Journal of Monetary Economics 57 (2010) 226–232

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Testing the transparency benefits of inflation targeting: Evidence from private sector forecasts ☆

Christopher Crowe

International Monetary Fund, USA

ARTICLE INFO

Article history:
Received 12 May 2008
Received in revised form
22 December 2009
Accepted 23 December 2009
Available online 4 January 2010

JEL classification:
D82
E52
E58
G14

Keywords:
Inflation targeting
Central bank transparency
Inflation forecasts
Propensity score matching

ABSTRACT

Using inflation forecast data for 11 IT adoption countries, IT adoption is found to promote convergence in forecast errors, suggesting that it enhances transparency. This result, which is subjected to several robustness checks and found to be robust, supports Morris and Shin's (2002) contention that better public information is most beneficial for forecasters with bad private information. However, it does not support their hypothesis that better public information could make private forecasts less accurate.

1. Introduction

The consensus view among both policymakers and academics is that the introduction of inflation targeting (IT) increases the transparency of monetary policymaking (Bernanke et al., 1999; Faust and Henderson, 2004; King, 1997; Mishkin and Schmidt-Hebbel, 2007; Svensson, 1999).[1] Following Geraats (2002), transparency is defined as "the removal of information asymmetries." This study offers a direct test of the transparency benefits of IT by assessing whether the introduction of IT helps the private sector to better forecast inflation. The underlying assumption is that one potential way in which IT enhances transparency is by reducing the informational asymmetry between the central bank and the private sector with respect to the former's private information about future inflation.

To assess this potential channel, this study employs the Consensus Economics dataset, focusing on medium-term inflation forecasts by individual forecasters. The use of individual forecasters (rather than averages per country) allows one to test whether the effect of IT differs across forecasters with different characteristics. This is useful, since the simple signal

☆ This work reflects the views of the author alone and does not reflect the views of the IMF, its Executive Board or Management. The author would like to thank, subject to the usual caveats, Paolo Mauro, Ellen Meade, Scott Roger, Philip Schellekens, Hyun Shin and participants at seminars at American University, the IMF and the Bank of England and the Royal Economic Society's 2008 Annual Conference, as well as the Editor and two anonymous Referees, for comments on earlier drafts of the paper, and Martin Minnoni for excellent research assistance.

E-mail address: ccrowe@imf.org

[1] See Crowe (2010) for a discussion of the literature on IT and transparency.

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2009.12.003

extraction model outlined in Section 2 predicts that IT's impact on forecast accuracy will be conditional on how accurate the forecasts are in the first place.

A critical issue in assessing the effect of IT is that IT adoption is likely to be non-random. To control for this potential source of bias, forecasters in IT adoption countries are matched with a control group of forecasters in non-IT adoption countries using propensity score matching. A number of papers in the literature on IT have started to adopt these techniques for dealing with non-random adoption of IT (Lin and Ye, 2007; Vega and Winkelried, 2005). However, this paper is the first to apply these techniques to the question of IT's transparency benefits. It is also the first to use forecaster-level data and therefore uses forecaster characteristics to undertake the match. A second econometric problem (endogeneity resulting from the inclusion of the initial forecast error) is dealt with using IV estimation. The effect of IT adoption is then estimated by comparing the behavior of forecasters' errors in the 12 months leading up to and following the adoption of IT in each country. The sample includes 139 forecasters across 11 IT-adoption episodes and a matched control group drawn from non-IT adoption countries.

The results strongly support the view that IT adoption leads to better private sector forecasts, conditional on forecasters' initial forecast errors: the effect is strongest for those whose initial forecast accuracy is worst, in line with the model. This finding is subjected to several robustness checks, and found to be robust. However, while IT adoption improves forecast accuracy most for the worst forecasters, there is no evidence that the best forecasters are harmed. Hence, the results are not supportive of Morris and Shin's (2002) concerns over transparency's potential downsides.

The rest of the paper is organized as follows. Section 2 outlines a simple signal extraction model with public and private information to motivate the discussion and provide some predictions. Section 3 describes the data used. Section 4 outlines the methodology, while Section 5 gives the results and Section 6 concludes.

2. Theoretical framework

This section outlines a simple signal extraction model to motivate the empirical analysis. Agents ("forecasters") seek to minimize the squared error of their inflation forecast f_i, given the actual inflation rate π:

$$u_i(f_i,\pi) \equiv -(f_i-\pi)^2. \tag{1}$$

The private sector agents observe the central bank's public signal (a combination of its public forecasts, statements and analysis), π_C, and also observe their own private signal π_i. Each signal is noisy:

$$\pi_C = \pi + \eta,$$

$$\pi_i = \pi + \varepsilon, \tag{2}$$

with the variance of the iid error terms denoted σ_η^2 and σ_ε^2, respectively; the precision of the two signals is denoted as

$$\alpha = \frac{1}{\sigma_\eta^2},$$

$$\beta = \frac{1}{\sigma_\varepsilon^2}. \tag{3}$$

Agents therefore optimally weight the signals according to their relative precision:

$$f_i^* = \frac{\alpha\pi_C + \beta\pi_i}{\alpha+\beta}. \tag{4}$$

Then the expected mean square forecast error is given by

$$\tilde{V} \equiv E[(f_i-\pi)^2] = \frac{1}{(\alpha+\beta)}. \tag{5}$$

Taking this relation to the data requires some identifying assumptions: (a) that the precision of the private signals is constant, for each forecaster i, over the two year time period covered in the empirical section ($\beta_t^i = \beta^i$); and (b) that the precision of the public signal depends, over the same two year period, only on some country-specific factor and whether the central bank of country j has adopted inflation targeting ($\alpha_t^j = \alpha^j(IT); IT = \{0,1\}$). Hence the statement that IT improves transparency is equivalent to the condition that $\alpha^j(1) > \alpha^j(0)$. The forecast error for a typical forecaster in a non-IT ($IT = 0$) country is then given by

$$(\tilde{V}^{ij}|IT=0) \equiv \tilde{V}_0^{ij} = \frac{1}{(\alpha^j(0)+\beta^i)} \quad \text{with} \quad \frac{\partial}{\partial\alpha^j}\tilde{V}^{ij} < 0 \quad \text{and} \quad \frac{\partial^2\tilde{V}^{ij}}{\partial\alpha^j\partial\tilde{V}_0^{ij}} = -2\tilde{V}_0^{ij} < 0. \tag{6}$$

The introduction of IT, if it does indeed enhance transparency, should therefore reduce the average forecast error for all forecasters, but particularly for those whose forecast errors ($\tilde{V}_0^{ij}$) are initially large. This interaction effect is itself non-linear. However, it is likely to be difficult to accurately estimate third derivatives with the limited data sample available. Hence, a linearized interaction effect (around $\tilde{V}_0^{ij}$) gives the approximate effect of IT on forecast errors (conditional on $\tilde{V}_0^{ij}$)