distribution of monthly price changes are labelled TM and the ones derived from annual price changes are labelled TY.

A vast majority of the trimmed mean measures were found to outperform Meanstd. In addition, measures derived from the distribution of monthly price changes were found to outperform those that use annual price changes. Overall, the results suggest that it is optimal to trim about 15 to 25 per cent of the lower and upper ends of the distribution. Since the marginal gain from trimming more than 20 per cent seems quite modest, the decision was made to adopt a 20 per cent trimmed mean.

For the weighted median, no decision needed to be made about how much to trim, since by definition the weighted median is a 50 per cent trimmed mean. However, the results indicated that the weighted median also benefits from using monthly rather than annual price changes in its computation. This modification was therefore made to the Bank's Wmedian as well.

Figure A1. Average deviation from total inflation (Y/Y, percentage points)

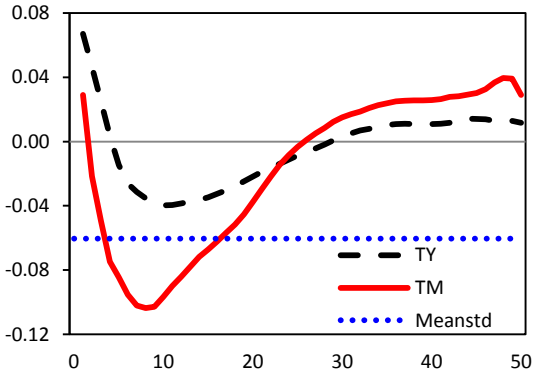

Figure A2. Volatility (standard deviation of Y/Y%)

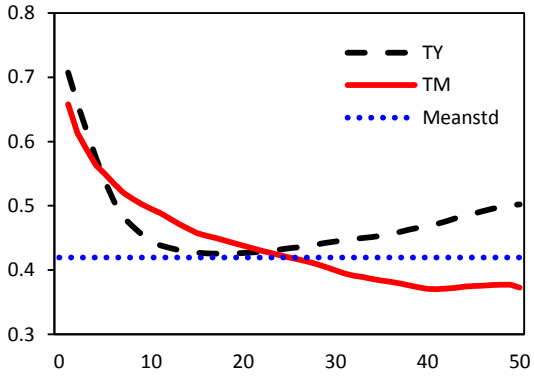

Figure A3. Persistence (sum of AR(5) coefficients on Q/Q%)

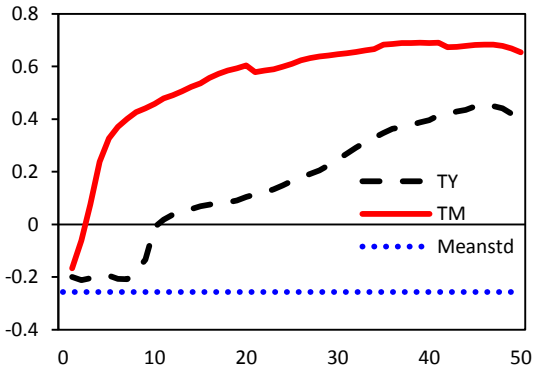

Figure A4. Responsiveness to economic slack (sum of coefficients on the output gap)