Working Paper | 491 | Outubro de 2018

FGV | SAO PAULO SCHOOL OF ECONOMICS

Where Did Inflation Targeting Matter?

Rodrigo S. Barbosa
Ricardo D. Brito
Vladimir K. Teles

Where Did Inflation Targeting Matter?

Rodrigo S. Barbosa

Sao Paulo School of Economics - FGV

Ricardo D. Brito

Insper Institute of Education and Research

Vladimir K. Teles

Sao Paulo School of Economics - FGV

Abstract

We investigate the effects of inflation targeting (IT) adoption on industrial economies by comparing each inflation targeter country (ITer) with its synthetic control, defined as the convex combination of non-IT countries that best reproduce the ITer counterfactual inflation trajectory. We show that most of the ITers enjoyed lower inflation and higher output growth than their synthetics in most of the 1990 years' IT experience. Although those gains could be transitory, they were economically important to justify IT Central Banks optimism with their regime choice, both case-by-case and on average.

Keywords: Inflation Targeting, Inflation-output growth short-run tradeoff, Synthetic Control, Case studies.

JEL Classification: E42, E52.

1. Introduction

Does inflation targeting (IT) matter? In an influential article, Ball and Sheridan (2005) address this question for seven pioneer IT countries (ITers) among twenty industrial economies through cross-section difference-in-difference regressions, finding that IT adoption did not show significant average effects on macro performance. Deepening in the control for observable variables, common time-trends and country fixed-effects, several authors use other comparative case study methods to produce sharper measures of the average effects of IT in similar samples without reaching a consensus. Lin and Ye (2007) conclude that IT had no significant effects on inflation measures matching propensity scores of IT adoption, while Brito (2011) shows that IT introduction reduced inflation and increased short-run output growth in dynamic panels with controls for IT covariates, for example.

Although better sets of control variables make treated and non-treated units more comparable for treatment evaluation purposes – i.e. make the conditional independence assumption (CIA) more plausible – given the presence of unmeasured factors can never be excluded, there remains reasonable uncertainty about the control group ability to reproduce the counterfactual outcome trajectories that treated units would have experienced in the absence of treatment. In the macro-policy context, this latter doubt is strengthened because the usual samples of countries are small for the idiosyncratic factors to get diversified away.

Under which circumstances would you trust the effectiveness of a medicine tested on seven out of twenty patients only? Is a measure of the average effect of treatment in this group a sufficient statistic? Unconditionally, it is insufficient because the intervention inferred effectiveness (or ineffectiveness) may result from a few patients that are successful (or unsuccessful) for other unobservable factors than the treatment. To increase confidence on small sample results, it is helpful to have supplementary disaggregated evidence of the similarities between treated and control patients' outcome variable trajectories. Disaggregation makes explicit how the treatment effect spreads among treated patients, and the trajectories' perspective warns about unobserved differences. Only patients that are originally alike are able to produce similar paths of the outcome variable over extended periods of time.

Therefore, more reliable than selecting non-ITers on the basis of subjective measures of affinity with ITers, it is to drive the selection process by the ability of the non-ITers to reproduce the ITers pre-intervention outcome trajectories. Additionally, given the small number of countries, it is assuring to evaluate the gains from the use of IT on a case-by-case basis.

In this article, we provide evidence on the IT regime performance from a case-by-case perspective that transparently addresses the comparability between ITers' and non-ITers' outcomes. We apply the synthetic control method, proposed by Abadie and Gardeazabal (2003) and Abadie, Diamond and Hainmueller (2010), that provides a systematic data-driven approach to choose comparison units and allows precise quantitative inference in small-sample comparative studies.

The idea behind this method is that, among a few countries, a convex combination of untreated countries, named synthetic control, usually better approximates the treated country characteristics than any single country alone. By selecting untreated countries that are comparable to the treated one and weighting the formers to best reproduce the latter outcome trajectory over extended periods of time, the synthetic control implicitly accounts for unobservable time-varying and heterogeneous effects of factors, thus being truly comparable and valid to evaluate the causal effects of IT. By systematizing the process of estimating the counterfactual, the method also enables to conduct falsification exercises, termed "placebo studies", that provide the building blocks for an alternative mode of quantitative inference in small samples.

To measure the effects of IT on ITers, we compare the levels of inflation and output growth of the 7 pioneer ITers with their respective synthetics, on a case-by-case and on average, during the 1990 decade of IT inception. We reveal that ITers' annual inflation rates were lower than their respective synthetic estimates in 36 out of 42 observations (the sum of the number of IT years among ITers until 1999). Given the synthetic interpretation of the counterfactual, we claim that IT lowered average annual inflation in 0.86 percentage points in Australia, in 0.91 percentage points in Canada, in 0.72 percentage points in New Zealand, in 0.95 percentage points in Sweden, and in 0.81 percentage points in the United Kingdom, without noticeable effects in Finland's and Spain's inflations (in Table 5.A). For the same ITer-synthetic pairs, ITers' annual output growth rates were higher than their respective synthetic estimates in 27 out of the 42 observations, in spite of pre-IT higher average growth of some synthetics.[1] Not accounting for these ITer-synthetic pre-IT growth differences, the ITers overcame their synthetics' annual average growth rates in 0.82 percentage points in Australia, in 0.61 percentage points in Canada, in 2.66 percentage points in Finland, in 1.44 percentage points in Spain, and in 0.94 percentage points in Sweden; and were overcome in 0.79 percentage points in New Zealand and in 0.76 percentage points in the United Kingdom (in Table 5.B). After accounting for these ITer-synthetic pre-IT growth differences, the IT relative improvement in average annual growth amounts to 2.20 percentage points in Australia, 4.44 in Finland and even turns into positive 2.51 percentage points in the United Kingdom (in Table 5.C).

When looked from the perspective of the average effect of IT on ITers – i.e. the average effect of treatment on treated (ATT) – in an annual AR(1) panel, the direct effect of IT on the annual inflation rates was a significant average reduction of 0.56, with a cumulative IT effect of 0.80 percentage points. Simultaneously, the direct effect of IT on the annual output growth rates was a significant average increase of 1.11, with a cumulative IT effect of 1.42 percentage points. [2]

Although the continuation of the enhancing IT effects was unclear from the 1999's pictures, the above numbers sketched a picture where IT adoption simultaneously caused inflation to lower and output growth to rise among ITers in general. Irrespective of their significance, such numbers seem economically important enough to justify the Central Bank(er)s optimism with the novel IT system in the beginning of the 2000s.[3]

The rest of this article is organized as follows. Section 2 briefly describes the synthetic control methodology and our adaptation to the IT evaluation context. The data sources and samples used are described in section 3. Section 4 reports and discusses the results. The conclusions are in section 5.

2. Methodology

This section provides an intuitive explanation of the synthetic control method applied to the IT evaluation context, compares it with the linear regression and PSM techniques already used in the IT literature and describes alternative inference procedures for small samples. For more details, we refer the reader to the original articles by Abadie and Gardeazabal (2003), Abadie, Diamond and Hainmueller (2010) (hereafter, ADGH), and Abadie, Diamond and Hainmueller (2011).

2.1. The synthetic control estimator

Studies that estimate the average effect of a treatment by comparative case between treated and non-treated groups, using regression or propensity score methods, implicitly assume the conditional independence assumption (CIA), i.e. that the differences between groups are limited to the post-treatment period, or that the statistical procedure controls for all relevant pre-treatment differences. If the treated and non-treated units are different in unobservable ways, then non-treated units are unable to reproduce the counterfactuals of treated, and the estimated impact of treatment may be biased.

ADGH's insight is that a synthetic unit constructed through a weighted average of non-treated units has better chances to reproduce the counterfactual trajectory of the treated unit than any single non-treated unit alone, mainly when the units of analysis are a few aggregate entities, like countries. By matching the pre-intervention outcomes of the IT treated country over extended periods of time, the synthetic country implicitly controls for confounding unobservable characteristics that vary with time, overcoming a major limitation of linear regression or propensity score matching techniques.

As Abadie, Diamond and Hainmueller (2011) intuitively explain, only units that are alike in both observed and unobserved determinants of the outcome variable should produce similar trajectories of the outcome variable over many time periods. Once the IT treated country and its synthetic pair have presented similar behavior over pre-IT adoption, a discrepancy in the outcomes post-IT adoption is interpreted as caused by the IT police itself.

Algebraically, suppose that we are interested in evaluating the impact of a treatment on the macro variable Y, for which we have a balanced panel for $j = 1, \ldots, (J + 1)$ countries and $t = 1, \ldots, T$ periods. The value that would be observed in country

j on date t without intervention, $Y_{j,t}^{N}$, is given by the general factor model $Y_{j,t}^{N} = f\left(\lambda_t \mu_j, \theta_t Z_{j,t}, \varepsilon_{j,t}\right)$, where μ_j represents country-fixed non-observables, $Z_{j,t}$ is for country observed covariates, $\varepsilon_{j,t}$ are unobserved transitory shocks with zero mean, and the coefficients λ_t and θ_t generalize the usual difference-in-difference (fixed-effects) model allowing the effects of confounding unobserved and observed characteristics to vary with time.[4] Without loss of generality, assume that only the first country (*j* = *1*) is treated by the program under assessment, and denote $Y_{1,t}^{I}$ the outcome under treatment.[5] Let $t = T_0$ be the treatment starting date and $t = (T_0 + 1)$ be the initial date of treatment potential effects. The one-period lag embodies the time needed for macro policies to transmit to the economy. For simplicity, assume that the program operates without interruption until the final date T and that it has no effect on the outcome before getting started, such that $Y_{j,t}^{I} = Y_{j,t}^{N}$ for $\forall\ j$ and $t \le T_0$. Finally, define X_1, the $(K \times 1)$ vector of country *1*'s pre-intervention characteristics that predict its post-intervention outcomes in the absence of treatment effects. In other words, X_1 comprises characteristic of country *1* realized up to time T_0, meant to predict post-intervention outcomes, but themselves not affected by the intervention.

In this context, the main variable of interest, the program's impact on country *1*:

$$\alpha_{1,t} = Y_{1,t}^{I} - Y_{1,t}^{N}, \qquad \forall\ t = (T_0 + 1), ..., T, \tag{1}$$

is not directly observable, and the value effectively observed is:

$$Y_{1,t} = Y_{1,t}^{N} + \alpha_{1,t} D_{1,t}, \tag{2}$$

with $D_{1,t}$ a dummy variable equal to 1 if $t > T_0$ or 0 otherwise.

Because $Y_{1,t} = Y_{1,t}^I$ for $t \geq (T_0 + 1)$ in equation (2), the value of $Y_{1,t}^I$ is observed, and to reach the objective of measuring the treatment's impact $\alpha_{1,t}$ in equation (1), we need to estimate $Y_{1,t}^N$.

The synthetic method estimates the counterfactual trajectory of the treated country, i.e. the synthetic control $\hat{Y}_{1,t}^N$, as the convex combination of the untreated countries that best resembles the pre-treatment predictors' characteristics of the treated outcome:

$$\min_W \|X_1 - X_0 W\|_V = \sqrt{(X_1 - X_0 W)' V (X_1 - X_0 W)} \quad (3)$$

subject to:

$$W = (w_2, ..., w_{J+1})' \; , \; \sum_{j=2}^{J+1} w_j = 1 \text{ , and } w_j \geq 0, \quad \forall \quad j = 2, ..., (J+1) \; ,$$

where X_0 is the $(K \times J)$ matrix with countries *j = 2, ..., (J + 1)* predictors characteristics, and V is a $(K \times K)$ diagonal and positive semidefinite matrix that reflects the relative importance of the different predictors of the outcome Y .

In more details, for each X_1's *k-th* row indicating the value of a country *1*'s pre-treatment characteristic that is correlated to its outcome, the coordinate (k, j) of matrix X_0 contains the pre-treatment value of the same characteristic for country j .

The vector of optimal weights W^* defines the synthetic control, a fictitious untreated country with predictors' characteristics $X_0 W^* = \hat{X}_1$ as close as possible to country *1*'s pre-treatment ones, X_1.[6] Similarly, the synthetic control estimator $\hat{Y}_{1,t}^N$ of $Y_{1,t}^N$ is given by:

$$\hat{Y}_{1,t}^N = Y_{0,t} W^* \text{ for } t = (T_0 + 1), ..., T , \tag{4}$$

with $Y_{0,t} = (Y_{2,t}, ..., Y_{J+1,t})$ the $(1 \times J)$ vector with the countries *j = 2, ..., (J + 1)* outcome variables.

Therefore, in case Y denotes inflation, the effect of IT on ITer's inflation is measured by the difference between the ITer's post-treatment inflation and its synthetic counterpart:

$$\hat{\alpha}_{1,t} = Y_{1,t}^I - \hat{Y}_{1,t}^N \text{ for } t = (T_0 + 1), ..., T . \tag{5}$$

Before we move on, a few points are in order to deepen intuition about the synthetic technique implementation.

Fist regarding the optimization problem (3), to generate a synthetic control that matches the pre-intervention variable of interest of the treated over extended periods of time, it is helpful to include in the vector X_1 many of its pre-intervention values $Y_{1,T_0 - j}$, for $0 \leq j < T_0$. Additionally with the same objective, ADGH suggest to make use of the degree of freedom on the choice of the matrix V to minimize the mean square error in predicting the variable of interest during the pre-intervention period.

Another aspect that deserves reflection is how the synthetic convex combination compares with other methods. Abadie, Diamond and Hainmueller (2011) prove that linear regression also implies a linear combination with coefficients that add to one, in spite of not restricting them to the [0; 1] interval and, therefore, allowing extrapolation outside the support of the data. Regression extrapolation could be detected if the regression weights were explicitly calculated. Because regression weights are not calculated in practice, however, the extent of the extrapolation produced by regression techniques is typically hidden from the analyst, as well as the similarities and differences between the units treated and the regressions' linear combination units. In contrast to regression analysis, and as in small sample comparative studies, the synthetic control method makes explicit the contribution of each comparison unit to the counterfactual of interest, allowing researchers to use quantitative and qualitative techniques to analyze the similarities between the treated unit and the synthetic control.

Third, while using weights that sum to one and fall in the [0; 1] interval prevents extrapolation biases, synthetic controls – like Mill's difference and other matching methods – may present severe interpolation biases if the control group contains units that are very different from the treated ones. Thus, it is important to restrict the control group to units with outcomes that are thought to be driven by the same structural process as the treated and that were not subject to idiosyncratic structural shocks during the period of study.

2.2. Inference

On a case-by-case basis, the synthetic method provides point estimates of the program's impact $\hat{\alpha}_{j,t} = \left(Y^{I}_{j,t} - \hat{Y}^{N}_{j,t}\right)$ for $t = \left(T_0 + 1\right), ..., T$, but does not offer confidence intervals to infer about its significance.

To address this issue, ADGH proposes an inference procedure analogous to the permutation test, useful in cases where the number of observations in the control group is small.[7] As explained by Abadie, Diamond and Hainmueller (2011), this alternative model of inference is based on the premise that our confidence that a particular treated-synthetic difference estimate reflects the impact of the intervention under study would be severely undermined if we obtained estimated effects of similar or even greater magnitudes in cases where the intervention did not take place.

In our notation terms, suppose there are $(J+1)$ individuals, of whom 1 will be randomly selected to receive treatment. Let Y_1^I be the observation of the treated patient and $Y_2^N, ..., Y_{J+1}^N$ be the observations of individuals who did not receive treatment. One possible measure of the impact of the treatment is the difference $\overline{\alpha} = \left(\overline{Y}^I - \overline{Y}^N\right)$, with $\overline{Y}^I = Y_1^I$ and $\overline{Y}^N = J^{-1}\sum_{i=2}^{J+1} w_i Y_i^N$, for given weights $\left(w_2, ..., w_{J+1}\right)$. Under the null hypothesis that the treatment does not affect the patient (H_0), the value of $\overline{\alpha}$ should be the same regardless of the individual selected for treatment, and there are $(J+1)$ different possibilities of selection, each with probability $(J+1)^{-1}$ of being observed. In the permutation method, we obtain the distribution of those $(J+1)$ differences and reject H_0 if $\overline{\alpha}$ is an extreme value of this distribution.

The inference procedure for the synthetic control suggested by ADGH is similar. The problem (3) is solved and the difference $\hat{\alpha}_{j,t} = \left(Y_{j,t}^I - \hat{Y}_{j,t}^N\right)$ is computed for every control unit $j = 2, ..., (J + 1)$, swapping the treated country by the control country j, called placebo study. The intervention is considered effective if the estimated impact on the country effectively subjected to treatment is extreme in relation to the distribution of the impacts on the placebos.

Because the choice of V influences the mean square error of the estimator, for each treated country, we compute its respective set of placebos with the same V matrix chosen for that treated country. And we exclude from the distribution the placebos that present a root mean square error (RMSE) larger that three times the RMSE of the treated country.

Additionally, given we have more than one treated country, it is also informative to compute the average effect of IT on ITers, that we implement through dynamic panel regressions of the inflation difference:

$$\hat{\alpha}_{n,t} = \rho^{\pi} \cdot \hat{\alpha}_{n,t-1} + \beta^{\pi} \cdot D^{IT}_{n,t-1} + \eta^{\pi}_{n} + \upsilon^{\pi}_{n,t} \quad \text{for } n = \textit{Australia, Canada, Finland, New Zealand, Spain, Sweden, and United Kingdom,} \tag{6}$$

where ρ^{π} and β^{π} are common coefficients, $D^{IT}_{n,t-1}$ is a dummy variable equal to *1* if ITer n has adopted IT in period *t-1* or *0* if it has not, η^{π}_{n} allows for cross-country fixed-effects, and $\upsilon^{\pi}_{n,t}$ is the disturbance; and of the output growth rate difference:

$$\hat{\alpha}^{g}_{n,t} = \rho^{g} \cdot \hat{\alpha}^{g}_{n,t-1} + \beta^{g} \cdot D^{IT}_{n,t-1} + \eta^{g}_{n} + \upsilon^{g}_{n,t}, \tag{7}$$

where $\hat{\alpha}^{g}_{n,t}$ is the growth rate difference between the ITer n and its synthetic (with the same weights found in problem (3) for inflation), ρ^{g} and β^{g} are common coefficients, η^{g}_{n} allows for cross-country fixed-effects, and $\upsilon^{g}_{n,t}$ is the disturbance.

In both equations (6) and (7), due to the $\hat{\alpha}$'s AR(1) structure, the direct impact of the IT policy, β, implies a cumulative effect of $[\beta/(1-\rho)]$, which accounts for subsequent decreasing over-time lagged $\hat{\alpha}$ on the right-hand side.

The small number of ITer-synthetic pairs (*N=7 or 6*) relative to number of years (*T=29*) and the likely weak exogeneity of the IT variable $D_{n,t-1}^{IT}$ drive us to use the pooled Ordinary Least Square (OLS) estimation, with ITer dummies (also known as Least Square Dummy Variable – LSDV), clustering standard errors by ITer for accurate inference purposes.

3. Data

As described in Table 1, to minimize concerns of interpolation bias, we choose to analyze a panel of 24 industrial countries that could be considered relatively homogeneous in terms of macroeconomic management and performance, as of the end of the 1980s. We restrict the analysis to countries that adopted the inflation-targeting regime prior to 1999. January of that year marked the end of national monetary policies for ten of those countries, due to the Euro adoption. During the 2000's, some of the IT regime's best practices had become widespread among non-IT central banks, what makes the measurement of the IT effect difficult in recent years. Therefore, we assess the impacts of the IT regime on inflation and output growth rates in seven IT pioneering countries: Australia, Canada, Finland, New Zealand, Spain, Sweden and the United Kingdom.[8]

< Insert **Table 1** around here >

The annual consumer price indices and per capita Gross Domestic Products series are from the World Economic Outlook (WEO) by the International Monetary Fund (IMF).

The literature diverges on when to date the adoption of IT and distinguishes between converging and constant IT (for examples of different IT calendars, see Johnson (2002), Corbo et al. (2002), Roger and Stone (2005), Ball and Sheridan (2006), and Mishkin and Schmidt-Hebbel (2007)). Mishkin and Schmidt-Hebbel's (2007) IT adoption dates are the closest to Bernanke's et al. (1999) country-case studies and IMF's (2006). They also better represent the common central banks operating procedure of initializing an IT regime with explicit targets for the coming year on, and not for the current year, described in Johnson (2002). The time lag also recognizes the stylized fact that it takes some time for the monetary policy to transmit to the economy, documented by Friedman (1961) and Batini and Nelson (2002), among others. This is an aspect we model below and, thus, the reason why we prefer to focus on Mishkin and Schmidt-Hebbel's (2007) IT adoption dates.

4. Results

This section first describes our selection of inflation predictors and the inflation model adjustment to data. Then, we analyze the results on a case-by-case and on average.

4.1. Model adjustment

To identify the impact of the IT adoption on the inflation rate via the synthetic control method, we first have to define a set of inflation predictors, whose pre-treatment values for the treated country X_1 will be approximated by a convex combination of the same predictors for units in the control group X_0W.

After a review of the theoretical and empirical literature on inflation prediction, we have parsimoniously chosen to use lagged inflation and lagged output growth as predictors only. A long tradition in Economics, well synthesized in the adaptive Phillips curve, states that current inflation relates to its past realizations and to some measure of economic activity. Although economic activity is often measured in terms of a gap from potential, like output gap or unemployment rate, what matters for monetary policy evaluation is how inflation and growth are being traded off in the short-run. Moreover, it is also our goal to match unobserved fixed effects more evident in levels than in gap measures.

In addition to those usual predictors, variables like money growth, central bank independence, fiscal balance, exchange rate regime, trade openness or terms of trade – shown to be significant in some inflation studies, like Ghosh et al. (2002) – have not been consistently significant to variations in the sample of periods and countries, or to inclusions of lagged inflation, common-time and country fixed effects, the reason why we leave them out.[9] As shown in Cecchetti et al. (2000) and Ang et al. (2006), among others, it is hard to find inflation indicators that consistently improve the simple autoregressive projections.

Therefore, given our selected predictor (X_1) for each of the seven ITers studied in Table 4.A: *(i)* the inflation rate in the year of IT adoption, *(ii)* the average inflation rate in the four years previous to the IT adoption year, *(iii)* the average growth rate in the five years previous to the IT adoption year, and *(iv)* the average inflation rate in the five-to-nine years previous to the IT adoption year; the Tables 2, 3 and 4.B respectively show the weights attributed to the inflation predictor variables (the main diagonal of the diagonal and positive semidefinite matrix V, in Table 2), the weights attributed to the non-IT countries forming the synthetic control (the vector W^*, in Table 3), and the pre-IT inflation predictors values of its synthetic (the vector X_0W, in Table 4.B) that result from solving problem (3) for each ITers.

Before analyzing the IT effects, it is interesting to get a perspective of how the ADGH's synthetic algorithm works in practice. In Table 2, we see that the relative importance of each inflation predictor varies a lot from country to country when let V free to better adjust pre-treatment inflation data. Although the inferential procedures are valid for any V, its choice influences the mean square error of the estimator. For this reason, the optimal V for each ITer (in Table 2) is going to be imputed in its respective placebo study below.

Table 3 illustrates the alleged transparency of the synthetic method, making explicit the contribution of each comparison country to the counterfactual of interest, and allowing the reader to qualitative analyze the similarities between the ITers and the synthetic controls. By comparing Tables 4.B with 4.A, the reader can also get a sense of the ITer-synthetic quantitative differences. The fit looks excellent for New Zealand and Spain, reasonably good for Canada and Sweden, and fair for Finland and United Kingdom. Not by coincidence, those latter four synthetics miss their ITers output growth rates, the predictor variable that gets almost zero importance in their V matrices. The fit for Australia seems poorer than for the others, and comparisons with its synthetic require more care.

4.2. The inflation targeting effects

As previously mentioned, we evaluate the impact of adopting the IT system by comparing the post-treatment inflation of the ITers with their synthetics. Given that the inflation rate is the adjusted variable Y in problem (3), the synthetic inflation intuition is that of the "counterfactual", i.e., the inflation performance that the ITer would have experienced had it not adopted IT.

Additionally, because of the inflation-output short-run tradeoffs involved in monetary policy, for a sensible benefit-cost analysis, we should also account for the simultaneous output growth rates of the ITers and their synthetics, with the same synthetic weights optimally defined to simulate the counterfactual inflation processes of the ITers. Because the weights that approximate the ITer pre-IT inflation rates might not approximate pre-IT output growth rates as well, however, the latter comparison is less straightforward. Due to pre-existing growth dissimilarities, the synthetic country output growth rate should not be interpreted as that of the counterfactual output process of ITer, but as a benchmark for the relative disinflation costs involved. The ITer-synthetic output growth difference tells us by how much the ITer grew more than its synthetic, while the output growth rate difference-in-difference (that controls for previously existing output growth dissimilarities) tells us by how much the output growth rate difference improved.

On case-by-case, Figures 1.1 to 1.7 plot the inflation trajectories of the seven ITer-synthetic pairs since 1971, allowing the reader to appreciate how well, and how long into the past, the synthetics replicate ITers. Additionally, because the large worldwide inflation reduction observed during the 1980s represents a change in magnitude that makes difficult to appreciate the low levels trajectories in the 1990s, Figures 2.1 to 2.7 provide zooms in of inflation and output growth trajectories around the times of IT adoptions. In all figures, the dotted vertical line indicates the year of IT adoption. The Table 4 and Figures 2.1.b-2.7.b show that some synthetics that did not match their ITer pre-IT output growth rates (Australia, Finland, and United Kingdom) presented higher pre-IT output growth than their ITers, what means these ITer-synthetic difference-in-difference measures are going to be even more favorable to ITers than the simple ITer-synthetic differences measures.

For each of the seven ITers studied, Table 5 shows the time-average post-treatment difference that measure the effect of IT on the ITer – i.e. the effect of treatment on treated, $\overline{\alpha}_n = (T - T_0)^{-1} \sum_{t=T_0+1}^{T} \hat{\alpha}_{n,t}$ – and its relative ranking position among the placebos in terms of inflation rates (in Table 5.A) and output growth rates (in Table 5.B). As just explained, because the synthetic control weights are defined to match the inflation trajectory, and not to match the output growth trajectory, we also present the post-treatment output growth differences adjusted to the pre-treatment growth differences in Table 5.C – i.e. output growth rate difference-in-difference instead of just output growth rate difference.

4.2.A. Canada and Sweden

First looking at Canada, Table 3 shows the countries composition of Canada's synthetic (W^*). It has approximately 56% of Denmark, 32% of Switzerland, 7% of Greece and 2% of Norway, with the other control countries contributing very little. Actually, Canada's inflation trajectory has some qualitative similarities with inflations in Denmark (from 1971 up to the late 1980s) and Switzerland (in the early 1970s and the early 1990s).[10] With such countries' weights, the synthetic of Canada is able to replicate Canada's pre-IT characteristics closely in Table 4 ($X_0 W \approx X_1$), and the synthetic's inflation path approximates Canada's inflation path before IT adoption (until 1991) reasonably in Figure 1.2 ($\hat{Y}_{1,t}^N = Y_{0,t} W^* \approx Y_{1,t}^N$ for $t = 1, ..., T_0$). It is noticeable how both inflations simultaneously go down during the 1980s.

After 1991, Canada and its synthetic are still highly correlated, although Canada presents systematically lower inflation rates. Looking at Figures 2.1.a and 2.1.b, which provide zooms in of the IT period, we see that Canada has better performance than its synthetic both in terms of average inflation and average output growth. Table 5 confirms the visual impression by showing the average Canada-synthetic difference in inflation and output growth after IT adoption. On average, during the eight year of IT (from 1992 to 1999), Canada presents annual inflations 0.91 percentage points lower than its synthetic. Although small, compared to the worldwide inflation reduction observed during the 1980s and early 1990s, it is economically important, given that the average inflation among industrial economies in the late 1990s was close to 2% per year. From the positive inflation-output tradeoff perspective of the Phillips curve, the fact that Canada presents annual output growth rates 0.60 percentage points higher than its synthetic is also noticeable as superior performance.

Attempting to infer the economic significance of the above numbers, we recur to the permutation method and rank the Canada-synthetic difference among the twelve placebo-synthetic differences that have RMSPE smaller than 3 times Canada's RMSPE. Canada-synthetic's inflation difference is the third more negative (in Table 5.A), and Canada-synthetic's output growth difference is the fourth more positive (in Table 5.B or 5.C). Among these twelve placebos, Norway is the only country that performs better than Canada both in terms of lower inflation and higher growth for the period.[11]

Similar to Canada, Sweden presents superior performance during the IT regime, with average annual inflation 0.95 percentage points lower and average annual output growth 0.94 higher than its synthetic (in Figures 2.6.a, 2.6.b, and column 6 of Table 5.A and 5.B). Noting that Sweden's synthetic presented higher pre-IT output growth than Sweden, and looking at the difference-in-difference output growth that controls for this pre-IT dissimilarity, Sweden's average annual output growth improved in 1.12 percentage points relative to its synthetic. Sweden ranks second in the inflation-reduction ranking among its fifteen placebos with RMSPE smaller than 3 times its own (in Table 5.A) and fifth in the output-growth ranking (in Table 5.B or 5.C), with no country performing better both in terms of inflation and output growth. Italy, the placebo country that performed better than Sweden in terms of inflation, presented simultaneous disappointing output growth rates, ranking number 13 out of 15.

4.2.B. Australia

Although Australia also seems to present enhancing IT effects from the view point of column 1 in Table 5, we must be cautious because of its synthetic inability to reproduce the Australian pre-treatment characteristics, shown in Table 4 and in Figure 1.1 around 1995 (the year of Australia IT adoption). This means that we have not been able to find a convex combination of industrial non-ITers that satisfactorily match Australia's pre-treatment characteristics (X_1), casting doubts on the resulting synthetic ability to reproduce the counterfactual trajectory of post-treatment inflation.

With this warning, we note that Australia ranks second best in inflation reduction (in Table 5.A) and third in output growth (in Table 5.B), with no country performing better, both in terms of inflation and output growth. If we ponder that Australia's synthetic presented higher pre-IT output growth than Australia, and we additionally take into account the difference-in-difference output growth, Australia becomes second in the output growth ranking (in Table 5.C).

4.2.C. New Zealand and United Kingdom

New Zealand and the United Kingdom provide the stylized Phillips curve experience, lowering inflation at the cost of lower growth relative to their synthetics. In the case of New Zealand, although the visual inspection of Figure 1.4 might suggest a poor fit, because of the freakish 1984-1990 inflation volatility, New Zealand's synthetic matches its pre-treatment characteristics very well in Table 4. We realize that New

Zealand lowered average annual inflation in -0.72 percentage points (in Table 5.A) at the cost of -0.79 percentage points lower average annual growth (in Table 5.B), not showing any obvious IT gain (nor any obvious IT loss).

The United Kingdom's experience seems similar to the New Zealand's, although with a worse match of the pre-treatment average growth. The United Kingdom also lowered average annual inflation in -0.81 percentage points (in Table 5.A) at the cost of -0.76 percentage points lower average annual growth (in Table 5.B). In this case, however, it is noticeable the United Kingdom used to grow less before IT than its synthetic. In fact, when the UK-synthetic growth difference is adjusted for pre-IT dissimilarities, it becomes clear that the UK-synthetic average annual growth difference actually improved in 2.51 percentage points post-IT adoption (in Table 5.C), suggesting that the UK's IT disinflation did not incur in output costs but in a bonus, like Canada and Sweden.

4.2.D. Finland and Spain

Finland and Spain are special cases of the IT regime, because they both adopted it as a mean to adjust their inflation to the Maastricht Treaty requirements.

Spain's synthetic also matches its pre-treatment characteristics well in Table 4. From Table 5, we realize that in spite of any difference in inflation under IT (in Table 5.A), Spain presents an annual average output growth 1.44 percent point higher than its synthetic (in Table 5.B).

Finland's experiences seem similar to Spain's, although with a worse match of the pre-treatment average growth. From Table 5, Finland presents annual average output growth rates 2.66 percent points higher than its synthetic. If additionally, we adjust to pre-IT differences, the relative average annual output growth improvement widens to

4.44 percentage points, putting Finland in the first position among its placebos in terms of growth.

4.2.E. Panel

We could summarize the individual ITers cases just presented as providing evidence that "pure" ITer (Australia, Canada, New Zealand, Sweden and United Kingdom) decreased inflation with relatively low output growth costs (not to say with output growth gains in 4 out of 5 countries), while the "transitory" ITers (Finland and Spain) grew more without bearing inflation increases. Until 1999, the ITers' annual inflation rates had been lower than their respective synthetic estimates in 36 out of 42 observations (the sum of the number of IT years among ITers). For same ITer-synthetic pairs, ITers' annual output growth rates were higher than their respective synthetic estimates in 27 out of the 42 observations, in spite of pre-IT higher average growth of some synthetics.

But what about the average effect of IT on ITers, i.e. the average effect of treatment on treated (ATT)? How this study results compare with previous analysis like Ball and Sheridan (2005), Lin and Ye (2009) or Brito (2011)?

To answer this question, we pool the treated-synthetic pairs and estimate the average effect of lagged IT in annual AR(1) models for inflation and output growth. In Table 6, we show the results for all seven ITers (in the Seven ITers columns) and for the ITers' group except Australia for suspicion of poor fit (in the Six ITers columns). Both for inflation and output, the IT effects are significant, no matter the sample or the estimation method, confirming Brito (2011). In column 2, for example, we show that the direct impact of IT on the average annual inflation of ITers was a significant 0.56 percentage points reduction, and the cumulative effect of IT on the annual inflation rate was a significant average reduction among ITers of 0.80 percentage points. Relative to output growth, without accounting for pre-IT growth differences in column 7, the direct impact of IT on the average annual growth of ITers was a significant 0.88 percentage points increase, and the cumulative effect of IT on the annual growth rate was an average increase among ITers of 1.32 percentage points. If accounted for pre-IT growth differences in this framework, including ITers fixed-effects in column 8, the direct impact of IT on the average annual growth of ITers was a significant 1.11 percentage points increase, and the cumulative effect of IT on the annual growth rate was an average increase among ITers of 1.42 percentage points.

Just to confirm Friedman's (1961) point about the sluggish effects of monetary policy, in column 1 of Table 6, we present the AR(1) panel without lagging the IT dummy variable – with $D_{n,t}^{IT}$ instead of $D_{n,t-1}^{IT}$ – and, in column 3, we use one IT dummy for each year of the IT program – instead of a single dummy for every year under the program. As expected, the contemporaneous IT dummy in column (1) has less impact than its lag in column (2), because the IT effects on inflation only become significant one year after IT adoption, in column (3).

5. Conclusions

Although industrial central banks that have primarily adopted IT seem happy with their choice, to the point of having influenced the IMF recommendations and many other countries in a subsequent wave of IT introductions after the mid 1990s, the existing evidence in terms of realized inflation and output growth does not seem to support their optimism.[12]

In this paper we point that the small sample available to study the IT regime adoption in industrial economies demands a case-by-case counterfactual, better represented by combinations of non-ITers countries able to replicate each ITer inflation trajectory over the pre-treatment period. We notice that the worldwide inflation reduction of the 1980s and early 1990s did not leave much room for similarly large-in-magnitude inflation reductions among industrial countries. Although adopted in a period when monetary management performance was already standing out, we have been able to show that IT gains were economically important and relatively well spread among the pioneer IT economies.

The results suggest that inflation targeting contributed to a lower inflation rate in Canada, Sweden, New Zealand, the United Kingdom, and Australia with reservations, especially after the first year of regime adoption. In Spain and Finland, we did not identify significant inflation reduction, but higher output growth. On average, the IT effects on ITers were significantly negative for inflation and significantly positive for output growth, at least during the inception years of the regime.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Synthetic control methods for comparative case studies: estimating the effect of California's tobacco control program," *Journal of the American Statistical Association* 105:490 (2010), 493-505.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Comparative politics and the synthetic control method," MIT working paper (2011).

Abadie, Alberto, and Javier Gardeazabal, "The economic costs of conflict: a case study of the Basque Country," *The American Economic Review* 93:1 (2003), 113-132.

Ang, Andrew, Geert Bekaert, and Min Wei, "Do macro variables, asset markets, or surveys forecast inflation better?" *Journal of Monetary Economics* 54:4 (2007), 1163-1212.

Ball, Laurence, and Niamh Sheridan, "Does inflation targeting matter?" (pp. 249-276), in Ben Bernanke and Michael Woodford (Eds.), *The inflation targeting debate* (Chicago, IL: University of Chicago Press, 2005).

Bank of Canada, *Renewal of the Inflation-Control Target: Background Information* (Ottawa: Bank of Canada, 2006).

Batini, Nicoletta, and Edward Nelson, "The lag from monetary policy actions to inflation: Friedman revisited," *International Finance* 4:3 (2001), 381-400.

Bernanke, Ben, Thomas Laubach, Frederic Mishkin, and Adam Posen, *Inflation targeting: lessons from the international experience* (Princeton, NJ: Princeton University Press, 1999).

Brito, Ricardo, "Inflation targeting did make a difference in industrial countries' inflation and output growth," Insper working paper (2011).

Cecchetti, Stephan, Rita Chu, and Charles Steindel, "The unreliability of inflation indicators," *Federal Reserve Bank of New York Current Issues in Economics and Finance* 6:4 (2000), 1-6.

Corbo, Vittorio, Oscar Landerretche, and Klaus Schmidt-Hebbel "Does Inflation Targeting Make a Difference?" (pp. 221-69), in Norman Loayza and Raimundo Soto (Eds.), *Inflation Targeting: Design, Performance, Challenges* (Santiago: Central Bank of Chile, 2002).

Ernst, Michael, "Permutation Methods: a basis for exact inference," *Statistical Science* 19:4 (2004), 676-685.

Friedman, Milton, "The lag in effect of monetary policy," *Journal of Political Economy* 69:5 (1961), 447-466.

Ghosh, Atish, Anne-Marie Gulde, and Holger Wolf, *Exchange rate regimes: choices and consequences* (Cambridge, MA: MIT Press, 2002).

International Monetary Fund (IMF), *Inflation targeting and the IMF* (Washington, DC: IMF, 2006).

Johnson, David, "The effect of inflation targeting on the behavior of expected inflation: evidence from an 11 country panel," *Journal of Monetary Economics* 49:8 (2002), 1521-1538.

King, Mervyn, *The inflation targeting ten years on*, Speech delivered to the London School of Economics. 19 November 2002, London, England.

Lin, Shu, and Haichun Ye, "Does inflation targeting really make a difference? Evaluating the treatment effect of inflation targeting in seven industrial countries," *Journal of Monetary Economics* 54:8 (2007), 2521-2533.

Mishkin, Frederic, and Klaus Schmidt-Hebbel, "Does inflation targeting make a difference?" NBER working paper no. 12876 (2007).

Roger, Scott, and Mark Stone, “On target? The international experience with achieving inflation targeting,” IMF working paper no. 05-163 (2005).

Walsh, Carl, “Inflation targeting: What have we learned?” *International Finance* 12:2 (2009), 195-233.

Tables and Figures

Table 1 - Industrial countries' classification in 1999 and dates of adoption

1.A - Inflation targeting countries	
Australia	1995
Canada	1991
Finland	1993
New Zealand	1990
Spain	1995
Sweden	1995
United Kingdom	1992

1.B - Non-inflation targeting countries	
Austria	Japan
Belgium	Netherlands
Denmark	Norway
France	Portugal
Germany	Singapore
Greece	South Korea
Iceland	Switzerland
Ireland	United States
Italy	

Note: Dates of "converging-IT" adoption from Mishkin and Schimdt-Hebbel (2007).

Table 2 - Inflation predictors weights that best reproduce pre-Inflation Targeting inflation trajectory

	Australia	Canada	Finland	New Zealand	Spain	Sweden	United Kingdom
Inflation(yr. of adopt.)	0.1008	1.9E-14	0.2090	0.0704	0.0026	0.8739	0.3404
Inflation(-1,-4)	0.7558	0.0786	1.9E-05	0.3685	0.0003	6.1E-09	0.5113
Output growth(-1,-5)	0.0342	0.3586	3.0E-07	0.1337	0.0603	4.6E-06	1.2E-08
Inflation(-5,-9)	0.1091	0.5628	0.7910	0.4274	0.9368	0.1261	0.1483

Notes: Each column shows the main diagonal of *V* in equation (1) for the respective ITer. *Inflation(yr. of adopt)* means annual inflation in the year of IT adoption, according to Table 1. *Variable(-a,-b)* means the variable average from *(yr. of adopt. - a)* to *(yr. of adopt.- b)* .

Table 3 - Non-inflation targeting country weights (in rows) in the synthetics of inflation targeting countries (in columns)

	Australia	Canada	Finland	New Zealand	Spain	Sweden	United Kingdom
Austria	0	0.002	0	0.031	0.020	0.003	0.015
Belgium	0	0.001	0	0.029	0.019	0.003	0.018
Denmark	0.192	0.562	0.279	0.051	0.017	0.005	0.025
France	0	0.001	0	0.034	0.015	0.005	0.018
Germany	0	0.003	0	0.031	0.015	0.002	0.011
Greece	0	0.071	0	0.030	0.067	0.099	0
Iceland	0	0	0.106	0.139	0.038	0.121	0.030
Ireland	0	0.001	0	0.025	0.023	0.003	0.014
Italy	0.144	0.001	0	0.033	0.656	0.007	0.014
Japan	0	0.001	0.288	0.022	0.015	0.002	0.013
Netherlands	0	0.002	0	0.030	0.023	0.003	0.015
Norway	0.664	0.023	0	0.427	0.019	0.004	0.026
Portugal	0	0	0	0.024	0.016	0.005	0.008
Singapore	0	0.001	0	0.020	0.016	0.002	0.011
South Korea	0	0	0	0.013	0.019	0.003	0.507
Switzerland	0	0.323	0.327	0.031	0	0.729	0.022
United States	0	0.007	0	0.030	0.023	0.006	0.254

Table 4 - Inflation predictors' characteristics before Inflation Targeting adoption

	Australia	Canada	Finland	New Zealand	Spain	Sweden	United Kingdom
			4.A. Inflation Targeter Country				
Inflation(yr. of adopt.)	4.64	5.60	2.20	6.22	4.64	2.53	4.72
Inflation(-1,-4)	1.98	4.55	4.64	7.64	5.29	4.61	6.36
Output growth(-1,-5)	0.98	1.50	-0.26	2.48	1.55	-0.49	1.85
Inflation(-5,-9)	7.53	5.80	4.99	14.24	6.49	6.23	4.50
			4.B. Synthetic Control Country				
Inflation(yr. of adopt.)	2.78	4.78	2.21	6.22	4.64	2.54	4.72
Inflation(-1,-4)	2.67	4.53	4.31	7.64	5.30	4.75	6.36
Output growth(-1,-5)	2.36	1.49	1.52	2.48	1.55	-0.31	5.12
Inflation(-5,-9)	5.73	5.77	5.00	14.25	6.50	6.24	4.51

Notes: Inflation(yr. of adopt) means annual inflation in the year of IT adoption, according to Table 1. *Variable(-a,-b)* means the variable average from *(yr. of adopt. - a)* to *(yr. of adopt.- b)* .

Table 5 - Average annual difference between the Inflation Targeting country and its synthetic (or the placebo and its synthetic)

	Australia	Canada	Finland	New Zealand	Spain	Sweden	United Kingdom
	5.A. Country-Synthetic Difference in Inflation						
Treated	-0.86	-0.91	-0.08	-0.72	-0.05	-0.95	-0.81
Placebos average	0.07	-0.20	-0.23	0.44	0.03	0.08	0.36
Treated ranking	2	3	8	4	7	2	4
	5.B. Country-Synthetic Difference in Per-capita GDP growth						
Treated	0.82	0.61	2.66	-0.79	1.44	0.94	-0.76
Placebos average	0.18	-0.29	-0.42	-0.34	0.18	0.36	0.10
Treated ranking	3	4	2	9	3	5	10
	5.C. Country-Synthetic Difference-in-Difference in Per-capita GDP growth						
Treated	2.20	0.60	4.44	-0.79	1.45	1.12	2.51
Placebos average	0.12	-0.25	-0.18	-0.48	0.26	0.81	0.61
Treated ranking	2	4	1	9	3	5	5
No. of countries	15	12	14	17	15	15	17

Notes: Treated is the annual average difference between the ITer (naming the column) and its synthetic during the IT years. *Placebos average* is the annual average difference between the placebos and their respective synthetic during the IT years of the ITer. *Treated ranking* is increasing for inflation and decreasing for per-capita GDP growth. *No. of countries* is the number of countries included in the rank, that presented RMSPE smaller than 3 times the ITer's RMSPE.

Table 6 - Panel estimates of the average inflation targeting effects on inflation and output growth (1971-1999)

Equation:	Inflation equation						Output growth equation		
Sample:	Seven ITers				Six ITers		Seven ITers		
Estimator:	OLS			FE-LS	OLS		OLS	FE-LS	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Inflation targeting	-0.36***								
	(0.09)								
L.Inflation targeting		-0.56***		-0.56***	-0.51**		0.88***	1.11***	
		(0.13)		(0.13)	(0.14)		(0.21)	(0.24)	
1st year of inflation targeting			0.67			0.47			-0.22
			(0.39)			(0.38)			(0.67)
2nd year of inflation targeting			-0.96**			-1.12**			0.12
			(0.35)			(0.37)			(0.73)
3rd year of inflation targeting			-0.93**			-0.69**			1.03**
			(0.33)			(0.27)			(0.32)
4th year of inflation targeting			-0.66*			-0.65*			1.74*
			(0.28)			(0.32)			(0.74)
5th year of inflation targeting			-0.18			-0.15			1.60**
			(0.13)			(0.14)			(0.46)
6th-10th year of inflation targeting			-0.21*			-0.20			1.10*
			(0.10)			(0.11)			(0.50)
L.Inflation	0.30***	0.30***	0.30***	0.29***	0.28***	0.28***			
	(0.07)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)			
L.GDP per-capita growth							0.33***	0.22**	0.20*
							(0.07)	(0.08)	(0.09)
Constant	0.04	0.07	0.04	0.07**	0.05	0.02	-0.41	-0.52***	-0.52***
	(0.07)	(0.07)	(0.08)	(0.03)	(0.09)	(0.09)	(0.27)	(0.07)	(0.08)
ICE of IT	-0.51***	-0.80***		-0.79***	-0.71***		1.32**	1.42***	
Observations	196	196	196	196	168	168	196	196	196
Adj. R2	0.09	0.09	0.07	0.09	0.08	0.05	0.13	0.09	0.08

Notes : Seven Iters includes all inflation targeters in Table 1, Six Iters excludes Australia. Pooled cross-section OLS (OLS) in columns (1)-(3) and (5)-(7), including country-fixed-effects (FE-LS) in columns (4) and (8)-(9), estimated with constant (omitted in the table). The data frequency is annual and $L.Z=Z_{t-1}$ means 1-year lag in relation to the dependent variable. Robust standard errors clustered by country in parentheses. *, **, *** indicate the significance level of 10%, 5%, and 1% respectively. ICE means Implied Cumulative Effect equal to $IT\ dummy_{t-1(or\ t)}/(1\text{-}Dependent\ Variable_{t-1})$, with significance tested by a nonlinear test.

Figure 1.1. Australia and Australia's synthetic inflation

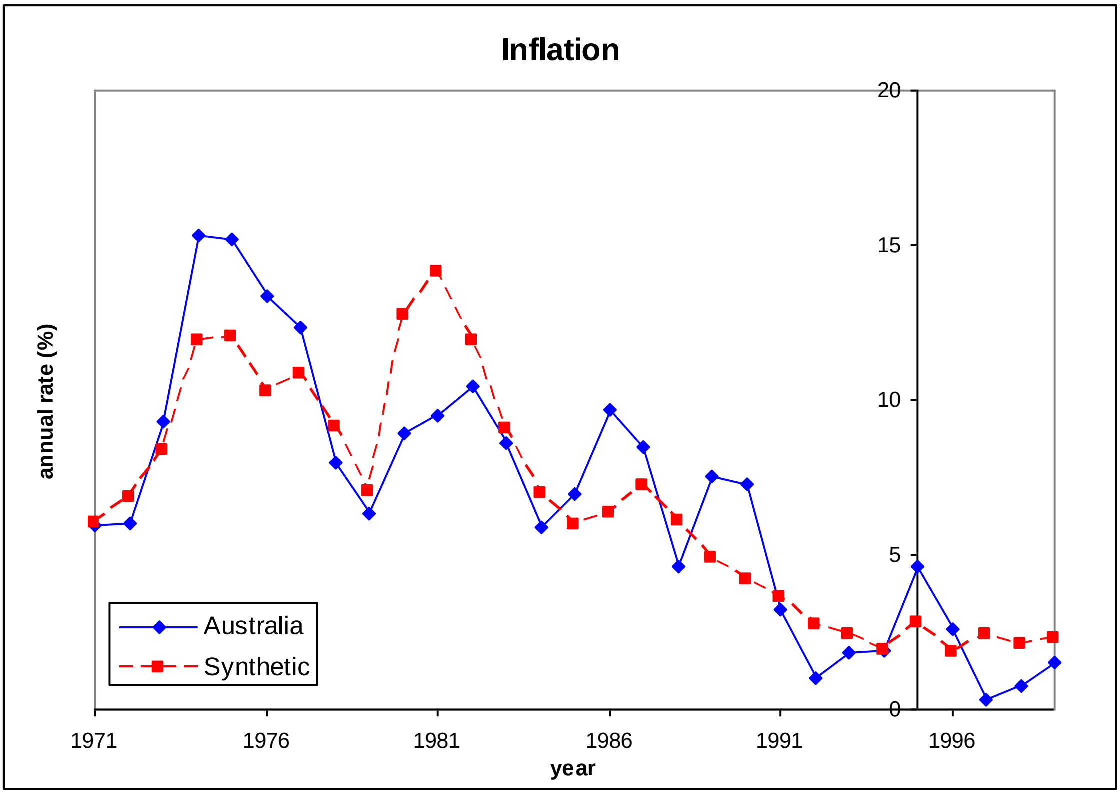

Figure 1.2. Canada and Canada's synthetic inflation

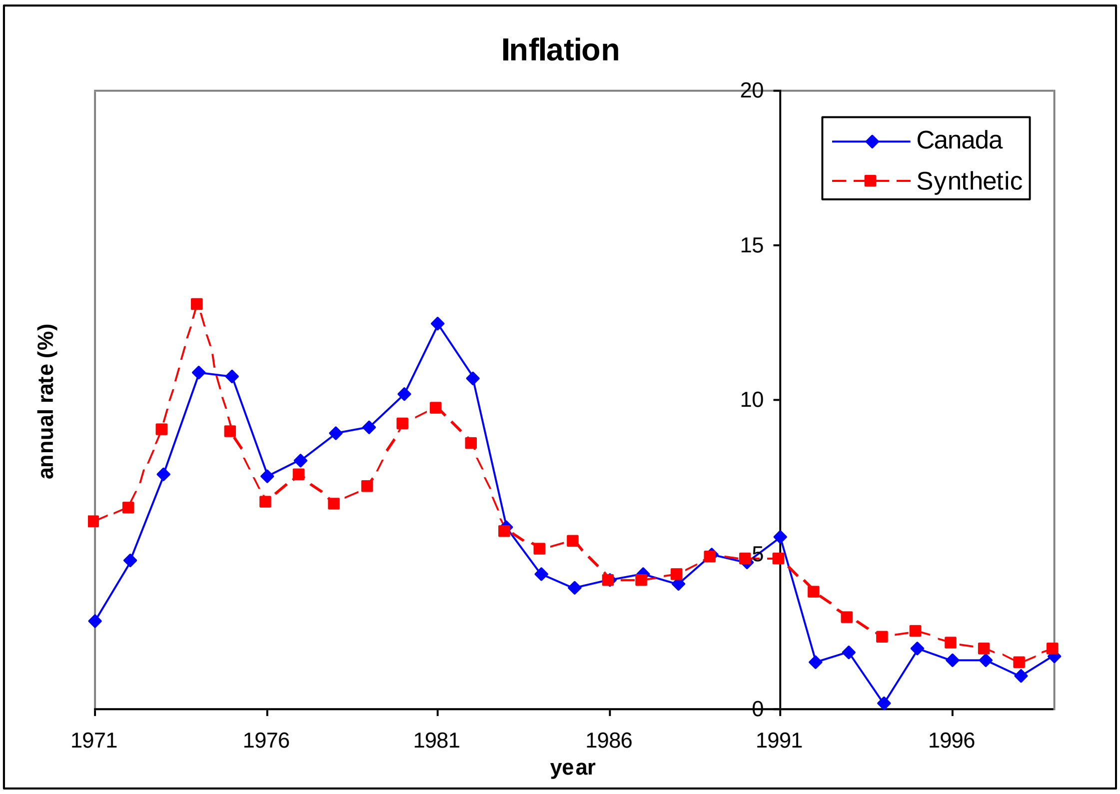

Figure 1.3. Finland and Finland's synthetic inflation

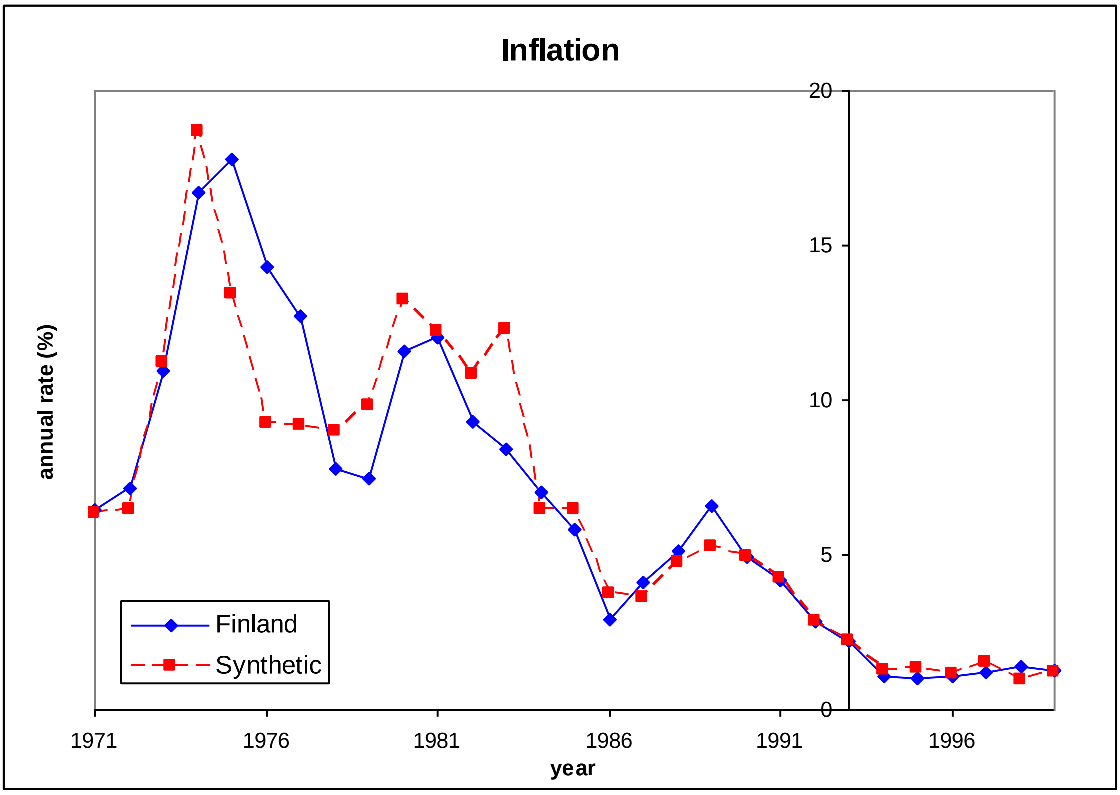

Figure 1.4. New Zealand and New Zealand's synthetic inflation

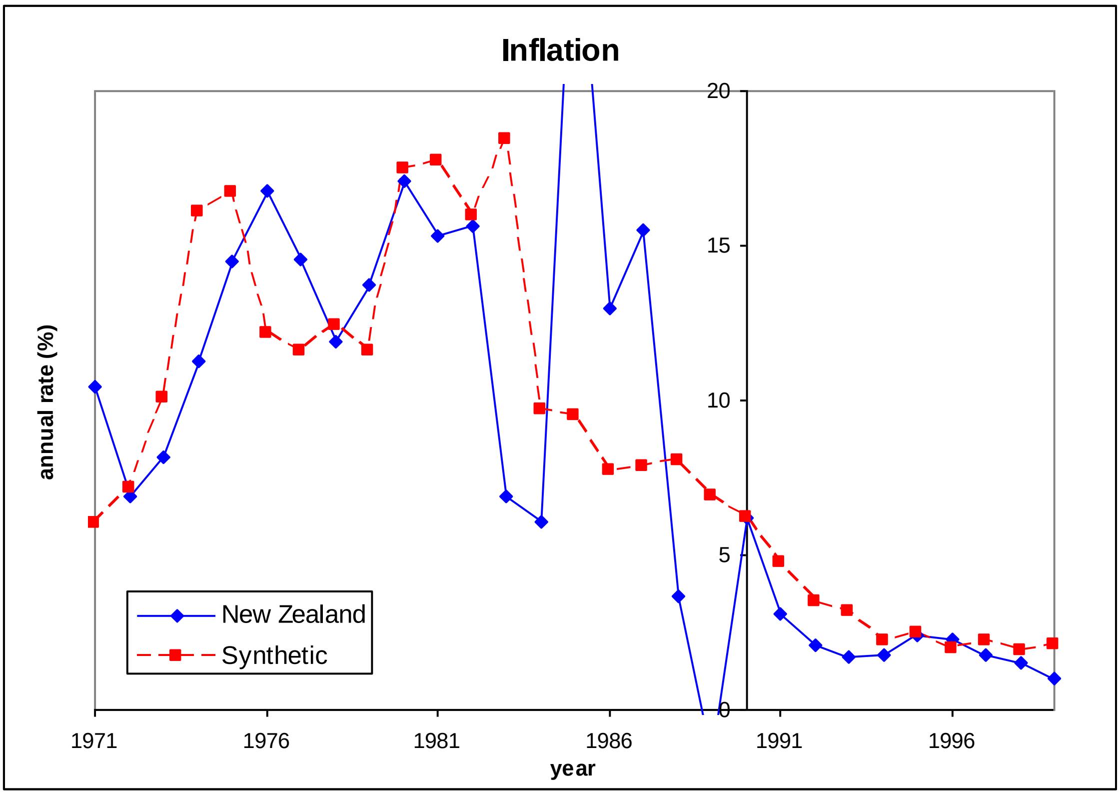

Figure 1.5. Spain and Spain's synthetic inflation

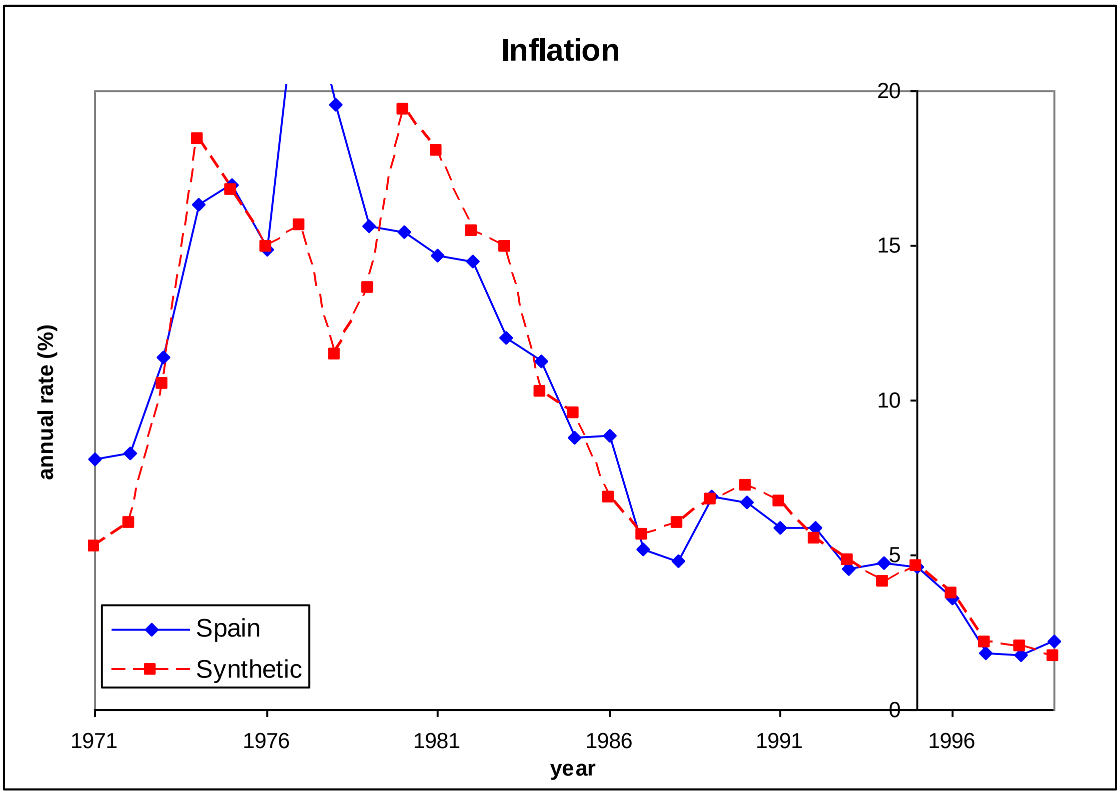

Figure 1.6. Sweden and Sweden's synthetic inflation

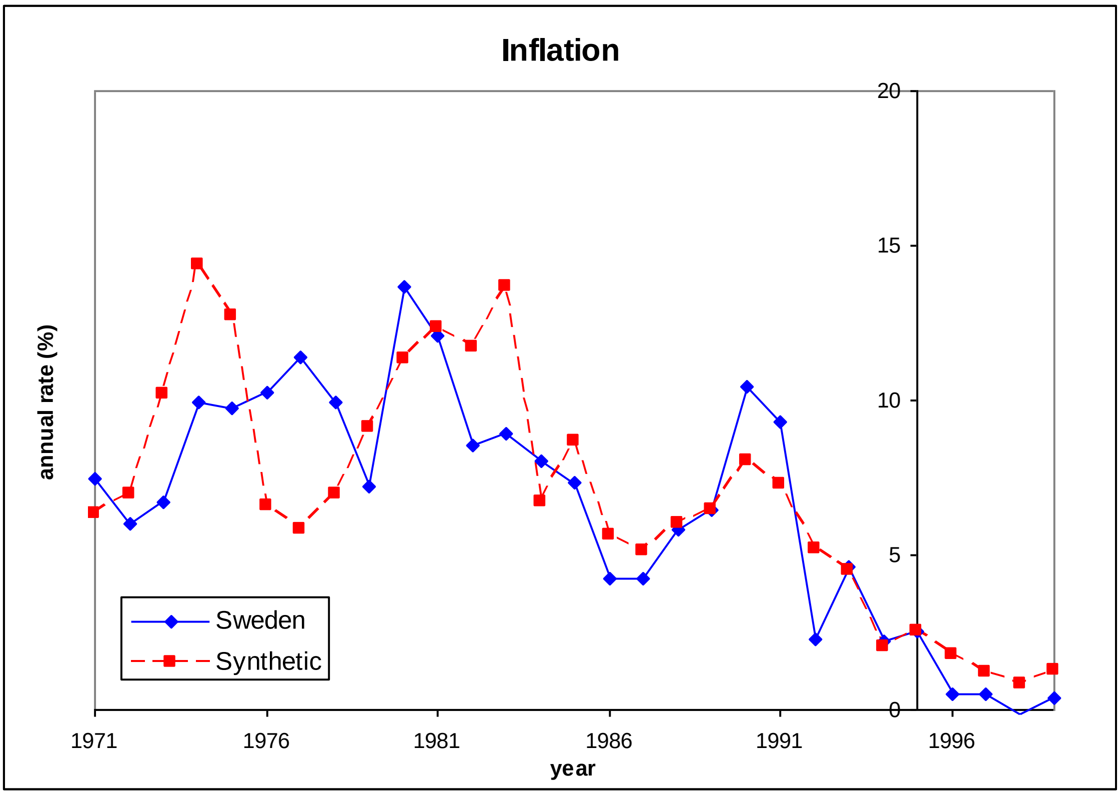

Figure 1.7. United Kingdom and United Kingdom's synthetic inflation

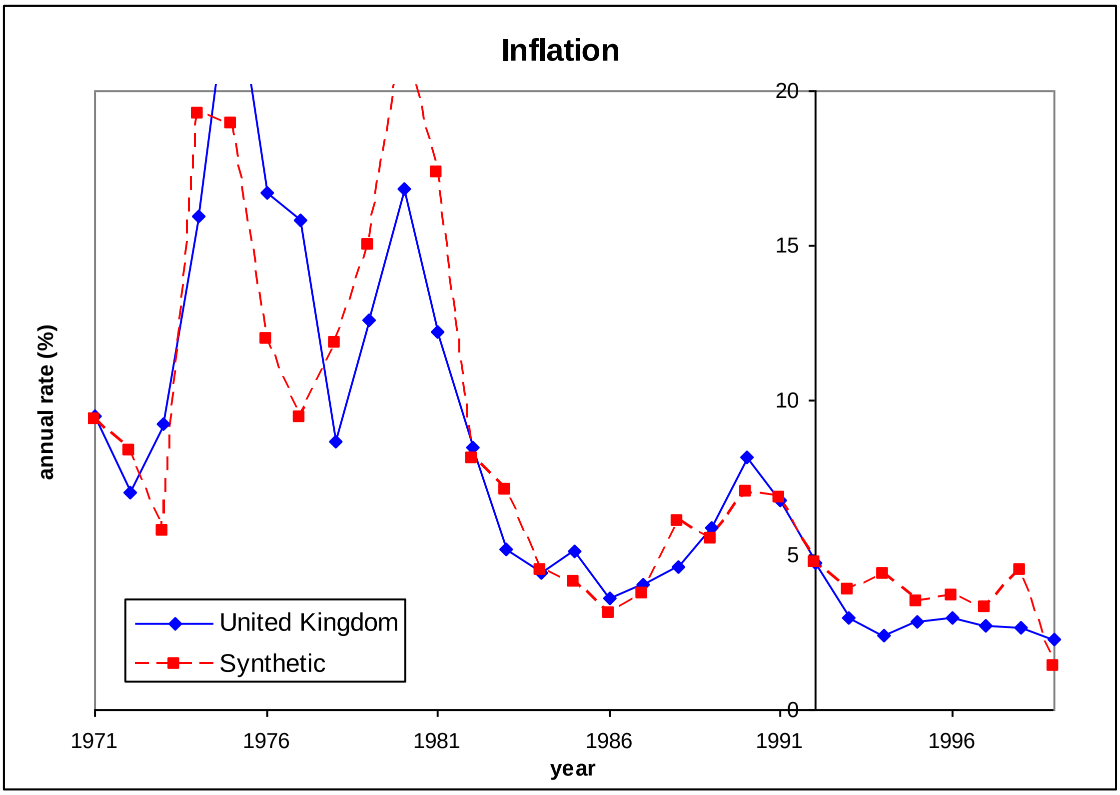

Figure 2. Zoom in of inflation and output growth rate trajectories around Inflation Targeting adoptions

Figure 2. *(cont.)* Zoom in of inflation and output growth rate trajectories around Inflation Targeting adoptions

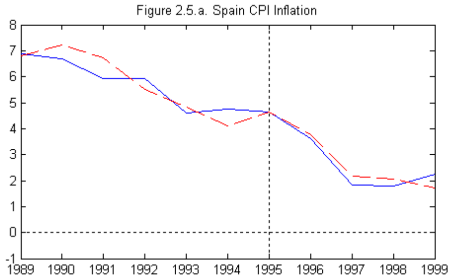

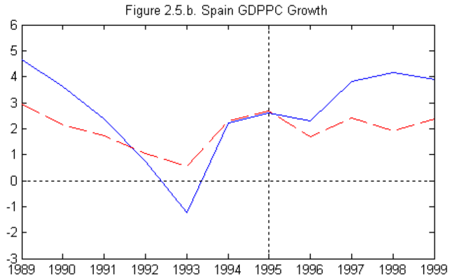

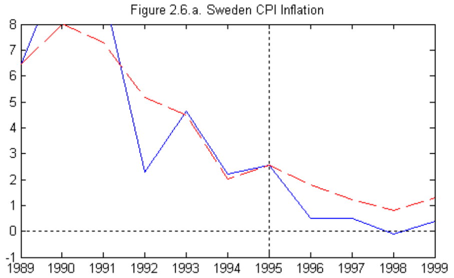

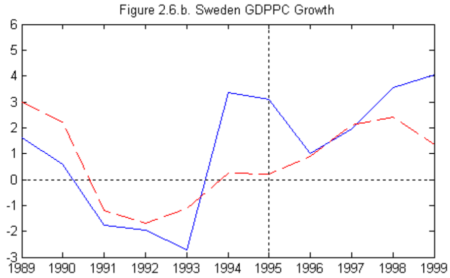

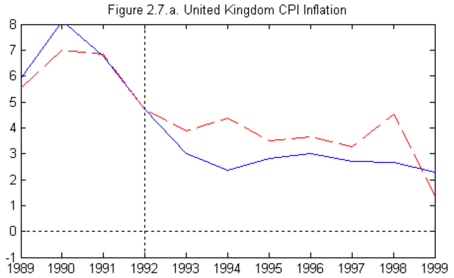

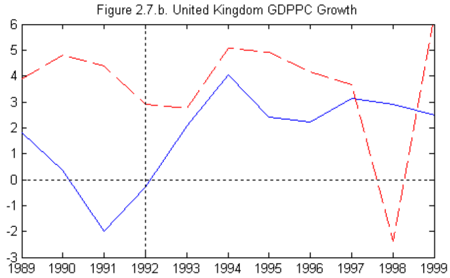

Appendix *(provided for the referee's benefit)*:

Figure A. Canada, Switzerland and Denmark inflations

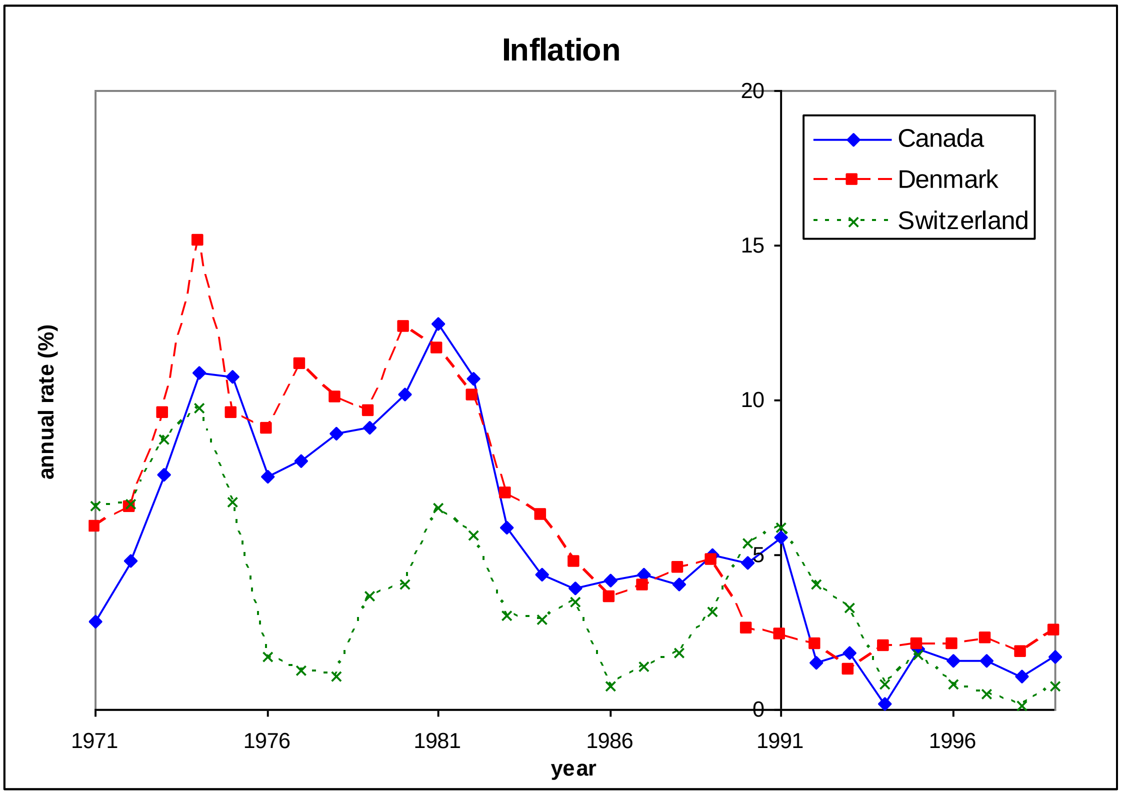

Table A - Average annual difference between the Inflation Targeting country and its synthetic (or the placebo and its synthetic)

	Australia	Canada	Finland	New Zealand	Spain	Sweden	United Kingdom
				A.A. Country-Synthetic Difference in Inflation			
Treated	-0.86	-0.91	-0.08	-0.72	-0.05	-0.95	-0.81
Placebos average	0.07	-0.20	-0.23	0.44	0.03	0.08	0.36
Standard error	(0.92)	(1.08)	(1.25)	(2.37)	(0.88)	(1.44)	(1.65)
t -statistic	-1.50	-1.30	0.21	-1.02	-0.14	-1.06	-1.32
p -value	0.07	0.10	0.58	0.15	0.45	0.14	0.09
				A.B. Country-Synthetic Difference in Per-capita GDP growth			
Treated	0.82	0.61	2.66	-0.79	1.44	0.94	-0.76
Placebo averages	0.18	-0.29	-0.42	-0.34	0.18	0.36	0.10
Standard error	(2.32)	(1.61)	(3.27)	(4.88)	(2.78)	(3.46)	(3.35)
t -statistic	0.42	1.10	1.65	-0.19	0.68	0.25	-0.48
p -value	0.34	0.14	0.05	0.58	0.25	0.40	0.68
				A.C. Country-Synthetic Difference-in-Difference in Per-capita GDP growth			
Treated	2.20	0.60	4.44	-0.79	1.45	1.12	2.51
Placebo averages	0.12	-0.25	-0.18	-0.48	0.26	0.81	0.61
Standard error	(2.32)	(1.41)	(3.22)	(2.49)	(2.78)	(3.46)	(2.87)
t -statistic	1.34	1.19	2.50	-0.26	0.63	0.13	1.23
p -value	0.09	0.12	0.01	0.60	0.26	0.45	0.11
No. of countries	15	12	14	17	15	15	17

Notes: Treated is the annual average difference between the ITer (naming the column) and its synthetic during the IT years. *Placebos average* is the annual average difference between the placebos and their respective synthetic during the IT years of the ITer. *t-statistic* adjusts the *standard error* for the random variable average, dividing it by the square root of the number of IT years and by a 0.4 autocorrelation along years. *p-value* is the one-sided probability with H_1 that the inflation difference is negative and the per-capita GDP growth difference is positve. *No. of countries* is the number of countries included in the rank, that presented RMSPE smaller than 3 times the ITer's RMSPE.

[1] Below, we explain that the synthetic control weights are defined to match the ITer inflation trajectory pre-IT adoption. For a sensible cost-benefit analysis, however, besides the post-IT inflation analysis, we also have to wonder how the contemporaneous output growth rates of these same ITer and synthetic country were compromised. Because the weights that approximate the ITer pre-IT inflation rates might not approximate pre-IT output growth rates as well, the latter comparison is less straightforward.

[2] The direct impact of lagged IT on inflation (or output growth) is given by β in equation (6) (or (7)); and its cumulative effect is given by $\beta/(1-\rho)$, which accounts for subsequent decreasing over-time lagged dependent variable on the right-hand side.

[3] King (2002) and Bank of Canada (2006) are two good examples of Central Bank(er)s optimism with IT. IMF (2006) is a document that pleads IT adoption.

[4] See Abadie, Diamond and Hainmueller (2010) for a formal proof that the synthetic method is able to control for effects of confounding unobserved and observed characteristics that vary with time.

[5] For multiple treated countries, the method can be applied to each treated separately, or to some aggregate of all treated.

[6] Note that the restriction of W^* to the subset of convex combinations, instead of simply linear combinations, makes $\hat{X}_1 = X_1$ infeasible when the characteristics of the case of interest do not belong to the convex hull of the potential controls' characteristics. Abadie, Diamond and Hainmueller (2011) show that the linear regression weights can always extrapolate to produce a perfect fit, even if X_1 is far from the convex hull of the columns of X_0.

[7] See Ernst (2004) for permutation methods, tests and exact inference.

[8] Iceland, Norway, Switzerland and South Korea, which adopted the IT strategy after 1999, are considered members of the control group.

[9] We have experimented with some of these variables, but got similar qualitative results to the ones being presented.

[10] See Figure A, in Appendix A, for a comparison of those three countries inflations.

[11] Instead of applying the permutation method idea of ranking the results, we could have used the traditional t-statistic test, in spite of the small number of elements. Those results are presented in Table A,

in the Appendix. In Canada's case, the *p*-values are 7% for the inflation difference and 8% for the output growth.

[12] See Walsh (2009) for a survey of the IT literature. He concludes that "… macroeconomic experiences among both inflation targeting and non-targeting developed economies have been similar …"