Flexible Inflation Forecast Targeting: Evidence from Canada*

Glenn D. Otto
University of New South Wales, Australia

and

Graham M. Voss
University of Victoria, Canada

Summary

We examine models of flexible inflation forecast targeting. The target criteria from these models restrict the conditionally expected paths of inflation and output targeted by the central bank. We estimate and test these criteria for Canada, an early adopter of inflation targeting. We show that the Bank of Canada systematically balances inflation and economic activity over the medium term policy horizon consistent with flexible inflation targeting. Further, our results establish the usefulness of flexible inflation targeting rules as an alternative to standard Taylor rules.

Keywords: monetary policy, inflation, inflation targeting.

Corresponding Author:

G Voss
Department of Economics
University of Victoria, Canada V8W 2Y2

Email: gvoss@uvic.ca
Telephone: +1 250 721 8545
Facsimile: +1 250 721 6214

---

*Financial support from the Australian Research Council (DP 1093363) is gratefully acknowledged. The second author also acknowledges the Hong Kong Institute for Monetary Research Fellowship programme for support on this project. Earlier versions of this paper benefitted from the comments of Maxym Chaban, Kenneth Kuttner, Eric Santor, Aarti Singh, Gregor Smith, and James Yetman, as well as seminar participants at the University of Victoria, CEA Conference 2009, HKIMR, the RBA Workshop on Monetary Policy in Open Economies, Griffith University, and ESAM 2008.

1. Introduction

Inflation targeting, the practice of specifying a numerical target for inflation and implementing forward-looking policy decisions to achieve the target, was initially developed by central banks as a transparent means of implementing credible monetary policy. Subsequent theoretical work by Svensson (1997, 1999), Woodford (2003, 2004), Svensson and Woodford (2005), and Woodford and Giannoni (2005), recasts inflation targeting as an optimal targeting rule. A key element of this theoretical work is *forecast* targeting, where the central bank uses its policy instrument to ensure that the bank's projections or forecasts for its target variables satisfy criteria consistent with minimizing social welfare loss.[1] A further important element is the idea of *flexible* targets that consider variables other than inflation, most notably economic activity.

We use the inflation forecast targeting framework as a means of investigating the operation of monetary policy in Canada over the period in which it has been formally targeting inflation. The target criteria from this framework provide restrictions on the conditionally expected paths of variables targeted by the central bank; they are in fact the Euler conditions from a linear quadratic optimization problem for the central bank. We estimate these conditions, providing a test of monetary policy under the maintained hypothesis that monetary policy has been implemented as if they were operating within the inflation forecast targeting framework. General specification tests then allow us to determine whether the conditions are satisfied. An attractive feature of the approach is that we need not concern ourselves with how the policy instrument is adjusted to achieve these conditions, which would require a structural representation of the entire economy.

Our primary motivation is to establish whether there is a close correspondence between inflation targeting as practiced and as prescribed by the theory of inflation forecast targeting. If actual behaviour is consistent with theory, then models of inflation forecast targeting

[1] Woodford (2007). Svensson (1997) is the seminal theoretical treatment of inflation forecast targeting. Earlier work by King (1994) discusses the idea as a practical description of monetary policy in the UK.

would arguably be useful tools for analysis, in the same way that policy instrument rules, such as the Taylor rule (Taylor, 1993), are used in policy analysis. More importantly, such behaviour would lend support to the arguments of Svensson (2003) and Woodford (2004, 2007) that central banks move to explicit statements and operation of inflation forecast targets, including the role for economic activity and how it is balanced against inflation.

An additional motivation is to examine directly the issue of flexible versus strict inflation targeting: the trade-off between inflation and output pursued by an inflation targeting central bank. Svensson (1999) defines pure inflation targeting as a regime where the target criteria involve only the projected path of inflation. Such a target arises when the central bank places no weight on variation in any variable other than inflation in its loss function. Flexible inflation targeting, in contrast, includes other variables in the target criteria, most commonly the projected path of the output gap.[2] The general consensus is that most inflation targeting central banks practise flexible inflation targeting.[3] Despite this consensus, there is not much direct empirical evidence in support of flexible inflation targeting nor, consequently, is there much evidence of the trade-offs central banks pursue. One reason for this is that most empirical descriptions of inflation targeting central banks are based upon policy instrument rules, which do not provide a direct means of discriminating between flexible and pure inflation targeting.[4] In contrast, our approach provides evidence on the actual balance between inflation and cyclical variation in output that central banks have pursued over the short term horizon.

In addition to the near term balance between inflation and output, there is a further im-

[2] Giannoni and Woodford (2005) consider in detail a variety of theoretical structures and their implications for target criteria, which in some instances include variables in addition to inflation and the output gap. While theoretically appealing, our focus here on inflation and output is more likely to be consistent with central bank practice.

[3] See for example Svensson (1999) and Bernanke *et al* (1999). Buiter (2006), however, is critical of the flexible inflation targeting approach, arguing that most central banks have mandates that are lexicographic in their targets. Price stability is ordered above other objectives. Thus economic activity is not to be traded-off against price stability, but considered only once inflation is at its target value.

[4] Policy instrument rules in most instances will include measures of output even if the loss function itself does not include output stabilization. See for example Svensson (2003).

plication of optimal forecast targeting that we can test. As Woodford (2004, 2007) notes, inflation forecast targets should be consistent across different horizons. So, for example, a pure inflation targeting regime should restrict the conditional expectations of inflation from the near term horizon out through to the end of the policy horizon. Similarly, the balance or trade-off between inflation and output variation under a flexible inflation forecast targeting rule should be the same across all horizons. The only relevant restriction is that the horizons must be ones for which monetary policy has some effect on the target variable. Our empirical framework allows us to examine this criterion and to our knowledge we are the first to do so.

To foreshadow our principal results, we do find evidence that for the period 1996–2007 the Bank of Canada has in effect operated a flexible inflation forecast target over the near term horizon of 12 to 18 months. The flexible target we identify weights both inflation and output growth with the relative weight on the latter between 0.2 and 0.4. The principal qualification is that there appears to be significant variation in the weights of the target across horizons. To support these results, we report some sensitivity analysis extending the model in various directions. One particular concern is weak instruments in our instrumental variables estimation, a common problem for empirical models of monetary policy. To address this, we demonstrate that our results are consistent with the identification robust methods based upon Anderson and Rubin (1949). Finally, and in keeping with the recommendations of Svensson (2003) and Woodford (2004, 2007), we demonstrate the practical use of our estimated forecast target based upon recent events in Canada.

Ours is not the first empirical study to consider the Euler conditions associated with optimal inflation forecast targeting. Favero and Rovelli (2003) estimate and test the Euler conditions associated with a particular structural model of central bank behaviour and the aggregate economy using US data. Their objective is to identify the preference parameters of the Federal Reserve, notably the targeted inflation rate, and determine whether there was a significant change in these preferences after the high inflation period of the 1970s.

Similarly, Dennis (2004, 2006) and Giannoni and Woodford (2005) also provide estimates for the United States for very general models of an optimizing central bank pursuing flexible inflation targeting.

Our approach is much simpler than these general models; we focus exclusively on the Euler conditions alone. The benefit of doing so is that ours is a limited information approach, imposing relatively little economic structure on the estimation. This can lead to less efficient estimation relative to full information methods applied to a complete structural model; however, we face much less danger of estimating a mis-specified model. This is particularly relevant in the current context because of well known difficulties with estimating different parts of the New Keynesian monetary model (see Henry and Pagan, 2004).

The other study that examines the forecast target conditions and which comes closest to ours in approach is Rowe and Yetman (2002). These authors examine whether the Bank of Canada has targeted either inflation or output in recent decades by asking whether there are predictable deviations from target values. The principles of our approach are identical to theirs; where we differ is in terms of the horizons we consider and the nature of the targets we examine. Rowe and Yetman (2002) focus on the long end of the policy horizon, 6 to 8 quarters, and a single target variable at a time. We focus on shorter horizons, where the trade-off between inflation and output is likely to be more evident, and we estimate and test flexible targets, weighted averages of inflation and output.[5]

Our study is also closely related to Kuttner (2004) though in this instance we share similar objectives rather than methods. As we do, he considers the Euler conditions restricting inflation and output in an optimal inflation forecasting framework. He uses data for a number of countries, inflation targeting and otherwise, to investigate the correlations between forecasts for inflation and output measures finding some support for forecast targeting. The

[5] A related paper, Ruge-Murcia (2009), also examines Canadian monetary policy, asking whether the Bank may be viewed as operating a price level rather than an inflation target. This analysis, however, focuses on the long-run properties of the price level whereas we focus on the short-run dynamics of inflation and in particular the nexus between inflation and economic activity.

critical difference between this study and Kuttner's is the basis of the forecasts underlying the inflation targets. He uses published central bank forecasts to examine the correlations between inflation and measures of output. These forecasts depend then upon the central bank's information set including its own assumptions about the path of monetary policy. In contrast, our estimation procedure constructs target forecasts conditional on publicly available information only and can be viewed as an assessment of monetary policy, at least with respect to inflation forecast targeting. The distinction between the two methods is subtle but of some importance. To see the difference, it is helpful to consider a simple situation of a central bank targeting only inflation and with a two year policy window. In this case, we would expect the two year inflation forecasts by the central bank to be a very good measure of its target rate of inflation. So the Bank's forecast here gives us a clear indication of its preferences for inflation. Our approach, which in essence asks whether any current variable systematically forecasts deviations of inflation from target, is assessing whether the Bank is meeting its inflation target. If there are predictable deviations then the Bank is not using all available information or is not implementing monetary policy consistent with forecast targeting.

Finally, although we limit our analysis to Canada in this paper, the results are of broader interest. Canada was an early adopter of inflation targeting and has had a stable and successful regime since at least 1996. In addition, Canada has monthly inflation and output data which is useful for our empirical methods. So in these regards, it provides an excellent case study for inflation targeting behaviour. But there is a further consideration of interest. The formal arrangements for inflation targeting in Canada are purely in terms of inflation; nonetheless, like most inflation targeting regimes, it is believed that the Bank of Canada has near term considerations for economic activity as well. Our analysis allows us to uncover this near term balance and provides evidence that IT, as practiced, is more nuanced than a single minded focus on inflation.

2. Inflation Targeting Conditions

2.1 Theoretical Framework

To motivate the conditions which we test, it is helpful to consider a simple environment consisting of a forward looking central bank with an objective function that only depends upon variation of inflation around a target. Given its model of the underlying economy and forecasts, the central bank will adjust the policy instrument to ensure that inflation does not deviate from target. Since in general the central bank's instrument only affects inflation with a lag, it will operate to ensure that expected inflation — at a horizon for which it can influence inflation — does not differ from target. If we suppose that relative to time t, the horizon under its control is $t+h,\ h \geq \bar{h}$, then optimal policy under strict inflation targeting requires;

$$E_t(\pi_{t+h} - \pi^*) = 0, \quad h \geq \bar{h} \tag{1}$$

where π_{t+h} is inflation at time $t+h$ and π^* is the target rate of inflation. An optimality condition or Euler equation like (1) can be derived using the standard New Keynesian model of optimal monetary policy for a central bank that is concerned only about inflation, Galí (2008). Rowe and Yetman (2002) is principally concerned with testing such a condition for Canada.

Few if any central banks claim to strictly target inflation and more general or *flexible* targets are likely to be a better characterization of central bank behaviour. The two flexible inflation targets commonly discussed in the literature on monetary policy are:

$$E_t\Big(\pi_{t+h} + \phi x_{t+h} - \pi^*\Big) = 0 \quad h \geq \bar{h} \tag{2}$$

$$E_t\Big(\pi_{t+h} + \phi(x_{t+h} - x_{t+h-1}) - \pi^*\Big) = 0 \quad h \geq \bar{h} \tag{3}$$

where x_t is the output gap, the difference between the logarithms of output and potential output. Both of these conditions are associated with models of monetary policy when the central bank's loss function depends upon variation in both inflation and output gaps but arise under different assumptions of central bank behaviour.

Condition (2) arises in a model with a forward-looking New Keynesian aggregate supply curve and the assumption that the central bank pursues discretionary monetary policy — that is, it re-optimizes monetary policy every period.[6] The condition has been referred to as *a leaning against the wind* since the bank trades off inflation higher than target against output below capacity.

Condition (3) also arises in a model with a forward-looking New Keynesian aggregate supply curve but in this case under the assumption that the central bank is able to commit to a time invariant optimal path for current and future monetary policy. With discretionary optimization, the central bank need not concern itself with the dynamic structure of the aggregate supply relation; it need only focus on the contemporaneous trade-off between the current output gap and inflation. Under commitment, the central bank's optimization problem does need to take into account the dynamic structure of the aggregate supply relation, in particular the dependence of current inflation on future inflation (with a forward-looking aggregate supply curve). Hence the trade-off between inflation and output will itself have a dynamic structure, as in condition (3).[7]

The parameter ϕ plays an important role in conditions (2) and (3) in that it captures the relative weight the output gap variable receives in the flexible inflation target. If we can obtain stable empirical values for this parameter then we have a description of monetary policy which is of general practical interest since it will capture the manner in which central banks balance their objectives over the near term. One could also attempt to use the underlying theoretical models to further interpret this parameter. Under certain assumptions about the structure of the underlying economy and preferences of the central bank, ϕ will be a simple function of the slope of the Phillips curve (α) and the weight on output in the

[6]Svensson (2003) provides an exact treatment of this condition when there are lags associated with monetary policy. The basic principle is outlined in Clardia, Galí, and Gertler (1999), among other places.

[7]For the derivation of the condition and a full discussion of the trade-offs, see Svensson (2003).

central bank loss function (λ); specifically, $\phi = \lambda/\alpha$ (see Svensson, 2003). The general message from this relationship is that a positive value for ϕ is consistent with a positive value for λ. Existing empirical work that directly estimates λ, for example Favero and Rovelli (2003) and Dennis (2004, 2006), find it to be very small or zero. While we cannot comment directly on the size of λ, we can at least infer from our estimates whether it is zero; that is, whether the central bank weights economic activity in its objectives.

There is a further issue concerning the sign of the ϕ parameter. In the presence of a forward looking aggregate supply curve, the sign is positive. However, if the aggregate supply curve is backward looking, then the same condition holds but ϕ would be less than zero.[8] Our estimates of ϕ are uniformly positive consistent with the forward looking aggregate supply assumption; however, as we do not directly estimate the aggregate supply relationship we are reluctant to draw to strong of inference from this aspect of our analysis.

2.2 Empirical Framework

Conditions (2) and (3) are the basis for our estimation. Due to a number of practical issues, however, they are not the exact conditions on which we focus our empirical work. The main hurdle we face is that these conditions are best estimated using monthly data. Monthly data is compelling for two reasons. First, it provides a reasonable sample size for estimation, which turns out to be critical for the generalized method of moments (GMM) used to identify and estimate the target conditions. Second, monthly data more closely approximates the frequency of monetary policy decision and provides a finer description of the conditional information set than does quarterly data, the natural alternative. The trade-off, however, is that monthly measures of the output gap (such as one constructed from a Hodrick-Prescott filter) are quite volatile and fail to provide reasonable estimates for conditions (2) and (3) (as reported below).

To proceed, we focus on the principal implication of either of the two conditions; specifically,

[8] Kuttner (2004) provides an intuitive discussion of this situation; Svensson (2003) provides complete details.

that the central bank operates monetary policy in a manner to balance off deviations of inflation from target against economic activity. From this perspective, the conditions can be approximated as:

$$E_t\Big(\pi_{t+h} + \phi(\Delta y_{t+h} - \Delta y) - \pi^*\Big) = 0, \quad h \geq \bar{h} \tag{4}$$

where π_{t+h} is inflation at time $t+h$ and Δy_{t+h} is a measure of economic growth at time $t+h$. The constant terms Δy and π^* are the target rates for inflation and output growth respectively.

This representation of the flexible inflation target has other attractive features. First, output growth has the advantage of being very transparent and consistent with central bank practice, which often focuses on output growth rather than measures of output gaps. Second, in broadest terms, we are interested in identifying what if any attempt is made to balance inflation and economic activity in the operation of monetary policy. In this regard, output growth serves as a useful and natural measure; even without appealing directly to theory, one might think that conditions such as (4) are a useful means of describing the objectives of the central bank. Finally, one can motivate this condition directly from condition (3) under the condition that $E_t(\Delta\bar{y}_{t+h} - \Delta y) = 0$, where $\bar{y}_t$ is potential output.

The other hurdle we face when estimating the flexible target conditions concerns instrument quality. When estimating moment conditions with GMM, it is well known that weak instruments can cause estimates and inference to be unreliable; see for example the survey article by Stock, Wright and Yogo (2002). For our purposes, to ensure reasonable instrument quality, it turns out to be necessary to normalize our moment conditions on output growth rather than inflation, since the latter is significantly more predictable. As Yogo (2004) notes, GMM with weak instruments is not invariant to such a transformation and in practice the implications of alternate transformations can be substantial. A similar strategy is pursued in Consolo and Favero (2009) when estimating Taylor rules.

In summary, then, the exact conditions we estimate are as follows:

$$E_t\Big(\Delta y_{t+h} + a_{1h}\pi_{t+h} - a_{0h}\Big) = 0, \quad h \geq \bar{h} \tag{5}$$

where $a_{1h} = 1/\phi_h$ and $a_{0h} = \pi_h^*/\phi_h + \Delta y$. Notice also that we allow the coefficients to vary across horizons, with the objective of testing whether or not this is the case. Recall Woodford's (2004) arguments that the trade-off in forecast targeting models should be invariant to the horizon under consideration.

3 Empirical Results

3.1 Data

Canada adopted its current inflation targeting framework in 1991. The current target band for inflation of 1–3 percent, however, was not put into operation until the end of 1995. So our sample is 1996:M1–2007:M12.[9] The Bank of Canada's inflation target is officially in terms of Total CPI inflation, measured monthly as the year on year percentage change. In practice, however, the Bank explicitly uses core inflation as an operational guide for monetary policy and has done so since the inception of inflation targeting. Following this practice our analysis focuses on core inflation. The current measure of core inflation, used since 2001, is constructed from the Total CPI by excluding eight of the most volatile components as well as the effects of changes in indirect taxes on the remaining components.[10] Prior to 2001, the measure of core inflation used by the Bank removed food and energy as well as the effect of indirect taxes from the CPI. The two series are closely correlated (the correlation is 0.83) so for simplicity we use the current core inflation measure for estimation. Some sensitivity analysis concerning different inflation measures is considered below.

For output, we use Statistics Canada's constant dollar monthly seasonally adjusted GDP. As discussed in the previous section, we use an output growth rate as a proxy for economic

[9] For details about the adoption of the inflation target, see Bank of Canada (1995).

[10] See Macklem (2001) for a more complete discussion of the role of core inflation in Canadian monetary policy.

activity. While there are a number of different growth rates one might use, for example month to month or year to year (i.e., first or twelfth log difference), for our purposes, we find that log third differences, comparable to quarterly growth, provides the most consistent and sensible results and this is what we focus on here. Some sensitivity analysis in this regard is provided below. Details on source and construction of all series used in estimation are provided in the data appendix.

For estimation, we require a set of instruments that are both valid and relevant. For the moment conditions that we estimate, there are a large number of potentially valid instruments: any variable known at time t is available as an instrument for our moment conditions. The only concern is to ensure that any series used as an instrument at time t has been released so that it is properly available. To ensure we have relevant instruments, we examined a wide range of valid instruments that might forecast either inflation or output growth over the horizons of interest to us (1–18 months as discussed below). As noted above, we are able to identify a small set of good quality instrument for inflation but not for output growth, which motivates the normalization in conditions (5).[11] For period t, the set is:

$$z_t = (1, \pi^c_{t-2}, \pi^{cm}_{t-2}, \Delta_{12} m_{t-2}) \tag{6}$$

where π^c_t is core Inflation, π^{cm}_t is commodity price inflation, and $\Delta_{12} m_t$ is the year on year growth rate of M2. (Specific details of the measures are provided in Table A1 of the data appendix.) All variables are lagged two months since, for each of these variables, the data for period $t-2$ is released within the month of period t or before.

A brief comment is warranted concerning the relatively small size of our instrument set, both in terms of the variables considered as well as the lags considered. Other studies that estimate similar moment conditions describing inflation targets, such as Favero and Rovelli

[11]Details are available from the authors. The strongest of our instruments for inflation is M2 growth. This is consistent with the Bank of Canada's own observations in the mid 1990s concerning inflation forecasting, see for example Bank of Canada (1995).

(2003), use much larger instrument sets with many more lags.[12] We found, however, that in order to satisfy Stock and Yogo's (2002) instrument quality criteria at both short and long horizons (we need instruments to predict up to 18 months ahead), a small instrument set was necessary.[13] The small instrument set is also helpful when we estimate systems of forecast equations as the number of moment conditions remains manageable.

3.2 Estimation

We first estimate the moment conditions as single equations for horizons $h = 1 \ldots 18$. As noted earlier, the conditions for the flexible targets should be satisfied after some $\bar{h}$, the horizon after which monetary policy has a significant impact on inflation and economic activity. *A priori*, we do not know $\bar{h}$, though conventional wisdom is that the lags of monetary policy are at least one year.[14] Since $\bar{h}$ is not known, we estimate the model over both short and medium term horizons with the expectation that the model should systematically fail at shorter horizons. This strategy allows us to empirically determine the appropriate horizons to consider. We are, however, limited as to how far ahead we look; instrument quality declines significantly for horizons greater than 18. The results are reported in Tables 1 and 2. Table 1 reports estimates for early horizons 1–6 while Table 2 reports estimates for horizons 7–18.

The equations are estimated using Generalized Method of Moments. We iterate the estimation using as the weighting matrix the inverse of successive estimates of the covariance matrix (iterations continue until the difference between successive objective functions are negligible). The covariance matrix is estimated following Newey and West (1987) with truncation parameters indicated in the table.[15] For each horizon estimated, we report the

[12]Similar moment conditions are also estimated in the empirical Phillips curve literature, such as Galí and Gertler (1999), where again much larger instrument sets are used.

[13]Rudd and Whelan (2007) in their examination of Phillips curve estimation also finds the need to reduce the instrument set in order to satisfy instrument quality conditions.

[14]The Bank of Canada's views on the lags of monetary policy are summarized by Macklem (2001, p3), "... it takes about a year before monetary policy actions even begin to significantly affect inflation (with most of the impact occurring in six to eight quarters)...".

[15]The choice of truncation parameter reflects the fact that for a forecast horizon of h, there is likely to be

direct parameters, a_0 and a_1, the recovered parameters, ϕ_h and π_h^*, Hansen's (1982) test for the over-identifying restrictions, and the F-statistic for the first stage regression: inflation at horizon h against the instrument set.

The F-statistic reported for each horizon is calculated from the appropriate Wald statistic, scaled by the number of instruments to construct its F-statistic equivalent, using the Newey and West (1987) robust covariance matrix estimator, with lag truncation parameter equal to $h-1$. This statistic provides an informal test for weak instruments, following the principles in Stock and Yogo (2002). The test is only informal since we are not assuming $i.i.d$ errors so their proposed tests and critical values are not strictly relevant. However, as Baum, Schaffer and Stillman (2007) argue, since there are no comparable critical values for non-$i.i.d$ errors it seems prudent to use the robust covariance estimator when constructing the test statistic and to use the critical values of Stock and Yogo (2002) as a guide.

The test for weak instruments that we use as a guide is Stock and Yogo's (2002) test for TSLS bias (relative to OLS). The null hypothesis is that the instruments are weak, with maximal TSLS relative bias of 5%, 10%, or 20%. The related 5% critical values for the F-statistic with our three instruments are: $\{13.91, 9.08, 6.46\}$ (see their Table 1). So, for example, if the F-statistic is greater than 10, say, then we can reject the weak instrument hypothesis (no more than 10% relative bias) at the 5% significance level. The same statistic can also be used to test for size distortion associated with Wald test statistics from the instrumental variables estimation. The 5% critical values for a desired maximal size of $\{0.10, 0.15, 0.20, 0.25\}$ are: $\{22.30, 12.83, 9.54, 7.80\}$ (see their Table 2). As an informal guide based upon these critical values, we consider an F-statistic of 10 or greater to be evidence against weak instruments.

The models for the near horizons, $h = 1 \ldots 6$, are reported in Table 1. For each horizon,

a moving average structure of $h-1$; Hansen and Hodrick (1980). For the single equation estimates, this is our choice of h. For the system estimates that follow, we choose h to be one less than the maximum horizon. All data and estimation programmes, in GAUSS, are available upon request.

there is no evidence that weak instruments are a significant problem. All of the reported F-statistics exceed the 5% critical values for the relative bias and for size distortions greater than 15%. However, the model itself is not consistent with the data at these horizons. For all but $h = 2$, we reject the over-identification restrictions at the 10 percent level; for $h = 2$, the marginal significance level is 12%, so there is no strong support for the model at this horizon either. This is as we would expect. One way to interpret Hansen's J statistic is that the moment conditions are not orthogonal to the instruments, suggesting that there are predictable deviations of the flexible inflation target. This would be consistent with the central bank being unable to control inflation and economic activity at these near horizons. (It is of course also consistent with more general mis-specification.)

The medium term horizons are reported in Table 2 and these are generally quite favourable to the flexible inflation targeting model, though not uniformly so. First note that, as we might anticipate, instrument quality is not as strong at these horizons, particularly at horizons 14–17 where the F-statistic is below, though not too far below, 10. This qualifies our point estimates and inference to some extent. We do have an additional advantage, however, as we are able to consider the variety of estimates over the different horizons. So, while the estimates and inference from the horizons with weak instruments are qualified, they can be compared to those horizons where instrument quality is stronger. And while there are some differences in the estimates across the horizons generally speaking the models for each horizon are telling a similar story. We return to the issue of instrument quality in the sensitivity analysis that follows.

As we found with the early horizons in Table 1, at horizons of 7 and 8, the model is rejected at the ten percent significance level; and at horizons 9 and 10, the marginal significance level for the J statistic is less than 15% so one might be cautious about accepting the model at these horizons. At horizons 11 through 18, however, we do not reject the model. These results are broadly consistent with the conventional wisdom that monetary policy begins to have a significant effect after a year. And in fact they provide a novel means of identifying the policy horizon.

The coefficient estimates in Table 2 are quite plausible. For each individual horizon, we obtain a positive ϕ_h coefficient somewhere between 0.2 and 0.7 and although there is significant variation all are statistically different from zero at conventional significance levels. (Recall that the directly estimated coefficients are a_{0h} and a_{1h}. These are used to recover the coefficients of interest, ϕ_h and π_h^*. For the latter, we require a target output growth rate, which we set to the mean output growth rate for the sample. Standard errors for the constructed coefficients ϕ_h and π_h^* are approximated using a standard first order approximation. For π^*, the standard errors are conditioned on the mean output growth rate.)

We interpret the results in Table 2 as evidence that the paths of inflation and output are consistent with a central bank that is successfully trading off inflation and output objectives over the medium term policy horizon; moreover, the ϕ_h coefficient gives us a specific measure of how the Bank balances inflation and economic activity. While we do not find this result too surprising, it does differ from results for the United States reported in Favero and Rovelli (2003) and Dennis (2004, 2006). These studies, which estimate central bank preferences directly, find either a very small or zero weight on output in the loss function, which would give rise to very small or zero values for our ϕ parameter.

We also obtain estimates for the target rate of inflation, π_h^*, though this measure must be interpreted with care as it is conditioned on the mean growth rate for output. Conditional on this growth rate, we get estimates for π_h^* that are typically near 1.75%, below the 2% mid-point of the Bank of Canada's target band. This result corresponds very closely to Rowe and Yetman (2002), which estimates a target inflation rate of 1.6% using a strict inflation targeting framework over the longer end of the policy horizon, 6 to 8 quarters.[16] One interpretation of our results, as well as Rowe and Yetman's, is that the Bank has been systematically biased downward in its control of inflation. However, we are dealing with a

[16] Rowe and Yetman (2002) considers a number of different samples; the 1.6% is for 1991-2001, the sample nearest ours.

finite sample of data and these estimates may be the outcome of negative shocks to inflation dominating positive shocks.

While individually each equation seems reasonable, taken as a whole there is a clear departure from the flexible inflation target framework; specifically, the coefficients vary substantially across the different horizons. And they do so in a systematic fashion: as we move out toward longer horizons, the coefficients on output growth ϕ_h tend to fall. While the decline is not monotonic, the trend is evident. One possibility is that these estimates accurately reflect the behaviour of the Bank and that at longer horizons the Bank is more focused on inflation than output. While not strictly consistent with the optimal flexible inflation targets, this seems at least plausible. The other possible explanation concerns falling instrument quality. At longer horizons, as noted above, the instrument quality does drop off. If there is a finite sample bias towards lower ϕ_h coefficients then this may explain the decline and one should be cautious about reading too much into the declining coefficients. We can investigate the first possibility more completely by considering system estimates that allow us, in a limited manner, to statistically test for equality across horizons. The second possibility is considered in more detail below when we look at identification robust methods.

Table 3 presents estimates from four systems of the moment conditions, in each case considering three moment conditions simultaneously. We restrict out attention to three conditions as larger sets of moment conditions are difficult to estimate. The four systems estimated are $h = 7 \ldots 9$, $h = 10 \ldots 12$, $h = 13 \ldots 15$ and $h = 16 \ldots 18$. Broadly speaking, the estimates are consistent with the flexible inflation targeting model. There are some notable differences, though, between the results in Table 2 and 3. First, in contrast to the single equation models, we do not reject the over-identifying restrictions for the $h = 7 \ldots 9$ model; the marginal significance level is 42%. This is inconsistent with the single equation models and points to either low power of the J test in the systems estimates or a tendency to over reject in the single equation estimation. On balance, we view the single equation results as likely to be more informative in this regard. The second difference of note concerns the ϕ_h parameters. They still show a declining trend as h increases but the coefficient value for any horizon h differs somewhat from its counterpart in Table 2. That said, the estimates are not that precisely estimated and lay within a standard deviation of each other.

Table 4 re-estimates these systems but under the restriction that the coefficients are the same across the three horizons of each model. At the bottom of the table is a test statistic of the restrictions. It compares the J statistic from the restricted model to that of the unrestricted model under the condition that both models are estimated using the same weighting function, in this case that of the unrestricted model. This statistic is distributed $\chi^2(4)$, see Davidson and MacKinnon (p. 368, 2004). In all cases, the restrictions are not rejected at conventional significance levels.

Consistent with our earlier discussion, the flexible inflation targeting model does not fit the data at horizons 7-9. The ϕ_h coefficient is statistically insignificant (p-value of 0.15). The models for the other horizon do support the model to the extent that the coefficients are statistically significant and of plausible values. Where they fail, again consistent with earlier results, is in terms of the apparent parameter inconstancy across the full set of horizons, particularly for ϕ_h, which takes on the values $\{0.57, 0.41, 0.27\}$ for $h = 10 \ldots 12$, $h = 13 \ldots 15$, and $h = 16 \ldots 18$ respectively. As a means of further exploring the parameter constancy, we consider models for $h = 10, 13, 16$, $h = 11, 14, 17$, and $h = 12, 15, 18$, reported in Table 5. Here we have clearer evidence against parameter constancy, with marginal significance levels for the test of the restrictions well below ten percent for the first two models. For the last model, which focuses on the longer of the horizons, the marginal significance level is 11%, providing limited support for the restriction.

To summarize, we find considerable support for the principles of flexible inflation targeting: there is a systematic trade-off being managed by the Bank of Canada between inflation and economic activity, the latter measured as quarter on quarter monthly output growth, over the medium term policy horizon. While our results are not as sharply focused as we would like concerning the policy horizon, a reasonable interpretation is that the model holds over the 12 to 18 month horizon. The significant departure from flexible inflation targeting principles is parameter inconstancy across different target horizons, with the results suggesting that the trade-off has less weight on output growth at longer horizons. The single equation estimates, which impose the least structure on the data, provide a range of weights on economic activity declining from 0.4 to 0.2 over horizons 12–18 and these are our preferred estimates. The restricted estimates for horizons $13-15$, $16-18$, and $12, 15, 18$, while slightly higher, are not too dissimilar and provide some additional support for this range of estimates. We further find a target rate of inflation estimate between 1.7 and 1.8 percent, somewhat lower than the Bank of Canada's target of 2%.

3.3 Sensitivity Analysis

Table 6 re-estimates the $h = 12$ single equation model using some alternative variables. We focus on the single equation estimates to conserve space. As there is a close correspondence between the single equation and system estimates there is no information lost. We choose the one year horizon model since it is reasonably well estimated and rests in the accepted control horizon for monetary policy; moreover, it is least likely to be affected by weak instruments. For ease of comparison, we re-report the results for $h = 12$ from Table 1.

The first consideration is the measurement of core inflation. The Bank of Canada changed the measure of core inflation it uses to guide monetary policy in May 2001. The results so far use the current measure for the whole sample. To examine this more completely, we splice the two series at May 2001. The estimates are reported in the second column of Table 6 and are essentially unchanged from the core measure reported in column one of the table. The only notable feature is that the instrument quality measure suggests a problem with weak instruments; however, the coefficients are largely unchanged.

The second consideration is the measurement of output. Columns three and four report estimates using month to month growth rates and year on year growth rates. Again, the results are essentially unchanged. Finally, the last two columns report estimates using a measure of the output gap, either in levels or in third (quarterly) differences. The output gaps are constructed using the Hodrick-Prescott filter; see Table A1 for details. In either case, the model fails. The estimated coefficients a_0 and a_1 are both statistically insignificant. This is most surprising with respect to the difference in the output gap, as the correlation between this variable and the output growth measure $\Delta_3 y_t$ is high, 0.89 over our sample. We suspect that the problem lies in the difficulty in constructing a reasonable measure of the output gap at the monthly frequency and leave this for future research.[17]

We next consider whether the models are stable across the sample. While we could use formal parameter stability tests problems arise because of instrument quality for the smaller samples. In particular, for the early part of the sample, the model suffers significantly from weak instruments. As a simple alternative, we report and examine estimates for the sample 2001:M6–2007:M12; these results are reported in Table 7 for horizons 10–12. In addition to providing some information concerning the parameter stability of the models, it also provides a further check on whether the change in the measure of core inflation by the Bank in May 2001 alters our conclusions.

Generally speaking, the results are still supportive of the flexible inflation targeting model. Instrument quality is fine; indeed, it appears to be somewhat better than the full sample estimates and one may wish to focus solely on these estimates as being the most reliable. For all horizons, the test for over-identifying restrictions is not rejected and, apart from $h = 18$, the coefficient estimates are all signed as before and statistically significant. Where there are some differences from the full sample estimates are the coefficient magnitudes. The inflation target estimate, π_h^*, tends to be somewhat larger, closer to the 2 percent mid-point of the Bank's official target. The ϕ_h coefficients also tend to be somewhat larger,

[17] In previous work, Otto and Voss (2009), where we use quarterly data, the results for quarterly growth rates and change in the output gap are very similar; however, these results are qualified by weak instruments.

particularly at the longer horizons. As noted, the one exception to this conclusion, and important qualification, is the estimate for ϕ_{18}, which is statistically insignificant. (Strictly speaking, the a_1 coefficient is statistically insignificant at the ten percent level so that the constructed ϕ coefficient is not defined.) The model is clearly mis-specified for this horizon, though this may be an artefact of the small sample. On balance, though, we see the results as still broadly supportive of the flexible inflation targeting model but indicating that there remains considerable uncertainty about the coefficient estimates and leaving open the possibility that the Bank of Canada has not been entirely consistent in its behaviour over the inflation targeting period.

Finally, we address possible concerns about weak instruments using identification robust methods based upon the Anderson and Rubin (1949) statistic. Here we follow the methods presented in Dufour, Khalaf, and Kichian (2006). See also the discussion in Dufour (2003).

Our primary concern is the estimation of the ϕ parameter in equation (4). This parameter is not very precisely estimated, certainly less so than the inflation target parameter π^*, and so there is naturally some concern about the instrumental variable estimator. Further, whether this variable differs from zero is of considerable interest since it lies at the heart of the flexible inflation targeting model, providing a clear distinction from the strict inflation targeting model. Consequently, we focus exclusively here on the ϕ parameter.

To proceed, it is helpful to modify our notation somewhat. First, re-write out basic model as a more standard instrumental variables problem:

$$\pi_{t+h} = -\pi^* - \phi(\Delta_3 y_{t+h} - \overline{\Delta_3 y}) + u_{t+h}, \quad h \geq \bar{h} \tag{7}$$

where we have conditioned on the mean output growth rate. Next partition our full instrument set z_t into $(1\ z_{2t})$. From an IV perspective, z_{2t} are the instruments for $\Delta_3 y_{t+h} - \overline{\Delta_3 y}$ while the unit column is simply an exogenous regressor providing the estimate of π^*. Notice that we have not re-normalized the equation in terms of output growth as we did in the GMM estimation. Since the Anderson-Rubin procedure is identification robust, weak instruments are not an issue.

The Anderson-Rubin regression is then, for a particular ϕ_{h0},

$$\pi_{t+h} + \phi_{h0}(\Delta_3 y_{t+h} - \overline{\Delta_3 y}) = -\pi^* + z_{2t}\gamma + \nu_{t+h}, \quad h \geq \bar{h} \tag{8}$$

where ν_{t+h} is a combination of the structural and reduced form errors. Under the null hypothesis that $\phi_h = \phi_{h0}$, the γ coefficients should be jointly zero and the Anderson-Rubin statistic is the F-statistic for the exclusion of z_{2t} from the above regression.

Since there are many possible values of ϕ_{h0}, we follow Dufour et al (2006) and others by searching over a suitable range of values and inverting the associated test statistic to obtain confidence intervals for the ϕ_h parameter. In addition to the confidence intervals using the traditional AR statistic, we also report confidence intervals using the analogous Wald statistic constructed using the Newey and West (1987) HAC robust covariance estimator in order to correct for the possible moving average structure of the error terms and possible heteroskedasticity. While we are not aware of any formal justification for this adjustment it seems a reasonable direction to pursue: the Anderson-Rubin statistic requires independent and identically distributed errors, which cannot be assumed in the current context.

The results for both statistics are reported in Table 8. In both cases, we consider a range of ϕ_h between zero and one with a step size of 1×10^{-4}. Extending the grid beyond these end points does not alter any conclusions since the p-values beyond this range are zero. The first column for each of the two statistics is the maximum p-value for the statistic over the range considered. Beside it is the associated value for ϕ_h, which is a Hodges-Lehmann type estimator and can be compared to our GMM estimates (see Dufour et al (2006) for detail). The third column for each statistic presents the 95% confidence intervals constructed for the ϕ_h parameters. The confidence intervals are constructed by finding the bounds for ϕ_h values for which the p-value exceeds 5%. In all cases, the confidence intervals are continuous. (To give an additional sense of these intervals, Figures A3 and A4 in the data appendix report the p-values over the range of ϕ_h values for each horizon. The figures also give a clear indication that the confidence intervals are continuous.)

Table 8 provides clear evidence in support of the estimates we obtained using GMM when we consider horizons greater than 12. The point estimates and confidence intervals are very similar to the GMM estimates. As we might expect, the confidence intervals for the HAC robust statistics are significantly wider. Notably, however, in no case for these horizons do we fail to reject the hypothesis that $\phi_h = 0$; in other words, the data are consistent with flexible rather than strict inflation targeting, as we found with the GMM estimation.

For horizons 10 and 11, however, the null hypothesis is always rejected, meaning there are no values of ϕ consistent with the flexible inflation targeting model at these horizons. This is evident from the empty confidence interval sets for these horizons for both sets of statistics (see also Figures A3 and A4 in the data appendix). These results are distinctly different from our GMM estimates. While for $h = 10$ (Table 1), there is some evidence against the model with a somewhat low J-statistic, for $h = 11$ this is not the case. Interestingly, for neither horizon does instrument quality appear to be a problem. Within the context of the inflation targeting model, the results in Table 8 provide further evidence that the relevant horizon for monetary policy is twelve months and beyond.

3.4 A Practical Exercise

The analysis above identifies flexible inflation targets, albeit with some variation in coefficients across horizons. To demonstrate how these targets can be usefully implemented and to give a sense of what the trade-off between inflation and output growth implies in practical terms, we consider a simple illustrative exercise designed to capture the basic elements of the monetary policy decision process.

The target we consider is a simplification of the results above, focusing on the medium term horizons $h = 12 \ldots 18$ and setting a value of ϕ to 0.30, the mid-point of our single equation estimates. The target is:

$$E_t \pi^c_{t+h} \;=\; 1.75 - 0.30 \cdot (E_t \Delta_3 y_{t+h} - \overline{\Delta_3 y}) \tag{9}$$

Here we have assumed that $\pi^* = 1.75$, consistent with our estimates, though lower than the Bank of Canada's target mid-level. Finally, we have written the targets in the same manner as Woodford (2007); this emphasizes that this is a flexible target for inflation, dependent upon projected economic activity.

On 1 June 2010, the Bank of Canada chose to raise interest rates in response to improvements in the Canadian economy and the need to begin to remove the monetary stimulus in place as a result of the financial crisis of 2008–09. To replicate this decision using our target, we use the data available in May 2010. Let $T =$ 2010:M4, i.e. April 2010; then the available data is, due to the various release lags: $\{\pi^c_t, \pi^{cm}_{t-1}, \Delta_{12} m_{t-1}, \Delta_3 y_{t-2} |\; t = 1...T\}$. This is possibly somewhat conservative, as the Bank of Canada almost surely has more recent information on the commodity price index and money supply but as this is an exercise the issue is not critical.

We estimate a simple VAR(4) for the four variables; other lag lengths were considered with the results largely unchanged. The VAR is used to provide forecasts for core inflation and output growth over $T + 1 \ldots T + 18$. The output growth forecasts are used to construct the target for core inflation. To provide a measure of uncertainty, we use a simple bootstrap (10000 replications) to construct $\pm$ 1 standard deviation bands for the target as well as for the core inflation forecast. The results are reported in Figure 1.

Our forecast for core inflation is fairly steady at about 2% or just below over the entire forecast horizon. For the target bands (the bands rather than the target itself are reported in the figure), these are initially quite low but rise steadily as we move out to the medium term horizons because our VAR is projecting above average near term growth followed by steadily declining growth rates over the eighteen month horizon. At the medium term horizon,

$h \geq 12$, the projections for core inflation are in the lower part of the flexible target band, suggesting little need for any change in monetary policy — at least conditional on these forecasts. Other considerations, such as household balance sheets and the very expansionary stance of monetary policy at the time, clearly influenced the decision to marginally tighten monetary policy in June 2010.

Setting aside any debate about whether a change in monetary policy is warranted, the figure highlights a key issue associated with flexible inflation targets. Because the flexible target is conditioned on forecasts of output (or some other measure of economic activity such as the output gap), the target itself comes with substantial uncertainty, which bears on the monetary policy decision. Of course, this uncertainty is implicit in all monetary policy decisions and the flexible targets simply make this uncertainty explicit. Nonetheless, as a means of discussing monetary policy this may increase difficulties of communication. It is also apparent, and consistent with our discussion concerning instrument quality, that there is greater uncertainty associated with the output forecast — and hence the forecast target — than there is for core inflation (as evident by the relative width of the relevant bands). Again, the implication is that communicating monetary policy decisions under flexible inflation targets is likely to be more difficult than simple inflation targets because of the associated uncertainty that becomes much more explicit with the flexible targets.

4. Conclusions

We estimate and test flexible inflation targets for Canada using monthly data over the inflation targeting period. The targets we estimate can be motivated from standard theoretical models of monetary policy in New Keynesian environments though they also have a fairly simple intuitive interpretation, representing the near term balance between inflation and cyclical variations in output that a central bank wishes to achieve.

Our results support flexible inflation forecast targets with a significant weight on output, between 0.2 and 0.4, in the target conditions over medium term horizons of 12-18 months based upon our single equation estimates. Our restricted systems estimates are slightly larger, 0.3 to 0.6. These targets are consistent with monetary policy that is successfully trading off inflation and economic activity. This result contrasts with other papers for the US that have not found support for an economically significant weight for economic activity in the objectives for monetary policy (e.g. Dennis, 2006). The principal qualification to our results is that the target coefficients are not stable across the policy horizon and that at longer horizons the weight on output growth is generally smaller. While this is not strictly consistent with theoretical models of optimal monetary policy, one possible justification may be the greater uncertainty associated with forecasts of output at longer horizons causing the bank to down weight output fluctuations at longer horizons.

One way to interpret our analysis is that it is analogous to the empirical literature on fitting Taylor rules as a description of monetary policy and a possible rule to which central banks might commit. Svensson (2003) and Woodford (2007) strongly advocate commitment by central banks to the sorts of forecast targets considered here. Our results demonstrate that the actual behaviour of the Bank of Canada is broadly consistent with meeting such targets and that a formal commitment would be a practical possibility. Our analysis also highlights an important practical issue in this regard: communicating a flexible target which comes with sampling uncertainty.

References

Anderson TW and Rubin H. 1949. Estimation of the parameters of a single equation in a complete system of stochastic equations. *Annals of Mathematical Statistics* **20**: 46–63.

Bank of Canada. 1995. *Monetary Policy Report.* Bank of Canada: Ottawa, Canada.

Baum CF, Schaffer ME, Stillman S. 2007. Enhanced routines for instrumental variables/GMM estimation and testing. Boston College Working Paper 667.

Bernanke BS, Laubach T, Mishkin FS, Posen AS. 1999. *Inflation Targeting: Lessons from the International Experience.* Princeton University Press: Princeton, NJ.

Buiter W. 2006. How Robust is the New Conventional Wisdom? The Surprising Fragility of the Theoretical Foundations of Inflation Targeting and Central Bank Independence. CEPR DP 5772. Centre for Economic Policy Research: London, UK.

Clarida R, Galí J, Gertler M. 1998. Monetary policy rules in practice: some international evidence. *European Economic Review* **42**: 1033–67.

Clarida R, Galí J, Gertler M. 1999. The Science of Monetary Policy: A New Keynesian Perspective. *Journal of Economic Literature* **37**: 1661–1707.

Consolo A, Favero CA. 2009. Monetary policy inertia: More fiction than fact?, *Journal of Monetary Economics* **56**: 900–6. DOI:10.1016/j.jmoneco.2009.06.007.

Davidson R, MacKinnon JG. 2004. *Econometric Theory and Methods.* Oxford University Press; Oxford, UK.

Dennis R. 2004. Inferring policy Inferring Policy Objectives from Economic Outcomes. *Oxford Bulletin of Economics and Statistics* **66** Supplement: 735–64.

Dennis R. 2006. The policy preferences of the U.S. Federal Reserve. *Journal of Applied Econometrics* **21**: 55–77.

Dufour J-M. 2003. Identification, weak instruments and statistical inference in econometrics. *Canadian Journal of Economics* **36**: 767–808.

Dufour, J-M, Khalaf L, Kichian M. 2006. Inflation dynamics and the New Keynesian Phillips Curve: An identification robust econometric analysis. *Journal of Economic Dynamics and Control* **20**: 1707–27.

Favero CA, Rovelli R. 2003. Macroeconomic stability and the preferences of the Fed: a formal analysis, 1961–98. *Journal of Money, Credit and Banking* **35**: 545–56.

Galí J. 2008. *Monetary Policy, Inflation and the Business Cycle.* Princeton University Press: Princeton, NJ.

Giannoni MP, Woodford M. 2005. Optimal inflation targeting rules. In *The Inflation Targeting Debate*, Bernanke BS, Woodford M (eds). University of Chicago Press: Chicago, IL.

Hansen LP. 1982. Large sample properties of generalized-method of moments estimators.

Econometrica **50**: 1029–54.

Hansen LP, Hodrick RJ. 1980. Forward exchange rates as optimal predictors of future spot rates: an econometric analysis. *Journal of Political Economy* **88**: 829–53.

Henry S, Pagan A. 2004. The Econometrics of the New Keynesian Policy Model: Introduction. *Oxford Bulletin of Economics and Statistics* **66** Supplement: 581–607.

King M. 1994. Monetary policy in the UK. *Fiscal Studies* **15**: 109–208.

Kuttner KN. 2004. The role of policy rules in inflation targeting. *Federal Reserve Bank of St Louis Review* **86**: 89–111.

Macklem T. 2001. A New Measure of Core Inflation. *Bank of Canada Review* Autumn: 3–12.

Newey W, West K. 1987. A simple, positive-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* **55**: 703–08.

Otto GD, Voss GM. 2009. Strict and Flexible Inflation Targets: An Empirical Tests of Inflation Forecast Targeting Models. HKIMR WP 20.

Rowe N, Yetman J. 2002. Identifying a policy-makers' target: an application to the Bank of Canada. *Canadian Journal of Economics* **35**: 239–56.

Rudd J, Whelan K. 2005. Modelling inflation dynamics: a critical review of recent research. Board of Governors of the Federal Reserve System FEDS Working Paper No 2005-06.

Ruge-Murcia FJ. 2009. Do Inflation-Targeting Central Banks Implicitly Target the Price Level? Mimeo, Department of Economics and CIREQ, University of Montreal.

Stock JH, Wright J, Yogo M. 2002. A survey of weak instruments and weak identification in generalised methods of moments. *Journal of Business and Economic Statistics* **20**: 518–29.

Stock JH, Yogo M. 2002. Testing for Weak Instruments in Linear IV Regression. NBER Technical Working Paper 284.

Svensson LEO. 1997. inflation forecast targeting: Implementing and monitoring inflation targets. *European Economic Review* **41**: 1111–46.

Svensson LEO. 1999. Inflation targeting as a monetary policy rule. *Journal of Monetary Economics* **43**: 607–54.

Svensson LEO. 2003. What is wrong with Taylor rules? Using judgement in monetary policy through targeting rules. *Journal of Economic Literature* **41**: 426–27.

Svensson LEO, Woodford M. 2005. Implementing optimal monetary policy through inflation forecast targeting. In *The Inflation Targeting Debate*, Bernanke BS, Woodford M (eds). University of Chicago Press: Chicago, IL.

Taylor JB. 1993. Discretion verses policy rules in practice. *Carnegie-Rochester Conferences Series on Public Policy* **39**: 195–214.

Woodford M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.*

Princeton University Press: Princeton, NJ.

Woodford M. 2004. Inflation targeting and optimal monetary policy. *Federal Reserve Bank of St Louis Review* 86(4), 15–41.

Woodford M. 2007. Forecast targeting as a monetary policy strategy: policy rules in practice. Mimeo, Columbia University.

Yogo M. 2004. Estimating the Elasticity of Intertemporal Substitution when Instruments are Weak. *Review of Economics and Statistics 86*: 797–810

Table 1: Single Equation Estimates for Flexible Inflation Target — Short Horizons

Sample: 1996:M1–2007:M12

Model: $E_t(\Delta_3 y_{t+h} + a_{1h}\pi^c_{t+h} - a_{0h})$; $\phi_h = 1/a_{1h}$ and $\pi^*_h = a_{0h}/a_{1h} - \overline{\Delta_3 y}/a_{1h}$.

Instruments: $z_t = \{1, \pi^c_{t-2}, \pi^{cm}_{t-2}, \Delta_{12} m_{t-2}\}$

h	1	2	3	4	5	6
a_{0h}	6.9751	6.4665	6.4368	5.1653	4.7840	5.3035
	(0.8574)	(1.0222)	(1.2038)	(1.2363)	(1.2553)	(1.3962)
	[0.0000]	[0.0000]	[0.0000]	[0.0001]	[0.0002]	[0.0002]
a_{1h}	2.1361	1.8681	1.8667	1.0840	0.8809	1.1088
	(0.4874)	(0.5703)	(0.6578)	(0.6345)	(0.6360)	(0.7103)
	[0.0000]	[0.0014]	[0.0053]	[0.0901]	[0.1685]	[0.1211]
ϕ_h	0.4681	0.5353	0.5357	0.9225	1.1351	0.9019
	(0.1068)	(0.1634)	(0.1888)	(0.5400)	(0.8195)	(0.5778)
	[0.0000]	[0.0014]	[0.0053]	[0.0901]	[0.1685]	[0.1211]
π^*_h	1.6892	1.6594	1.6447	1.6592	1.6088	1.7468
	(0.0811)	(0.1198)	(0.1410)	(0.2786)	(0.3835)	(0.2595)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0001]	[0.0000]
J	4.9182	4.1997	4.6710	6.8313	7.0274	6.3682
	[0.0855]	[0.1225]	[0.0968]	[0.0329]	[0.0298]	[0.0414]
F	36.3744	23.3334	15.9099	16.1688	16.1465	16.6552

Standard errors are reported in ( ). For ϕ_h and π^*_h, standard errors are calculated using a first order approximation. The estimate and standard errors for π^*_h are conditional on the mean of output growth: $\overline{\Delta_3 y} = 3.37$. Marginal significance levels, p-values, are reported in [ ]. For coefficient estimates, p-values are based upon a two-sided t-test. J is Hansen's (1982) test statistic for over-identification. Covariance estimation is Newey and West (1987) HAC with smoothing parameter set to $h-1$.

F is the Wald statistic (scaled by number of instruments) for the slope coefficients of the reduced form regression of π^c_{t+h} on the instruments z_t using the Newey and West (1987) HAC with smoothing parameter set to $h-1$. As a point of reference, these may be compared to the critical values for the null of weak instruments (with respect to TSLS relative bias) from Stock and Yogo (2002), Table 1: $\{13.91, 9.08, 6.46\}$ for 5%, 10%, and 20% bias.

Table 2: Single Equation Estimates of Flexible Inflation Target

Sample: 1996:M1–2007:M12

Model: $E_t(\Delta_3 y_{t+h} + a_{1h}\pi^c_{t+h} - a_{0h})$; $\phi_h = 1/a_{1h}$ and $\pi^*_h = a_{0h}/a_{1h} - \overline{\Delta_3 y}/a_{1h}$.

Instruments: $z_t = \{1, \pi^c_{t-2}, \pi^{cm}_{t-2}, \Delta_{12} m_{t-2}\}$

h	7	8	9	10	11	12	13	14	15	16	17	18
a_{0h}	6.5626	6.8492	5.9357	5.9047	6.1559	7.4825	8.0566	10.6671	11.2253	10.9839	10.6925	10.6874
	(1.5678)	(1.5682)	(1.4036)	(1.3095)	(1.2497)	(1.5536)	(1.7810)	(2.3434)	(2.2705)	(2.1758)	(2.0317)	(1.8517)
	[0.0001]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
a_{1h}	1.8359	1.9498	1.4943	1.5006	1.6457	2.3595	2.6590	4.0550	4.4020	4.2946	4.1551	4.1651
	(0.8145)	(0.8151)	(0.7271)	(0.6798)	(0.6373)	(0.8175)	(0.9397)	(1.2763)	(1.2524)	(1.2102)	(1.1345)	(1.0562)
	[0.0260]	[0.0183]	[0.0420]	[0.0291]	[0.0110]	[0.0046]	[0.0054]	[0.0019]	[0.0006]	[0.0005]	[0.0004]	[0.0001]
ϕ_h	0.5447	0.5129	0.6692	0.6664	0.6076	0.4238	0.3761	0.2466	0.2272	0.2328	0.2407	0.2401
	(0.2417)	(0.2144)	(0.3256)	(0.3019)	(0.2353)	(0.1468)	(0.1329)	(0.0776)	(0.0646)	(0.0656)	(0.0657)	(0.0609)
	[0.0260]	[0.0183]	[0.0420]	[0.0291]	[0.0110]	[0.0046]	[0.0054]	[0.0019]	[0.0006]	[0.0005]	[0.0004]	[0.0001]
π^*_h	1.7408	1.7861	1.7192	1.6913	1.6948	1.7444	1.7638	1.8003	1.7853	1.7736	1.7631	1.7576
	(0.1545)	(0.1273)	(0.1717)	(0.1742)	(0.1611)	(0.1227)	(0.1156)	(0.0922)	(0.0860)	(0.0861)	(0.0852)	(0.0790)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
J	5.2560	4.9901	4.4760	3.8055	2.7580	2.0270	1.5581	0.5499	0.0841	0.0204	0.0671	0.2051
	[0.0722]	[0.0825]	[0.1067]	[0.1492]	[0.2518]	[0.3629]	[0.4588]	[0.7596]	[0.9588]	[0.9898]	[0.9670]	[0.9025]
F	19.2474	19.3945	20.3103	18.8478	13.9137	11.1325	10.6483	9.0292	7.8001	8.0693	8.7389	10.0234

Standard errors are reported in ( ). For ϕ_h and π^*_h, standard errors are calculated using a first order approximation. The estimate and standard errors for π^*_h are conditional on the mean of output growth: $\overline{\Delta_3 y} = 3.37$. Marginal significance levels, p-values, are reported in [ ]. For coefficient estimates, p-values are based upon a two-sided t-test. J is Hansen's (1982) test statistic for over-identification. Covariance estimation is Newey and West (1987) HAC with smoothing parameter set to $h-1$.

F is the Wald statistic (scaled by number of instruments) for the slope coefficients of the reduced form regression of π^c_{t+h} on the instruments z_t using the Newey and West (1987) HAC with smoothing parameter set to $h-1$. As a point of reference, these may be compared to the critical values for the null of weak instruments (with respect to TSLS relative bias) from Stock and Yogo (2002), Table 1: $\{13.91, 9.08, 6.46\}$ for 5%, 10%, and 20% bias.

Table 3 Unrestricted System Estimates of Flexible Inflation Target

Sample: 1996:M1–2007:M12

Model: $E_t(\Delta_3 y_{t+h} + a_{1h}\pi^c_{t+h} - a_{0h})$; $\phi_h = 1/a_{1h}$ and $\pi^*_h = a_{0h}/a_{1h} - \overline{\Delta_3 y}/a_{1h}$.

Instruments: $z_t = \{1, \pi^c_{t-2}, \pi^{cm}_{t-2}, \Delta_{12} m_{t-2}\}$

	System: $h = 7 \ldots 9$			System: $h = 10 \ldots 12$			System: $h = 13 \ldots 15$			System: $h = 16 \ldots 18$		
h	7	8	9	10	11	12	13	14	15	16	17	18
a_{0h}	5.5688	5.7774	5.6129	5.8521	5.9054	6.6716	7.1776	7.0418	7.4225	11.5137	11.1051	10.4100
	(1.4382)	(1.4880)	(1.4060)	(1.0700)	(1.1170)	(1.2526)	(1.2076)	(1.3009)	(1.4444)	(1.6648)	(1.6096)	(1.6234)
	[0.0002]	[0.0002]	[0.0001]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
a_{1h}	1.2582	1.3852	1.3052	1.4924	1.5291	1.9556	2.1990	2.1308	2.3531	4.4615	4.2398	3.9054
	(0.7464)	(0.7767)	(0.7342)	(0.5556)	(0.5708)	(0.6424)	(0.6434)	(0.6859)	(0.7542)	(0.9793)	(0.9328)	(0.9142)
	[0.0945]	[0.0771]	[0.0780]	[0.0083]	[0.0084]	[0.0029]	[0.0009]	[0.0024]	[0.0023]	[0.0000]	[0.0000]	[0.0000]
ϕ_h	0.7948	0.7219	0.7662	0.6701	0.6540	0.5113	0.4547	0.4693	0.4250	0.2241	0.2359	0.2561
	(0.4715)	(0.4048)	(0.4310)	(0.2495)	(0.2441)	(0.1680)	(0.1330)	(0.1511)	(0.1362)	(0.0492)	(0.0519)	(0.0599)
	[0.0945]	[0.0771]	[0.0780]	[0.0083]	[0.0084]	[0.0029]	[0.0009]	[0.0024]	[0.0023]	[0.0000]	[0.0000]	[0.0000]
π^*_h	1.7503	1.7404	1.7210	1.6654	1.6603	1.6900	1.7330	1.7248	1.7236	1.8261	1.8252	1.8035
	(0.1898)	(0.1710)	(0.1812)	(0.1572)	(0.1660)	(0.1393)	(0.1115)	(0.1197)	(0.1121)	(0.0727)	(0.0715)	(0.0683)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
J		6.0187			4.8935			1.8457			1.9902	
		[0.4211]			[0.5575]			[0.9333]			[0.9206]	

Standard errors are reported in ( ). For ϕ_h and π^*_h, standard errors are calculated using a first order approximation. The estimate and standard errors for π^*_h are conditional on the mean of output growth: $\overline{\Delta_3 y} = 3.37$. Marginal significance levels, p-values, are reported in [ ]. For coefficient estimates, p-values are based upon a two-sided t-test. J is Hansen's (1982) test statistic for over-identification. Covariance estimation is Newey and West (1987) HAC with smoothing parameter set to one less than the maximum horizon for the system.

Table 4 Restricted System Estimates of Flexible Inflation Target

Sample: 1996:M1–2007:M12

Model: $E_t(\Delta_3 y_{t+h} + a_{1h}\pi^c_{t+h} - a_{0h})$; $\phi_h = 1/a_{1h}$ and $\pi^*_h = a_{0h}/a_{1h} - \overline{\Delta_3 y}/a_{1h}$.

Instruments: $z_t = \{1, \pi^c_{t-2}, \pi^{cm}_{t-2}, \Delta_{12} m_{t-2}\}$

	System: $h = 7\ldots9$	System: $h = 10\ldots12$	System: $h = 13\ldots15$	System: $h = 16\ldots18$
a_{0h}	4.8624	6.4816	7.4955	9.7725
	(1.1975)	(1.0109)	(1.1209)	(1.7417)
	[0.0001]	[0.0000]	[0.0000]	[0.0000]
a_{1h}	0.9140	1.7638	2.4230	3.6431
	(0.6375)	(0.5219)	(0.6065)	(0.9765)
	[0.1542]	[0.0010]	[0.0001]	[0.0003]
ϕ_h	1.0941	0.5670	0.4127	0.2745
	(0.7631)	(0.1678)	(0.1033)	(0.0736)
	[0.1542]	[0.0010]	[0.0001]	[0.0003]
π^*_h	1.6365	1.7660	1.7040	1.7583
	(0.2462)	(0.1205)	(0.0952)	(0.0690)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]
J	6.1430	6.0061	2.3840	3.2312
	[0.8031]	[0.8148]	[0.9925]	[0.9754]
$J^R - J^{UR}$	0.6840	2.5853	1.8156	4.2216
	[0.9533]	[0.6294]	[0.7696]	[0.3768]

Standard errors are reported in ( ). For ϕ_h and π^*_h, standard errors are calculated using a first order approximation. The estimate and standard errors for π^*_h are conditional on the mean of output growth: $\overline{\Delta_3 y} = 3.37$. Marginal significance levels, p-values, are reported in [ ]. For coefficient estimates, p-values are based upon a two-sided t-test. J is Hansen's (1982) test statistic for over-identification. Covariance estimation is Newey and West (1987) HAC with smoothing parameter set to one less than the maximum horizon for the system.

$J^R - J^{UR}$ is the difference of the criterion functions for the restricted and unrestricted model, both estimated using the weighting function from the unrestricted model (this is not the same as the difference of the Js reported for the two models in the tables). Under the null the restrictions are true, this difference is distributed $\chi^2(4)$; see Davidson and MacKinnon (p 368, 2004).

TABLE 5 RESTRICTED SYSTEM ESTIMATES OF FLEXIBLE INFLATION TARGET CONTINUED

Sample: 1996:M1–2007:M12

Model: $E_t(\Delta_3 y_{t+h} + a_{1h}\pi^c_{t+h} - a_{0h})$; $\phi_h = 1/a_{1h}$ and $\pi^*_h = a_{0h}/a_{1h} - \overline{\Delta_3 y}/a_{1h}$.

Instruments: $z_t = \{1, \pi^c_{t-2}, \pi^{cm}_{t-2}, \Delta_{12} m_{t-2}\}$

	System: $h = 10, 13, 16$	System: $h = 11, 14, 17$	System: $h = 12, 15, 18$
a_{0h}	6.6161	6.8590	7.8411
	(0.5382)	(0.5980)	(0.9647)
	[0.0000]	[0.0000]	[0.0000]
a_{1h}	1.9085	2.0550	2.5952
	(0.2899)	(0.3081)	(0.5269)
	[0.0000]	[0.0000]	[0.0000]
ϕ_h	0.5240	0.4866	0.3853
	(0.0796)	(0.0730)	(0.0782)
	[0.0000]	[0.0000]	[0.0000]
π^*_h	1.7026	1.6994	1.7241
	(0.0773)	(0.0818)	(0.0710)
	[0.0000]	[0.0000]	[0.0000]
J	5.4591	3.9753	3.5846
	[0.8585]	[0.9485]	[0.9641]
$J^R - J^{UR}$	16.8289	9.6716	7.3819
	[0.0021]	[0.0463]	[0.1170]

Standard errors are reported in ( ). For ϕ_h and π^*_h, standard errors are calculated using a first order approximation. The estimate and standard errors for π^*_h are conditional on the mean of output growth: $\overline{\Delta_3 y} = 3.37$. Marginal significance levels, p-values, are reported in [ ]. For coefficient estimates, p-values are based upon a two-sided t-test. J is Hansen's (1982) test statistic for over-identification. Covariance estimation is Newey and West (1987) HAC with smoothing parameter set to one less than the maximum horizon for the system.

$J^R - J^{UR}$ is the difference of the criterion functions for the restricted and unrestricted model, both estimated using the weighting function from the unrestricted model (this is not the same as the difference of the Js reported for the two models in the tables). Under the null the restrictions are true, this difference is distributed $\chi^2(4)$; see Davidson and MacKinnon (p 368, 2004).

Table 6: Sensitivity Analysis

Sample: 1996:M1–2007:M12

Model: $E_t(X_{t+h} + a_{1h}\pi_{t+h} - a_{0h})$; $\phi_h = 1/a_{1h}$ and $\pi_h^* = a_{0h}/a_{1h} - \overline{x}/a_{1h}$.

Instruments: $z_t = \{1, \pi_{t-2}^c, \pi_{t-2}^{cm}, \Delta_{12} m_{t-2}\}$

Horizon: $h = 12$.

Specification	$\pi = \pi^c$ $X = \Delta_3 y$	$\pi = \pi^{sp}$ $X = \Delta_3 y$	$\pi = \pi^c$ $X = \Delta y$	$\pi = \pi^c$ $X = \Delta_{12} y$	$\pi = \pi^c$ $X = \Delta_3 x$	$\pi = \pi^c$ $X = x$
a_0	7.4825	6.9226	8.4776	8.6836	1.8308	0.5735
	(1.5536)	(1.4051)	(2.0365)	(1.8479)	(1.4385)	(0.7218)
	[0.0000]	[0.0000]	[0.0001]	[0.0000]	[0.2055]	[0.4284]
a_1	2.3595	2.1107	2.7935	3.1526	0.9445	0.3639
	(0.8175)	(0.7495)	(1.0623)	(1.0186)	(0.7153)	(0.3720)
	[0.0046]	[0.0057]	[0.0096]	[0.0024]	[0.1891]	[0.3298]
ϕ_h	0.4238	0.4738	0.3580	0.3172	1.0587	2.7478
	(0.1468)	(0.1682)	(0.1361)	(0.1025)	(0.8017)	(2.8085)
	[0.0046]	[0.0057]	[0.0096]	[0.0024]	[0.1891]	[0.3298]
π_h^*	1.7444	1.6848	1.8318	1.7107	1.8991	1.8446
	(0.1227)	(0.1238)	(0.1258)	(0.0677)	(0.2387)	(0.3075)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
J	2.0270	1.7448	1.7442	2.8602	0.0604	2.4873
	[0.3629]	[0.4180]	[0.4181]	[0.2393]	[0.9703]	[0.2883]
F	11.1325	9.9421	11.1325	11.1325	11.1325	11.1325

Spliced core inflation π^{sp} is defined in Table A1. x refers to the output gap constructed as annual percentage deviation from HP filtered trend. See Table A1 for details.

Standard errors are reported in ( ). For ϕ_h and π_h^*, standard errors are calculated using the delta method. The estimate and standard errors for π_h^* are conditional on the mean of output growth. Marginal significance levels, p-values, are reported in [ ]. For coefficient estimates, p-values are based upon a two-sided t-test. J is Hansen's (1982) test statistic for over-identification. Covariance estimation is Newey and West (1987) HAC with smoothing parameter set to $h - 1$.

F is the Wald statistic (scaled by number of instruments) for the slope coefficients of the reduced form regression of π_{t+h}^c on the instruments z_t using the Newey and West (1987) HAC with smoothing parameter set to $h - 1$. As a point of reference, these may be compared to the critical values for the null of weak instruments (with respect to TSLS relative bias) from Stock and Yogo (2002), Table 1: $\{13.91, 9.08, 6.46\}$ for 5%, 10%, and 20% bias.

Table 7: Further Sensitivity Analysis — Later Sample

Sample: 2001:M5–2007:M12

Model: $E_t(\Delta_3 y_{t+h} + a_{1h}\pi^c_{t+h} - a_{0h})$; $\phi_h = 1/a_{1h}$ and $\pi^*_h = a_{0h}/a_{1h} - \overline{\Delta_3 y}/a_{1h}$.

Instruments: $z_t = \{1, \pi^c_{t-2}, \pi^{cm}_{t-2}, \Delta_{12} m_{t-2}\}$

h	10	11	12	13	14	15	16	17	18
a_{0h}	5.1853	6.7828	7.4086	7.8483	6.8614	7.3682	7.7326	5.4628	4.7909
	(0.8958)	(1.0612)	(1.5564)	(1.4371)	(1.0203)	(0.7835)	(1.1451)	(1.5319)	(1.9748)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0007]	[0.0183]
a_{1h}	1.2141	1.9830	2.3674	2.6052	2.1360	2.4356	2.6353	1.4273	1.0899
	(0.4585)	(0.5729)	(0.8477)	(0.7913)	(0.5793)	(0.4603)	(0.6428)	(0.7991)	(1.0374)
	[0.0104]	[0.0010]	[0.0070]	[0.0017]	[0.0005]	[0.0000]	[0.0001]	[0.0792]	[0.2977]
ϕ_h	0.8236	0.5043	0.4224	0.3838	0.4682	0.4106	0.3795	0.7006	0.9175
	(0.3110)	(0.1457)	(0.1513)	(0.1166)	(0.1270)	(0.0776)	(0.0926)	(0.3923)	(0.8733)
	[0.0104]	[0.0010]	[0.0070]	[0.0017]	[0.0005]	[0.0000]	[0.0001]	[0.0792]	[0.2977]
π^*_h	2.0196	2.0422	1.9750	1.9634	1.9327	1.9030	1.8971	1.9125	1.8879
	(0.1654)	(0.1142)	(0.1070)	(0.1043)	(0.0941)	(0.0827)	(0.0776)	(0.0887)	(0.1252)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
J	0.1102	0.7797	0.0008	0.0477	0.9375	1.7992	2.0953	1.4068	0.7703
	[0.9464]	[0.6772]	[0.9996]	[0.9764]	[0.6258]	[0.4067]	[0.3508]	[0.4949]	[0.6803]
F	24.1595	39.0492	28.3352	17.4491	15.3254	11.6954	20.5183	16.8623	31.1465

Standard errors are reported in ( ). For ϕ_h and π^*_h, standard errors are calculated using the delta method. The estimate and standard errors for π^*_h are conditional on the mean of output growth. Marginal significance levels, p-values, are reported in [ ]. For coefficient estimates, p-values are based upon a two-sided t-test. J is Hansen's (1982) test statistic for over-identification. Covariance estimation is Newey and West (1987) HAC with smoothing parameter set to $h-1$.

F is the Wald statistic (scaled by number of instruments) for the slope coefficients of the reduced form regression of π^c_{t+h} on the instruments z_t using the Newey and West (1987) HAC with smoothing parameter set to $h-1$. As a point of reference, these may be compared to the critical values for the null of weak instruments (with respect to TSLS relative bias) from Stock and Yogo (2002), Table 1: $\{13.91, 9.08, 6.46\}$ for 5%, 10%, and 20% bias.

Table 8: Identification Robust Methods

Sample: 1996:M1–2007:M12

	Anderson-Rubin			Anderson-Rubin HAC Robust		
h	Max p-value	ϕ_h	95% C.I.	Max p-value	ϕ_h	95% C.I.
10	0.0007	0.3459	–	0.0013	0.4905	–
11	0.0129	0.3186	–	0.0268	0.2107	–
12	0.1893	0.2785	[0.1976, 0.4137]	0.2991	0.2248	[0.1216, 0.5037]
13	0.4830	0.2634	[0.1645, 0.4565]	0.5188	0.2369	[0.1143, 0.5702]
14	0.8066	0.2478	[0.1435, 0.4594]	0.9087	0.2373	[0.0973, 0.5643]
15	0.9868	0.2311	[0.1264, 0.4385]	0.9940	0.2265	[0.0858, 0.5044]
16	0.9987	0.2315	[0.1256, 0.4443]	0.9993	0.2323	[0.0874, 0.5053]
17	0.9919	0.2360	[0.1279, 0.4605]	0.9950	0.2404	[0.0947, 0.5768]
18	0.9671	0.2457	[0.1360, 0.4840]	0.9803	0.2373	[0.1024, 0.6985]

Anderson-Rubin refers to the Anderson-Rubin F-statistic, distributed as $F(k_2, T - k_1 - k_2)$ where $T = 126$, $k_1 = 1$, $k_2 = 3$. Anderson-Rubin HAC Robust from the analogous Wald statistic, distributed $\chi^2(k_2)$ using the Newey and West (1987) HAC robust covariance matrix estimator with smoothing parameter set to $h - 1$. Maximum p-values are constructed from inverting the test statistics for values of ϕ_h ranging from 0 to 1. ϕ_h is the value associated with the maximal p-value. The 95% confidence intervals (C.I.) are constructed from the inverted p-statistics for all values of ϕ_h.

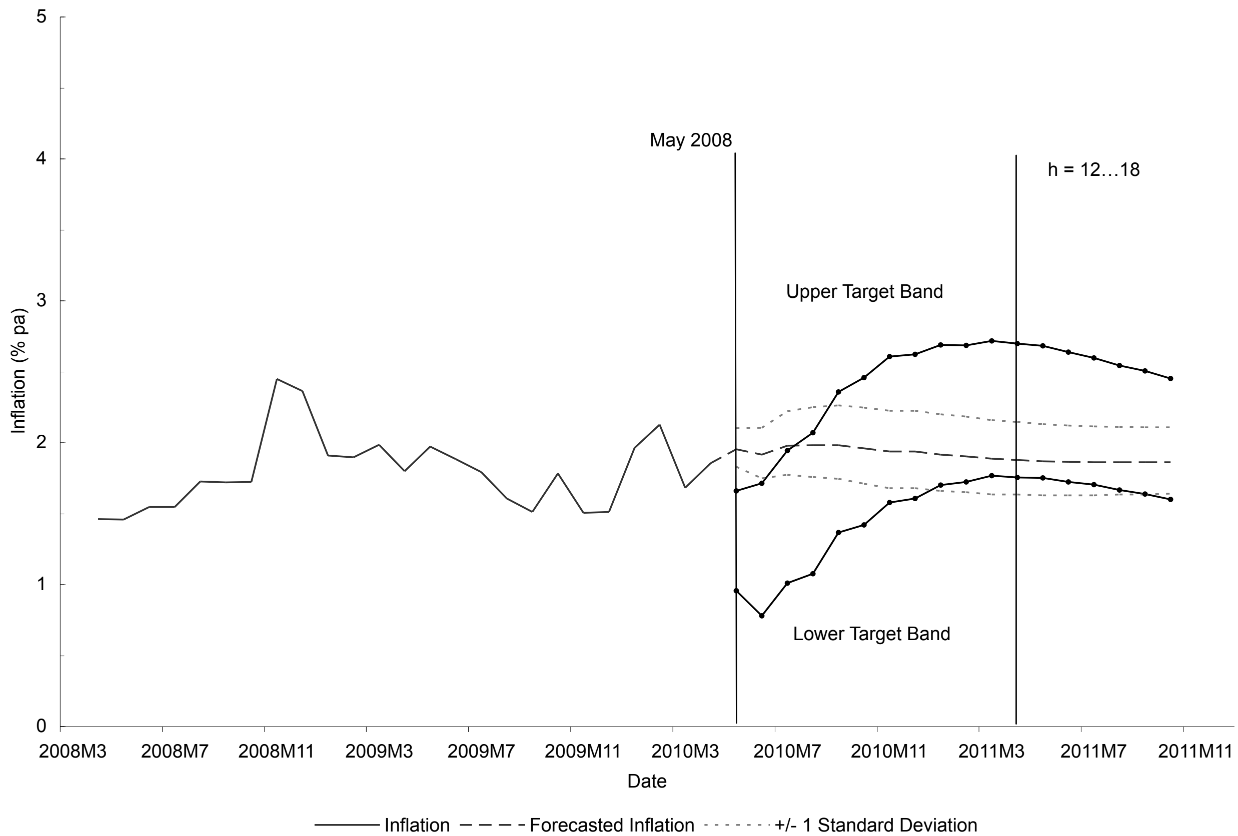

Figure 1: Core Inflation Forecasts with Flexible Inflation Target Bands