A PERSISTENCE-WEIGHTED MEASURE OF INFLATION FOR CANADA

by Mikael Khan

An essay submitted to the Department of Economics
in partial fulfillment of the requirements
for the degree of Master of Arts

Queen's University
Kingston, Ontario, Canada
July 2011

Acknowledgments

I am grateful to Gregor Smith for his patience, support, and guidance. I also thank Patrick Sabourin, Marc-Andre Gosselin, and Louis Morel for encouragement and helpful comments. Financial support from Queen's Economics Department (QED) and the Social Sciences and Humanities Research Council (SSHRC) is greatly appreciated.

Table of Contents

1. Introduction

The ability of the central bank to achieve price stability rests crucially on distinguishing persistent movements in inflation from transitory ones. This task is complicated by the high frequency noise in conventional measures of inflation, which provides the basis for focusing on "core inflation" to guide monetary policy. Measures of core inflation attempt to capture price movements that are most relevant for the forward-looking nature of monetary policy. This is typically achieved by excluding highly volatile prices. For example, the Bank of Canada targets a 2% annual increase in the consumer price index (CPI) but operationally refers to a measure that excludes 8 of its most volatile components.

Though a host of alternative measures of inflation have been suggested, and are indeed monitored by central banks, none of them explicitly measure what the central bank should be most concerned with: inflation persistence.[1] In fact, focusing on volatility can be misleading. For illustrative purposes, consider the variance of a first-order autoregressive process, $\sigma^2/(1-\rho^2)$, where σ^2 is the error variance. In this simple univariate framework, where persistence is measured by the autoregressive parameter ρ, it is clear that for a given σ^2 a highly persistent process may also be highly volatile. Hence, ignoring volatile prices can amount to ignoring those that contain important information for the conduct of monetary policy.

Since the implementation of inflation targeting in Canada, there has been a well-documented decline in aggregate inflation persistence (Benatti (2008), Mendes and Murchison (2010)). Nevertheless, there remains the possibility that at a disaggregated

[1] For an exhaustive account of the literature on inflation persistence, see Fuhrer (2009)

level certain prices display a greater degree of persistence than others. In fact, preventing persistent movements in inflation requires a clear understanding of persistence in the individual prices that comprise it. Yet, despite persistence being an important property of inflation, no detailed study in Canada has been conducted. Furthermore, no alternative measure of inflation monitored by the Bank of Canada explicitly accounts for it.

In this essay, I measure the persistence in disaggregated price data and study its evolution since the implementation of inflation targeting in Canada. This allows for a determination of the price movements that most warrant policy attention. Furthermore, it allows for an analysis of whether this answer is stable over time, despite the lack of regime change over the sample period. Finally, I exploit estimates of persistence to reweight the consumer price index according to the relative persistence of its components. This persistence-weighted measure of inflation provides a summary of underlying inflationary pressure in a more direct and appropriate manner than prevailing measures of core inflation.

There is a tendency in the literature on core inflation to evaluate core inflation measures on the basis of how accurately they forecast total inflation. The intuition is that these measures are intended to capture the underlying trend of inflation, which inflation itself should converge to in the long run. However, this approach is problematic under inflation targeting. One should expect the information content of core inflation measures to be reflected in the monetary policy decisions that keep inflation at (or near) target. If that is indeed the case, then such measures of inflation should not help forecast inflation at the horizon at which the central bank targets it (24 to 36 months, in Canada). The more appropriate evaluation criterion is to examine the relationship between core inflation measures and monetary policy decisions. That being said, there remains the possibility that alternative inflation measures contain useful information about future inflation that the central bank ought to exploit, but does not currently do so. This would be revealed if such measures were found to have useful forecasting properties for total inflation. Therefore, I evaluate the persistence-weighted measure of inflation on the basis of its relationship with the central bank's key policy rate, as well as its role in forecasting total inflation at various horizons.

2. Related Research

Persistence is an important property of inflation, yet despite consensus regarding its definition, there is no such consensus as to how it should be measured. Defining persistence as the speed of decay of a shock to inflation, one can infer that the concept is related to the impulse response function (IRF) of inflation. However, the IRF itself does not provide a useful measure of persistence, since it is an infinite length vector. This has motivated authors to focus on related scalar measures of persistence.

Pivetta and Reis (2007) argue that persistence is a univariate property and hence, the appropriate way to measure it is through univariate methods. This entails estimating an autoregression for inflation, and obtaining a measure of persistence based on the estimated parameters of the equation. Common measures are the sum of autoregressive coefficients, the largest autoregressive root, and the half-life of inflation. Pivetta and Reis apply all three of these methods to U.S. data, concluding that inflation persistence has been largely unchanged over the past 3 decades. Similarly, Levin and Piger (2004) employ the sum of autoregressive coefficients to show that, allowing for breaks in the intercept, U.S. inflation exhibits little persistence.

The univariate approach is also prevalent in central bank research. Lunneman and Matha (2004) estimate the sum of autoregressive coefficients for disaggregated price data in the Euro area, concluding that disaggregated prices do not display a high degree of persistence. Cutler (2001) estimates an AR(1) process for each component of the consumer price index in the UK, reweighting the CPI according to the relative magnitude of the estimated coefficients. Bilke and Stracca (2008) replicate this study for the Euro area, but allow for persistence to be measured by an AR(p) process instead. They find their persistence-weighted measure to be highly correlated with monetary policy decisions.

In Canada, Mendes and Murchison (2010) estimate an AR(1) process for inflation both before and after the implementation of inflation targeting. Not surprisingly, they document a substantial decline in inflation persistence since the inflation-targeting framework has been in place. This result is also evident in the work of Benatti (2008), who shows that Canadian inflation followed a unit root prior to inflation targeting but can currently not be distinguished from white noise. On disaggregated prices, Wilkinson (2011) examines statistical properties of CPI components at the aggregate and provincial level. Energy, shelter, and tobacco are shown to have the most volatile prices, while the magnitudes of certain large price movements are shown to be province-specific. The author suggests a "trimmed mean" approach to measuring core inflation may be preferable for capturing underlying inflationary pressure.

The dominance of univariate approaches to measuring persistence is motivated, at least in part, by the relative simplicity of computing them. However, the appropriateness of such measures has been questioned. For instance, Cogley, Primiceri, and Sargent (2007) compare univariate and multivariate approaches to examining inflation-gap persistence in the U.S. Measuring persistence based on predictability (an issue that I later discuss in detail), they argue that conditioning persistence on an information set better captures movements in persistence. This is evident in Cogley and Sargent (2005), where the authors document changes in the correlation of inflation with lags of itself as well as those of other macroeconomic variables. Hence, their preferred measure of persistence is obtained in a vector autoregression (VAR) that includes inflation. The central bank is ultimately concerned with persistence when formulating policy decisions, so it seems reasonable to condition estimates of persistence on the central bank's information set. Despite this argument for a multivariate approach, I consider both univariate and multivariate methods in this paper. The following section details my method.

3. Measuring Persistence

I follow Cogley, Primiceri, and Sargent (2007) in measuring persistence based on predictability. The intuition is that past shocks to inflation contribute to its future predictability, while future shocks cause forecast errors. Therefore, the continuing influence of past shocks can be measured by the proportion of predictable variation in inflation. The R^2 of j-step ahead forecasts is their preferred measure of persistence. For the purposes of this paper, the R^2 measure is attractive for two reasons. Firstly, it does not present a dramatic deviation from prevailing methods when used in the univariate context. To see why, consider the expression for the R^2 of an AR(1) process with time-varying parameters,

$$R_{jt}^2 = 1 - \left[\frac{\left(1-\rho_t^{2j}\right)\sigma_{\varepsilon t}^2/(1-\rho_t^2)}{\sigma_{\varepsilon t}^2/(1-\rho_t^2)} \right] \tag{1}$$

A little algebra shows the above expression simplifies to ρ_t^{2j}. Hence, persistence is only a function of the autoregressive coefficient ρ, which would still be the case were I to use other, more common univariate methods.

Now, consider a vector autoregession (VAR), also with time-varying parameters,

$$y_t = X'_{t-1}\theta_{t-1} + \varepsilon_{yt} \tag{2}$$

where y_t is a vector of endogenous variables, X_{t-1} includes a constant and lagged values of y_t, and ε_{yt} is the innovation vector. Any VAR(p) can be expressed in its companion form (see Enders (2003) for a simple derivation) as,

$$z_{t+1} = \mu_t + A_t z_t + \varepsilon_{zt+1} \tag{3}$$

The vector z_t contains current and lagged values of y_t, μ_t contains the intercepts, and the matrix A_t includes the autoregressive coefficients. Obtaining a scalar measure of persistence in this context is more problematic because persistence depends on a matrix of coefficients, A_t, instead of a single parameter ρ. One possible approach is to focus on the largest autoregressive root. However, simply ignoring additional roots may be misleading. The R^2 gets around this complexity by providing a scalar measure of persistence without discarding potentially useful information. Hence, all measures of persistence reported in this paper will be based on $R_{j.}^2$ Furthermore, all results in this paper are based on j=1. Naturally, persistence decays as the forecast horizon increases, the speed of decay being slower for more persistent components. Most persistence estimates found in this study are low, and hence converge rapidly to zero as the horizon increases. These results are less relevant for the purposes of this paper, and are available upon request.

To allow for changes in persistence over time, all estimates will be obtained using a 10-year rolling window. On the one hand, the stable monetary regime over the sample period I examine should not be associated with fundamental changes in the inflation process. However, it is still possible that sector-specific structural factors could have altered the persistence in certain disaggregated prices. Rolling estimates account for this possibility, while also measuring persistence in a manner that replicates the real-time inflation monitoring process of the central bank.

While measuring persistence is an important first step in determining the relative importance of prices for monetary policy, it is not sufficient. A persistent shock to a CPI component should only warrant policy attention if it comprises a significant portion of consumer spending. Hence, a true measure of policy relevance must account for both persistence and expenditure shares. This can be achieved by applying the following weighting method,

$$w_{it} = \frac{e_{it}\rho_{it}}{\sum_{i=1}^{n} e_{it}\rho_{it}} \tag{4}$$

where, for the i-th CPI component at time t, w_{it} is the weight it is assigned, ρ_{it} is the estimated persistence, and e_{it} is its corresponding expenditure share in the CPI basket. The weights sum to 1, reflecting the relative policy importance of each item in the CPI basket. These persistence-based weights are then used to construct a measure of inflation that explicitly accounts for the relative persistence in each component of the CPI.

To obtain univariate estimates of persistence, I estimate the following p-th order autoregressive process,

$$\pi_{it} = \alpha + \sum_{k=1}^{p_i} \rho_{p_i} \pi_{it-k} + \varepsilon_t \tag{5}$$

where π_{it} is the first-difference of the log of the i-th CPI component at time t. The order of the regression, p, is determined in each case by the Schwartz criterion. All results in this paper are robust to using the Bayesian Information Criterion (BIC) instead.

Multivariate persistence is estimated in a vector autoregression (VAR), given by equation (2). The vector y_t contains total inflation, its 8 components, the unemployment rate, and the three-month Treasury bill rate. Including unemployment and a short-term interest rate in the VAR follows Cogley and Sargent (2005) and Cogley *et al* (2007). Of course, the central bank's actual information set contains hundreds, if not thousands, of variables. For practical reasons and degrees of freedom restrictions, a wider information set cannot be incorporated. However, it is a reasonable approximation to view the central bank's information set in terms of real activity, inflation, and the stance of monetary policy. Furthermore, to measure persistence in disaggregated CPI components, it is important to allow for feedback between these components. For example, a positive shock to the shelter component of the CPI could well manifest itself in higher prices for household furniture and other household items. To that end, the VAR conditions estimates of persistence on appropriate variables.

4. Data

Inflation is measured as the log-difference of each component of the CPI. This data is available at various different levels of aggregation. While the univariate approach places no restriction on the level of aggregation used, this is not true of the VAR where I am restricted by degrees of freedom. Hence, I focus my analysis on 8 components that comprise the CPI. The CPI expenditure-based weights vary according to systematic updates to the basket conducted by Statistics Canada. For the period covered in this paper, these changes occur in 1995, 1998, 2003, 2004, and 2007.

For the VAR, the unemployment rate is for all individuals in Canada, 15 years and above. The short-term nominal interest rate is the return on three-month Treasury bills. The inflation and unemployment data is seasonally adjusted. The choice to focus on the inflation-targeting regime necessitates using data starting in 1992. Most recent data is for 2011m03. All data is published on a monthly basis by Statistics Canada. It is important to use data at a monthly frequency, as this mirrors the frequency at which the central bank becomes aware of new inflation data.

5. Predictability Findings

The choice of using a 10-year rolling window with data beginning in 1992 means all persistence estimates begin in 2002. Univariate persistence estimates are shown in Figure 1. Confidence intervals are based on the approximate standard error of R^2 (see Cohen *et al* (2003) for details). Inflation, by this measure, exhibits little persistence in the case of each CPI component. The overall range of the R^2 statistics is between 0 and 0.18. Food and shelter inflation rates exhibit the least persistence, followed closely by those of household items and transportation. The R^2 for these components remains close to zero throughout the sample. On average, inflation in clothing and footwear, and health and personal care exhibit the highest degree of persistence.

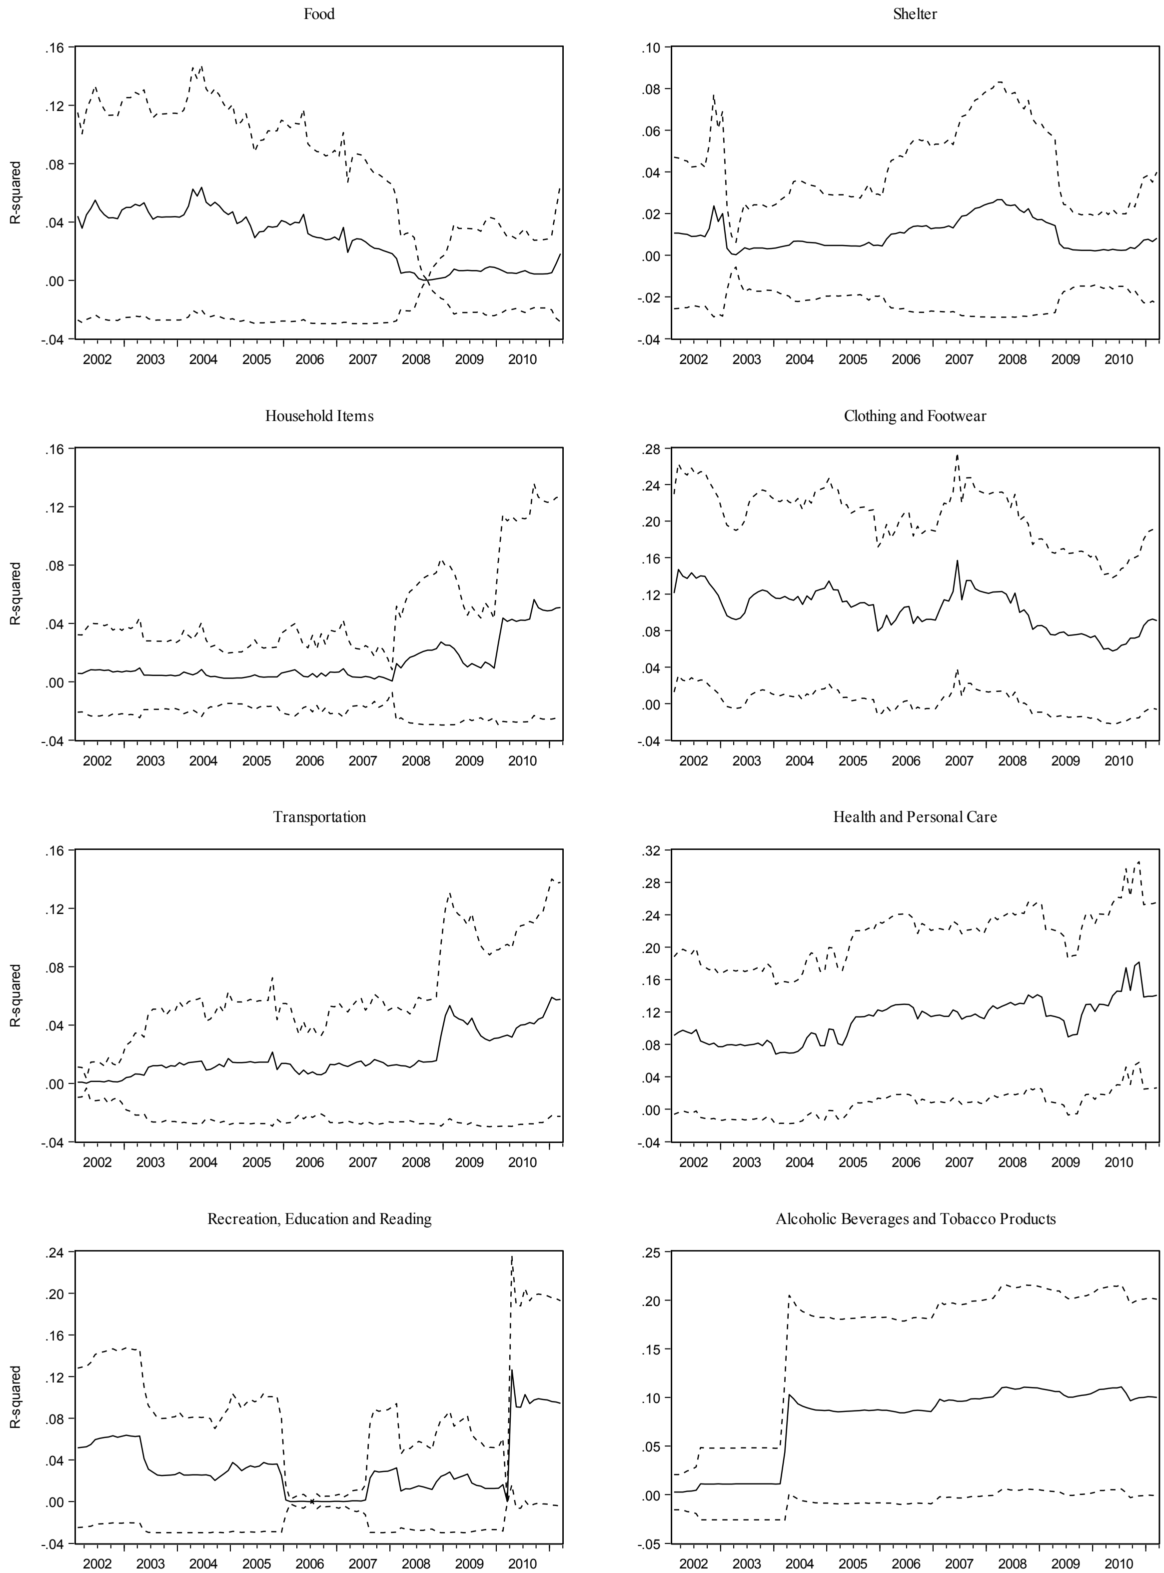

Figure 1. Univariate persistence estimates. Dotted lines are 95% confidence intervals

Univariate estimates of persistence appear to be relatively stable over the inflation-targeting regime. While there is some time variation in each case, it occurs within a narrow range. This is evidenced by the standard deviations of the R^2 statistics, which range from 0.01 to 0.04. However, there are certain changes in persistence that warrant further attention. The inflation rate of the transportation component experiences a sudden sustained increase in persistence in 2008m12. Similar changes occur for recreation, education and reading, and alcoholic beverages and tobacco products in 2010m4 and 2004m4, respectively.

To analyze whether these sudden pronounced changes are attributable to structural breaks, I conduct Chow breakpoint tests at the dates at which they occur. I test for both individual and joint breaks in the intercept and autoregressive coefficient. Results are shown in Table 1.

Table 1. Chow breakpoint test p-values (Univariate)

	Intercept	AR coefficient	All coefficients
Transportation	0.63	0.48	0.73
Recreation, education and reading	0.73	0.06	0.16
Alcohol and tobacco products	0.78	0.71	0.92

For inflation in transportation and alcoholic beverages and tobacco products, the null hypothesis of no breaks at the specified breakpoints cannot be rejected. In the case of recreation, education and reading, there is some evidence of a break in the autoregressive coefficient. The null cannot be rejected at the 5% level of significance but is rejected at 10% (the p-value is 0.06).

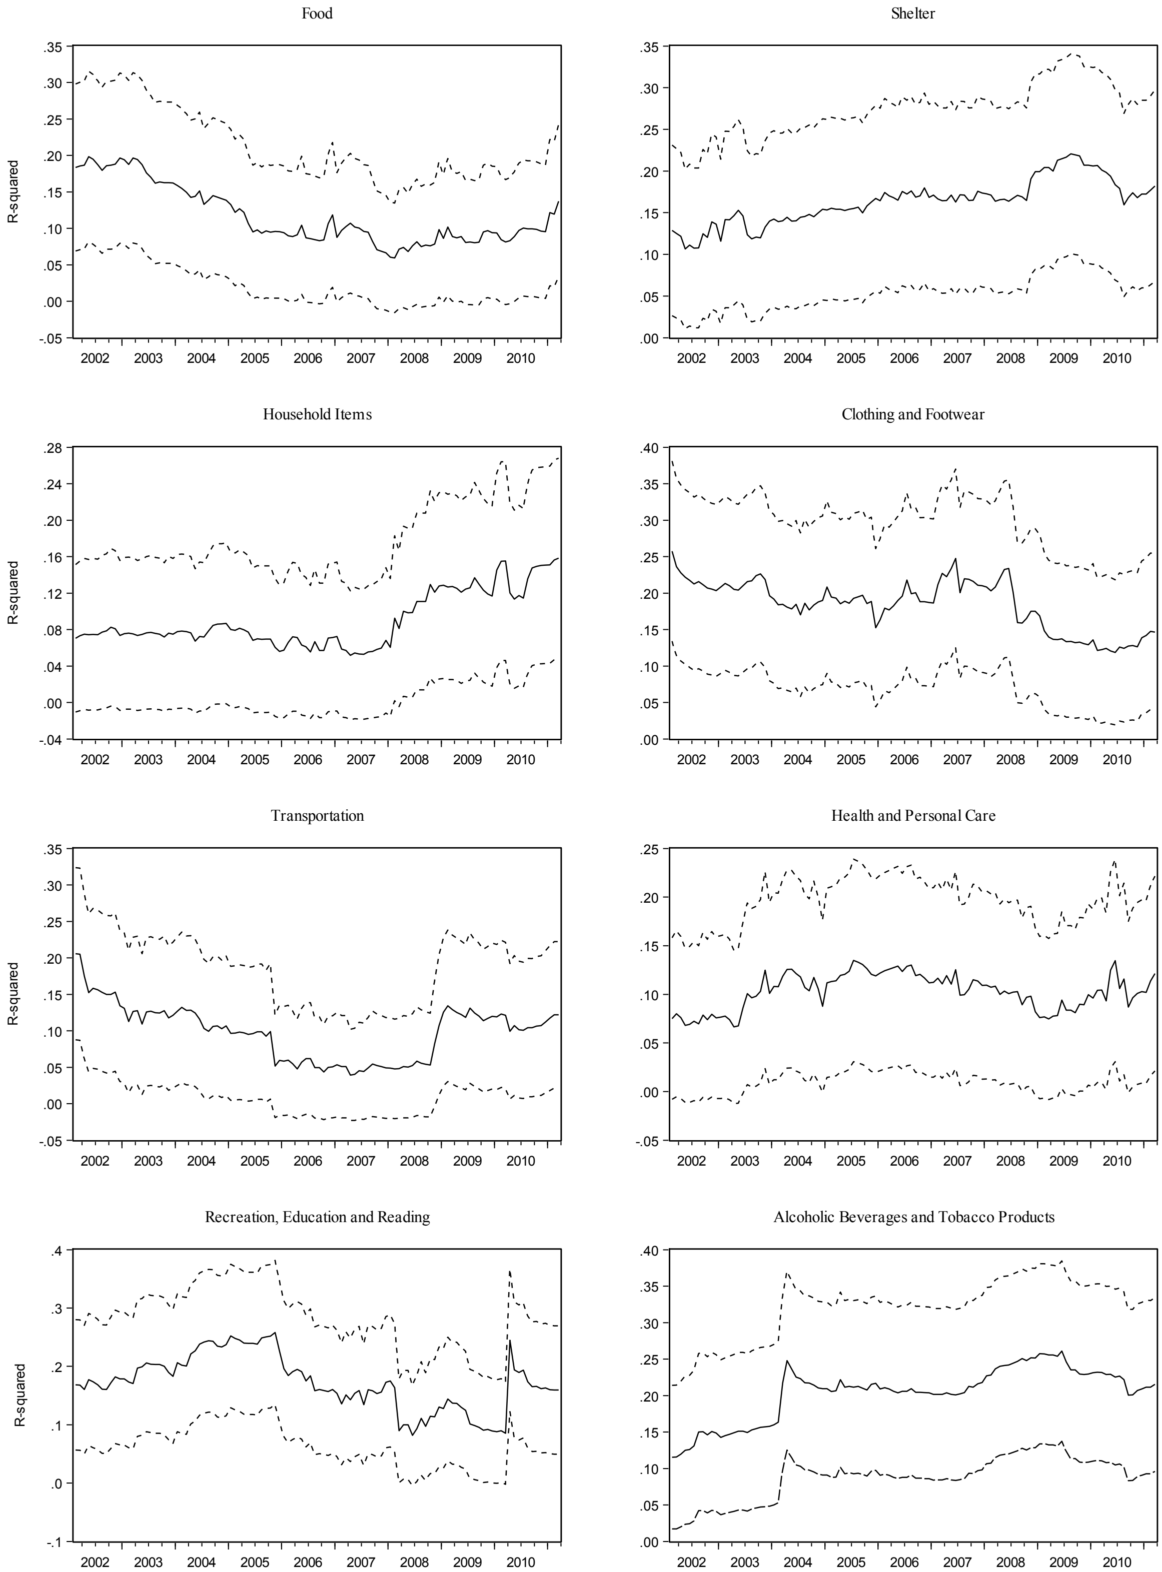

Figure 2. Multivariate persistence estimates. Dotted lines are 95% confidence intervals

Multivariate estimates of persistence differ from the univariate case. One result, evident in Figure 2, is that the inflation rate of each CPI component exhibits greater persistence when measured in a VAR. The most significant departure from the univariate case is that food and shelter inflation have average R^2 statistics of 0.12 and 0.16 respectively, having remained close to zero when measured using the univariate method. Clothing and footwear, and alcoholic beverages and tobacco products have the most persistent rates of inflation. Inflation in household items exhibits the least persistent. However, it is worth noting that there is little variation in the multivariate persistence estimates between components. While in the univariate case some components were substantially less persistent than others, all multivariate R^2 measures, on average, lie within a narrow range of 0.1 to 0.2.

Despite differences in magnitude, the evolution of the multivariate persistence estimates is similar to univariate case. There is modest time variation in the estimates. The standard deviations range from 0.02 to 0.05. The previously documented increase in the inflation persistence of transportation, recreation, education and reading, and alcoholic beverages and tobacco products, is also evident in the multivariate context. Once again I conduct Chow tests to examine whether these changes are the result of structural breaks. Results are detailed in Table 2.

Table 2. Chow breakpoint test p-values (Multivariate)

	Intercept	AR coefficient	All coefficients
Transportation	0.45	0.59	0.68
Recreation, education and reading	0.25	0.01*	0.07
Alcohol and tobacco products	0.01*	0.11	0.34

* the null of no break at the specified breakpoint is rejected at the 5% level of significance

Once again, there is no evidence of a structural break in the transportation series. However, there is now strong evidence of a structural break in the autoregressive coefficient for inflation in recreation, education and reading. There is also evidence of a break in the intercept for inflation in alcoholic beverages and tobacco products. In these cases, the null hypothesis of no break is rejected at the 5% level.

It is hence evident that examining inflation persistence in a VAR has important implications for both the absolute and relative persistence of CPI components. In absolute terms, each component displays greater persistence than in the univariate case. This is due to the gain in predictability from using additional relevant variables. In relative terms, the differences in persistence between components are not as pronounced as under the univariate approach. This will have important implications in the following section, where the estimates are exploited to reweight the CPI.

6. Reweighting the CPI

As described in section 3, I use equation (4) to generate persistence-based weights. These weights reflect the relative policy importance of each component in the CPI by explicitly accounting for their expenditure shares and persistence estimates. Naturally, since I estimate persistence using a rolling window, these weights also vary over time. This section presents a description of the persistence-based weights and compares them to the original CPI expenditure-based weights. The CPI expenditure-based weights shown in all tables are from the 2005 basket, but these weights have varied only trivially over the sample period.

Table 3. Univariate persistence-based weights

	Std. dev	Min	Max	Mean	CPI*
Food	0.11	0.00	0.35	0.17	0.17
Shelter	0.07	0.00	0.25	0.09	0.27
Household Items	0.04	0.00	0.18	0.04	0.11
Clothing and Footwear	0.05	0.08	0.31	0.20	0.05
Transportation	0.08	0.00	0.34	0.12	0.20
Health and Personal Care	0.04	0.10	0.26	0.17	0.05
Recreation, Education and Reading	0.09	0.00	0.38	0.12	0.12
Alcoholic Beverages and Tobacco Products	0.05	0.00	0.16	0.09	0.03

*2005 basket

Table 3 shows that the univariate weights are highly variable. This is particularly evident in the weight assigned to food, which has a standard deviation of 11 percentage points. 6 of the 8 components, at some point in time, are assigned a weight of zero. Clothing and footwear, health and personal care, and alcoholic beverages and tobacco products are (on average) assigned substantially greater weights than their corresponding expenditure-based weights in the CPI. The most heavily down-weighted component is shelter. It receives an average weight of just 0.09 in the persistence-weighted index, compared to its current expenditure weight of 0.27 in the CPI. On average, the highest weight is assigned to clothing and footwear.

Table 4. Multivariate persistence-based weights

	Std. dev	Min	Max	Mean	CPI*
Food	0.04	0.09	0.25	0.15	0.17
Shelter	0.06	0.20	0.42	0.33	0.27
Household Items	0.02	0.05	0.13	0.07	0.11
Clothing and Footwear	0.02	0.05	0.11	0.08	0.05
Transportation	0.04	0.07	0.25	0.14	0.20
Health and Personal Care	0.01	0.02	0.05	0.04	0.05
Recreation, Education and Reading	0.04	0.08	0.23	0.15	0.12
Alcoholic Beverages and Tobacco Products	0.01	0.03	0.07	0.05	0.03

*2005 basket

On the other hand, each weight obtained using the multivariate approach (Table 4) has a lower standard deviation than its corresponding univariate weight. Furthermore, no component is assigned a weight of zero at any point in time. On average, shelter is the highest weighted component. This is in stark opposition to the univariate case where shelter is amongst the lowest weighted. Hence, the two approaches have dramatically different policy implications.

A striking result is that, on average, the multivariate weights are remarkably similar to the original CPI expenditure-based weights. This is driven by the previously documented fact that in relative terms, multivariate estimates of persistence do not vary substantially between components. Hence, when adopting the double-weighting scheme outlined in (4) the expenditure-based weights have a more dominant impact on determining the final weights than do the persistence estimates.

Having constructed persistence-based weights of CPI components, I reweight the CPI basket accordingly. For both the univariate and multivariate cases, I construct an index with the base year 2002. This allows for a comparison with prevailing measures of inflation. Figure 3 plots the year-over-year evolution of these measures along with total inflation, while Figure 4 compares them to the Bank of Canada's official core inflation measure (CPIX). Table 5 contains descriptive statistics. I also include CPIW in the table. CPIW is an alternative measure of inflation closely monitored by the Bank of Canada. Instead of excluding certain volatile components, this measure weights components by the inverse of their volatility.

The average year-over-year growth rate of the univariate persistence-weighted measure is 1.35. This is lower than the average growth rates of total and core inflation over the sample period, and below the Bank of Canada's 2% inflation target. It is, however, well within the 1 to 3 percent target range. This measure is also smoother than total inflation, but substantially more volatile than core inflation. This is not surprising, as core inflation is constructed using a volatility criterion, and as previously discussed, persistence and volatility are different concepts.

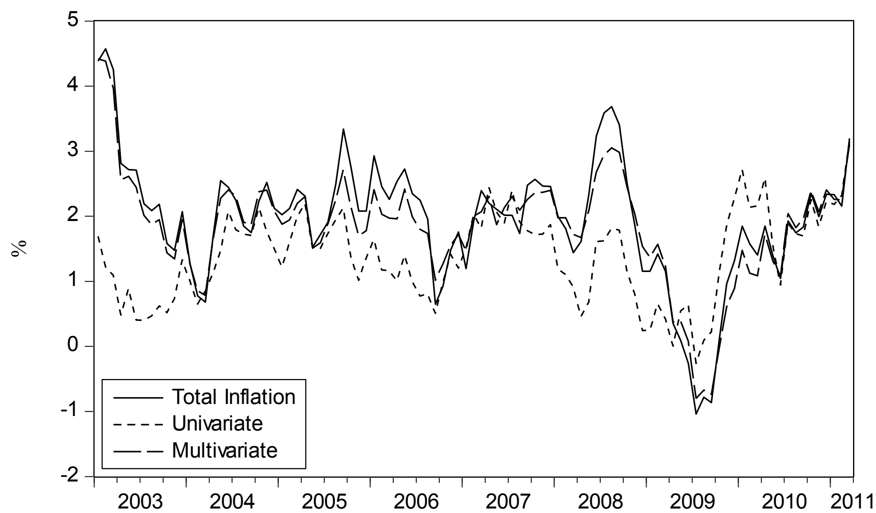

Figure 3. Comparison of year-over-year growth of persistence-weighted measures with total inflation

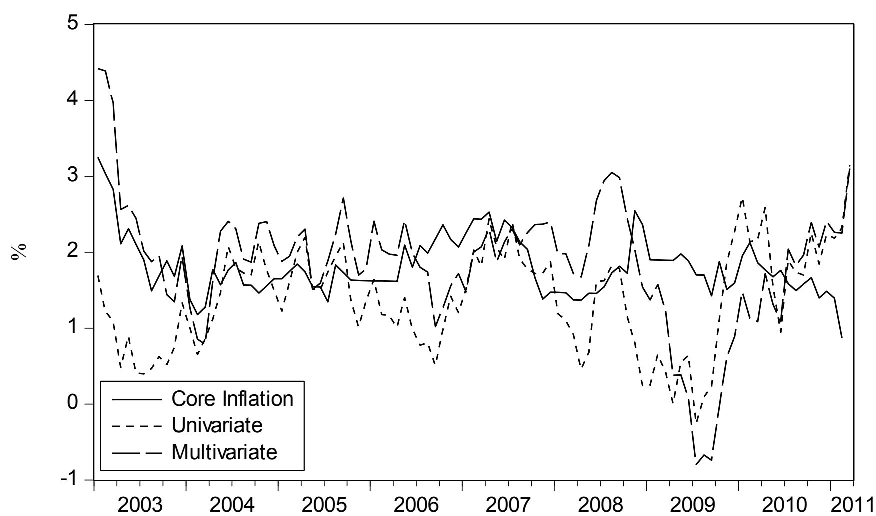

Figure 4. Comparison of year-over-year growth of persistence-weighted measures with core inflation

Table 5. Descriptive statistics of alternative measures of inflation
year-over-year percentage change

	Mean	Median	Maximum	Minimum	Std. Dev.
Univariate	1.36	1.44	2.72	-0.26	0.66
Multivariate	1.85	1.96	4.42	-0.80	0.86
Total inflation	1.94	2.02	4.57	-1.04	0.97
Core inflation	1.81	1.73	3.25	0.87	0.39
CPIW	1.88	1.85	2.90	1.10	0.42

The univariate persistence-based measure implies lower inflationary pressure than total inflation throughout much of the sample period. This begins to change during the recent economic downturn that started towards the end of the 2008. The persistence-weighted measure does not decelerate as sharply as total inflation during the recession, but accelerates more during the recovery. Since mid-2010 this measure closely tracks the movements in total inflation. This implies that the recent spike that has seen inflation rise above the 2% target may require close scrutiny from policymakers. This is even more evident when the measure is compared to the Bank of Canada's core inflation measure. The current gap between core inflation and the persistence-weighted measure is the largest it has been over the sample period.

It is evident that the measure based on the multivariate approach is very similar to total inflation. In fact, in many periods the two measures are indistinguishable. This is, of course, a result of the similarity between the weights used to compute both measures. The persistence-based measure is marginally less volatile than total inflation, but substantially more volatile than core inflation. The average growth rate of the persistence-weighted measure is 1.85, which is about the mid-point of the average growth rates of total and core inflation over the same period.

The multivariate persistence-weighted measure also implies lower inflationary pressure throughout a majority of the sample period. However, this difference is less pronounced than in the univariate case since the measure closely tracks total inflation. It also implies much stronger inflationary pressure than what is currently implied by the official core inflation measure.

7. Evaluation

The next step is to evaluate the usefulness of constructing a measure of inflation that explicitly accounts for the relative persistence of its components. First, I examine whether the persistence-weighted measures are correlated with monetary policy decisions. This correlation does not imply that the central bank could have set past policy in response to the measures derived in this paper. However, it is entirely possible that the central bank could have been reacting indirectly to the information content of these measures. Hence, a strong correlation with policy decisions suggests these measures contain important information for the conduct of monetary policy and that the central bank has been responding to this information. If these measures are not correlated with policy decisions, they may still contain information that the central bank ought to respond to. This is later ascertained by examining whether the persistence-weighted measures help forecast future total inflation.

Figure 5 plots the univariate and multivariate persistence-weighted measures of inflation along with the Bank of Canada's target for the overnight rate. Table 6 displays the correlation coefficient of each measure with the policy rate. The univariate persistence-weighted measure does not appear to track movements in the policy rate. The correlation coefficient of just 0.04 is indicative of this absence of a relationship. However, the multivariate measure is much more highly correlated with monetary policy decisions. In this case the correlation coefficient is 0.47, which is similar to that of total inflation with the policy rate.

Surprisingly, of the official measures, this correlation is weakest for core inflation. The strongest correlation is with CPIW. The coefficient of 0.49 is just above that of the multivariate persistence-weighted measure. Hence, conditioning persistence on the central bank's information set results in a measure that better reflects its decision making process.

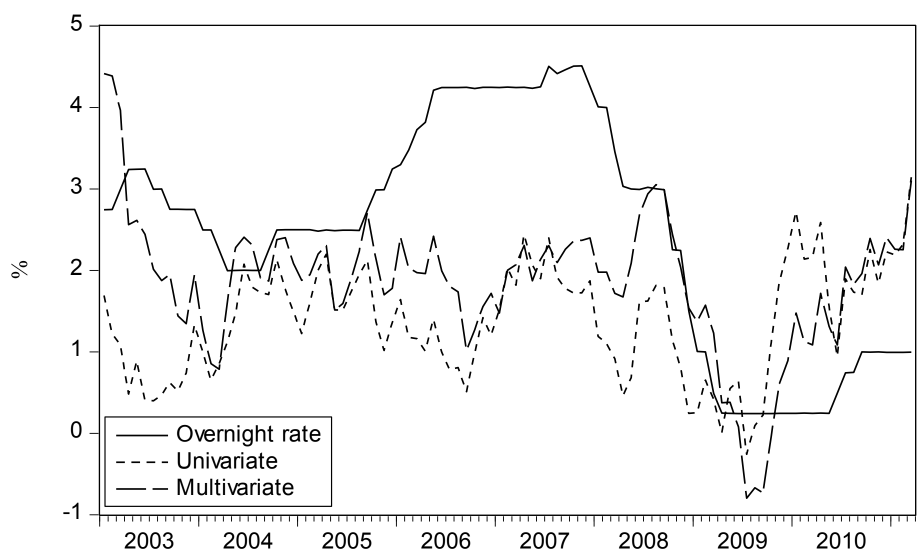

Figure 5. Persistence-weighted measures of inflation with the key policy rate

Table 6. Correlations with overnight rate

Univariate	0.04
Multivariate	0.47
Total inflation	0.46
Core inflation	0.26
CPIW	0.49

The absence of a relationship between the univariate measure and monetary policy decisions raises the possibility that the measure contains information about the future path of inflation, which the central bank ought to include in its information set. Though the evidence suggests this is less likely for the multivariate measure, I examine whether both measures are unbiased predictors of future inflation. I follow the method proposed by Cogley (2005) for examining the forecasting performance of alternative measures of inflation. It entails estimating the following equation,

$$\pi_{t+h} - \pi_t = \alpha_h + \beta_h\left(\pi_t - \pi_{core_t}\right) + \mu_{t+h} \tag{6}$$

where h is the forecast horizon, $\pi_t - \pi_{core_t}$ is the gap between total and core inflation, and $\pi_{t+h} - \pi_t$ is the subsequent change in total inflation. The intuition behind this approach is that the core deviation should be inversely related to the change in inflation, for a sufficiently large h. If the restriction $\alpha = 0$ and $\beta = -1$ holds, then equation (6) simply collapses to,

$$\pi_{t+h} = \pi_{core_t} + \mu_t \tag{7}$$

and the core measure is an unbiased predictor of future inflation. The restriction on α follows from the fact that both left and right-hand side variables should be mean zero. The restriction on β reflects whether the core deviation correctly captures the degree of temporary movements. A value of β that is negative but has an absolute value less than 1 implies the core deviation will overstate both current transients changes and the future change in inflation. If β is negative but greater than 1 in absolute value, these will be understated by the core deviation.

Figure 6 reports estimates of β at various forecast horizons, along with the R^2 obtained from estimating equation (6) at each forecast horizon. Each equation is estimated from 2003m1 to 2011m3-h. Confidence intervals are based on HAC standard errors. Ideally one should evaluate forecast performance out of sample, but the sample size prohibits this. Table 7 (see appendix) reports p-values from the joint test that $\alpha = 0$ and $\beta = -1$. I conduct the test only for the persistence-weighed measures of inflation. That CPIX and CPIW pass this test is well documented elsewhere (see Lafleche and Armour (2006)).

Using the univariate persistence-weighted measure in equation (6) results in β estimates that are not statistically different from -1, at the 6 and 12-month horizons. The measure also appears to have reasonable explanatory power at these horizons, evidenced by the R^2s of 0.6 and 0.5. However, it also evident that CPIX and CPIW outperform the univariate measure at these horizons. Furthermore, the hypothesis that $\alpha = 0$ and $\beta = -1$ is rejected at the 5% level at each forecast horizon. Hence, by this criterion, the univariate measure is not useful for forecasting total inflation. The failure of this measure to perform well in both evaluation criteria examined supports the possibility that univariate methods are not appropriate for measuring persistence.

Turning to the multivariate measure, it is clear that it does not perform at all well in the forecasting exercise. The β estimates vary substantially from -1 at most horizons, and the confidence intervals around these estimates are large. The R^2 statistics are well below those of other measures, and the test that $\alpha = 0$ and $\beta = -1$ is also rejected at each forecast horizon. However, this is not necessarily a negative result. Unlike the univariate measure, the multivariate measure is correlated with monetary policy decisions. Hence, the poor forecasting performance may be attributed to the information content of the measure already being reflected in policy that seeks to keep inflation low and stable.

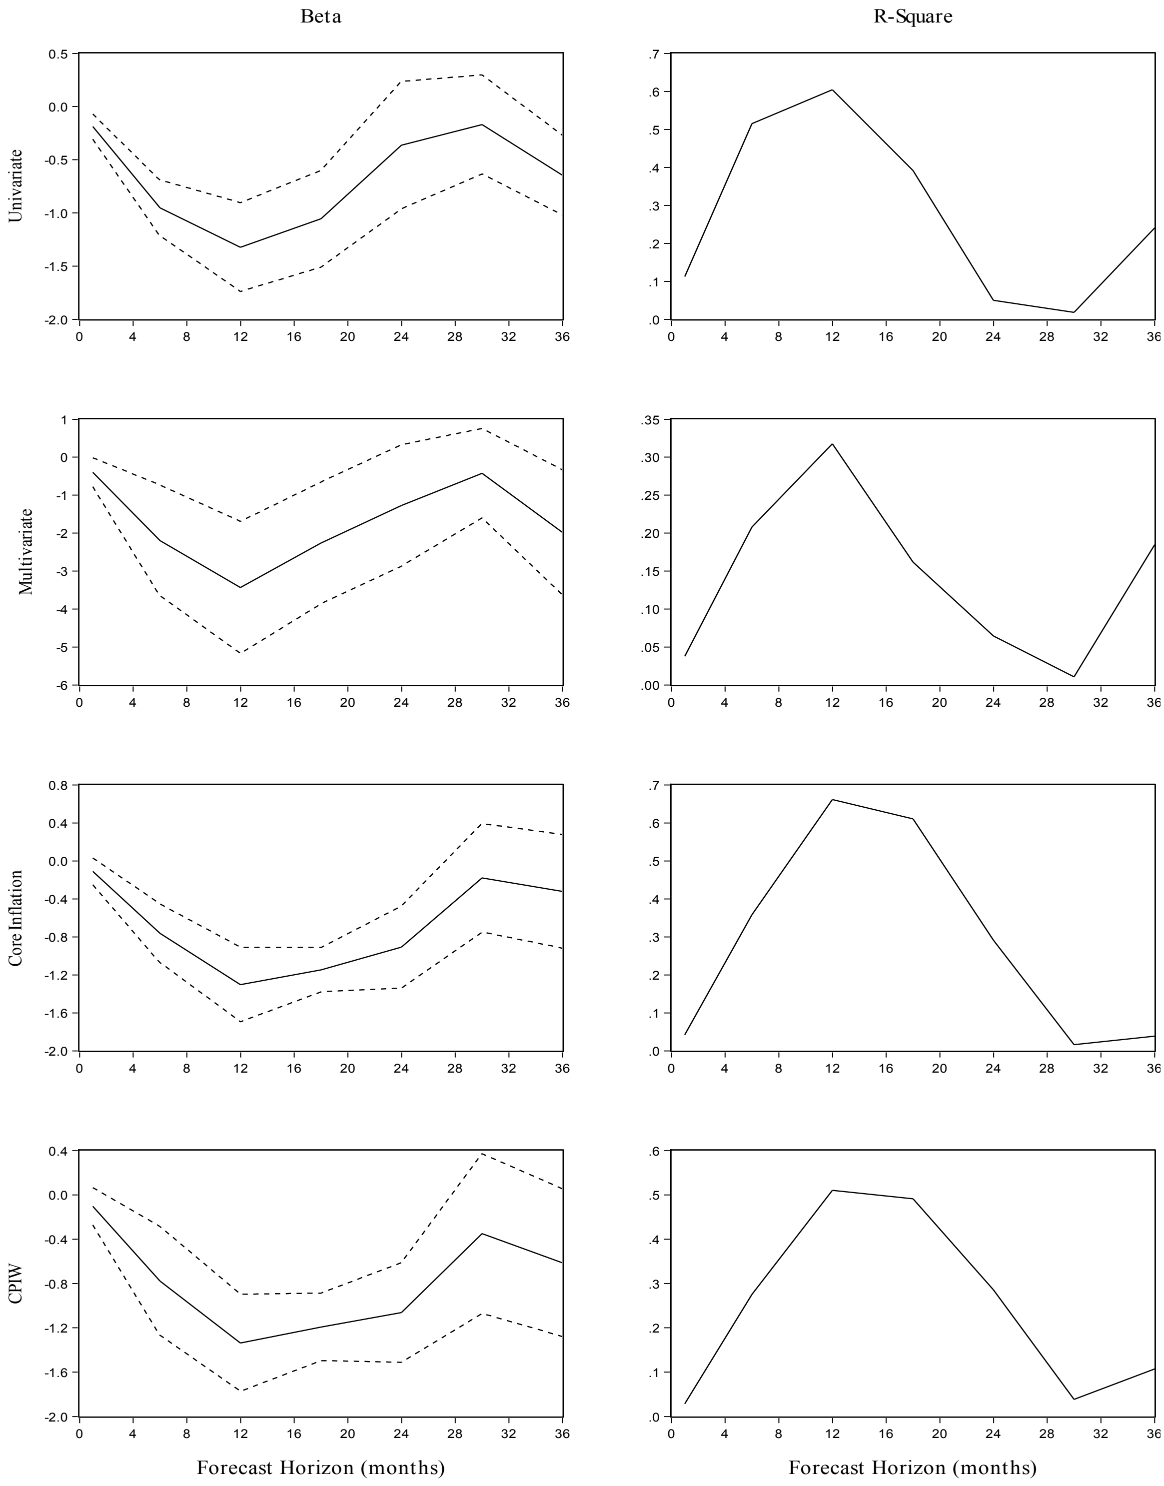

Figure 6. Estimates and R-square measures of equation (6) for various measures of inflation

8. Disaggregation

So far, results indicate that measuring inflation in a multivariate context can lead to quite different estimates than those implied by more common univariate methods. Furthermore, when these estimates are exploited to generate a persistence-weighted measure of inflation, the multivariate measure is shown to have more desirable properties. However, one departure of the approach in this paper from other attempts at reweighting the CPI basket is the reliance on a very high level of aggregation. Using a lower level of aggregation may reveal additional useful information that is not captured by focusing on just 8 components. Hence, I extend my analysis to 54 CPI components, which is the level of aggregation used by the Bank of Canada in constructing measures of core inflation.

For the univariate method, using more disaggregated data is not complicated. I simply estimate equation (5) for 54 components instead of 8. However, as noted in previous sections, estimating persistence in a VAR prohibits the use of a large set of regressors. This degrees of freedom issue arose from the desire to include lags of each CPI component in each equation. An alternative approach suited to more disaggregated data is to estimate reduced-form equations for each CPI component as follows,

$$\pi_{it} = \alpha + \sum_{k=1}^{p_i} \rho_{p_i} \pi_{it-k} + X_{t-1} + \varepsilon_t \tag{8}$$

which is just equation (5) augmented with a set of lagged explanatory variables. As in the VAR, this includes total inflation, the unemployment rate, and the three-month Treasury bill rate. The only difference is lagged values of other CPI components are excluded. While it is desirable to allow for feedback between CPI components, the role of shocks to other components may be well proxied by lagged total inflation. Hence, I apply each step of this study to 54 CPI components. This data was provided by the Bank of Canada, and is adjusted for changes in indirect taxes.

For brevity, I do not show graphs of rolling estimates for each of the 54 components. Table 8 reports the average persistence-based weights obtained from both univariate and multivariate methods, and compares them to the original CPI weights.

Table 8 reveals a striking result regarding the univariate persistence-based weights. The three highest weighted components, rented accommodation, mortgage interest cost, and replacement cost are all subcomponents of shelter. Aggregating these weights reveals that shelter receives a weight of over 50%. This is a significant result, as shelter received a near-zero weight when univariate persistence was measured in the context of 8 CPI components. Hence, the level of aggregation appears to be important for determining the relative importance of individual prices. The same methodology leads to drastically different conclusions when the level of aggregation is altered.

Aside from shelter, most CPI components are assigned a weight within 2 percentage points of the original expenditure-based weights. Notable exceptions are passenger vehicles, food at restaurants, education, and communications. These components are down-weighted by 3 to 5 percentage points when the basket accounts for persistence.

Turning to the multivariate estimates, it is evident from Table 8 that they do not vary substantially from univariate estimates. Once again, rented accommodation, mortgage interest cost, and replacement cost are assigned the highest weights of the 54 components. These weights are marginally lower than those computed from the univariate method. Though shelter receives an overall weight of over 50% this does not present as dramatic a deviation from the 8-component case as did the univariate method. Shelter received an average weight of over 30% when the multivariate method was applied to just 8 components.

Table 8. Weights of 54 CPI components

Percentage share in the basket	Univariate	Multivariate	CPI (2005 basket)
Meat	0.81	1.17	2.35
Fish	0.04	0.08	0.34
Dairy products and eggs	0.46	0.73	1.88
Bakery and cereal products	0.50	0.53	1.84
Fruit, fruit preparations and nuts	0.18	0.48	1.27
Vegetables and vegetable preparations	0.82	0.78	1.23
Other food and non alcoholic beverages	2.26	1.99	2.97
Food purchased from restaurants	0.54	1.90	5.15
Rented accommodation	11.86	11.54	5.36
Mortgage interest cost	33.71	28.16	5.66
Replacement cost	10.22	8.79	3.27
Property tax	0.12	0.58	3.31
Homeowners' home and mortgage insurance	0.65	0.66	1.15
Homeowners' maintenance and repairs	1.99	1.81	1.51
Other owner accommodation expenses	5.25	4.46	1.58
Electricity	4.97	4.48	2.51
Water	0.03	0.08	0.51
Natural gas	0.07	0.28	1.33
Fuel oil and other fuels	0.13	0.28	0.42
Communications	0.44	0.94	2.95
Childcare and domestic services	0.02	0.41	0.98
Household chemical products	0.21	0.24	0.51
Paper, plastic and foil supplies	0.52	0.46	0.59
Other household goods and services	0.08	1.04	2.02
Furniture	1.63	1.66	1.56
Household textiles	0.60	0.54	0.43
Household equipment	0.08	0.33	1.64
Services related to furniture and equipment	0.01	0.05	0.27
Clothing	4.06	3.77	3.53
Footwear	0.82	0.76	0.88
Clothing accessories, watches and jewelry	0.09	0.16	0.61
Clothing materials, notions and services	0.15	0.38	0.34
Passenger vehicles	1.53	2.40	6.25
Lease rent	0.03	0.44	1.46
Gasoline	3.77	4.11	4.92
Passenger vehicle parts, maintenance and repairs	0.62	0.78	1.86
Other passenger vehicle operating expenses	3.63	3.51	3.56

Local transportation	0.07	0.09	0.63
Intercity transportation	0.54	0.50	1.09
Healthcare goods	0.15	0.21	1.35
Healthcare services	0.15	0.16	1.12
Personal care supplies and equipment	1.23	1.16	1.27
Personal care services	0.04	0.23	0.99
Recreational equipment and services	0.09	0.25	1.76
Recreational vehicles	0.04	0.33	0.95
Operation of recreational vehicles	0.29	0.33	0.48
Home entertainment equipment, parts, and services	0.13	0.13	1.19
Travel services	1.11	1.18	2.17
Other cultural and recreational services	1.86	2.15	2.40
Education	0.05	0.15	2.67
Reading material and other printed matter	0.03	0.17	0.60
Alcoholic beverages served in establishments	0.01	0.07	0.55
Alcoholic beverages purchased from stores	0.58	1.31	1.17
Tobacco products	0.73	0.85	1.35

There are hence two important conclusions to be drawn from this analysis. The first is that univariate and multivariate estimates of persistence are similar when more disaggregated data is used. This is not the case when a more aggregate approach is adopted. Secondly, multivariate estimates appear to be more robust to the level of aggregation than univariate estimates which are particularly sensitive to it.

As in the 8-component case, I reweight the CPI basket by the persistence-based weights. The resulting measures of inflation are graphed with total inflation in Figure 7, while Figure 8 compares them to core inflation. Table 9 reports descriptive statistics. The univariate and multivariate measures now are extremely similar. Both measures imply lower inflationary pressure from 2003-2004, and closely track total inflation from 2004 to mid-2006. Subsequently, these measures imply higher inflationary pressure till 2009, reaching a peak of approximately 5 percent in mid-2008. This corresponds with a period of excess demand in Canada. In response to the latest economic downturn, both persistence-weighted measures decline sharply in tandem with total inflation. However, they both remain below total inflation during the ongoing recovery. This story is largely true in relation to core inflation as well. The main difference is that core inflation does not react as strongly to the business cycle, so the persistence-weighted measures are well below core inflation during the recent economic downturn.

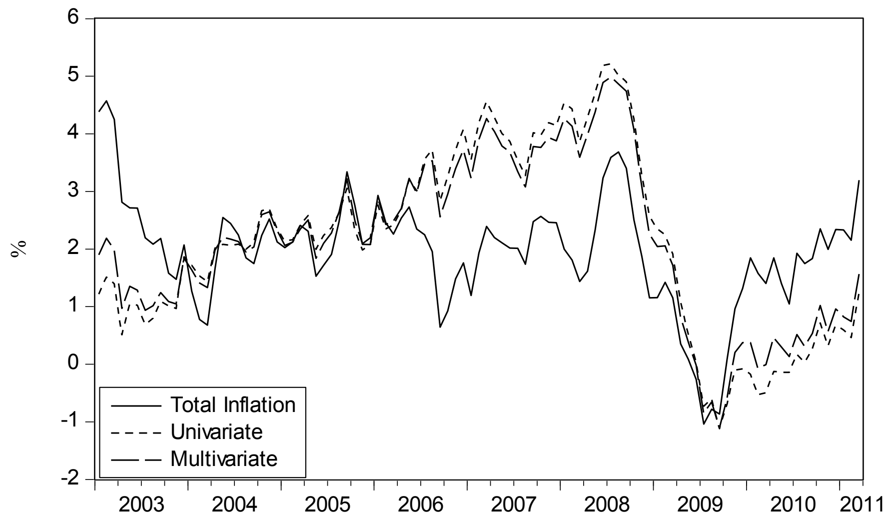

Figure 7. Comparison of year-over-year growth of persistence-weighted measures with total inflation

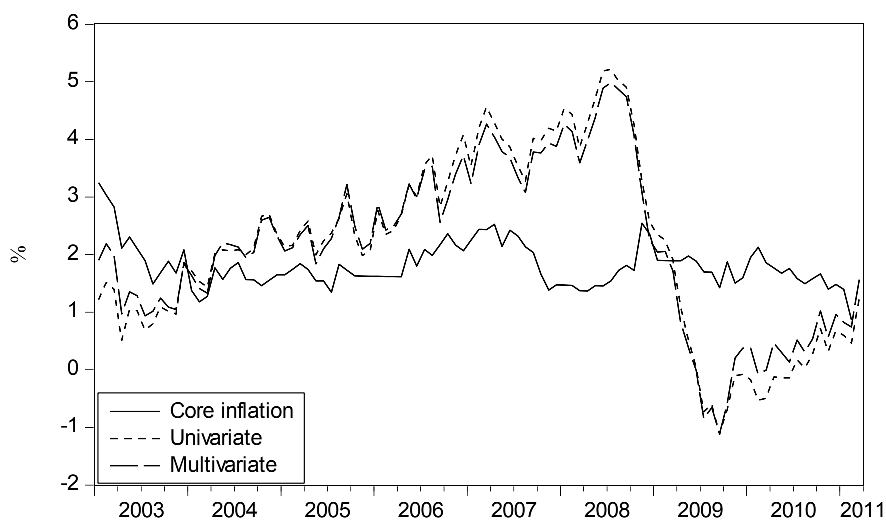

Figure 8. Comparison of year-over-year growth of persistence-weighted measures with core inflation

Table 9. Descriptive statistics of alternative measures of inflation
year-over-year percentage change

	Mean	Median	Maximum	Minimum	Std. Dev.
Univariate	2.14	2.15	5.21	-1.12	1.61
Multivariate	2.15	2.15	4.99	-1.10	1.44
Total inflation	1.94	2.02	4.57	-1.04	0.97
Core inflation	1.81	1.73	3.25	0.87	0.39
CPIW	1.88	1.85	2.90	1.10	0.42

Table 9 reveals that both persistence-weighted measures have average growth rates of 2.1 percent over the sample period. This is in line with the 2 percent inflation target, and the average growth rates of other official measures reported in the table. In opposition to other measures of core inflation, both persistence-weighted measures are more volatile than total inflation. Of the two, the multivariate measure has a lower standard deviation.

Figure 9 displays the persistence-weighted measures of inflation with the Bank of Canada's key policy rate. There appears to be a striking correlation between these measures and the policy rate. As reported in Table 10, the correlation coefficients of the univariate and multivariate measures with the policy rate are 0.79 and 0.80, respectively. These far exceed the correlation of any other measure of inflation (official measures and persistence-weighted measures based on 8 components) with monetary policy decisions.
Hence, there is strong evidence that the central bank has reacted to the information content of these measures. However, despite this favourable result, there is one caveat. Both persistence-weighted measures assign high weights to mortgage interest cost. This component is almost entirely determined by the Bank's policy rate itself, so attaching a high weight to it will naturally reveal a strong correlation of the resulting measure with

the policy rate. Furthermore, for a persistence-weighted measure to serve as a useful operational guide, it may not be desirable to highly weight a component that is largely endogenous to monetary policy. To address these concerns, I exclude mortgage interest cost from the persistence-weighted measures and reexamine their relationship with the policy rate.

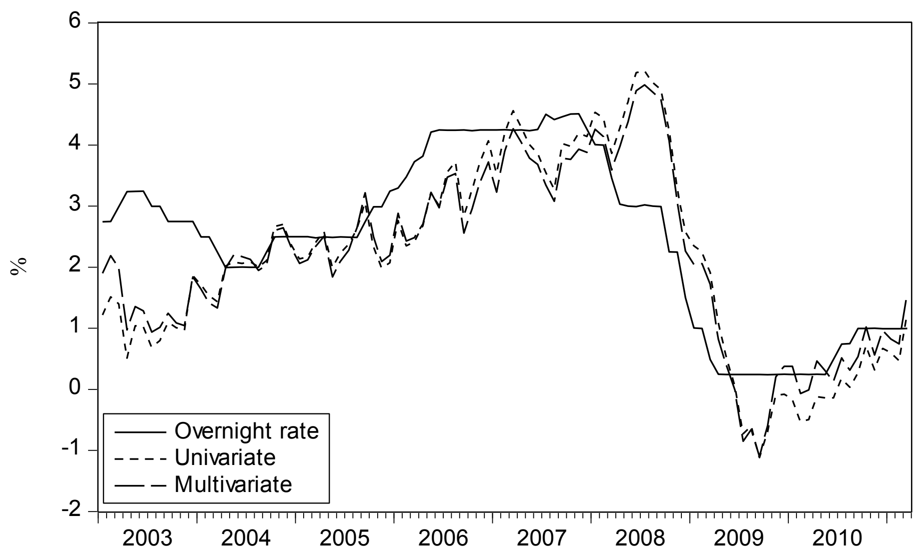

Figure 9. Persistence-weighted measures of inflation with the key policy rate

Table 10. Correlations with overnight rate

Univariate	0.79
excl. mortgage interest cost	0.58
Multivariate	0.80
excl. mortgage interest cost	0.58
Total inflation	0.46
Core inflation	0.26
CPIW	0.49

As reported in Table 10, the correlation coefficients do decrease once mortgage interest cost is excluded. However, they are still high (0.58 in both cases) and indeed still higher than that of any other measure of inflation. Therefore, this positive result is robust to whether or not mortgage interest cost is included in the persistence-weighted measures. It is also worth noting that both total inflation and CPIW do not exclude mortgage interest cost, so the relative performance of the persistence-weighted measures in this regard is likely even more superior than Table 10 suggests.

As in the 8-component case, I investigate whether the persistence-weighted measures contain useful information about future inflation that is not exploited by the central bank. Given the high correlation of these measures with the policy rate, it seems unlikely that the forecasting exercise will yield positive results. Nevertheless, I estimate equation (6) for both univariate and multivariate measures constructed using 54 CPI components. Figure 10 contains estimates of β as well as the R^2 of both regressions at each forecast horizon. Table 7 (see appendix) reports p-values of the joint test that $\alpha = 0$ and $\beta = -1$ in each regression.

It is evident that neither measure meets the proposed criteria to be considered an unbiased predictor of total inflation. The estimates of β are significantly above -1 at each forecast horizon, while the explanatory power is also negligible. The joint test that $\alpha = 0$ and $\beta = -1$ is rejected at each horizon for both persistence-weighted measures. The strong correlation of these measures with policy decisions and the lack of predictive content for future inflation may indicate that the relative persistence of CPI components is fully reflected in monetary policy. A more sophisticated approach to investigating this possibility is to examine the performance of the persistence-weighted measures in the central bank's policy rule. I undertake this exercise is the following section, proceeding only with the measures derived using 54 CPI components.

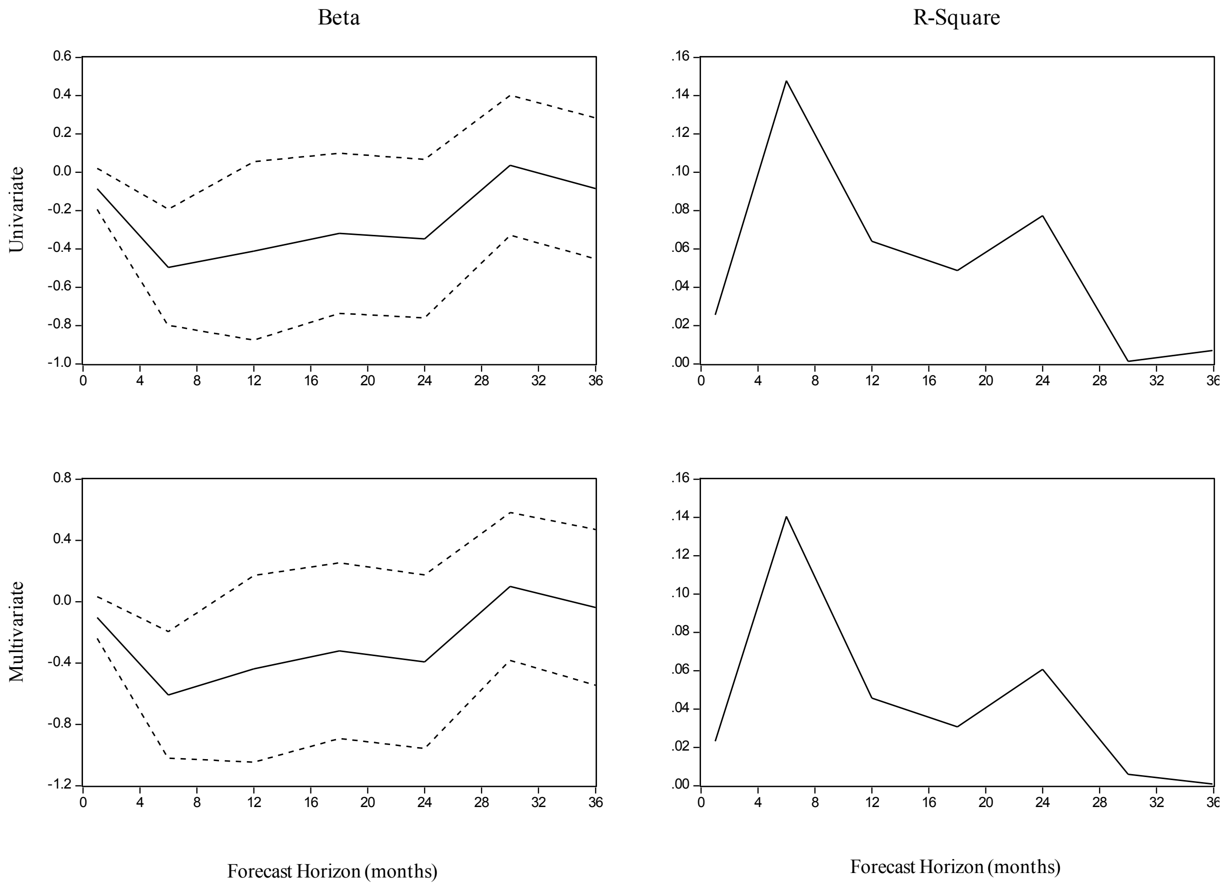

Figure 10. Estimates and R-square measures of equation (6) for persistence-weighted measures of inflation

9. Persistence-Weighted Measures in the Central Bank's Policy Rule

A monetary policy rule expresses the central bank's policy instrument as a function of the key variables to which it reacts. The concept was popularized by Taylor (1993) and is now widely employed by central banks. The policy rule in the Bank of Canada's model of the Canadian economy, ToTEM, is a Taylor-type rule (see Murchison and Rennison (2006)). It expresses the target for the overnight rate as a function of the lagged policy rate (to capture the fact that policy reacts gradually to economic developments), the long-run neutral interest rate, the deviation of inflation from target, and the output gap. More precisely, the rule is expressed as follows,

$$R_t = \rho R_{t-1} + (1-\rho)\left[R^* + \phi_\pi(E_t\pi_{t+k} - \pi^*) + \phi_y(y_t)\right] \tag{9}$$

where R_t is the overnight rate, R^* is the neutral overnight rate, $E_t\pi_{t+k}$ is the expected inflation rate k periods ahead, π^* is the inflation target, and y_t is the output gap. My goal is to estimate this rule on a monthly basis, in a manner than closely replicates the decision making process of the central bank. Before doing so, there are a few practical issues that require consideration.

The ToTEM policy rule is forward-looking, so it includes the deviation of future expected inflation from the inflation target. I do not have access to a historical series of the Bank of Canada's inflation forecasts, nor do such forecasts exist for the persistence-weighted measures computed in this essay. Hence, when estimating (9) I use the contemporaneous deviation of inflation from target.

Another important issue relates to the measurement of the output gap. The output gap series made available by the Bank of Canada is a quarterly measure and is therefore not suitable for estimating the policy rule on a monthly basis. Furthermore, constructing a monthly output gap measure in a real-time manner is problematic. Common filters such as the Hodrick-Prescott (HP) filter and Baxter-King (BK) bandpass filter are two-sided, so they use data that was not available to policymakers at the time policy decisions were made. To resolve these difficulties, I propose three alternative measures of the output gap. The first of these is to simply construct a monthly series from the Bank of Canada's quarterly series. I do so using quadratic average interpolation. The output gap is a measure that is unlikely to have many high-frequency movements, so interpolating the quarterly series may provide a good approximation. The second method is to apply the filter proposed by Christiano and Fitzgerald (2003). This filter is a finite-sample approximation to the ideal bandpass filter. It also has a one-sided representation, so it can be used in real-time. Finally, I construct a simple measure of the output gap by taking the deviation of output from a 5-year moving average. Note that these methods themselves replicate the real-time policymaking process, but the GDP data I use is revised. Ideally, one would want to use real-time data as well. Unfortunately, no such real-time database exists.

One final issue in estimating the monthly policy rule is that the Bank of Canada does not make policy decisions every month. Each year, policy decisions are made on 8 predetermined dates. Therefore, I include a dummy variable in the policy rule, which equals 1 in each month a policy decision was made, and 0 otherwise.

Table 11 displays the estimated parameters of interest in the policy rule using the persistence-weighted measures, as well as total and core inflation. For each measure, the equation is estimated with the 3 different versions of the output gap. The persistence-weighted measures exclude mortgage interest cost.

The coefficient on the deviation of the univariate persistence-weighted measure from the 2 percent inflation target is significant at the 5% level in two of the three equations. It is not statistically significant when the output gap based on the 5-year moving average is used. The same is true of the multivariate measure, but the coefficient on its deviation from the inflation target is in each case higher than that of the univariate measure. Furthermore, significance occurs at the 1% level as opposed to 5% in the univariate case. Total inflation is also significant at the 5% level in the same two equations. However, in each case the estimated coefficient is lower than when the persistence-weighted measures are used. The results are very similar to those obtained using the univariate measure, but not as favourable as when the multivariate measure is used. Interestingly, core inflation is not statistically significant in any equation. Despite the moving average-based output gap being statistically significant each time it is used in the policy rule, it introduces some counter-intuitive results that may point to its inappropriateness in measuring economic slack. For example, its inclusion in the policy rule results in negative estimates of the long-run neutral rate of interest.

Hence, there is further evidence that the central bank has implicitly responded to the information content of the persistence-weighted measures. Furthermore, the estimated coefficients on these measures in the policy rule imply that this response has been more pronounced than for other measures of inflation. This result holds true in every estimated equation, so it is robust to alternative measures of the output gap. However, the comparison of the performance of the persistence-weighted measures in the policy rule to the official measures of inflation should be qualified. The Bank of Canada is known to use expected rather than current inflation in its policy rule, so this result could simply be driven by the alternative specification adopted in this paper. Nevertheless, the contemporaneous significance of the persistence-weighted measures in the policy rule is important in itself. Together with prior evidence reported in this study, it reinforces the possibility that these measures contain useful information for the future path of inflation, and that the central bank responds to this information to meet its inflation objectives.

Table 11. Estimated Parameters of Policy Rule, 2003-2010

	R_{t-1}	R^*	$\pi_t - \pi^*$	${y_t}^{BOC}$	${y_t}^{CF}$	${y_t}^{MA}$	$\overline{R^2}$
Univariate							
-	0.959**	0.084	0.048*	0.014			0.982
	(0.029)	(0.089)	(0.023)	(0.028)			
-	0.965**	0.064	0.050*		0.001		0.983
	(0.020)	(0.063)	(0.020)		(0.002)		
-	0.901**	-0.098	-0.009			0.006**	0.985
	(0.022)	(0.056)	(0.022)			(0.001)	
Multivariate							
-	0.962**	0.073	0.062**	0.006			0.982
	(0.028)	(0.088)	(0.024)	(0.028)			
-	0.966**	0.060	0.063**		0.000		0.983
	(0.020)	(0.062)	(0.022)		(0.002)		
-	0.903**	-0.083	0.002			0.006**	0.985
	(0.022)	(0.057)	(0.024)			(0.001)	
Total							
-	0.946**	0.152	0.047*	0.034			0.982
	(0.028)	(0.083)	(0.023)	(0.024)			
-	0.966**	0.097	0.049*		0.002		0.982
	(0.021)	(0.064)	(0.022)		(0.002)		
-	0.902**	-0.081	0.004			0.006**	0.985
	(0.022)	(0.055)	(0.023)			(0.001)	
Core							
-	0.934**	0.197**	0.070	0.055*			0.981
	(0.030)	(0.092)	(0.054)	(0.024)			
-	0.956**	0.140*	0.079		0.004*		0.982
	(0.023)	(0.076)	(0.055)		(0.002)		
-	0.890**	-0.058	0.083			0.006**	0.986
	(0.022)	(0.048)	(0.046)			(0.001)	

standard errors in parentheses
*p<0.05, **p<0.01

10. Conclusion

In this essay, I have applied both univariate and multivariate techniques to measuring an important property of disaggregated price data in Canada: persistence. My results indicate that, in absolute terms, inflation in most CPI components exhibit little persistence. However, I also find that the relative persistence of inflation in these components varies quite dramatically depending on the method used to measure persistence, as well as the level of aggregation at which it is measured. The difference between univariate and multivariate estimates of persistence is particularly pronounced at the 8-component level of aggregation. For instance, univariate estimates imply monetary policy should place a trivial weight on movements in inflation in shelter, while multivariate estimates indicate these movements warrant serious policy attention. This divergence diminishes as I decrease the level of aggregation to 54 CPI components, with multivariate estimates proving to be more robust to the level of aggregation. I also find the multivariate method to capture structural breaks in certain prices series that are not reflected in the univariate context.

I have then exploited these estimates to reweight the CPI in a manner that explicitly accounts for the relative persistence (and hence, policy relevance) of its components. Conceptually, I believe this approach is more in line with the objectives of central banks when constructing measures of core inflation. Indeed, an evaluation of my persistence-weighted measures of inflation raises the possibility that the central bank has been implicitly setting policy in response to the information content of these variables. This is evidenced by their strong correlation with the key policy rate, as well as their performance in the central bank's policy rule. In fact, by these criteria the persistence-weighted measures outperform both total inflation and the official core inflation measure. I also find the persistence-weighted measures are not useful for predicting total inflation. This, however, is precisely what one would expect if their information content were already reflected in the conduct of an inflation-targeting central bank.

11. References

Benati, L., 2008, Investigating Inflation Persistence Across Monetary Regimes, *Quarterly Journal of Economics* 123, 1005–1060.

Bilke, L. and L. Stracca, 2008, A Persistence-Weighted Measure of Core Inflation in the Euro Area, ECB working paper no. 905.

Christiano, L.J. and T.J. Fitzgerald, 2003, The Bandpass Filter, *International Economic Review* 44, 435-465.

Cogley, T., 2002, A Simple Adaptive Measure of Core Inflation, *Journal of Money, Credit and Banking* 34, 99-113.

Cogley, T. and T.J. Sargent, 2005, Drifts and Volatilities: Monetary Policies and Outcomes in the Post War U.S., *Review of Economic Dynamics* 8, 262-302.

Cogley, T., Primiceri, G.E., and T.J. Sargent, 2007, Inflation-Gap Persistence in the U.S., Unpublished Manuscript.

Cutler, J., 2001, A New Measure of Core Inflation in the UK, MPC discussion paper no. 3.

Enders, W., 2003, *Applied Econometric Time Series*, Wiley Series in Probability and Statistics, John Wiley & Sons, Inc.

Fuhrer, J.C., 2009, Inflation Persistence, Federal Reserve Bank of Boston working paper no. 14

Lafleche, T. and J. Armour, 2006, Evaluating Measures of Core Inflation, Bank of Canada Review, summer.

Levin, A.T. and J.M. Piger, 2004, Is Inflation Persistence Intrinsic in Industrial Societies?, ECB working paper no. 334.

Lunneman, P. and T.Y. Matha, 2004, How Persistent is Disaggregate Inflation? An Analysis Across EU15 Countries and HICP Sub-Indices, ECB working paper no. 415.

Mendes, R. and S. Murchison, 2010, Declining Inflation Persistence in Canada: Causes and Consequences, Bank of Canada Review, winter.

Murchison, S. and A. Rennison, 2006, ToTEM: The Bank of Canada's New Quarterly Projection Model, Bank of Canada technical report no. 97.

Pivetta, F. and R. Reis, 2007, The Persistence of Inflation in the United States, *Journal of Economic Dynamics and Control* 31, 1326-58.

Taylor, J., 1993, Discretion Versus Policy Rules in Practice, *Carnegie-Rochester Conference Series on Public Policy* 39, 195–214.

Wilkinson, G., 2011. The Behavior of Consumer Prices Across Provinces, Bank of Canada discussion paper no. 2.

12. Appendix

Table 7. P-values from joint test of $\alpha = 0$ and $\beta = -1$ in equation (6)

	8 CPI Components		54 CPI Components	
Horizon (months)	Multivariate	Univariate	Multivariate	Univariate
1	0.00	0.00	0.00	0.00
6	0.00	0.00	0.00	0.00
12	0.00	0.00	0.00	0.00
18	0.00	0.00	0.00	0.00
24	0.00	0.00	0.00	0.00
36	0.00	0.00	0.00	0.00