The Effect of Inflation Targeting: A Mean-Reverting Mirage?*

Petra M. Geraats†

University of Cambridge

February 2014

Abstract

Inflation targeting has become a popular monetary policy strategy during the last two decades. This has given rise to a lively debate about the empirical effects of the adoption of inflation targeting. Some influential empirical studies have argued that the apparent improved performance of inflation targeters is merely regression to the mean, and controlling for the initial condition they conclude that inflation targeting does not matter. This paper challenges these findings that the apparent benefits of inflation targeting have basically been a mean-reverting mirage. It formally establishes that controlling for the initial condition generally leads to biased estimates of the 'treatment effect' of inflation targeting. In addition, it uses simulations to illustrate that tests based on such a specification have low power to detect the effectiveness of inflation targeting. As a result, prominent empirical findings that inflation targeting does not matter due to regression to the mean are misleading as the estimated treatment effects are biased and their tests lack power to distinguish an oasis from a mirage.

KEY WORDS: monetary policy; inflation targeting

JEL CODES: E52, E31, C52

---

*This paper is based on my discussion of Ball (2010) at the ECB/Handbook of Monetary Economics Conference "Key Developments in Monetary Economics" in Frankfurt. I thank participants of this conference as well as those of the Rimini Conference in Economics & Finance in Toronto, and seminar participants at the Nottingham School of Economics, for their comments.

†Faculty of Economics, University of Cambridge, Cambridge, CB3 9DD, United Kingdom. Email: Petra.Geraats@econ.cam.ac.uk.

1 Introduction

Inflation targeting has become a popular monetary policy strategy during the last two decades. This has given rise to a lively debate about the empirical effects of the adoption of inflation targeting. In a highly influential paper, Ball and Sheridan (2004) argue that the apparent improved performance of inflation targeters is merely 'regression to the mean'. They use a difference-in-differences specification that includes the initial condition to control for this, and conclude that inflation targeting does not matter. Ball (2010) uses a similar specification in the Handbook of Monetary Economics and also finds little evidence that inflation targeting has been beneficial.

This paper challenges these findings that the apparent benefits of inflation targeting have basically been a mean-reverting mirage. It finds that tests of the effects of inflation targeting using the Ball-Sheridan (BS) specification generally yield a biased estimate of the 'treatment effect' and tend to have low power to detect the effectiveness of inflation targeting. As a result, their specification could fail to find any effect of inflation targeting even if it has been highly effective.

It is shown analytically that the inclusion of the initial condition in the BS specification alters the interpretation of the coefficients. In particular, the coefficient estimate of the inflation-targeting indicator variable in their specification generally does not capture the treatment effect, i.e. the difference due to inflation targeting. Instead, if persistence is small, it reflects the difference in performance between inflation targeters and the control group in the period after the adoption of inflation targeting. So, if inflation targeting is actually effective at reducing inflation, but to an average level similar to others, then the BS specification gives the incorrect impression that inflation targeting has been ineffective.

In addition to showing analytically that estimates of the effect of inflation targeting tend to be biased in the BS specification, the low power of tests based on it is illustrated using Monte Carlo simulations. For instance, for plausible parameter values the paper finds that there would be no evidence of a significant effect of inflation targeting in 63% of replications even if there has in fact been a statistically and economically significant reduction in inflation of 2 percent point. The low power of the BS specification extends to cases in which their estimated treatment effect is actually unbiased and regression to the mean is relevant. Even in such a case it is shown that the BS specification fails to detect a significant reduction in inflation due to inflation targeting in 34% of replications.

Thus, tests of the effect of inflation targeting based on the BS specification tend to be unreliable.

The remainder of this paper is organized as follows. Section 2 sets up the framework for the analysis and provides a simple illustrative example that shows how the Ball-Sheridan specification could yield misleading results. This example is generalized to allow for persistence in section 3. The issue of regression to the mean is analyzed in section 4 and section 5 concludes.

2 Model

The effect of inflation targeting could be estimated using a difference-in-differences approach by comparing the change in a variable X_i (e.g. inflation in country i) before and after the adoption of inflation targeting (the 'treatment') to the change in X_i for others (the 'control group'). This leads to the specification

$$\Delta X_i = a_0 + b_0 I_i + \varepsilon_i \tag{1}$$

where $\Delta X_i \equiv X_{i2} - X_{i1}$ denotes the change in X_i from period 1 to 2, I_i is an indicator variable for country i adopting inflation targeting in period 2, ε_i is white noise. The coefficient a_0 captures the average change in X in the control group, and b_0 the effect of the treatment of inflation targeting on X.

However, suppose that countries with higher initial inflation are more likely to adopt inflation targeting (as is observed empirically), so that X_{i1} and I_i are positively correlated. In particular, Ball and Sheridan (2004) argue that X_{i1} may be high because of temporary shocks. If countries with high X_{i1} decide to adopt inflation targeting, then X_{i2} would be expected to be lower because of 'regression to the mean', even if inflation targeting were completely ineffective. So, estimation of (1) using ordinary least squares (OLS) would lead to a downward bias in b_0 because of a negative correlation between I_i and ε_i, and thereby overestimate the reduction in X due to the treatment effect.

To overcome this problem, Ball and Sheridan (2004) suggest to include the initial condition X_{i1}, so

$$\Delta X_i = a + bI_i + cX_{i1} + \varepsilon_i \tag{2}$$

If there is regression to the mean for X, the coefficient c for the initial condition X_{i1} would be expected to be negative, so a higher initial value X_{i1} reduces ΔX_i, leading to a relatively lower level of X_{i2}. The coefficient b is meant to capture the treatment effect of inflation targeting on X, corrected for regression to the mean. In the special case of $c = 0$, this specification reduces to (1) with $a = a_0$ and $b = b_0$.

To better understand the properties of the Ball-Sheridan (BS) specification (2), we first consider a simple illustrative example.

2.1 Illustrative Example

Assume that X_{it} is described by

$$X_{it} = \begin{cases} \mu_{Ot} + \varepsilon_{it} & \text{for} \quad I_i = 0 \\ \mu_{It} + \varepsilon_{it} & \text{for} \quad I_i = 1 \end{cases} \tag{3}$$

where μ_{It} and μ_{Ot} denote the average level of X in period t for inflation targeters and others, respectively, and ε_{it} is i.i.d. white noise with $\mathrm{E}\left[\varepsilon_{it}\right] = 0$ and $\mathrm{Var}\left[\varepsilon_{it}\right] = \sigma_{it}^2 \geq 0$ for all i and t, so X_{i1} and X_{i2} are independent. Suppose that inflation targeters are effective at achieving the inflation target X^* on average in period 2 so that $\mu_{I2} = X^*$, while other countries have an average of $\mu_{O2} = \mu_O$. So,

$$X_{i2} = \begin{cases} \mu_O + \varepsilon_{i2} & \text{for} \quad I_i = 0 \\ X^* + \varepsilon_{i2} & \text{for} \quad I_i = 1 \end{cases} \tag{4}$$

Note that the BS specification (2) can also be written as

$$X_{i2} = a + bI_i + (1 + c)\, X_{i1} + \varepsilon_i \tag{5}$$

This means that

$$X_{i2} = \begin{cases} a + (1 + c)\, X_{i1} + \varepsilon_i & \text{for} \quad I_i = 0 \\ a + b + (1 + c)\, X_{i1} + \varepsilon_i & \text{for} \quad I_i = 1 \end{cases}$$

Matching coefficients with (4) yields $c = -1$ and $\varepsilon_i = \varepsilon_{it}$, as the result should hold for any realization of X_{i1} and ε_{it}. Focusing on $I_i = 0$ and $I_i = 1$ then gives $a = \mu_O$ and $a + b = X^*$, respectively, which implies $b = X^* - \mu_O$. As a result, the BS specification (2) yields $a = \mu_O$, $b = X^* - \mu_O$ and $c = -1$.

This result also follows from the estimation of (5) by ordinary least squares (OLS). Let N be the number of observations in the sample, including $N_I \in \mathbb{N}$ adopting inflation targeting in period 2 and $N_O \in \mathbb{N}$ without inflation targeting, where $N = N_O + N_I$.

The observations X_{it} are described by (3). For analytical convenience, assume that $\sum_{i \in R} \varepsilon_{it} = 0$, where R denotes the monetary policy regime (with $I_i = 0$ or $I_i = 1$), so the sample average $\bar{X}_{Rt}$ of X_{it} equals $\bar{X}_{Ot} = \mu_{Ot}$ and $\bar{X}_{It} = \mu_{It}$ for $I_i = 0$ and $I_i = 1$, respectively.[1] Assume also that $\sum_{i \in R} \varepsilon_{i1} \varepsilon_{i2} = 0$, so the OLS estimate $\hat{\boldsymbol{\beta}}$ of $\boldsymbol{\beta} \equiv (a, b, 1 + c)'$ satisfies $\hat{\boldsymbol{\beta}} = \boldsymbol{\beta}$ exactly.[2] Then appendix A.1 shows that the OLS estimate for (5) equals

$$\hat{\boldsymbol{\beta}} = \begin{pmatrix} \mu_{O2}, & \mu_{I2} - \mu_{O2}, & 0 \end{pmatrix}' \tag{6}$$

So, again $a = \mu_{O2} = \mu_O$, $b = \mu_{I2} - \mu_{O2} = X^* - \mu_O$ and $1 + c = 0$, so $c = -1$. The same outcome is obtained for OLS estimation of (2).

This result has important implications for the interpretation of the coefficients in the BS specification. When the data are described by (3), the intercept a equals the average period 2 level of X for countries in the control group without inflation targeting, rather than the average change in X in the control group. Furthermore, the coefficient b does not capture the *average change* in X due to the treatment of inflation targeting (i.e. $(\mu_{I2} - \mu_{I1}) - (\mu_{O2} - \mu_{O1})$), but the *difference in the average level* of X between the treatment and control group in period 2 (i.e. $\mu_{I2} - \mu_{O2}$). Finally, the variable X_{i1} capturing the initial condition has a negative coefficient with a magnitude of one, or a zero coefficient in the specification (5) in levels. The latter result is intuitive since X_{i1} and X_{i2} are assumed to be independent according to (3).

This illustrative example shows how the coefficients in the BS specification could be completely misinterpreted. In particular, consider the plausible case in which countries that adopted inflation targeting initially had a structurally higher level of inflation than others ($\mu_{I1} > \mu_{O1}$) and after the adoption of inflation targeting successfully reduced it to their inflation target which is set at $X^* = \mu_O$, whereas those without inflation targeting experienced no change in average inflation ($\mu_{O1} = \mu_{O2} = \mu_O$). Then a regression using the BS specification (2) would give a treatment coefficient $b = 0$, giving the incorrect impression that inflation targeting has been ineffective!

The same result holds if there was also a (smaller) decline in average inflation for those without inflation targeting, such that $\mu_{O1} > \mu_{O2} = \mu_O$. No matter how high average inflation (μ_{I1}) initially was before inflation targeting, whenever the inflation target is set close to the average level of inflation of others ($X^* \approx \mu_O$), the estimated

[1] This assumption is relaxed in section 4.

[2] This presumes that $N \geq 3$ and $\exists \varepsilon_{i1} \neq 0$ to ensure that the three parameters in $\boldsymbol{\beta}$ can be estimated.

treatment effect is close to zero ($b \approx 0$), despite the fact that inflation targeting has successfully reduced inflation.

Clearly, b in the BS specification produces a biased estimate of the treatment effect of inflation targeting. However, in the special case of $\mu_{I1} = \mu_{O1}$, b yields the treatment effect, as the latter is reduced to $\mu_{I2} - \mu_{O2}$. But for $\mu_{I1} = \mu_{O1}$ there is no structural difference in initial inflation between countries that adopt inflation targeting and others, which appears to be at odds with the facts.

3 Persistence

The example above is based on the strong assumption that X_{it} is independent over time, which is not realistic when focusing on inflation or many other macroeconomic variables. In particular, although inflation targeters tend to show little inflation persistence, for other countries inflation tends to be quite persistent (Benati 2008). So it is important to allow for persistence in X, in particular X_{Ot}.

Before analyzing a more general case below, suppose now that X_{it} follows a random walk for countries without inflation targeting, so $X_{i2} = X_{i1} + \varepsilon_{i2}$ for $I_i = 0$. In particular, assume that X_{i1} is still given by (3) in period 1, but that now for period 2

$$X_{i2} = \begin{cases} \mu_{O1} + \varepsilon_{i1} + \varepsilon_{i2} & \text{for} \quad I_i = 0 \\ \mu_{I2} + \varepsilon_{i2} & \text{for} \quad I_i = 1 \end{cases} \tag{7}$$

where ε_{i2} is i.i.d. white noise. So, the effect of ε_{i1} is persistent for countries without inflation targeting, whereas inflation targeters manage to break with the past and are no longer affected by ε_{i1}. Assume again that $\sum_{i \in R} \varepsilon_{it} = 0$ and $\sum_{i \in R} \varepsilon_{i1}\varepsilon_{i2} = 0$, and denote $\sum_{i \in R} \varepsilon_{it}^2 = S_{Rt}$ and $S_t = S_{Ot} + S_{It}$, where R denotes the regime (with $I_i = 0$ or $I_i = 1$). Then appendix A.2 shows that the OLS estimate for (5) equals

$$\hat{\boldsymbol{\beta}} = \left(\tfrac{S_{I1}}{S_1}\mu_{O1}, \quad \mu_{I2} - \mu_{I1} + \tfrac{S_{I1}}{S_1}\left[\mu_{I1} - \mu_{O1}\right], \quad \tfrac{S_{O1}}{S_1} \right)' \tag{8}$$

The interpretation of the estimated coefficients is again quite different from what may be expected for the BS specification. The intercept a does not capture the average change in X in the control group (which equals zero here), but a fraction S_{I1}/S_1 of μ_{O1}, where S_{I1} captures the volatility of the shocks in period 1 for countries that subsequently adopt inflation targeting, with $0 < S_{I1}/S_1 < 1$.[3] Furthermore, the coefficient b does not

[3]The strict inequalities presume that $\exists \varepsilon_{i1} \neq 0$ for each regime R.

equal the average change in X due to the inflation targeting treatment, which is equal to $\mu_{I2} - \mu_{I1}$ in this case. Instead, if inflation targeting is effective at breaking with the past and reducing average inflation from $\mu_{I1} > \mu_{O1}$ to $\mu_{I2} < \mu_{I1}$, then the estimate for the 'treatment' coefficient b is *smaller* in magnitude than the actual effect. Thus again, the *estimated treatment effect is generally biased*, unless $\mu_{I1} = \mu_{O1}$. Note that this bias is increasing in S_{I1}. So, if inflation targeters experienced relatively high initial volatility (which is plausible since they tend to be small open economies), the bias in the estimated treatment effect would be exacerbated. Finally, the estimate for the 'mean-reversion' coefficient c is equal to $S_{O1}/S_1 - 1 = -S_{I1}/S_1 < 0$, so its magnitude is also increasing in the initial volatility for inflation targeters.

The bias for b makes it likely that OLS estimation of the BS specification would fail to find that inflation targeting has been effective. This can be illustrated by a Monte Carlo simulation. Suppose that inflation X_{i1} is described by (3) and X_{i2} by (7), where $\mu_{O1} = 2$ and $\mu_{I1} = 4 > \mu_{I2} = 2$; ε_{it} is normally distributed, $\varepsilon_{it} \sim N(0, \sigma^2)$ for all i and t with $\sigma^2 = 1$; and $N_O = N_I = 10$, so $N = 20$.[4] These parameter values imply a treatment effect of -2, which is significant in size, and a 95% confidence interval around the inflation target $\mu_{I2} = 2$ of $X_{i2} \in (0, 4)$ for $I_i = 1$, which appears empirically plausible. Then the OLS estimates for (2) are $\hat{a} = 1.00$ (0.64), $\hat{b} = -1.00$ (0.73) and $\hat{c} = -0.50$ (0.27), based on 100,000 replications (with standard errors in parentheses). It is straightforward to check that these coefficient estimates are consistent with the analytical result in (8). The magnitude of the estimated treatment effect $\hat{b}$ is clearly biased downward. Furthermore, the null hypothesis that inflation targeting is ineffective $H_0 : b = 0$ cannot be rejected in 73% of replications (using a significance level of 5%), despite the fact that inflation targeting has successfully achieved a sizeable reduction in average inflation compared to the control group. So, the BS specification has low power to detect the effect of inflation targeting.

To check to what extent the result would be better for a (much) smaller variance of the shocks, assume now that $\sigma^2 = 1/4$ instead (implying a 95% confidence interval of $X_{i2} \in (1, 3)$ under inflation targeting). Then the simulations yield $\hat{a} = 1.00$ (0.57), $\hat{b} = -1.00$ (0.59) and $\hat{c} = -0.50$ (0.27), so the coefficient estimates remain the same (in line with (8)) while the standard errors are reduced, but $H_0 : b = 0$ can still not be rejected in 63% of replications using the BS specification (again using 5% significance). In sharp contrast, using the specification in differences (1) without the initial condition

[4]Ball and Sheridan (2004) and Ball (2010) also use a sample size of 20 for their regressions.

X_{i1} (i.e. restricting $c = 0$), OLS estimation yields the unbiased result $\hat{a}_0 = 0.00$ (0.19) and $\hat{b}_0 = -2.00$ (0.27), and rejects $H_0 : b_0 = 0$ in all replications, using the same simulation. This is despite the fact that X_{i1} and I_i are highly correlated with a coefficient of 0.90. Clearly, a strong correlation between X_{i1} and I_i need not imply that OLS estimation of (1) is biased.

So far, the results in this section have been based on the assumption that X_{it} is independent over time for inflation targeters. However, it is probably optimistic to presume that inflation targeting implies a complete break with the past, so it is important to also allow for some persistence for inflation targeters. Nevertheless, assuming a random walk for inflation is problematic under inflation targeting. First of all, from a theoretical perspective, an effective inflation targeter is able to achieve an inflation target X^* on average regardless of past shocks, so ε_{i1} should not have a permanent effect. Furthermore, empirical evidence (Benati 2008) shows that inflation persistence is very low for inflation targeters, which is inconsistent with a random walk. So, a more general specification is used to model persistence.

Assume that X_{i1} is still given by (3), except that now the assumption of independence between ε_{i1} and ε_{i2} is relaxed. Instead, let $\varepsilon_{i2} = \rho_R \varepsilon_{i1} + \eta_{i2}$, where ρ_R denotes the persistence parameter for regime R, with $0 \leq \rho_R \leq 1$, and η_{i2} is i.i.d. white noise. This means that

$$X_{i2} = \begin{cases} \mu_{O2} + \rho_O \varepsilon_{i1} + \eta_{i2} & \text{for} \quad I_i = 0 \\ \mu_{I2} + \rho_I \varepsilon_{i1} + \eta_{i2} & \text{for} \quad I_i = 1 \end{cases} \tag{9}$$

This convenient hybrid specification nests the previous two data generation processes. In particular, $\rho_O = \rho_I = 0$ gives (3), while $\rho_I = 0$ and $\rho_O = 1$ with $\mu_{O2} = \mu_{O1}$ yields (7). Assume again that $\sum_{i \in R} \varepsilon_{i1} = 0$, $\sum_{i \in R} \varepsilon_{i1}^2 = S_{R1}$ and $S_1 = S_{O1} + S_{I1}$, as well as $\sum_{i \in R} \eta_{i2} = 0$ and $\sum_{i \in R} \varepsilon_{i1} \eta_{i2} = 0$, where R denotes the regime ($I_i = 0$ or $I_i = 1$). Then appendix A.3 shows that the OLS estimate for (5) equals

$$\hat{\boldsymbol{\beta}} = \begin{pmatrix} \mu_{O2} - \bar{\rho}\mu_{O1} \\ \mu_{I2} - \mu_{O2} - \bar{\rho}\left[\mu_{I1} - \mu_{O1}\right] \\ \bar{\rho} \end{pmatrix} \tag{10}$$

where $\bar{\rho} \equiv \frac{1}{S_1}\left(\rho_O S_{O1} + \rho_I S_{I1}\right)$ is a weighted average of ρ_R, with the weight S_{R1}/S_1 reflecting the initial relative volatility in regime R.

For the special case in which $\rho_I = \rho_O = 0$, $\bar{\rho} = 0$ so (10) reduces to (6). In addition, in the case of $\rho_I = 0$ and $\rho_O = 1$ with $\mu_{O2} = \mu_{O1}$, it is straightforward to check that $\bar{\rho} = S_{O1}/S_1 = 1 - S_{I1}/S_1$, so (10) is equal to (8). It is clear from (10) that the bias in the estimated treatment effect is not specific to these two cases but holds more generally for the BS specification. In particular, the true treatment effect equals $(\mu_{I2} - \mu_{I1}) - (\mu_{O2} - \mu_{O1})$, so the bias is $(1 - \bar{\rho})(\mu_{I1} - \mu_{O1})$, which is positive for $\mu_{I1} > \mu_{O1}$. So, if inflation targeters initially had structurally higher inflation than others, but then managed to reduce it, the magnitude of the estimated treatment effect is biased downward. This means that the BS specification underestimates the magnitude of the treatment effect, making it likely to incorrectly conclude that inflation targeting has been ineffective.

Again, there is no bias in the estimated treatment effect if $\mu_{I1} = \mu_{O1}$, but empirically initial inflation appears to have been structurally higher for inflation targeters than for others. Furthermore, the estimates for the intercept and treatment effect in the BS specification are both unbiased in the special case of $\bar{\rho} = 1$, which requires $\rho_O = \rho_I = 1$, so X follows a random walk for both inflation targeters and others. But, as mentioned before, a random walk in inflation is incompatible with a successful inflation targeter who manages to break with the past and achieve an inflation target X^* on average. Therefore, if inflation targeting is indeed effective, then $\bar{\rho} \neq 1$ and the estimated treatment effect of inflation targeting using the BS specification is biased, making it less likely to find a reduction in inflation.

Note that this bias in the estimated treatment effect is due to the BS specification that includes the initial condition X_{i1} as explanatory variable in an attempt to control for regression to the mean. In the specification in differences (1) without the initial condition (i.e. restricting $c = 0$), there is no bias and the OLS estimates for a_0 and b_0 are $(\bar{X}_{O2} - \bar{X}_{O1}) = (\mu_{O2} - \mu_{O1})$ and $(\bar{X}_{I2} - \bar{X}_{I1}) - (\bar{X}_{O2} - \bar{X}_{O1}) = (\mu_{I2} - \mu_{I1}) - (\mu_{O2} - \mu_{O1})$, respectively.[5]

4 Regression to Mean

The analysis so far has allowed for initial differences between inflation targeters and others based on structural factors, such as $\mu_{I1} > \mu_{O1}$ or $S_{I1} > S_{O1}$, but it has not considered selection into inflation targeting based on transitory shocks ε_{i1}, thereby precluding the issue of regression to the mean.

[5] This is derived in appendix A.4.

Suppose now that X_{it} is again described by (3), but that inflation targeting is completely ineffective and that there is no structural difference between inflation targeters and others, so $\mu_{Ot} = \mu_{It} = \mu_t$ and $\sigma^2_{Ot} = \sigma^2_{It} = \sigma^2_t$. Instead, countries that happen to have high inflation in period 1 with $X_{i1} > \mu_1$ decide to adopt inflation targeting, whereas others do not. This selection into inflation targeting means that $\bar{X}_{I1} > \mu_1 > \bar{X}_{O1}$. In period 2, however, $\bar{X}_{I2} = \bar{X}_{O2} = \mu_2$. Then OLS regression of difference-in-difference specification (1) yields an estimate $\hat{b}_0$ of $\left(\bar{X}_{I2} - \bar{X}_{I1}\right) - \left(\bar{X}_{O2} - \bar{X}_{O1}\right) = \bar{X}_{O1} - \bar{X}_{I1} < 0$, suggesting that inflation targeting has reduced X_{it}, although the true treatment effect b_0 is zero since inflation targeting has been assumed to be ineffective with $\mu_{Ot} = \mu_{It}$. Clearly, the estimated treatment effect $\hat{b}_0$ is biased; the reduction in X_{It} is simply regression to the mean. The bias is caused by the violation of the assumption that $\sum_{i \in R} \varepsilon_{i1} = 0$ as $\sum_{i \in I} \varepsilon_{i1} > 0 > \sum_{i \in O} \varepsilon_{i1}$ due to sample selection, so that $\bar{X}_{I1} > \mu_{I1} = \mu_{O1} > \bar{X}_{O1}$. The BS specification includes the initial condition X_{i1} in an attempt to control for regression to the mean.[6]

Ball (2010) argues that the BS specification gives an unbiased estimate $\hat{b}$ of the true treatment effect b_0. In his derivation, Ball (2010, appendix 1.1) assumes that selection into inflation targeting is based on $I_i = u_0 + u_1 X_{i1} + \eta_i$, where η_i is assumed to be independent of X_{i1}. However, $I_i \in \{0, 1\}$ is an indicator variable, so η_i must depend on X_{i1} and thereby on ε_{i1}, which means that his argument for unbiasedness breaks down. Nevertheless, for the special case in which $\mu_{O1} = \mu_{I1}$, the estimated treatment effect $\hat{b}$ is unbiased (as has been shown before). Intuitively, for $\mu_{O1} = \mu_{I1}$ variation in X_{i1} is entirely due to ε_{i1}, so controlling for X_{i1} removes the effect of regression to the mean. But for $\mu_{O1} \neq \mu_{I1}$ variation in X_{i1} also reflects structural differences that do not disappear over time, so including X_{i1} distorts the estimated treatment effect. However, even if $\mu_{O1} = \mu_{I1}$, the BS specification could yield misleading results and fail to find a significant effect of inflation targeting, as is illustrated by the following Monte Carlo simulations.

Suppose that inflation X_{it} is described by (3), where $\mu_{I1} = \mu_{O1} = \mu_{O2} = 2$ (so $\hat{b}$ is unbiased), $\mu_{I2} = 1$ (so the true treatment effect is -1), $\varepsilon_{it} \sim N\left(0, \sigma^2\right)$ with $\sigma^2 = 1$, and $N = 20$. Assume that country i adopts inflation targeting if $X_{i1} > 2$ (so regression to the

[6]Ball (2010, p. 1307) claims that this addresses the problem of endogeneity of I_i. However, simply including the selection variable (i.c. X_{i1}) in the regression generally does not solve endogeneity problems. Instead, the difference-in-differences approach could be combined with propensity score matching to obtain a suitable, comparable control group, as in Vega and Winkelried (2005).

mean applies). Then the OLS estimates for (2) are $\hat{a} = 2.0$ (0.60), $\hat{b} = -1.0$ (0.79) and $\hat{c} = -1.0$ (0.41), based on 100,000 replications (with standard errors in parentheses). Although these are consistent with (6) and the estimated treatment effect $\hat{b}$ is unbiased, the null hypothesis that inflation targeting is ineffective, $H_0 : b = 0$, cannot be rejected in 77% of replications (using a significance level of 5%), despite the fact that inflation targeting has successfully reduced average inflation. In contrast, using (1) instead, $H_0 : b_0 = 0$ is rejected in nearly all replications.[7] So, even when the BS specification (2) yields an unbiased estimate of the treatment effect, it may have much lower power to detect the effectiveness of inflation targeting than (1).

Now consider the same setup, but with less variable shocks or a stronger treatment effect, which both amount to a reduction in the relative importance of regression to the mean compared to the inflation targeting treatment. For instance, assuming that $\sigma^2 = 1/4$ or that $\mu_{I2} = 0$ (so the inflation target μ_{I2} now lies outside the 95% confidence interval for X_{i1}) produces the result that $H_0 : b = 0$ cannot be rejected in 34% of replications (again using 5% significance).[8] Ironically, for this case in which the BS specification is unbiased, its poor performance actually improves precisely when the relative importance of the regression to the mean that it aims to correct for declines compared to the treatment effect. However, despite the fact that inflation targeting has significantly reduced structural inflation, the BS specification still fails to detect that inflation targeting has been effective in about one third of replications.

The presence of regression to the mean presumes that the effect of the shock ε_{i1} is temporary, so it does not hold if X_{it} follows a random walk, which would yield a permanent effect of ε_{i1}. In addition, regression to the mean does not apply if the higher level of X_{i1} for $I_i = 1$ is a structural feature due to a higher mean μ_{I1}. In the latter case, the difference-in-difference specification (1) yields an unbiased estimate $\hat{b}_0$ of the treatment effect, whereas the magnitude of $\hat{b}$ in the BS specification (2) is generally biased downwards, as shown in section 3.

The crucial question is whether X_{i1} and I_i are correlated because of temporary shocks ε_{i1} or fundamental factors μ_{R1}. This may be hard to distinguish and it is likely to

[7]To be precise, the OLS estimates for (1) are $\hat{a}_0 = 0.80$ (0.37) and $\hat{b}_0 = -2.60$ (0.53), which reflects $\bar{X}_{I1} = 2.8$, $\bar{X}_{O1} = 1.2$, $\bar{X}_{I2} = 1.0$ and $\bar{X}_{O2} = 2.0$, while $H_0 : b_0 = 0$ is not rejected for only 0.4% of replications, using the same simulation.

[8]The OLS estimates of (2) are $\hat{a} = 2.0$ (0.68), $\hat{b} = -1.0$ (0.39) and $\hat{c} = -1.0$ (0.41) for $\sigma^2 = 1/4$, and $\hat{a} = 2.0$ (0.60), $\hat{b} = -2.0$ (0.79) and $\hat{c} = -1.0$ (0.41) for $\mu_{I2} = 0$.

depend on the context. For instance, if X_{it} is the rate of inflation in one year, then a high level of X_{it} could plausibly be due to a temporary positive shock ε_{it}. But if X_{it} is the (average) rate of inflation over a period of half a decade, then it is more likely to reflect a high structural factor μ_{it}. In the latter case, one would not expect X_{it} to automatically go down due to regression to the mean.[9] In particular, some countries may have structural features (e.g. small open economy, weak institutions) that make it more difficult to control inflation. They may suffer from structurally high inflation that is unlikely to subside unless measures are taken to mitigate the problem in some way (e.g. inflation targeting).

In Ball and Sheridan (2004), the pre-targeting sample period is at least 5 years and even up to 30 years. So, X_{i1} is a longer run average that is unlikely to exhibit much regression to the mean as it is already close to the mean μ_{i1}. This means that $\bar{X}_{I1} > \bar{X}_{O1}$ is mostly due to $\mu_{I1} > \mu_{O1}$. If countries with high μ_1 decide to adopt inflation targeting, X_{i1} and I_i are correlated, but OLS estimation of (1) is unbiased, whereas the BS specification (2) is biased (as $\mu_{I1} \neq \mu_{O1}$), unless $\rho_I = \rho_O = 1$ (as shown in section 3). But in the latter case, X_{it} follows a random walk and the effect of ε_{i1} is permanent, so there cannot be regression to the mean, which was the motivation for the BS specification.

To summarize, the presence of regression to the mean does not mean that the BS specification (2) is more suitable than the usual difference-in-differences specification (1) to test the effectiveness of inflation targeting. In particular, the BS specification generally yields a biased estimate of the treatment effect, unless $\mu_{I1} = \mu_{O1}$. But even in the latter case, it may have much lower power to detect the effectiveness of inflation targeting than (1), although its performance improves when regression to the mean becomes relatively less important. However, when X_{it} is a multi-year average, it exhibits little regression to the mean, so $\bar{X}_{I1} > \bar{X}_{O1}$ reflects $\mu_{I1} > \mu_{O1}$, which means that the estimated treatment effect using the BS specification is biased.

[9]Following Ball and Sheridan's (2004) baseball analogy, when I have a low batting average in a few games, one may think it is just temporary (perhaps due to an injury). But when my low batting average persists over time, the problem is more likely to be structural (e.g. reflecting poor hand-eye coordination), so one would not expect my batting average to go up.

5 Concluding Remarks

In influential contributions to the literature on the empirical effects of inflation targeting, Ball and Sheridan (2004) and Ball (2010) suggest that apparent improvements, such as a reduction in inflation, simply reflect regression to the mean after countries with temporarily high inflation decided to adopt inflation targeting. Using a modified difference-in-differences specification that aims to control for this by including the initial condition, they find little evidence that inflation targeting has been beneficial.

This paper exposes the shortcomings of their empirical approach by showing that their specification generally yields a biased estimate of the 'treatment effect' of inflation targeting and that their test of the effectiveness of inflation targeting has low power. It is shown analytically that their inclusion of the initial condition in a difference-in-differences specification alters the interpretation of the coefficients. In particular, if the persistence in the variable of interest is sufficiently small, the coefficient estimate of the inflation-targeting indicator variable in their regression does not capture the treatment effect of inflation targeting compared to the control group, but rather the difference in performance post-inflation targeting. So, when inflation targeting has succeeded in reducing inflation to the level of others, the Ball-Sheridan (BS) specification suggests it has been ineffective.

There are two special cases in which the estimate of the treatment effect is unbiased in the BS specification: (i) the structural average is the same for inflation targeters and others in the pre-targeting period, which appears to be at odds with empirical facts for inflation; and (ii) the variable follows a random walk for both inflation targeters and others, which is inconsistent with the central tenet of inflation targeting that it makes inflation stationary around the inflation target.

Some Monte Carlo simulations are used to illustrate the low power of tests based on the BS specification. For instance, if inflation targeters manage to significantly reduce average inflation by two percent points, while inflation of others follows a random walk, there is no evidence of a significant effect of inflation targeting in their specification for 63% of replications. In another simulation in which their estimated treatment effect is actually unbiased and regression to the mean is present, the BS specification fails to detect a significant reduction in inflation due to inflation targeting in 34% of replications.

However, when the BS specification is applied to multi-year averages, as in Ball and Sheridan (2004) and Ball (2010), there is unlikely to be much regression to the mean as the variable is already close to its mean. This means that the observed differences in average inflation between inflation targeters and others in the pre-targeting period mostly reflect differences in the structural average, which implies that the BS specification produces a biased estimate of the treatment effect.

To conclude, influential empirical findings that inflation targeting does not matter due to regression to the mean are misleading as their estimated treatment effects are biased and their tests lack power to distinguish an oasis from a mirage.

A Appendix

This appendix provides a formal derivation of the analytical results in this paper.

A.1 Illustrative Example

This section shows that estimation of $\boldsymbol{\beta} = (a, b, 1 + c)'$ in (5) using OLS yields (6) when the sample satisfies (3) with $\sum_{i \in R} \varepsilon_{it} = 0$ and $\sum_{i \in R} \varepsilon_{i1}\varepsilon_{i2} = 0$, where R denotes the regime (with $I_i = 0$ or $I_i = 1$).

Without loss of generality, order the observations i such that $I_i = 0$ for $i = 1, ..., N_O$ and $I_i = 1$ for $i = N_O + 1, ..., N$. Then the $N \times 3$ matrix of explanatory variables and the $N \times 1$ vector of the dependent variable are given by

$$\mathbf{X} = \begin{bmatrix} \mathbf{1}_{N_O} & \mathbf{0}_{N_O} & \mu_{O1}\mathbf{1}_{N_O} + \boldsymbol{\varepsilon}_{O1} \\ \mathbf{1}_{N_I} & \mathbf{1}_{N_I} & \mu_{I1}\mathbf{1}_{N_I} + \boldsymbol{\varepsilon}_{I1} \end{bmatrix} \text{ and } \mathbf{y} = \begin{pmatrix} \mu_{O2}\mathbf{1}_{N_O} + \boldsymbol{\varepsilon}_{O2} \\ \mu_{I2}\mathbf{1}_{N_I} + \boldsymbol{\varepsilon}_{I2} \end{pmatrix}$$

where $\mathbf{1}_N \equiv (1, ..., 1)'$ and $\mathbf{0}_N \equiv (0, ..., 0)'$ are $N \times 1$ vectors of ones and zeros, respectively; and $\boldsymbol{\varepsilon}_{Ot} \equiv (\varepsilon_{1t}, ..., \varepsilon_{N_O t})'$ and $\boldsymbol{\varepsilon}_{It} \equiv (\varepsilon_{N_O+1,t}, ..., \varepsilon_{Nt})'$ are $N_O \times 1$ and $N_I \times 1$ vectors of ε_{it} for $I_i = 0$ and $I_i = 1$, respectively. Note that $\mathbf{1}'_{N_O}\boldsymbol{\varepsilon}_{Ot} = \sum_{i=1}^{N_O} \varepsilon_{it} = 0$ and $\mathbf{1}'_{N_I}\boldsymbol{\varepsilon}_{It} = \sum_{i=N_O+1}^{N} \varepsilon_{it} = 0$. In addition, $\boldsymbol{\varepsilon}'_{R1}\boldsymbol{\varepsilon}_{R2} = 0$ for $R \in \{O, I\}$. For ease of notation, let $\boldsymbol{\varepsilon}'_{Ot}\boldsymbol{\varepsilon}_{Ot} = \sum_{i=1}^{N_O} \varepsilon_{it}^2 = S_{Ot}$ and $\boldsymbol{\varepsilon}'_{It}\boldsymbol{\varepsilon}_{It} = \sum_{i=N_O+1}^{N} \varepsilon_{it}^2 = S_{It}$, and denote $S_t = S_{Ot} + S_{It}$ for $t \in \{1, 2\}$.

The OLS estimate equals $\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{y}$. To compute $\hat{\boldsymbol{\beta}}$, start with straightforward matrix multiplication to get

$$\mathbf{X}'\mathbf{X} = \begin{bmatrix} N_O + N_I & N_I & N_O\mu_{O1} + N_I\mu_{I1} \\ N_I & N_I & N_I\mu_{I1} \\ N_O\mu_{O1} + N_I\mu_{I1} & N_I\mu_{I1} & N_O\mu_{O1}^2 + N_I\mu_{I1}^2 + S_1 \end{bmatrix}$$

using the fact that $\mathbf{1}'_{N_R}\mathbf{1}_{N_R} = N_R$, $\mathbf{1}'_{N_R}\boldsymbol{\varepsilon}_{R1} = 0$ and $\boldsymbol{\varepsilon}'_{R1}\boldsymbol{\varepsilon}_{R1} = S_{R1}$ for $R \in \{O, I\}$. Note that $\det(\mathbf{X}'\mathbf{X}) = N_O N_I S_1 > 0$, so $\mathbf{X}'\mathbf{X}$ is nonsingular. Its inverse equals

$$(\mathbf{X}'\mathbf{X})^{-1} = \frac{1}{N_O N_I S_1} *$$
$$\begin{bmatrix} N_I\left(N_O\mu_{O1}^2 + S_1\right) & -N_I\left(N_O\mu_{O1}^2 + S_1 - N_O\mu_{O1}\mu_{I1}\right) & -N_O N_I \mu_{O1} \\ -N_I\left(N_O\mu_{O1}^2 + S_1 - N_O\mu_{O1}\mu_{I1}\right) & N_O N_I\left(\mu_{O1} - \mu_{I1}\right)^2 + N S_1 & -N_O N_I\left(\mu_{I1} - \mu_{O1}\right) \\ -N_O N_I \mu_{O1} & -N_O N_I\left(\mu_{I1} - \mu_{O1}\right) & N_O N_I \end{bmatrix}$$

Postmultiplying this expression by $\mathbf{X}'$ and simplifying gives

$$\left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}' = \frac{1}{N_O N_I S_1} * \tag{11}$$

$$\begin{bmatrix} N_I S_1 \mathbf{1}'_{N_O} - N_O N_I \mu_{O1} \boldsymbol{\varepsilon}'_{O1} & -N_O N_I \mu_{O1} \boldsymbol{\varepsilon}'_{I1} \\ -N_I S_1 \mathbf{1}'_{N_O} - N_O N_I \left(\mu_{I1} - \mu_{O1}\right) \boldsymbol{\varepsilon}'_{O1} & N_O S_1 \mathbf{1}'_{N_I} - N_O N_I \left(\mu_{I1} - \mu_{O1}\right) \boldsymbol{\varepsilon}'_{I1} \\ N_O N_I \boldsymbol{\varepsilon}'_{O1} & N_O N_I \boldsymbol{\varepsilon}'_{I1} \end{bmatrix}$$

Using the fact that $\mathbf{1}'_{N_R}\boldsymbol{\varepsilon}_{R2} = 0$ and $\boldsymbol{\varepsilon}'_{R1}\boldsymbol{\varepsilon}_{R2} = 0$, postmultiplying by $\mathbf{y}$ yields

$$\left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'\mathbf{y} = \frac{1}{N_O N_I S_1}\begin{pmatrix} N_I S_1 N_O \mu_{O2} \\ -N_I S_1 N_O \mu_{O2} + N_O S_1 N_I \mu_{I2} \\ 0 \end{pmatrix}$$

Simplifying produces (6):

$$\hat{\boldsymbol{\beta}} = \begin{pmatrix} \mu_{O2} \\ \mu_{I2} - \mu_{O2} \\ 0 \end{pmatrix}$$

A.2 Random Walk

Now assume instead that X_{i2} is given by (7), so that

$$\mathbf{y} = \begin{pmatrix} \mu_{O1}\mathbf{1}_{N_O} + \boldsymbol{\varepsilon}_{O1} + \boldsymbol{\varepsilon}_{O2} \\ \mu_{I2}\mathbf{1}_{N_I} + \boldsymbol{\varepsilon}_{I2} \end{pmatrix}$$

Then premultiplying this by (11) gives

$$\left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'\mathbf{y} = \frac{1}{N_O N_I S_1}\begin{pmatrix} N_I S_1 N_O \mu_{O1} - N_O N_I \mu_{O1} S_{O1} \\ -N_I S_1 N_O \mu_{O1} - N_O N_I \left(\mu_{I1} - \mu_{O1}\right) S_{O1} + N_O S_1 N_I \mu_{I2} \\ N_O N_I S_{O1} \end{pmatrix}$$

Simplifying yields (8):

$$\hat{\boldsymbol{\beta}} = \begin{pmatrix} \frac{S_{I1}}{S_1}\mu_{O1} \\ \mu_{I2} - \left[\frac{S_{O1}}{S_1}\mu_{I1} + \frac{S_{I1}}{S_1}\mu_{O1}\right] \\ \frac{S_{O1}}{S_1} \end{pmatrix} = \begin{pmatrix} \frac{S_{I1}}{S_1}\mu_{O1} \\ \mu_{I2} - \mu_{I1} + \frac{S_{I1}}{S_1}\left[\mu_{I1} - \mu_{O1}\right] \\ \frac{S_{O1}}{S_1} \end{pmatrix}$$

A.3 General Persistence

Now assume instead that X_{i2} is given by (9), so that

$$\mathbf{y} = \begin{pmatrix} \mu_{O2}\mathbf{1}_{N_O} + \rho_O \boldsymbol{\varepsilon}_{O1} + \boldsymbol{\eta}_{O2} \\ \mu_{I2}\mathbf{1}_{N_I} + \rho_I \boldsymbol{\varepsilon}_{I1} + \boldsymbol{\eta}_{I2} \end{pmatrix}$$

where $\boldsymbol{\eta}_{Ot} \equiv \left(\eta_{1t}, ..., \eta_{N_O t}\right)'$ and $\boldsymbol{\eta}_{It} \equiv \left(\eta_{N_O+1,t}, ..., \eta_{Nt}\right)'$ are $N_O \times 1$ and $N_I \times 1$ vectors of η_{it} for $I_i = 0$ and $I_i = 1$, respectively. Note that $\mathbf{1}'_{N_R}\boldsymbol{\eta}_{R2} = 0$ and $\boldsymbol{\varepsilon}'_{R1}\boldsymbol{\eta}_{R2} = 0$ for $R \in \{O, I\}$. Then premultiplying $\mathbf{y}$ by (11) gives

$$\left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'\mathbf{y} = \frac{1}{N_O N_I S_1} *$$

$$\begin{pmatrix} N_I S_1 N_O \mu_{O2} - N_O N_I \mu_{O1} \rho_O S_{O1} - N_O N_I \mu_{O1} \rho_I S_{I1} \\ -N_I S_1 N_O \mu_{O2} - N_O N_I \left(\mu_{I1} - \mu_{O1}\right) \rho_O S_{O1} + N_O S_1 N_I \mu_{I2} - N_O N_I \left(\mu_{I1} - \mu_{O1}\right) \rho_I S_{I1} \\ N_O N_I \rho_O S_{O1} + N_O N_I \rho_I S_{I1} \end{pmatrix}$$

Simplifying produces (10):

$$\hat{\boldsymbol{\beta}} = \begin{pmatrix} \mu_{O2} - \frac{1}{S_1}\left(\rho_O S_{O1} + \rho_I S_{I1}\right)\mu_{O1} \\ -\mu_{O2} + \mu_{I2} - \frac{1}{S_1}\left(\rho_O S_{O1} + \rho_I S_{I1}\right)\left(\mu_{I1} - \mu_{O1}\right) \\ \frac{1}{S_1}\left(\rho_O S_{O1} + \rho_I S_{I1}\right) \end{pmatrix} = \begin{pmatrix} \mu_{O2} - \bar{\rho}\mu_{O1} \\ \mu_{I2} - \mu_{O2} - \bar{\rho}\left(\mu_{I1} - \mu_{O1}\right) \\ \bar{\rho} \end{pmatrix}$$

where $\bar{\rho} \equiv \frac{1}{S_1}\left(\rho_O S_{O1} + \rho_I S_{I1}\right)$.

A.4 Differences-in-Differences Result

Assume that X_{i2} is still given by (9), but consider now the differences-in-differences specification (1) without the initial condition X_{i1} (i.e. restricting $c = 0$). So, the $N \times 2$ matrix of explanatory variables and the $N \times 1$ vector of the dependent variable are given by

$$\mathbf{X}_R = \begin{bmatrix} \mathbf{1}_{N_O} & \mathbf{0}_{N_O} \\ \mathbf{1}_{N_I} & \mathbf{1}_{N_I} \end{bmatrix} \quad \text{and } \mathbf{y}_R = \begin{pmatrix} \left(\mu_{O2} - \mu_{O1}\right)\mathbf{1}_{N_O} + \left(\rho_O - 1\right)\boldsymbol{\varepsilon}_{O1} + \boldsymbol{\eta}_{O2} \\ \left(\mu_{I2} - \mu_{O1}\right)\mathbf{1}_{N_I} + \left(\rho_I - 1\right)\boldsymbol{\varepsilon}_{I1} + \boldsymbol{\eta}_{I2} \end{pmatrix}$$

Then

$$\mathbf{X}'_R\mathbf{X}_R = \begin{bmatrix} N_O + N_I & N_I \\ N_I & N_I \end{bmatrix} \quad \text{and } \left(\mathbf{X}'_R\mathbf{X}_R\right)^{-1} = \begin{bmatrix} \frac{1}{N_O} & -\frac{1}{N_O} \\ -\frac{1}{N_O} & \frac{N_O+N_I}{N_I N_O} \end{bmatrix}$$

so

$$\left(\mathbf{X}_R'\mathbf{X}_R\right)^{-1}\mathbf{X}_R' = \begin{bmatrix} \frac{1}{N_O}\mathbf{1}_{N_O}' & \mathbf{0}_{N_I}' \\ -\frac{1}{N_O}\mathbf{1}_{N_O}' & -\frac{1}{N_O}\mathbf{1}_{N_I}' + \frac{N_O+N_I}{N_I N_O}\mathbf{1}_{N_I}' \end{bmatrix}$$

Postmultiplying this by $\mathbf{y}_R$, using $\mathbf{1}_{N_R}'\boldsymbol{\varepsilon}_{R1} = 0$ and $\mathbf{1}_{N_R}'\boldsymbol{\eta}_{R2} = 0$:

$$\left(\mathbf{X}_R'\mathbf{X}_R\right)^{-1}\mathbf{X}_R'\mathbf{y}_R = \begin{pmatrix} (\mu_{O2}-\mu_{O1}) \\ -(\mu_{O2}-\mu_{O1}) + \left[-\frac{N_I}{N_O} + \frac{N_O+N_I}{N_O}\right](\mu_{I2}-\mu_{O1}) \end{pmatrix}$$

Therefore, the OLS estimate of $\boldsymbol{\beta}_0 \equiv (a_0, b_0)'$ in (1) is equal to

$$\hat{\boldsymbol{\beta}}_0 = \begin{pmatrix} \mu_{O2}-\mu_{O1} \\ (\mu_{I2}-\mu_{I1}) - (\mu_{O2}-\mu_{O1}) \end{pmatrix}$$

References

Ball, L. (2010), The performance of alternative monetary regimes, *in* B. M. Friedman and M. Woodford, eds, 'Handbook of Monetary Economics', Vol. 3, Elsevier, chapter 23, pp. 1303–1343.

Ball, L. M. and Sheridan, N. (2004), Does inflation targeting matter?, *in* B. S. Bernanke and M. Woodford, eds, 'The Inflation-Targeting Debate', Vol. 32 of *NBER Studies in Business Cycles*, University of Chicago Press, chapter 6, pp. 249–276.

Benati, L. (2008), 'Investigating inflation persistence across monetary regimes', *Quarterly Journal of Economics* **123**(3), 1005–1060.

Vega, M. and Winkelried, D. (2005), 'Inflation targeting and inflation behavior: A successful story?', *International Journal of Central Banking* **1**(3), 153–175.