How Flexible Can Inflation Targeting Be and Still Work?

Kenneth N. Kuttner and Adam S. Posen*

September 30, 2011

Abstract

This paper takes up the issue of the flexibility of inflation targeting regimes, with the specific goal of determining whether the monetary policy of the Bank of England, which has a formal inflation target, has been any less flexible than that of the Federal Reserve, which does not have such a target. The empirical analysis uses the speed of inflation forecast convergence, estimated from professional forecasters' predictions at successive forecast horizons, to gauge the perceived flexibility of the central bank's response to macroeconomic shocks. Based on this criterion, there is no evidence to suggest that the Bank of England's inflation target has compelled it to be more aggressive in pursuit of low inflation than the Federal Reserve.

JEL codes: E42, E58, E65.

Keywords: Inflation targeting, inflation expectations, monetary policy.

*Williams College Economics Department and the NBER, kenneth.n.kuttner@williams.edu; and Peterson Institute for International Economics and Monetary Policy Committee, Bank of England, aposen@piie.com. Prepared for *Current Topics in Monetary and Fiscal Policy International Conference on the U.S. Economy, A Conference to Honor Ben Friedman*, Federal Reserve Bank of Boston, April 22-23, 2011. The paper has benefitted from the comments of the discussant, Gauti Eggertsson, as well as those of the conference participants. Marilyne Tolle provided expert research assistance. The views expressed herein and any errors are solely those of the authors, and not of the Bank of England, the Monetary Policy Committee, or the PIIE.

1 *Introduction*

After years of apparent success, inflation targeting (IT) has recently come under fire for its inability to prevent (or even its contribution to) the global financial crisis of 2008–09. The recent debate has largely centered on whether inflation targeting prevented central banks from responding sufficiently to asset price bubbles.[1]

But IT has always had its skeptics, even before the latest wave of scrutiny. One longstanding question is whether the policy framework could actually matter for outcomes as much as, and in the way that, its advocates said it should. IT proponents, such as Bernanke *et al.* (1999) and Svensson (1997, 1999a), portrayed it as a flexible monetary framework that would allow central banks to pursue countercyclical monetary policy while maintaining medium- and long-run price stability. (The phrase "disciplined discretion" has sometimes been used to convey this idea.)

Critics found this claim too good to be true. Neumann & von Hagen (2002), Friedman (2003), and Ball & Sheridan (2004) expressed doubts that merely declaring a numerical inflation target, while reserving the right to respond to shocks, would offer any tangible advantage over conventional discretionary frameworks. The second question is whether IT's touted advantages would come at the cost of greater output instability, due to a stricter adherence to a rigid policy rule. This concern has often been expressed by Federal Reserve policymakers, such as Kohn (2003) and Meyer (2004), and by some academics, notably Friedman & Kuttner (1996) and Friedman (2004). The essential issue, in other words, was that of credibility versus flexibility: how IT could be constraining enough to have an effect, and yet flexible enough to allow for meaningful output stabilization.

The challenge facing IT proponents was to resolve this paradox — to explain how central banks could have their cake and eat it too. King (1997) and Kuttner & Posen (1999) argued informally that flexible inflation targeting allowed central banks to anchor expectations while approximating the optimal state contingent rule. Theoretical rationales for this view include Walsh's (1995, 2003) contracting interpretation of IT, and models, such as that of Faust & Svensson (2001), in which the transparency associated with IT mitigated the problem of private information. Empirical studies,

[1] Contributions to this debate include Walsh (2009), Nelson (2009), Giavazzi & Giovannini (2010), Clarida (2010), Dale *et al.* (2010), McCallum (2010), Orphanides (2010), Svensson (2010b), Mishkin (2011), and Issing (2011).

including Siklos (1999) and Kuttner & Posen (2001), documented a reduction in both the level and persistence of inflation in countries that adopted IT. However these studies found little difference between implicit (e.g., U.S.) and explicit (e.g., U.K.) inflation targeters. Consequently, much of the extant evidence fails to address the argument put forth in Friedman (2003) and Ball & Sheridan (2004) that most central banks, not just inflation targeters, enjoyed better outcomes in recent years. One reason for this is that the widespread adoption of IT happened to coincide with a period of historically low inflation and output volatility (the "great moderation") and mild macroeconomic shocks. As Kuttner (2004) noted, IT had not yet been tested by a situation involving meaningful tradeoffs between low inflation and stable output.

This paper is an effort to assess IT's flexibility, and to evaluate the critique that the policy is excessively constraining during periods of large economic shocks. Section 2 reviews the theoretical literature bearing on the flexibility-credibility issue as it applies to inflation targeting. Of particular interest are the conditions under which a relaxation or outright suspension of the inflation target might be warranted, and the question of whether such an action would undermine confidence in the target. Section 3 proposes a novel gauge of flexibility: the speed with which medium-term inflation expectations, measured by professional forecasts, converge to the target following a disturbance. The analysis entails a comparison between the U.K., which has an explicit inflation target, and the U.S., which does not. The past few years' experience are especially relevant on this dimension, as many countries, including the U.S. and the U.K., have experienced much larger inflation and output fluctuations than over the preceding two decades.

The main empirical findings are twofold. The first is that in both countries, long-run inflation expectations, measured by the spreads between nominal and indexed bonds, are largely unresponsive to economic news, suggesting that inflation expectations have remained anchored in spite of the recent inflation volatility. The second and more novel finding is that there is no discernible difference across countries in the perceived speed of inflation adjustment. This result holds for positive as well as negative deviations from the inflation target, and for large as well as small deviations. If anything, the Bank of England may be slightly *more* flexible (in the sense of allowing more gradual adjustment) than the Federal Reserve. Taken together, the results indicate that whatever inflation stability IT has been able to confer on U.K. monetary policy, it has not come at the

expense of diminished flexibility.

2 *Theoretical perspectives on the flexibility-credibility tradeoff*

Having grown out of practical imperatives rather than a fully-articulated model, mapping inflation targeting into a consistent theoretical framework presents a challenge. Theoretical analyses of IT typically draw on the well-known Barro & Gordon (1983) model and its descendants. While at best a crude characterization of the policymaking environment, it is nonetheless useful for illustrating the tension between flexibility and credibility.

The model is so familiar that the briefest of summaries will suffice. The conventional if unrealistic assumption is that the central bank directly controls inflation, and output is then determined by a Lucas aggregate supply (AS) curve,

$$y = y^* + \beta(\pi - \pi^e) + \varepsilon \ , \tag{1}$$

where y^* is equilibrium output and ε is a supply (output) shock. The second key assumption is that the central bank minimizes an objective function of the form

$$L_{cb} = (\pi - \pi^*)^2 + \lambda(y-k)^2 \ . \tag{2}$$

where k represents the central bank's desired output gap, and π^* is the inflation target. Minimizing this loss function taking π^e as given (i.e., under discretion) yields the following targeting rule (TR),

$$\pi = \pi^* + \lambda\beta(k-y) \ . \tag{3}$$

Two implications follow. One is the classic inflation bias result: if $k > y^*$, then in equilibrium inflation will exceed the central bank's target, while output will remain at y^*. The second is that the optimal central bank response will be partially to offset adverse output shocks by choosing a higher inflation rate.

The model's solution is readily depicted in (y, π) space as the intersection of the AS and TR curves, equations 1 and 3. This is shown in figure 1a, assuming for simplicity that $\pi^* = y^* = 0$.

Figure 1: Graphical representation of the Barro-Gordon and Rogoff models

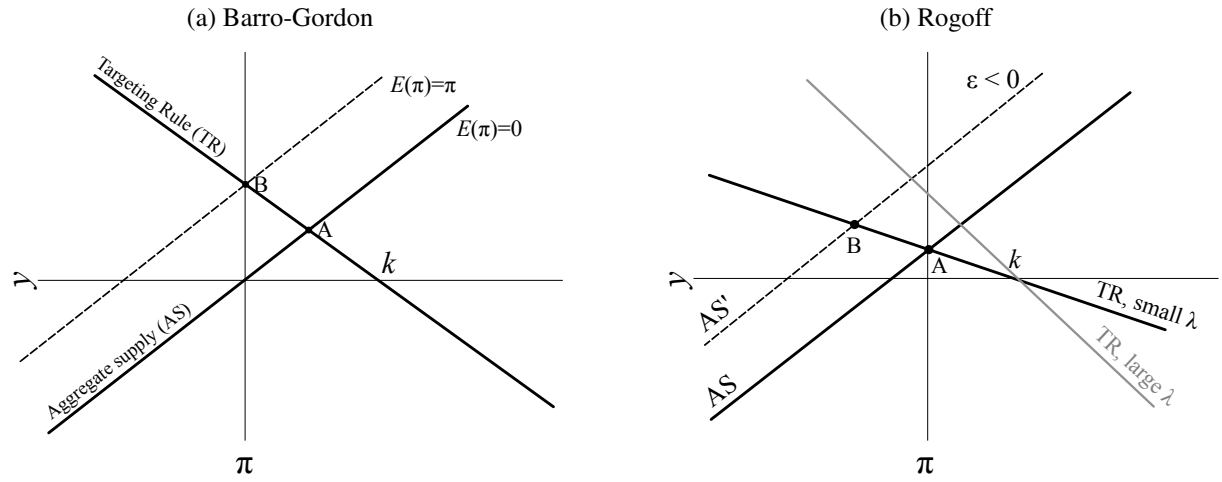

If expected inflation were zero, the equilibrium would occur at point A. Recognizing the central bank's incentives, however, the public will correctly anticipate the central bank's policy and set $\pi^e = \lambda\beta k$. In the absence of any shock, equilibrium would occur at point B, at the intersection of the TR curve with the vertical axis. Output shocks horizontally displace the AS curve, with negative shocks shifting the curve to the left. The degree of accommodation will depend on the slope of the TR curve, $\lambda\beta$.

Given that inflation in excess of the target is costly, in the absence of shocks the discretionary outcome is plainly suboptimal. It would therefore be desirable (assuming it were feasible) to "pre-commit" to a zero-inflation policy, in which case the TR curve in figure 1a would be replaced by a horizontal line at $\pi = 0$. Such a policy would prevent the central bank from responding appropriately to economic shocks, however. In the Barro-Gordon framework, the flexibility-versus-credibility tension manifests itself as a choice between a rigid zero-inflation rule and a fully discretionary policy. A central bank minimizing *ex ante* expected loss would choose the rule if the variance of the output shocks were below a certain threshold, and discretion if the shocks' variance exceeded that threshold.

Rogoff's (1985) twist on the Barro-Gordon framework was to appoint a "conservative" central banker whose weight on output, λ, was less than that of society. This corresponds to a shallower TR curve, as depicted in figure 1b. Inflation is lower than it would otherwise have been (the TR

Figure 2: The Flood-Isard state-contingent targeting rule

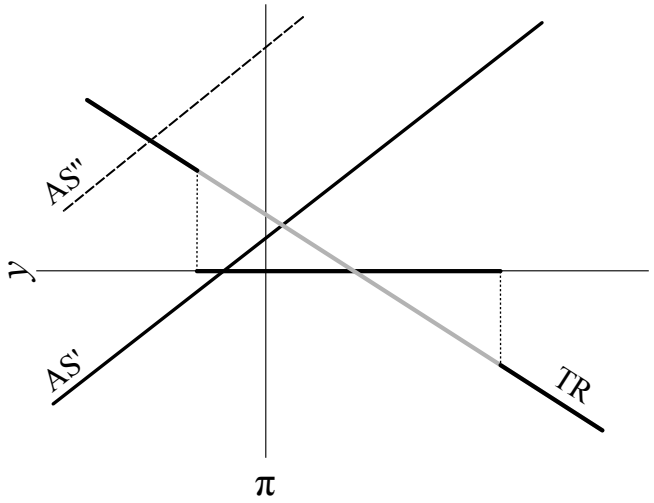

line intersects the vertical axis at a lower level), but output shocks will not be offset as much as they would have been under discretion. There is again a tradeoff between credibility (interpreted as conservatism), which reduces the inflation bias; and flexibility, which allows for an appropriately accommodative response to shocks.

Neither the Barro-Gordon nor the Rogoff model allows the central bank's policy response to depend on the magnitude of the shocks. Flood and Isard's (1989) insight was that the cost in terms of suboptimal inflation stabilization of strict adherence to a rule is proportional to the square of the shock size. Consequently, there are gains from conditioning the policy response on the size of the shock, with more accommodation of larger shocks. Allowing the policymaker to switch from an inflation-only rule to full discretion for large shock realizations would give rise to a TR curve like that depicted in figure 2.[2] In this case, the central bank would obey a strict zero-inflation rule under normal circumstances, when inflation bias made the larger contribution to the loss function, but revert to discretion when the cost of inflation was outweighed by the cost of output volatility.[3] This insight would naturally generalize to the decision to deviate from a Rogoff-style conservative (small λ) policy in favor of one of full accommodation.

[2]This is only a rough caricature of the full Flood-Isard model, which considers mixed strategies, and the choice between full discretion and a commitment to a partially state-contingent base growth rule.

[3]Lohmann (1992) showed how such a policy framework could result from a game between the government and the central bank, in which the central bank resorts to the flexible policy under the threat of a political override.

The most straightforward interpretation of IT in the context of this class of models is as "inflation only" targeting. While simplistic, it is easily motivated by the Barro & Gordon (1983) model, as it corresponds to the "precommitment" case. This interpretation of IT as inflation-only targeting has been strenuously contested by King (1997) and Bernanke *et al.* (1999), among many others. Svensson (2010a) in particular has insisted that *flexible* inflation targeting is fundamentally a targeting rule, like equation (1) above, expressing the optimal tradeoff between output and inflation fluctuations:[4]

> "Therefore, flexible inflation targeting can be described as "forecast targeting": "the central bank chooses a policy-rate path so that the forecast of inflation and resource utilization stabilizes both inflation around the inflation target and resource utilization around a normal level or achieves a reasonable compromise between the two."

Svensson's interpretation raises three important questions, however. One is what distinguishes IT from the conventional theory, laid out in Tinbergen (1952), of optimal policy in a situation with more targets than instruments. As Benjamin Friedman once observed when discussing a conference paper on IT:

> "[D]epiction of inflation targeting as maximizing a utility function including both an inflation term AND a term in the output gap... That's what people like me had been doing for decades... the idea that inflation targeting was how one maximized an objective function in output and inflation is what led me to suggest that, like Molière's Monsieur Jordaine, I must have been talking inflation targeting all along without realizing it."[5]

Friedman's (2004) more substantive objection was that committing to a specific numerical target for inflation obliges central banks to achieve that target at all costs, even when doing so would compromise the output stabilization objective.[6] The objection applies with particular force to a situation in which a large adverse shock threatened to breach the inflation target. Concern

[4] Pace Svensson, IT is often modeled as a Taylor-style instrument rule: see, for example, Galí (2002) and McCallum (2002).

[5] Email communication with Benjamin M. Friedman, April 5, 2011. The reference is to the central character in Molière's *Le Bourgeois Gentilhomme* (1670), Act II, sc. iv, stating "Par ma foi, il y a plus de quarante ans que je dis de la prose, sans que j'en susse rien." ("Good heavens! For more than forty years I have been speaking prose without knowing it.")

[6] Friedman (2003) also pointed out that stabilizing output while talking only about the inflation objective would be inconsistent with targeters' professed emphasis on transparency.

Figure 3: The Friedman Critique

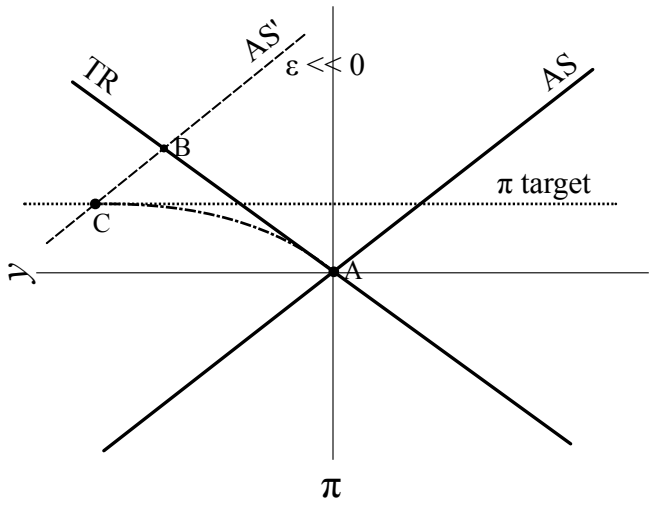

with overshooting the target would then require the central bank to behave more conservatively (depicted as a flattening of the TR curve) in the face of large shocks: in figure 3, the economy would move to point C rather than the optimal point B. The central bank becomes a prisoner of its own anti-inflationary rhetoric.

Faust & Henderson (2004) presented a fully worked-out version of Friedman's (2004) critique. They observed that optimal policy involves a commitment to return inflation to the target gradually and asymptotically, and not necessarily at some fixed horizon. When shocks are small, the two- to three-year horizon typically specified by inflation-targeting central banks is sufficient to bring inflation reasonably close to the target. A longer horizon would be appropriate for large shocks, however, and hitting the target over an abbreviated period of time may be insufficiently accommodative.

One rejoinder to this critique is that inflation targeting frameworks typically include an "escape clause," which allows the central bank to defer the attainment of the target in order to respond to large economic shocks with an appropriate degree of accommodation. The rationale for such a policy follows naturally from the Flood & Isard (1989) and Lohmann (1992) analyses.

A third issue is exactly how merely specifying a targeting rule anchors expectations and prevents central banks from opportunistically pursuing a discretionary policy — a question that is especially pertinent when central banks have the option to exercise an escape clause. Presum-

ably, the release of information by the inflation targeting central bank is the basis for anchoring expectations, and thus any additional flexibility, rather than anything distinctive about the inflation targeting central bank's policy approach to stabilization in and of itself.

The information provided by inflation targeting takes two forms. One is the public announcement of a numerical target. The second consists of detailed communication about the state of the economy and the outlook for inflation and other macroeconomic variables.[7] The target comes into play by removing (or at least reducing) uncertainty about the level of inflation around which the central bank is stabilizing price movements. This should enable the public and markets to hold the central bank accountable for whether the central bank is living up to its commitment. In Walsh's (1995, 2003) view, a specific target allows for the implementation of a contract that aligns the central banker's incentives with those of the public. An alternative interpretation, discussed below, is that the target, and the procedures associated with missing the target (e.g., writing letters to elected officials) provide an opportunity to communicate to the public the special circumstances that might require the temporary relaxation of the target. The purpose of detailed communication is clearly to share with the public the central bank's assessment of the shocks facing the economy, and to dispel the perception that the central bank might be concealing information about the nature of those shocks.

Understanding exactly how these mechanisms work requires appealing to a more complex set of theoretical models which, while related to the Barro & Gordon (1983) framework, include two additional features. Following Backus & Driffill (1985), one is that central bankers differ either according to their desired output level (k in the context of the model sketched above), or their degree of conservatism (the λ parameter in the same model). This is a plausible assumption, given that Kuttner & Posen (2010) found a tendency for unanticipated central bank governor turnover to generate significant financial market reactions.[8]

The second feature is that unobserved shocks prevent the private sector from distinguishing fluctuations in output and/or inflation that are caused by macroeconomic disturbances from those caused by the central bank deviating from its stated rule. This can be modeled as in Faust &

[7] Bernanke *et al.* (1999), Posen (2002), Kuttner (2004) and Svensson (2010a) discussed the operational aspects of transparency under inflation targeting.

[8] The evidence presented in the paper suggests that differences in λ may be more relevant than differences in k.

Svensson (2001) by assuming that rather than having perfect control over inflation, the central banker sets an "intended" inflation rate, and the actual rate is the intended rate plus some shock, η. This is in turn the sum of two shocks, the observable ξ and the unobservable ν.

The inference problem can be illustrated schematically using the AS and TR curves as in figure 4a. Suppose the economy starts at point A, and experiences an adverse output shock that shifts the AS curve to AS$'$, moving the economy to point C. This could happen in one of two ways. One is that the central bank acts like a dove, and follows a rule with $k > 0$, as in the dashed line. Alternatively, the central bank can claim that it was shooting for the "intended" inflation rate (which would have put the economy at point B), but because of a positive ν realization, inflation turned out to be higher than planned. Similar logic applies to the case of an unknown λ, shown in figure 4b: with a high-λ central bank, an adverse supply shock would move the economy to point C; but this is observationally equivalent to the response of a low-λ central bank, an "intended" inflation rate consistent with point B, and the realization of a positive ν shock. Essentially, the presence of the unobserved control error ν allows the central bank to behave less conservatively and get away with it.

With no way to distinguish between observationally equivalent outcomes, the public will resort to signal extraction to assess the plausibility of the central banks' excuse taking into account the relevant relative variances. If the variance of ν is small (i.e., there is greater transparency of intentions), then the central bank will not be able to get away with its ruse. More specifically, it will have the incentive to "behave itself" to reduce the degree to which higher-than-expected inflation causes people to revise (upward) their estimate of the unobserved k.[9]

In the Faust-Svensson framework, a more transparent central bank is one in which the variance of ν is small relative to the variance of ξ. The disclosure of information by inflation-targeting central banks presumably reduces the noise-to-signal ratio, thus enforcing more responsible behavior. Therefore, to the extent that it promotes transparency, inflation targeting strengthens central banks' incentive to build a reputation for low inflation by refraining from opportunistic behavior. In some sense, this is the inverse of the Cukierman & Meltzer (1986) result that ambiguity allows central banks to pursue a discretionary policy.

[9] In the Faust-Svensson model, it turns out that simply revealing k is the worst of all possible worlds, as that would remove the incentive for the central bank to preserve its reputation.

We conjecture that analogous logic would apply to a model with measurement error in the output gap, i.e. $y = \hat{y} + \zeta$, where $\hat{y}$ is the true output gap and ζ is measurement error, and unobserved central bank preferences.[10] This would also create an inference problem: the equilibrium at point A in figure 5 could either be the outcome of a small shock (a shift to AS$'$), a dovish (large λ) central banker, and no measurement error; or the result of a large shock (a shift to AS$''$), a conservative central banker, and a positive realization of the output gap measurement error, ζ. As in Faust & Svensson (2001) and Cukierman & Meltzer (1986), the private sector has to make inferences about the output shock and the central bank's preferences based on what it knows about the relative variances of the unobservables. Once again, this gives the central banker an incentive to maintain a reputation for low inflation.

A property common to all models of this type is that the shocks' relative variances — the signal-to-noise ratio — determine the optimal policy. In the pre-crisis "great moderation" world with small output shocks, the measurement error associated with the output gap would loom large. The inference problem would be nontrivial in this environment, and building a reputation would require conservative policymaking. The post-crisis reemergence of large output shocks has surely altered that calculus: to the extent that the uncertainty associated with output gap measurement has become (relatively) less relevant, it has become easier to distinguish opportunistic behavior from an appropriate response to macroeconomic conditions. This suggests that policymakers can afford to respond more aggressively to adverse shocks during dire economic circumstances: for a very large shock (a shift to AS$'''$), as in point C in figure 5, the central bank can respond optimally to the output shock, without worrying that doing so will undermine the institution's hard-earned reputation. The tension in inflation targeting between credibility and flexibility should therefore be less of an issue under these circumstances, especially if the overt exercise of the escape clause or the public explanation of the inflation overshoot allows the central bank to communicate the extraordinary nature of the shock.

[10] Kuttner (1994) proposed a method for quantifying this error, and Orphanides (2001) attributed much of the inflation of the 1970s to inaccurate estimates of the output gap.

Figure 4: Models with private information

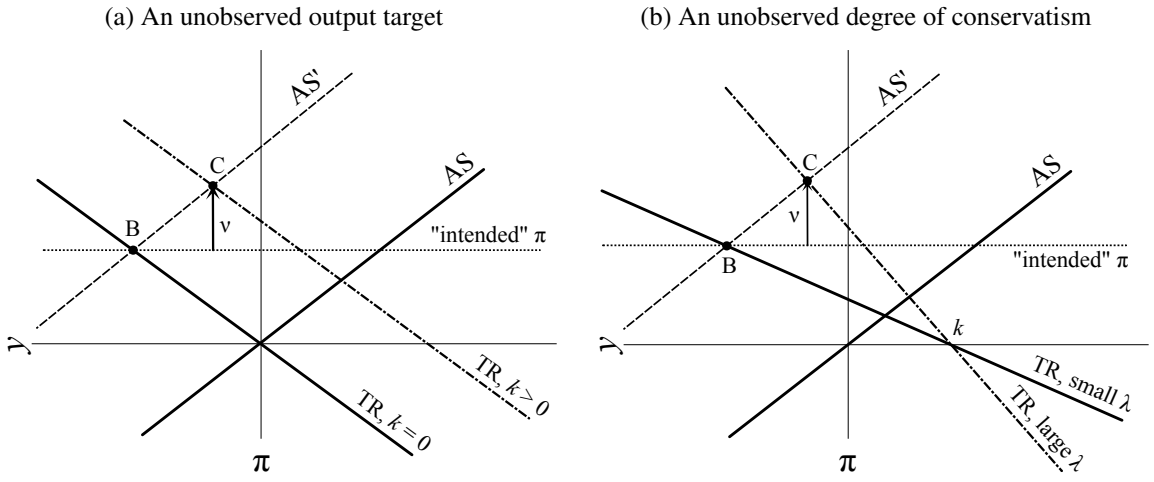

Figure 5: Conservatism with an escape clause

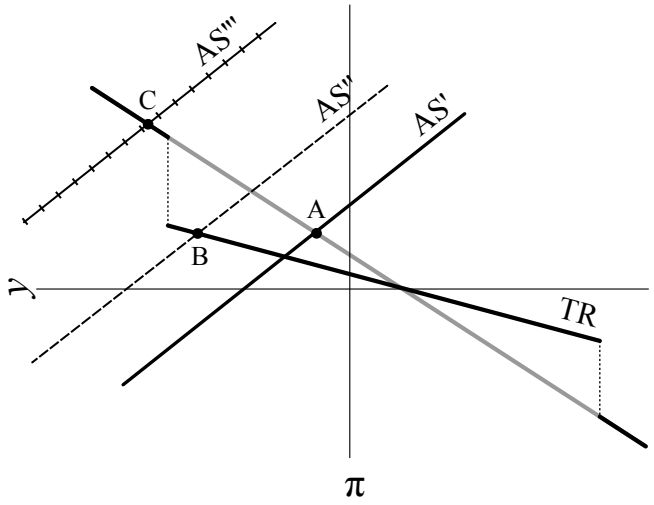

3 *Credibility, flexibility, and the behavior of inflation forecasts*

The goals of this section are twofold. The first is to assess the degree to which inflation expectations are anchored in the U.S. and the U.K. The second is to determine whether inflation targeting has created the expectation that the Bank of England will respond aggressively to return inflation to its target within a fixed period of time. Of particular interest is whether the Bank of England is perceived to behave more conservatively in this regard than the U.S. Federal Reserve, which operates without an explicit inflation target.

With rising inflationary pressures and a steep recession, the experience of the past few years should be especially informative with regard to these two goals. As shown in figures 6 and 7, inflation in both countries has been noticeably more volatile since 2007 than during the preceding eight years: 12-month inflation exceed the Fed's presumed implicit CPI objective of 2.5 percent by over three percentage points in the third quarter of 2008, and dropped below that target by 3.6 percentage points one year later. With the exception of a single quarter, the U.K. has consistently exceed its 2 percent CPI target since the end of 2007, even as the unemployment rate climbed to over 7.5 percent. The tension between credibility and flexibility therefore has been particularly acute in the U.K.

3.1 Are long-run inflation expectations well anchored?

Two alternative approaches exist for gauging inflation expectations: surveys of professional forecasters or households, and financial market indicators, such as breakeven inflation rates based on the spreads between nominal and inflation-indexed bonds. A clear advantage of market-based measures over surveys is their high frequency and real-time availability, making it possible (in theory) to gauge the market's instantaneous reaction to news. Moreover, in some countries, including the U.K., breakeven rates are available for longer horizons (e.g., ten years) than survey inflation expectations data. In addition, there is a presumption that because bond yields are determined in markets where "real money" is at stake, the implicit forecasts would be more meaningful than those from survey respondents. However a serious problem with the breakeven rates is their contamination with liquidity and term premia, which Pflueger & Viceira (2011) concluded account for much of the variability in TIPS yields in the U.S. Breakeven inflation rates are therefore flawed indicators

Figure 6: Inflation and unemployment in the U.S.

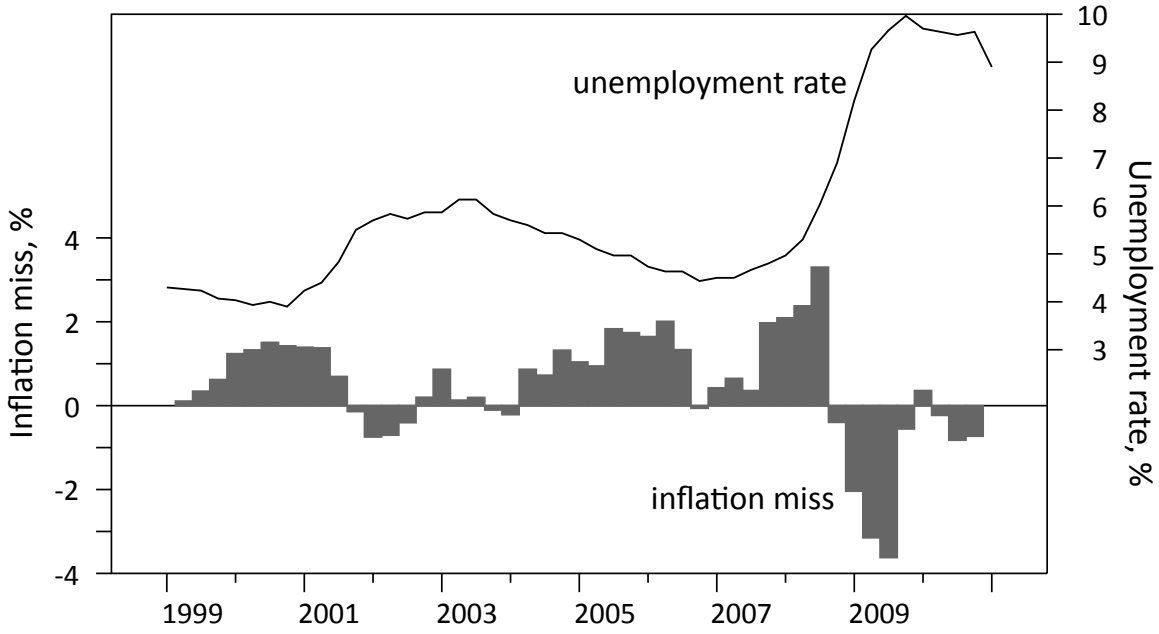

Figure 7: Inflation and unemployment in the U.K.

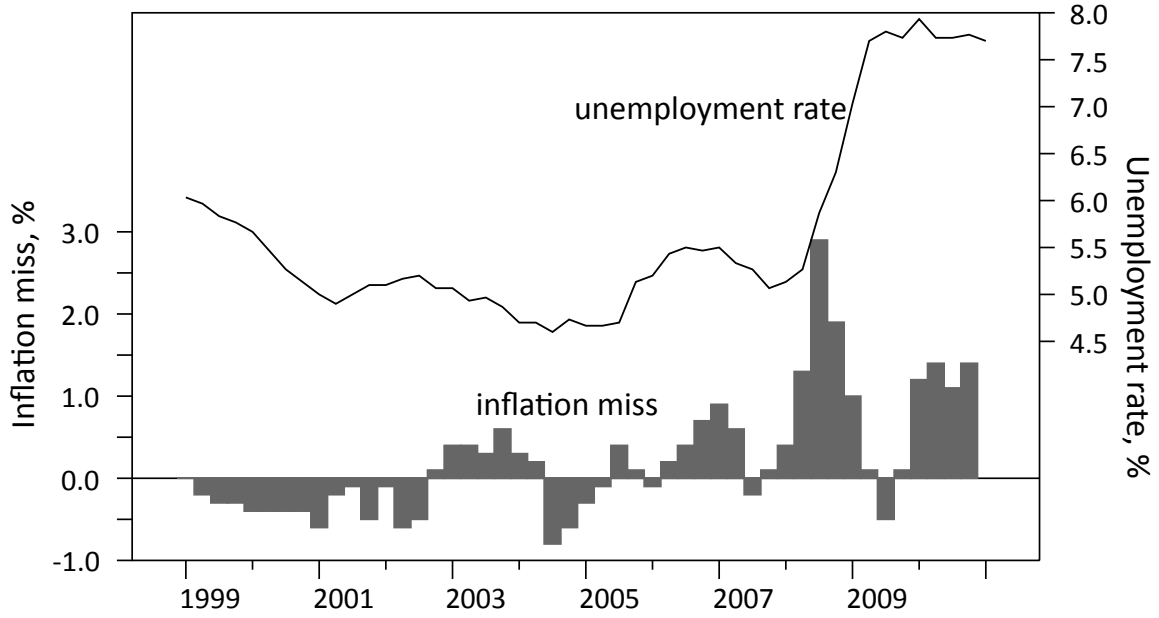

of expected inflation.[11] While only available at lower frequencies, survey-based inflation forecasts have the clear advantage of being uncontaminated by term and liquidity premia.

As shown in Figure 8, the 10-year CPI inflation forecasts for the U.S. from the Survey of Professional Forecasters (SPF) have remained remarkably stable at 2.5% since the late 1990s, with a standard deviation of only 7 basis points.[12] Only recently, since the onset of the financial crisis and ensuing recession, have those forecasts declined appreciably. At no time does there appear to have been an "inflation scare." The figure also shows that in contrast with the SPF forecasts, the long-term inflation forecasts obtained from the nominal-TIPS spread have been quite volatile, with standard deviations of 28 and 32 basis points for the 5-year-ahead 5-year forward and 10-year-ahead breakeven rates. Not surprisingly, the correlation coefficients between these two measures and the 10-year SPF forecast are low, and indeed have the wrong sign: -0.05 and -0.31 respectively.[13]

Defects notwithstanding, the long-term breakeven rates have been used in event study analyses by Gürkaynak *et al.* (2005), Gürkaynak *et al.* (2007), and Beechey *et al.* (2011) to gauge the impact of macroeconomic news on long-run inflation expectations. The studies report that U.S. breakeven rates exhibit a statistically significant response to economic news, whereas they generally do not in inflation-targeting countries.

The results reported in Table 1 confirm these studies' findings for the U.S. Regressions of changes in various expected inflation proxies (the 10-year nominal forward rate, the 10-year breakeven inflation rate, and the 5-year forward 5-year breakeven rate) on the surprise component of the core CPI and payroll employment releases show responses that are highly statistically significant. The coefficients' magnitudes are quite small in economic terms, however. A one-standard-deviation core CPI surprise would lead to a 1.6 basis point response in the 10-year forward breakeven inflation rate, for example. The maximum in-sample surprise of -0.5 percentage point would have generated a 8.6 basis point response. The effects of employment surprises are

[11] Gürkaynak *et al.* (2010) extracted an estimate of expected inflation, but their estimates are conditional on the questionable assumption that long run expected inflation is nonstationary, while the liquidity premium follows an AR(1).

[12] This statistic, and those reported subsequently, omit the four quarters from 2008Q3 through 2009Q2 (shaded in the plot), the height of the financial crisis.

[13] In terms of differences, the correlations are -0.18 and -0.12.

even smaller, with a one-standard deviation surprise producing a 1.3 basis point response and the maximum in-sample surprise of 330 thousand leading to a 4.8 basis point change. While these results suggest that macroeconomic news may have some impact on bond market participants' long-run inflation expectations, the effect is inconsequential in economic terms. These findings, along with the lack of any meaningful fluctuations in the long-run SPF forecasts, suggest that inflation expectations are reasonably well anchored in the U.S.

Table 2 presents an analogous set of results for the U.K., for inflation and GDP news (relative to the initial, "advance" report). Here, the only statistically significant responses are in the nominal forward rate, but even these are quite small in economic terms as the largest inflation and GDP surprises observed in the sample would have generated responses of only three basis points. The reaction of the 10-year forward breakeven and the 5-year forward 5-year breakeven rates are negligible in magnitude, and statistically insignificant.[14] The lack of any measurable reaction of market-based expected inflation proxies suggests that, at least on this metric, U.K. inflation expectations are well anchored.

3.2 How fast do inflation expectations converge?

The evidence presented above, and elsewhere in the literature, suggests that inflation targeting effectively anchors long-run expectations. We turn now to the question of whether stabilizing expectations comes at the cost of an excessively inflexible commitment to the inflation objective. This is the essence of the Friedman (2004) and Faust & Henderson (2004) concern that a target framed only in terms of inflation could compel the central bank to behave more conservatively than it would under the optimal state-contingent rule. Alternatively, to the extent that it allows the central bank to communicate its policies and reassure the public that it is not behaving opportunistically, an explicit objective may allow an inflation targeting central bank to respond *more* flexibly to adverse economic shocks than one that was concerned with preserving its reputation.

An informative indicator of a central bank's flexibility is the speed with which inflation is expected to converge to its target following an economic shock. As shown by Svensson (1997), this speed of adjustment is directly related to the central bank's degree of conservatism, represented by the λ in section 2. The idea is that a central bank committed to hitting its target at a specific

[14] Similar non-responses are obtained with other GDP releases, and with jobless claims data.

Table 1: High-frequency financial market response to macro news, U.S.

	Regressor				
Dependent variable	Core CPI	Employment	Std dev	Min	Max
Δ 10-year nominal forward	12.9**	22.0***	7.18	−19.3	25.0
	(6.0)	(7.3)			
Δ 10-year forward breakeven	17.2***	14.5***	4.64	−13.2	23.3
	(4.0)	(4.3)			
Δ 5-year/5-year forward breakeven	10.8**	8.8*	5.11	−13.5	24.1
	(4.4)	(5.2)			
Standard deviation	0.095	0.086			
Minimum	−0.5	−0.33			
Maximum	0.4	0.19			

Notes: The coefficients reported are from a regression of the change in the indicated inflation expectations proxy, measured in basis points, on the surprise component of either the core CPI inflation rate, measured as a monthly percent change, or the change in payroll employment, measured in millions. The expected changes in the macro variables are obtained from Money Market Services. The sample runs from June 1999 through June 2011, excluding the financial crisis period from September 2008 through June 2009. The CPI surprise regressions include 138 observations, and the employment surprise regressions include 136.

Table 2: High-frequency financial market response to macro news, U.K.

	Regressor				
Dependent variable	Inflation	Advance GDP	Std dev	Min	Max
Δ 10-year nominal forward	5.2**	6.3**	4.59	−17.0	17.1
	(2.3)	(3.0)			
Δ 10-year forward breakeven	0.0	1.1	3.02	−12.5	18.5
	(4.0)	(1.6)			
Δ 5-year/5-year forward breakeven	−2.5	−1.7	5.11	−13.4	21.0
	(1.5)	(1.8)			
Standard deviation	0.164	0.24			
Minimum	−0.5	−0.1			
Maximum	0.4	0.5			

Notes: The coefficients reported are from a regression of the change in the indicated inflation expectations proxy, measured in basis points, on the surprise component of either the inflation rate (RPIX pre-2004, CPI post-2004), measured as a monthly percent change, or the rate of GDP growth, measured in percentage points. The expected changes in the macro variables are obtained from Bloomberg. The sample runs from June 1999 through June 2011, excluding the financial crisis period from September 2008 through June 2009. The inflation surprise regressions include 135 observations, and the GDP surprise regressions include 45.

Figure 8: Alternative Measures of Long-Run Inflation Expectations, U.S.

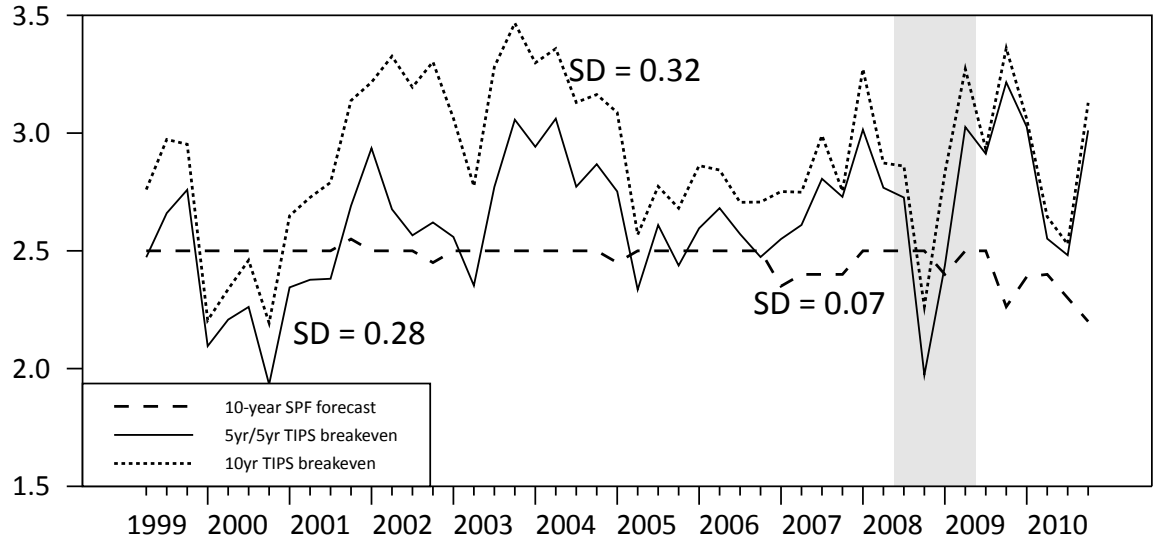

Notes: The SPF data are from the Federal Reserve Bank of Philadelphia, and the breakeven inflation rates from the Board of Governors of the Federal Reserve System. The shaded area corresponds to the 2008Q3 to 2009Q2 financial crisis period.

horizon (e.g., two years) would have to *increase* its speed of adjustment when there was a large shock. An unconstrained central bank, on the other hand, would choose not to accelerate the adjustment process, even if that meant overshooting the target at the two-year horizon.

One way to gauge this speed of convergence would be to estimate a time series model for the inflation rate. In a simple AR(1) model, for example, gradual convergence would be manifested in a large autoregressive parameter. An approach similar to this has been used in other work, such as that of Siklos (1999) and Kuttner & Posen (2001), to assess the long-run impact of inflation targeting. However this approach is more appropriate for the analysis of long time series, and less well-suited to shorter samples considered here.

3.2.1 A look at the medium-term inflation forecasts

Our approach to assessing flexibility relies instead on medium-term, one- to four-quarter-ahead inflation forecasts. Comparing forecasts at different horizons, it is possible to estimate the speed with which inflation expectations converge to their long-run mean, which should correspond to the central bank's inflation objective. Our method does not require estimating a time series process for

Figure 9: Median Inflation Forecasts, U.S.

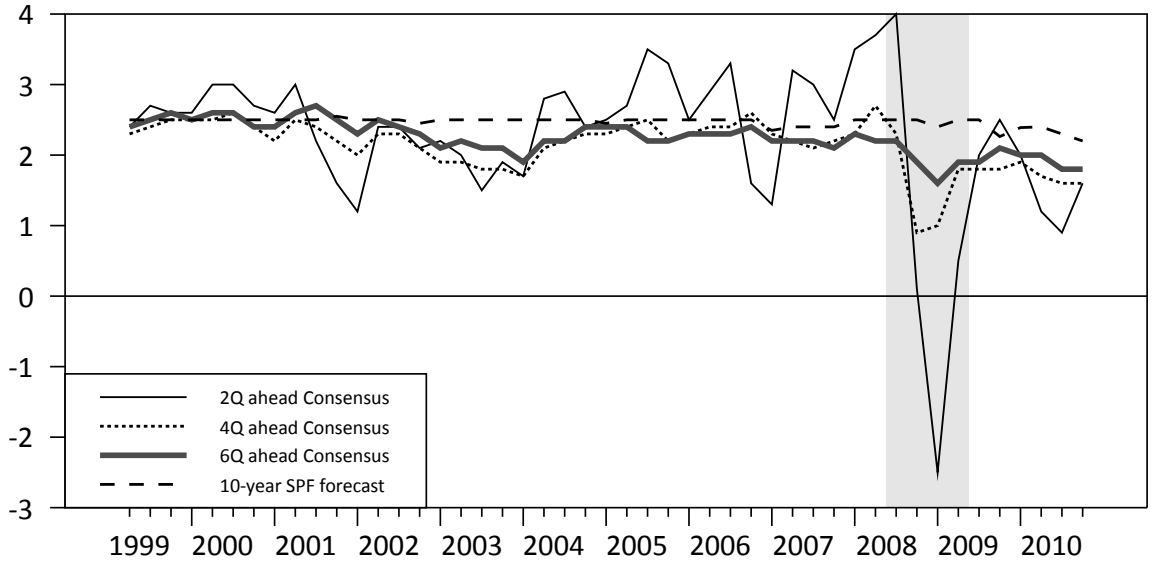

Notes: The forecast data come from Consensus Economics. The shaded area corresponds to the 2008Q3 to 2009Q2 financial crisis period.

Figure 10: Median Inflation Forecasts, U.K.

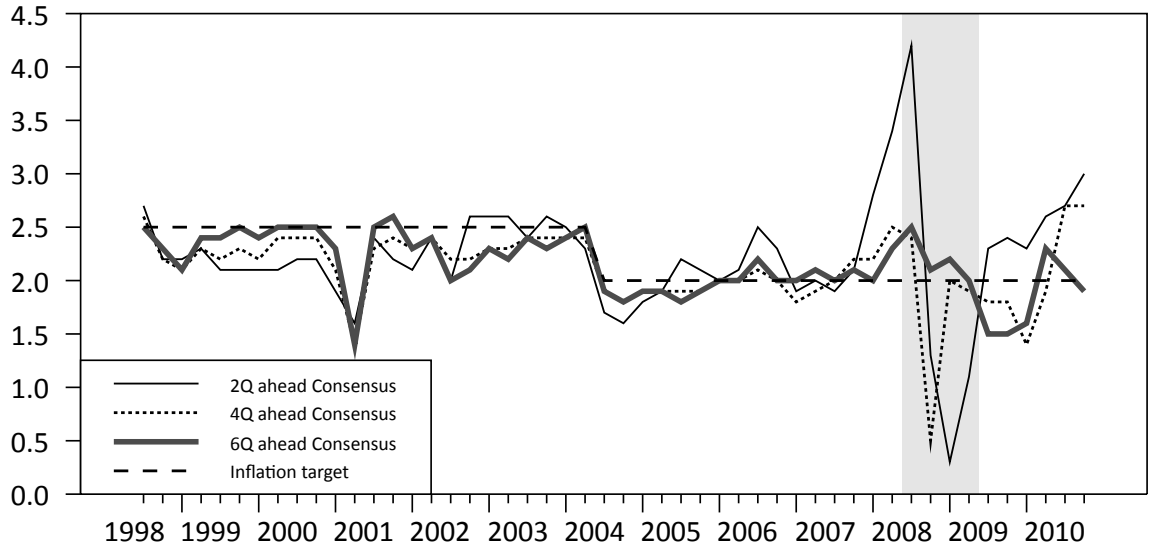

Notes: The forecast data come from Consensus Economics. The shaded area corresponds to the 2008Q3 to 2009Q2 financial crisis period.

inflation, and it avoids the complications associated with inference in the presence of a unit root. Moreover, our concern is less about actual inflation, which is apt to be very noisy, and more about the behavior of inflation expectations.

We use the year-over-year (i.e., change over four quarters) inflation forecasts provided by Consensus Economics, which collects quarterly forecasts from a panel currently consisting of 27 professional economists in the U.S., and 21 in the U.K.[15] The forecast horizons extend from one to six quarters. Forecasts for the two-, four- and six-quarter ahead forecasts are depicted in Figures 9 and 10, along with the 10-year SPF forecast for the U.S., and the official inflation target for the U.K.

In neither the U.S. nor the U.K. are even the six-quarter inflation forecasts equal to the long-run rate (the SPF 10-year forecast for the U.S. or the target for the U.K.). In the U.S., the six-quarter-ahead forecast fell well below 2.5 percent in 2003-04 and again in the most recent recession. In the U.K., the medium-term inflation forecasts were generally slightly less than the official 2.5 percent pre-2004 RPIX target, and exhibited noticeable fluctuations. After 2004, the medium-term forecasts deviated quite a bit from the 2 percent target, particularly just prior to the financial crisis as inflation rose, and as the economy fell into recession after the crisis. But what is important for our empirical method is the fact that the longer-horizon forecasts all tend to be closer to the long-run tendency than the shorter-term forecasts; our estimates of the speed of convergence are based on these differences.

Table 3 documents the variance of the U.S. inflation forecasts at the various horizons. Not surprisingly, the mean of the forecasts corresponds closely to the widely held belief that the Fed implicitly targets a CPI inflation rate of roughly 2.5 percent.[16] The key observation is that the standard deviation, shown in the second row of the table, falls as the forecast horizon increases. In the U.S., for example, the standard deviation of the one-quarter-ahead forecast is 83 basis points, falling to 29 basis points at the four-quarter horizon, and 22 basis points at six quarters. The third line of the table shows the standard deviation in the difference in adjacent forecasts, e.g.,

[15]The SPF compiles similar forecast data, but we use the Consensus data for both countries in order to maximize comparability.

[16]The Fed in recent years has shifted its emphasis towards the PCE deflator, whose growth rate since 1999 has averaged 0.32 percentage point less than that of the CPI.

between the three- and four-quarter ahead forecasts made at a given point in time. Here too, the standard deviation falls with forecast horizon. The standard deviation of the difference between the three- and four-quarter-ahead forecasts is 25 basis points, compared with 42 basis points for the difference between the two- and one-quarter ahead forecasts. Somewhat puzzlingly, the standard deviation falls by more than half going from the four- to five-quarter horizon, suggesting that many forecasters in the Consensus panel may simply be extending their four-quarter-ahead forecasts to the longer horizons.

Table 4 shows a comparable set of descriptive statistics for the U.K. data, allowing for the mean to change with the switch from the RPIX to the CPI in 2004. The statistics exhibit the same patterns as in the U.S., the means corresponding closely to the inflation target, and the standard deviations declining with horizon. Interestingly, the shorter-horizon standard deviations are considerably less than those for the U.S.: 43 basis points, versus 83. The four-quarter-ahead standard deviations are very similar, however.

3.2.2 A statistical model of expectations convergence

The basis for the model is a decomposition of the current inflation rate π_t into the sum of "underlying" inflation $\hat{\pi}_t$ and a transitory term v_t, which could represent measurement error, temporary inflation shocks, and short-term inflation dynamics,

$$\pi_t = \hat{\pi}_t + v_t \quad . \tag{4}$$

A plausible assumption is that underlying inflation, $\hat{\pi}_t$, is expected to revert to the (implicit or explicit) inflation target π^*,

$$E_t \hat{\pi}_{t+1} = \phi(\pi^* - \hat{\pi}_t) + \hat{\pi}_t \quad , \tag{5}$$

a partial adjustment equation in which ϕ represents the speed of convergence. This can also be written as

$$E_t(\hat{\pi}_{t+1} - \pi^*) = (1-\phi)(\hat{\pi}_t - \pi^*) \quad , \tag{6}$$

Table 3: Medium-term inflation forecasts, U.S.

	Forecast horizon, quarters					
	1	2	3	4	5	6
Forecast mean	2.53	2.43	2.31	2.28	2.20	2.27
Forecast standard deviation	0.83	0.68	0.45	0.29	0.25	0.22
Standard deviation of change	0.58	0.42	0.39	0.25	0.10	0.08

Notes: The statistics are for the four-quarter change in the CPI inflation rate, at the horizon indicated in each column. Data are from Consensus Economics.

Table 4: Medium-term inflation forecasts, U.K.

	Forecast horizon, quarters					
	1	2	3	4	5	6
Forecast mean, RPIX period	2.27	2.27	2.17	2.28	2.30	2.32
Forecast mean, CPI period	2.31	2.25	2.12	2.01	1.98	1.95
Forecast standard deviation	0.43	0.36	0.29	0.26	0.23	0.23
Standard deviation of change	0.23	0.21	0.19	0.18	0.16	0.13

Notes: The statistics are for the four-quarter change in the RPIX pre-2004 or the CPI inflation rate post-2004, at the horizon indicated in each column. Data are from Consensus Economics.

in which $1-\phi$ represents the degree of gradualism. This implies that the expected k-step ahead gap between inflation and its long-run rate decays geometrically at rate $1-\phi$,

$$E_t(\hat{\pi}_{t+k} - \pi^*) = (1-\phi)^k(\hat{\pi}_t - \pi^*) = (1-\phi)E_t(\hat{\pi}_{t+k-1} - \pi^*) \ . \tag{7}$$

The expected *observed* inflation includes the expectations of both the underlying inflation term, and the transitory v term. Assuming underlying inflation is given by equation (7), this can be written as

$$E_t(\pi_{t+k} - \pi^*) = (1-\phi)^k(\hat{\pi}_t - \pi^*) + E_t v_{t+k} \ ; \tag{8}$$

or as

$$E_t(\pi_{t+k} - \pi^*) = (1-\phi)E_t(\hat{\pi}_{t+k-1} - \pi^*) + E_t v_{t+k} \ . \tag{9}$$

Interpreting the time t survey-based k-period ahead forecast, $\pi^f_{t+k|t}$ as $E_t\pi_{t+k}$, equation (9) suggests a regression model in which the $k+1$-period-ahead forecast is regressed on the k-period-ahead forecast,

$$\pi^f_{t+k|t} = \phi\pi^* + (1-\phi)\pi^f_{t+k-1|t} + e_{k,t} \quad , \tag{10}$$

in which the composite error term includes a conventional error term $\varepsilon_{k,t}$, plus the unobserved expectations of the future transitory inflation terms.

$$e_{k,t} = E_t v_{t+k} - (1-\phi)E_t v_{t+k-1} + \varepsilon_{k,t} \ . \tag{11}$$

The slope coefficient can be interpreted as the gradualism parameter, $1-\phi$; equivalently, one minus the slope corresponds to the speed of convergence.[17] The implicit inflation objective can be calculated as the intercept divided by one minus the slope coefficient. The Consensus data used below are year-over-year inflation, but an analogous relationship holds between adjacent four-quarter forecasts.

The problem with this approach is that the unobserved $E_t v_{t+k-1}$ and $E_t v_{t+k}$ terms will be correlated with the forecast appearing on the right-hand side of the equation. This is because the correct

[17] Alternatively, the regression could be estimated with the difference in adjacent periods' forecasts, $\pi^f_{t+k|t} - \pi^f_{t+k-1|t}$, as the dependent variable.

specification is in terms of underlying inflation; using the inflation forecast instead introduces an errors-in-variables problem. However if the forecastability of the transitory component diminishes over time, then using longer-horizon forecasts will mitigate the problem. As an illustration, suppose v followed an MA process so that $E_t v_{t+1} = \theta_1 v_t$ and $E_t v_{t+2} = \theta_2 v_t$. In this case, the regression equation for the three-period- and two-period-ahead forecasts would be

$$\pi^f_{t+3|t} = \phi\pi^* + (1-\phi)\pi^f_{t+2|t} - (1-\phi)\theta_2 v_t + \theta_3 v_t + \varepsilon_{3,t} \quad . \tag{12}$$

In general, both of the transitory inflation terms will be correlated with the regressor. If v follows an MA(1) process, however, then $\theta_2 = \theta_3 = 0$, and the bias problem will disappear. Regressing $\pi^f_{t+2|t}$ on $\pi^f_{t+1|t}$ will yield a biased estimate of $1-\phi$ in the MA(1) case, however, and the direction of the bias will depend on the sign of θ_1.

3.2.3 Results

The results from estimating equation (9) on U.S. data from 1999Q2 to 2010Q4 (omitting the four periods corresponding to the financial crisis) for the one- to four-quarter forecast horizons appear in Table 5.[18] With R-squareds in the range of 0.53 to 0.74, the model describes the forecast data quite well. As expected, the mean inflation rate implied by the estimated slope and intercept terms is close to the Fed's presumed target. The slope coefficients themselves, which can be interpreted in terms of the perceived degree of gradualism, range between 0.47 and 0.71 depending on the horizon. Using GLS to estimate the equations for the two-, three- and four-quarter-ahead forecasts gives an estimate of 0.64, and the cross-equation restrictions are not rejected. This implies that 36 percent of the gap between underlying inflation and the objective is expected to close each quarter. If inflation were one percentage point above its mean, for example, the convergence parameter would imply that in one year the gap would be reduced to under 0.18 percent.[19]

The results for the U.K., shown in Table 6 for the 1998Q3 to 2010Q4 sample, are strikingly similar to those for the U.S. The equations' fit is similar, with R-squareds in the 0.7 range. As

[18] The longer-horizon five- and six-quarter forecasts are not used because the changes in adjacent forecasts are small, on the same order of magnitude as the one significant digit precision with which the forecasts are reported. Note that the right-hand-side variable in the regression involving the one-quarter-ahead forecast is the *current* quarter's forecast (the survey is conducted before all three months of inflation data are available).

[19] Not surprisingly, the estimated speed of convergence is much slower before the late 1990s.

expected, the implied objective closely matches the official Bank of England RPIX target pre-2004, and its CPI target post-2004. More importantly, the slope coefficients are very similar to those for the U.S., implying a comparable degree of perceived gradualism. The jointly estimated parameter of 0.71 says that 29 percent of any inflation gap is expected to be eliminated each quarter, and that a one percent deviation of inflation from its target is expected to shrink to a quarter of a percentage point over four quarters. Over the whole sample, therefore, there is no evidence to suggest that the Bank of England's inflation target has led it to attack inflation any more aggressively than the U.S. Federal Reserve.

The full sample results reported in Tables 5 and 6 may conceal state-dependent differences in the central banks' perceived behavior, however. A central bank concerned with maintaining a hawkish reputation may respond more aggressively during high inflation periods than it would when inflation was low. Similarly, a bank that felt constrained to hit its target at a fixed horizon might be unwilling to (optimally) lengthen that horizon when hit with a large adverse shock. Inversely, as argued above in section 2 , large shocks may allow central banks to act *less* conservatively to the extent that the legitimate response to macroeconomic shocks is more easily distinguished from "cheating." We investigate these possibilities by estimating the speed of convergence (alternatively, the degree of gradualism) on two pairs of mutually exclusive subsamples: whether inflation is above or below its target, and whether inflation is significantly either above or below target.[20]

Table 7 presents the results for the U.S. What is striking is how little the estimates depend on the subsample. The estimated degree of gradualism ranges between 0.61 and 0.69 (speeds of convergence between 0.39 and 0.31) regardless of whether the four-quarter inflation rate is above or below 2.5 percent, or whether the rate falls outside of a ± 0.75 percentage point range around the target.

As in the U.S., the estimates for the U.K., reported in Table 8, also seem largely independent of the prevailing rate of inflation. The degree of gradualism is 0.78 when inflation is above target, and 0.80 when below, a small difference both economically and statistically. The only noticeable difference in adjustment speeds is in the case of large versus small deviations from the target.[21]

[20] In the context of the U.S., "target" refers to an assumed 2.5 percent CPI inflation objective.

[21] Given the lower inflation volatility in the U.K. vis à vis the U.S., a ± 0.5 threshold was used instead of the ± 0.75

Here, inflation convergence appears to be somewhat *more* gradual when there are large shocks: 0.80 for large deviations, versus 0.64 for small deviations, which is consistent with the hypothesis that inflation targeting has allowed the Bank of England to respond slightly more flexibly to large shocks. Given the relatively small number of observations (12), one would not want to push this conclusion too far. Still, the result shows that the Bank of England's inflation target has not compelled it to behave more like an "inflation nutter" when faced with large deviations of inflation from its target.

4 Conclusions

This paper has revisited the flexibility-versus-credibility tension inherent in inflation targeting generally, with the specific goal of determining whether the monetary policy of the Bank of England, which has a formal inflation target, has been any less flexible than that of the Federal Reserve, which does not have such a target. The empirical analysis brought to bear a novel method for assessing flexibility: the speed of inflation forecast convergence, estimated by comparing professional forecasters' predictions at successively longer horizons.

We found no evidence to suggest that the Bank of England's inflation target compelled it to fight inflation any more aggressively than the Fed. Inflation forecasts converge at comparable rates in both countries. The lack of any detectable asymmetry in either country between the responses to positive versus negative deviations suggest that overshooting the target has not damaged the central banks' credibility; nor has above-target inflation elicited a disproportionately aggressive policy response. Similar results hold when comparing large versus small deviations from the target. Overall, the Bank of England seems to have been no less flexible than the Fed.

This leaves open the question of whether IT has had the desired effects on inflation expectations, given that the policy does not seem to lead to less flexibility, at least in the case of the Bank of England. Numerous studies have shown that in the U.K. inflation and its volatility have declined since the target's adoption, and expectations have stabilized — but this has happened in most other countries as well, including the U.S. The unconditional volatility of inflation is less than it is in the U.S., but that could be due to a variety of other reasons. The limited evidence presented in the

range for the U.S.

paper suggests that inflation expectations are better anchored in the U.K. than in the U.S., although the quantitative differences are slight. A definitive answer to this question remains elusive.

Given that IT seems not to have imposed an inflexible commitment on the Bank of England, it is tempting to ask whether IT is *too* flexible. Perhaps there are gains to be had from moving further towards the credibility end of the credibility-flexibility spectrum. Price level targeting is one frequently mooted framework that would move in that direction. In many countries, IT has successfully stabilized inflation expectations, but a credible commitment to a price path would, under some circumstances, generate stability-enhancing *changes* in inflation expectations.[22] The cost of such a policy in terms of diminished flexibility cannot be dismissed, however. Indeed, the U.K.'s recent experience is instructive in that regard. Having overshot its inflation target for more than three years, a price level target would perversely, given the current weak state of the British economy, oblige the Bank of England to keep inflation *below* two percent for an extended period of time. The issues raised in this paper therefore go beyond the choice of IT versus discretion, and also apply to the assessment of other proposed alternatives to IT.

[22] See Svensson (1999b) for a theoretical analysis, and Gaspar *et al.* (2010) for a simulation-based study.

Table 5: Estimated expectations convergence model, U.S.

	Forecast horizon, quarters (k)				
	1	2	3	4	2–4
Intercept	0.86	0.62	0.93	1.12	0.75
Implied inflation target	2.50	2.13	2.14	2.10	2.08
Slope $(1-\hat{\phi})$	0.65	0.71	0.57	0.47	0.64
Standard error	(0.07)	(0.06)	(0.06)	(0.07)	(0.03)
R-squared	0.70	0.74	0.70	0.53	
DW statistic	1.07	2.15	2.09	1.40	1.74

Notes: The sample runs from 1999Q2 through 2010Q4, excluding the financial crisis period from 2008Q3 through 2009Q2. The total number of usable observations is 43. The results reported in the 2–4 quarter horizon are from the GLS estimation of the equations for the 2-, 3- and 4-quarter horizons, constraining the coefficients to be equal across equations. The p-value for the likelihood ratio test of equal coefficients is 0.27.

Table 6: Estimated expectations convergence model, U.K.

	Forecast horizon, quarters (k)				
	1	2	3	4	2–4
Intercept	0.42	0.62	0.67	0.62	0.66
2% target dummy	0.06	−0.04	−0.14	−0.16	−0.12
Implied RPI target	2.07	2.27	2.27	2.30	2.28
Implied CPI target	2.34	2.11	1.81	1.72	1.87
Slope $(1-\hat{\phi})$	0.80	0.73	0.70	0.73	0.71
Standard error	(0.07)	(0.06)	(0.06)	(0.08)	(0.04)
R-squared	0.76	0.76	0.76	0.71	
DW statistic	1.74	1.64	1.49	1.68	1.64

Notes: The sample runs from 1998Q3 through 2010Q4, excluding the financial crisis period from 2008Q3 through 2009Q2. The total number of usable observations is 46. The results reported in the 2–4 quarter horizon are from the GLS estimation of the equations for the 2-, 3- and 4-quarter horizons, constraining the coefficients to be equal across equations. The p-value for the likelihood ratio test of equal coefficients is 0.60.

Table 7: Estimated degree of gradualism across subsamples, U.S.

	Forecast horizon, quarters (k)				Obser-	LR test
	2	3	4	2–4	vations	p-value
Above 2.5%	0.53	0.55	0.39	0.61	23	0.21
	(0.15)	(0.08)	(0.11)	(0.05)		
Below 2.5%	0.67	0.45	0.40	0.63	20	0.35
	(0.13)	(0.13)	(0.14)	(0.06)		
Outside $\pm$ 0.75% band	0.64	0.56	0.61	0.63	20	0.07
	(0.07)	(0.06)	(0.07)	(0.03)		
Inside $\pm$ 0.75% band	0.85	0.55	0.38	0.68	23	0.42
	(0.13)	(0.11)	(0.11)	(0.05)		

Notes: The sample runs from 1999Q2 through 2010Q4, excluding the financial crisis period from 2008Q3 through 2009Q2. The results reported in the 2–4 quarter horizon are from the GLS estimation of the equations for the 2-, 3- and 4-quarter horizons, constraining the coefficients to be equal across equations. The p-values for the likelihood ratio test of equal coefficients are reported in the last column.

Table 8: Estimated degree of gradualism across subsamples, U.K.

	Forecast horizon, quarters (k)				Obser-	LR test
	2	3	4	2–4	vations	p-value
Above target	0.78	0.84	0.77	0.78	20	0.25
	(0.11)	(0.12)	(0.10)	(0.06)		
Below target	0.73	0.63	0.86	0.80	26	0.35
	(0.15)	(0.08)	(0.19)	(0.05)		
Outside $\pm$ 0.5% band	0.82	0.82	0.72	0.80	12	0.09
	(0.09)	(0.13)	(0.13)	(0.06)		
Inside $\pm$ 0.5% band	0.70	0.53	0.79	0.64	34	0.12
	(0.10)	(0.06)	(0.15)	(0.04)		

Notes: The sample runs from 1998Q3 through 2010Q4, excluding the financial crisis period from 2008Q3 through 2009Q2. The results reported in the 2–4 quarter horizon are from the GLS estimation of the equations for the 2-, 3- and 4-quarter horizons, constraining the coefficients to be equal across equations. The p-values for the likelihood ratio test of equal coefficients are reported in the last column.

References

Backus, David, & Driffill, John. 1985. Inflation and Reputation. *The American Economic Review*, **75**(3), pp. 530–538.

Ball, Laurence M., & Sheridan, Niamh. 2004. Does Inflation Targeting Matter? *Pages 249–282 of:* Bernanke, Ben S., & Woodford, Michael (eds), *The Inflation Targeting Debate*. University of Chicago Press for the National Bureau of Economic Research.

Barro, Robert J., & Gordon, David B. 1983. A Positive Theory of Monetary Policy in a Natural-Rate Model. *Journal of Political Economy*, **91**, 589–610.

Beechey, Meredith J., Johannsen, Benjamin K., & Levin, Andrew T. 2011. Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area Than in the United States? *American Economic Journal: Macroeconomics*, **3**(2), 104–29.

Bernanke, Ben S., Laubach, Thomas, Mishkin, Frederic S., & Posen, Adam S. 1999. *Inflation Targeting*. Princeton, NJ: Princeton University Press.

Clarida, Richard. 2010. *What Has — and Has Not — Been Learned about Monetary Policy in a Low Inflation Environment? A Review of the 2000s*. Paper presented at "Revisiting Monetary Policy in a Low Inflation Environment" at the Federal Reserve Bank of Boston's 55th economic conference.

Cukierman, Alex, & Meltzer, Allan H. 1986. A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information. *Econometrica*, **54**(5), 1099–1128.

Dale, Spencer, Proudman, James, & Westaway, Peter. 2010. The inflation-targeting regime in the United Kingdom: a view from Threadneedle Street. *Oxford Review of Economic Policy*, **26**(1), 3–14.

Faust, Jon, & Henderson, Dale W. 2004. Is Inflation Targeting Best-Practice Monetary Policy? *Federal Reserve Bank of Saint Louis Review*, **86**(4), 117–144.

Faust, Jon, & Svensson, Lars E. O. 2001. Transparency and Credibility: Monetary Policy with Unobservable Goals. *International Economic Review*, **42**(2), 369–397.

Flood, Robert P., & Isard, Peter. 1989. Monetary Policy Strategies. *IMF Staff Papers*, **36**(3), 612–632.

Friedman, Benjamin M. 2003. The Use and Meaning of Words in Central Banking: Inflation Targeting, Credibility, and Transparency. *In:* Mizen, Paul (ed), *Essays in Honour of Charles Goodhart*, vol. 1. Cheltenham, U.K. and Northampton, MA: Elgar.

Friedman, Benjamin M. 2004. Why the Federal Reserve Should Not Adopt Inflation Targeting. *International Finance*, **7**(1), 129–136.

Friedman, Benjamin M., & Kuttner, Kenneth N. 1996. A Price Target for U.S. Monetary Policy? Lessons from the Experience with Money Growth Targets. *Brookings Papers on Economic Activity*, **1996**(1), 77–146.

Galí, Jordi. 2002. Targeting Inflation in an Economy with Staggered Price Setting. *Pages 271–294 of:* Loayza, Norman, & Soto, Raimundo (eds), *Inflation Targeting: Design, Performance, Challenges*. Series on Central Banking Analysis, and Economic Policies, vol. 5. Central Bank of Chile.

Gaspar, Vitor, Smets, Frank, & Vestin, David. 2010. Is the Time Ripe for Price-Level Path Stability? *Chap. 2, pages 22–51 of:* Siklos, Pierre L., Bohl, Martin T., & Wohar, Mark E. (eds), *Challenges in Central Banking*. Oxford University Press.

Giavazzi, Francesco, & Giovannini, Alberto. 2010 (July). *The Low-interest-rate Trap*. VoxEU, http://www.voxeu.org/index.php?q=node/5309.

Gürkaynak, Refet S., Sack, Brian, & Swanson, Eric. 2005. The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models. *The American Economic Review*, **95**(1), 425–436.

Gürkaynak, Refet S., Levin, Andrew T., Marder, Andrew N., & Swanson, Eric T. 2007. Inflation Targeting and the Anchoring of Inflation Expectations in the Western Hemisphere. *Pages 415–465 of:* Mishkin, Frederic S., & Schmidt-Hebbel, Klaus (eds), *Monetary Policy under Inflation Targeting*. Series on Central Banking Analysis, and Economic Policies, vol. 11. Central Bank of Chile.

Gürkaynak, Refet S., Wright, Jonathan H., & Swanson, Eric T. 2010. The TIPS Yield Curve and Inflation Compensation. *American Economic Journal: Macroeconomics*, **2**(1), 70–92.

Issing, Otmar. 2011 (7–8 March). *Lessons for Monetary Policy*. Remarks at the IMF Conference on Macro and Growth Policies in the Wake of the Crisis.

King, Mervyn. 1997. Changes in UK Monetary Policy: Rules and Discretion in Practice. *Journal of Monetary Economics*, **39**, 81–97.

Kohn, Donald L. 2003 (January). *Comments on Marvin Goodfriend's "Inflation Targeting in the United States?"*. Remarks delivered at the National Bureau of Economic Research conference on Inflation Targeting, Bal Harbor, Florida.

Kuttner, Kenneth N. 1994. Estimating Potential Output as a Latent Variable. *Journal of Business & Economic Statistics*, **12**(3), 361–368.

Kuttner, Kenneth N. 2004. A Snapshot of Inflation Targeting in its Adolescence. *Pages 6–42 of:* Kent, Christopher, & Guttmann, Simon (eds), *The Future of Inflation Targeting*. Sydney: Reserve Bank of Australia.

Kuttner, Kenneth N., & Posen, Adam S. 1999. *Does Talk Matter All? Inflation Targeting and Central Bank Behavior*. Working Paper 99-10. Peterson Institute for International Economics.

Kuttner, Kenneth N., & Posen, Adam S. 2001. Beyond Bipolar: A Three-Dimensional Assessment of Monetary Frameworks. *International Journal of Finance and Economics*, **6**, 369–387.

Kuttner, Kenneth N., & Posen, Adam S. 2010. Do Markets Care Who Chairs the Central Bank? *Journal of Money, Credit and Banking*, **42**(2–3), 347–372.

Lohmann, Susanne. 1992. Optimal Commitment in Monetary Policy: Credibility versus Flexibility. *The American Economic Review*, **82**(1), 273–286.

McCallum, Bennett T. 2002. Inflation Targeting and the Liquidity Trap. *Pages 395–437 of:* Loayza, Norman, & Soto, Raimundo (eds), *Inflation Targeting: Design, Performance, Challenges*. Series on Central Banking Analysis, and Economic Policies, vol. 5. Central Bank of Chile.

McCallum, Bennett T. 2010 (October). *Targets for Monetary Policy: Inflation, Exchange Rates, and Others*. Unpublished manuscript, Shadow Open Market Committee.

Meyer, Laurence H. 2004. Practical Problems and Obstacles to Inflation Targeting. *Federal Reserve Bank of St. Louis Review*, **86**(1), 151–160.

Mishkin, Frederic S. 2011 (February). *Monetary Policy Strategy: Lessons from the Crisis*. Working Paper 16755. National Bureau of Economic Research.

Nelson, Edward. 2009. An Overhaul of Doctrine: The Underpinning of UK Inflation Targeting. *The Economic Journal*, **119**(538), F333–F368.

Neumann, Manfred J. M., & von Hagen, Jürgen. 2002. Does Inflation Targeting Matter? *Federal Reserve Bank of Saint Louis Review*, **84**(4), 127–148.

Orphanides, A. 2001. Monetary policy rules based on real-time data. *The American Economic Review*, **91**(4), 964–985.

Orphanides, Athanasios. 2010. Reflections on Inflation Targeting. *Chap. 3, pages 13–24 of:* Cobham, David, Eitrheim, Øyvind, Gerlach, Stefan, & Qvigstad, Jan F. (eds), *Twenty Years of Inflation Targeting: Lessons Learned and Future Prospects*. Cambridge University Press.

Pflueger, Carolin E., & Viceira, Luis M. 2011 (March). *Inflation-Indexed Bonds and the Expectations Hypothesis*. Working Paper 16903. National Bureau of Economic Research.

Posen, Adam S. 2002. Commentary on 'Does It Pay To Be Transparent? International Evidence from Central Bank Forecasts', by Chortareas et al. *Federal Reserve Bank of Saint Louis Review*, **84**(4), 119–126.

Rogoff, Kenneth. 1985. The Optimal Degree of Commitment to an Intermediate Monetary Target. *Quarterly Journal of Economics*, **100**, 1169–89.

Siklos, Pierre L. 1999. Inflation-target design: Changing inflation performance and persistence in industrial countries. *Federal Reserve Bank of Saint Louis Review*, **81**, 47–58.

Svensson, Lars E. O. 1997. Inflation forecast targeting: Implementing and monitoring inflation targets. *European Economic Review*, **41**, 1111–1146.

Svensson, Lars E. O. 1999a. Inflation targeting as a monetary policy rule. *Journal of Monetary Economics*, **43**(3), 607–654.

Svensson, Lars E. O. 1999b. Price-Level Targeting versus Inflation Targeting: A Free Lunch? *Journal of Money, Credit and Banking*, **31**(3), 277–295.

Svensson, Lars E. O. 2010a. Inflation Targeting. *Chap. 22 of:* Friedman, Benjamin M., & Woodford, Michael (eds), *Handbook of Monetary Economics*, vol. 3B. North Holland.

Svensson, Lars E. O. 2010b (February). *Inflation Targeting after the Financial Crisis*. Speech delivered to the International Research Conference, "Challenges to Central Banking in the Context of the Financial Crisis".

Tinbergen, J. 1952. *On the Theory of Economic Policy*. North-Holland.

Walsh, Carl E. 1995. Optimal Contracts for Central Bankers. *The American Economic Review*, **85**(1), 150–167.

Walsh, Carl E. 2003. Accountability, Transparency, and Inflation Targeting. *Journal of Money, Credit and Banking*, **35**(5), 829–849.

Walsh, Carl E. 2009. Inflation Targeting: What Have We Learned? *International Finance*, **12**(2), 195–233.