

Discussion of “Targeting Inflation from Below: How Do Inflation Expectations Behave?”

Eric T. Swanson
University of California, Irvine

The adoption of inflation targeting by central banks around the world has been nothing short of spectacular. In the twenty-five years since the Reserve Bank of New Zealand first adopted the inflation-targeting (IT) framework, it has become the *de facto* standard against which all other monetary policy frameworks are compared. While some countries, such as New Zealand, the United Kingdom, Sweden, and Norway, are very explicit about how their monetary policy is tied to IT, other countries, such as the United States, are not explicit followers of the framework and can only be considered inflation targeters “in spirit,” if at all.

Given the widespread success of inflation targeting as an idea, it is perhaps surprising that there is not more empirical evidence documenting its success as a monetary policy framework. In fact, the empirical evidence on this question is decidedly mixed, as can be seen in the nice literature review in Michael Ehrmann’s paper. While many studies have found that countries that adopted IT have lower inflation and/or better-anchored inflation expectations than they did before adopting IT, the difficulty lies in showing that those countries have lower inflation and/or better-anchored inflation expectations than a *control group* of countries that did not adopt IT. Ideally, the control group should also control for the initial level of inflation, since IT countries as a whole have tended to start from a higher initial inflation level, which was what drove them to adopt IT in the first place (Ball and Sheridan 2005). This correlation in the data makes it hard to separate the effects of IT from those of potential mean reversion (Ball and Sheridan 2005).

Thus, the difficulties in the empirical literature seem to be not so much documenting an improved performance of IT countries *per se*, but rather documenting a performance in IT countries that is *significantly better* than in other countries that didn’t adopt IT.

In “Targeting Inflation from Below: How Do Inflation Expectations Behave?” Michael Ehrmann looks at a somewhat different question. Rather than ask whether IT has affected the behavior of inflation and inflation expectations relative to a control group of *non-IT* countries, Michael investigates whether the behavior of inflation and inflation expectations differs across IT countries with *low* vs. *average* or *high* inflation. The comparison that’s being studied by Michael in this paper is thus substantially different from the comparisons that have been analyzed before, and may be more cleanly identified in the data.

1. Data and Sample

Of course, for any empirical study, it’s important to have good data. The first point to note about Michael’s analysis is the wealth of data that is brought to bear on the issue. The Consensus Economics data set consists of monthly observations for about ten countries over twenty or more years. Moreover, the Consensus Economics data includes not just the level of inflation but also inflation expectations for the current year, next year, and individual forecasters’ predictions, which allows the researcher to construct measures of cross-sectional forecast dispersion, or disagreement. The richness of the data set gives Michael a reasonable chance of identifying the difference that he is interested in—the differential behavior of inflation and inflation expectations in IT countries with low vs. high inflation.

Although the Consensus Economics inflation and mean inflation expectations data go back to at least the beginning of 1990 for all the countries in the sample, Michael deliberately chooses to focus only on the period during which each listed country was an inflation targeter. Thus, the “start dates” listed in the first column of table 1 of his paper are not due to data availability, but rather correspond to the date that inflation targeting in each country began. The reason for this choice is the paper’s different focus from the previous literature—rather than compare IT vs. non-IT countries, Michael’s paper focuses on the comparison between IT countries with low vs. average or high inflation. This results in a somewhat reduced sample size of about 1,872 country-month observations, relative to the maximum possible sample size of about 3,000 country-month observations that would be available using all of the Consensus Economics data

(not to mention the additional non-IT countries of France, Germany, Italy, and the Netherlands, which could be included in the analysis prior to 1999 along with Spain.) The somewhat reduced sample size that results from this choice is made up for by the clearer focus of the regressions on the comparison of interest—the additional observations discussed above would not help to identify the difference between IT countries with low vs. high inflation.

2. Correlation over Time and across Countries

As impressive as the large sample size in the data set is, it's important to bear in mind that those 1,872 (or even 3,000+) observations are not independent. Inflation is a persistent, slow-moving process in every country, so the residuals in each of Michael's regressions are likely to be serially correlated. Although each of those regressions includes country fixed effects, those fixed effects control only for differences in the average *level* of inflation across countries, and do not fully remove the serial correlation that is present within each country. A few of Michael's regressions include time fixed effects as well, but those also will not eliminate the problem—time fixed effects remove the average level of inflation residuals across countries each month but do not correct for the fact that countries with idiosyncratically high inflation one month will also tend to have idiosyncratically high inflation for several subsequent months.

Inflation is also likely to be correlated across countries due to changes in global commodity prices, global business-cycle conditions, and other global factors, but here the time fixed effects that Michael includes in a few of the regressions will typically be sufficient to soak up the correlation. Regressions without those time fixed effects, however, will be subject to this additional source of correlation as well.

A separate, but closely related, issue is that inflation has trended downward over time in virtually every country. After removing country fixed effects, these downward trends will appear as serial correlation in the residuals within each country, independent of whether inflation deviations from trend are serially correlated or not.

For all of these reasons, there are effectively far fewer than 1,872 independent observations in each of Michael's regressions. Ordinary least squares standard errors will be severely downward biased (and

t-statistics upward biased), even when country and/or time fixed effects are included. Thus, the use of the Driscoll-Kraay (1998) panel standard error correction in each regression is crucial for ascertaining the true statistical significance of the results. This correction (or something like it) is not performed nearly as often as it should be in panel studies of inflation targeting. However, as with the standard Newey-West (1987) procedure for a single time series, it is important to specify the number of lags of potential serial correlation in the Driscoll-Kraay (1998) correction. A moderate number of monthly lags, such as twelve, may be sufficient if serially correlated inflation deviations from trend are the only problem, but if the inflation data suffer from downward *trends* over the sample, then twelve lags may not be sufficient, since the residual serial correlation will persist for several years. (Of course, the best correction in case of downward trends would be to explicitly allow for these trends in the regression specification itself, rather than try to control for them in the standard error-correction procedure, but given that Michael's regressions do not include time trends, we should bear in mind the possibility of biased standard errors here.) It would be helpful if Michael offered some guidance as to the appropriate lag length for the standard error correction for the Consensus Economics inflation and inflation forecast data, and considered a specification with time trends as a check on the results.

3. Regression Results

The results in Michael's regressions are quite robust. IT countries with persistently low inflation (below target for nine months or more) seem to have substantially lower inflation expectations than IT countries with average or high inflation (table 3 in his paper), and those expectations are more sensitive to the level of inflation itself (also table 3). IT countries with persistently low inflation also seem to have significantly greater inflation forecast dispersion (table 4 in his paper). The statistical significance of these results holds across a wide variety of specifications in both tables, which is remarkable given the difficulty the literature has had in finding robust differences between IT vs. non-IT countries. Apparently, the difference between low- vs. average- or high-inflation IT countries is easier to identify in the data than the difference between IT vs. non-IT countries.

The results in tables 3 and 4 suggest that inflation expectations in IT countries with persistently low inflation are not as well anchored as in average- or high-inflation IT countries. Inflation expectations are below target, they are more below target for countries with lower inflation, and disagreement across forecasters is greater for countries that are more below target. The effects are not symmetric—they apply only to countries that are running below target, and not above. The regressions do not shed light on *why* inflation expectations seem to become unanchored when inflation is below rather than above target, but they do suggest that there is a cost of inflation running below target. To the extent that the benefits of inflation targeting come from a firmer anchoring of inflation expectations, those benefits seem to decline substantially if inflation remains below target for more than a few months. Of course, this raises the intriguing possibility that the benefits may decline to the point where non-IT countries could actually perform *better* than IT countries in a low-inflation environment, in terms of inflation and inflation expectations being closer to target and better anchored.

4. Interpretation of the Results

Let me now speculate as to *why* Michael finds the results that he does. In other words, why does inflation below target seem to de-anchor inflation expectations in IT countries, while inflation above target does not have the same de-anchoring effect?

A natural explanation seems to be one of perceived impotence of the central bank. When inflation is above target, it's natural for the public and the media to interpret the higher inflation as an optimal choice of the central bank. After all, the central bank's alternative was to raise interest rates further and thereby reduce real economic activity.

However, when inflation is running *below* target, it's more difficult for the public and the media to interpret the outcome as an optimal choice. Presumably, central banks prefer more real economic activity to less, as long as inflation does not rise substantially above target, and yet the central bank did not stimulate the economy. It's hard to imagine why the central bank wouldn't have done so unless it was somehow unable to stimulate the economy. Thus, inflation target misses on the downside naturally seem to suggest an explanation

based on central bank impotence rather than optimization. This perception of impotence is probably strengthened by the obvious zero lower bound constraint faced by many central banks in recent years.

It's important to note that Michael's results suggest a problem of *perceptions* of central bank impotence by forecasters, rather than actual impotence. Much theoretical work (e.g., Reifschneider and Williams 2000; Eggertsson and Woodford 2003) and empirical evidence (e.g., Gurkaynak, Sack, and Swanson 2005; Swanson and Williams 2014) suggests that central banks can work around the zero lower bound constraint without too much difficulty as long as they have some ability and willingness to commit to policy actions in the future. Nevertheless, central banks may suffer from a perception of impotence to the extent that these workarounds are not understood by the public and the Consensus Economics forecasters.¹

However, if Michael's results are driven by the zero lower bound and private-sector perceptions of central bank impotence, then this does call into question whether the inflation-targeting criterion is necessary or even relevant for the analysis. In other words, do we need to restrict attention to inflation-targeting countries to obtain the same results that Michael finds? Or would we find very similar results if we looked at a broader sample of low-inflation vs. high-inflation countries that included both IT and non-IT central banks? Some additional research on this question seems like it would be warranted.

5. Conclusions

In summary, Michael asks a question in this paper that is different from what has typically been considered in the inflation-targeting literature. As a result, he gets stronger and more robust results than is typical for that literature. Apparently, the difference between IT countries with low inflation vs. average or high inflation is better identified in the data than is the difference between IT and non-IT countries themselves.

¹ Of course, it's also possible that central banks do suffer from impotence to some extent at the zero lower bound, which would help to explain why several of them allowed inflation to run below target for several months or even years without doing more to stimulate their economies.

IT countries with persistently low inflation seem to suffer from a “de-anchoring” problem. Their inflation expectations are lower (relative to IT countries with average or high inflation), their inflation expectations are more sensitive to the level of inflation itself, and there is more disagreement across forecasters about the future path of inflation.

A natural explanation for these findings seems to be the zero lower bound and private-sector perceptions of central bank impotence. However, this raises the question whether the restriction to IT countries is necessary to obtain the same result. To the extent that *all* central banks face a problem of impotence (or perceived impotence) at the zero lower bound, then we should expect to see similar results even for non-IT central banks. Future research into this question seems warranted.

References

- Ball, L., and N. Sheridan. 2005. “Does Inflation Targeting Matter?” In *The Inflation-Targeting Debate*, ed. B. Bernanke and M. Woodford, 249–76. Chicago: University of Chicago Press.
- Driscoll, J. C., and A. C. Kraay. 1998. “Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data.” *Review of Economics and Statistics* 80 (4): 549–60.
- Eggertsson, G. B., and M. Woodford. 2003. “The Zero Bound on Interest Rates and Optimal Monetary Policy.” *Brookings Papers on Economic Activity* (Spring): 139–211.
- Gurkaynak, R. S., B. Sack, and E. T. Swanson. 2005. “Do Actions Speak Louder than Words? The Response of Asset Prices to Monetary Policy Actions and Statements.” *International Journal of Central Banking* 1 (1): 55–93.
- Newey, W., and K. West. 1987. “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix.” *Econometrica* 55 (3): 703–8.
- Reifschneider, D., and J. C. Williams. 2000. “Three Lessons for Monetary Policy in a Low-Inflation Era.” *Journal of Money, Credit and Banking* 32 (4): 936–66.
- Swanson, E. T., and J. C. Williams. 2014. “Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates.” *American Economic Review* 104 (10): 3154–85.