

The Road Ahead

for Canadian Inflation Targeting

Christopher Ragan

Department of Economics
McGill University
christopher.ragan@mcgill.ca

and

Special Advisor
Bank of Canada
cragan@bankofcanada.ca

April 25, 2005

Abstract

This paper reviews the evidence regarding the success of inflation targeting, with a particular emphasis on Canada. It examines the level and volatility of inflation, the volatility of output growth, and the anchoring of inflation expectations. The paper comes out in favour of renewing the inflation targets beyond 2006, not least because of the logical coherence of the overall policy framework. It also discusses two substantive ways in which the inflation targets might be amended in the future, but about which more research needs to be done: the adoption of price-level targeting and the reduction of the inflation target to one percent. Finally, the paper provides a simple framework for thinking about the need for central-bank communications, and makes a case for the Bank of Canada placing more emphasis on longer-term education and less on the short-term signalling of its policy intentions.

This paper was prepared for the Bank of Canada's conference on "Issues in Inflation Targeting" held in Ottawa, April 28-29, 2005. It was written while the author was the visiting Special Advisor at the Bank of Canada. The views expressed here should not be attributed to the Bank.

The Road Ahead for Canadian Inflation Targeting

Christopher Ragan
April 25, 2005

When I was first asked to write a paper for this conference, I must admit to having less than my usual enthusiasm. Quite apart from the fact that I was unfamiliar with the large and growing literature on inflation targeting, my sense was that there were few unresolved issues. Inflation targeting had been with us in Canada for fourteen years, and it was widely viewed as being a successful approach to monetary policy. Whatever issues remained were likely to be less interesting than the ones in the late 1980s when the battle for “price stability” was first joined, or the ones in the early and mid 1990s when the main challenges were how to operationalize and communicate an effective inflation-targeting regime.

A few months’ of reading has changed my view. I was right that the remaining policy issues and related debates are less exciting than the ones from fifteen years ago. This is probably a good thing. But several interesting issues still need to be resolved, and there is still room to improve inflation targeting in Canada.

This paper offers my assessment of Canada’s inflation-targeting regime as well as my view of some outstanding issues. The current set of Canadian inflation targets expire at the end of 2006; this expiration leads to the following three questions which motivate the paper:

1. Should the inflation targets be extended beyond 2006?
2. If yes, should the targets be amended in any substantive ways?
3. Can the Bank of Canada improve its overall communications strategy?

The first two questions are obvious ones to ask when examining the potential renewal of the inflation targets. The third question is less obvious, except when one recognizes that central-bank communication is an essential part of inflation targeting, not least because of the emphasis inflation-targeting regimes place on the management of inflation expectations. Moreover, as the

other more “technical” aspects of inflation targeting have been worked out over the past decade, more emphasis has naturally fallen on the role of communications.

The following sections of the paper provide my views on these three questions. The first section reviews the evidence on the performance of inflation targeting, in Canada and elsewhere. I examine the level and volatility of inflation, the volatility of output growth, and the anchoring of inflation expectations. Aside from objective measures of macroeconomic performance, I also discuss why an important benefit of inflation targeting lies in its logical coherence, and in the relative ease with which the policy goals can be communicated to the public. I come out clearly in favour of extending Canada’s inflation targets.

The overall concern for the ongoing costs of inflation and for uncertainty regarding the long-run path of the price level suggests two ways to amend the inflation targets. Section 2 offers a brief review of the case for price-level targeting and the case for lowering the inflation target. In both cases, important questions remain to be answered regarding the optimal policy choice. Both issues are part of the Bank’s overall research program, and both are addressed in papers at this conference. I do not offer detailed discussions; I merely lay out the key issues and raise a few questions.

Section 3 explores the future of the Bank’s communications strategy, a strategy which has made impressive advances over the past decade. Despite these advances, however, there is more work to be done. I argue for a switch in emphasis, away from “short-term signalling” and toward more “longer-term education”. There is a sense in which the two types of communication are substitutes, but improvements in longer-term education are likely to generate more benefits — for the Bank and for society — than improvements in short-term signalling.

The paper’s final section addresses some of the likely opposition to my recommendation for better longer-term educational communications. Education is difficult to do correctly, and it requires considerable time and effort. But there are large benefits, including clearer thinking both inside and outside the Bank about economics and monetary policy. In addition, better communication of this type inevitably leads to a richer dialogue between the Bank and the public, an important part of central-bank transparency.

Section 1. Has Inflation Targeting Been Successful?

Determining the success of Canada's inflation-targeting (IT) regime requires that some reference point be chosen. We can evaluate aspects of Canada's economic performance in the pre-IT period and compare it with the performance following the adoption of inflation targeting. In this case, we need to establish how much of any change in observed performance is the result of the adoption of inflation targeting. A second approach is to compare Canada's performance during inflation targeting with the contemporaneous performance in other countries which were not using IT regimes. In this case, however, the implied counterfactual experiment is not clear: if Canada had not been inflation targeting, what would it have been doing?

When trying to think about the success of Canada's inflation-targeting regime, I continually find myself facing this quandary of the unclear counterfactual. If the Bank of Canada adopts the view that the long-run control of inflation is the only practical policy goal, and if it continues to take seriously the benefits of flexible exchange rates, then isn't any monetary policy it adopts going to look very much like inflation targeting? Perhaps some details in approach would differ, but their overall importance is likely to be small. Supporting evidence for this conjecture is presented by Collins and Siklos (2004) who find that the estimated monetary-policy reaction functions in Canada, New Zealand, and Australia, all formal inflation targeters, do not differ in significant ways from the reaction function in the United States, a country that does not formally target inflation but is nonetheless committed to maintaining low inflation and a flexible exchange rate.

The following review of the empirical evidence displays a little schizophrenia. At times I compare Canada's performance before and after the adoption of inflation targeting. At other times I compare the performance in IT countries with that in non-IT countries. In both cases I consider the extent to which inflation targeting can be credited with the observed differences in economic performance.

A. The Level and Volatility of Inflation

Figure 1 shows Canada's annual rates of CPI and core inflation from 1975 to 2004. As is readily apparent, inflation has been much lower on average in the years following the 1991 adoption of

inflation targeting than in the previous fifteen. It has also been less volatile. Table 1 shows inflation's mean and standard deviation in the pre-IT and IT years. It is unclear how to define appropriately the relevant periods for the comparison of economic outcomes. Some might argue that the genuine regime change began with John Crow's Hanson Lecture (Crow, 1988), even though formal inflation targets were not adopted until 1991. Others might argue that the early years of inflation targeting were a transition period that ought to be excluded from any comparison. Table 1 therefore shows three alternative definitions of the earlier and later periods: 1975-88 and 1989-04; 1975-91 and 1992-04; and 1975-90 and 1995-04. With all three definitions, the average level of inflation (both headline CPI and core inflation) and its standard deviation are lower in the later period than in the earlier period.¹

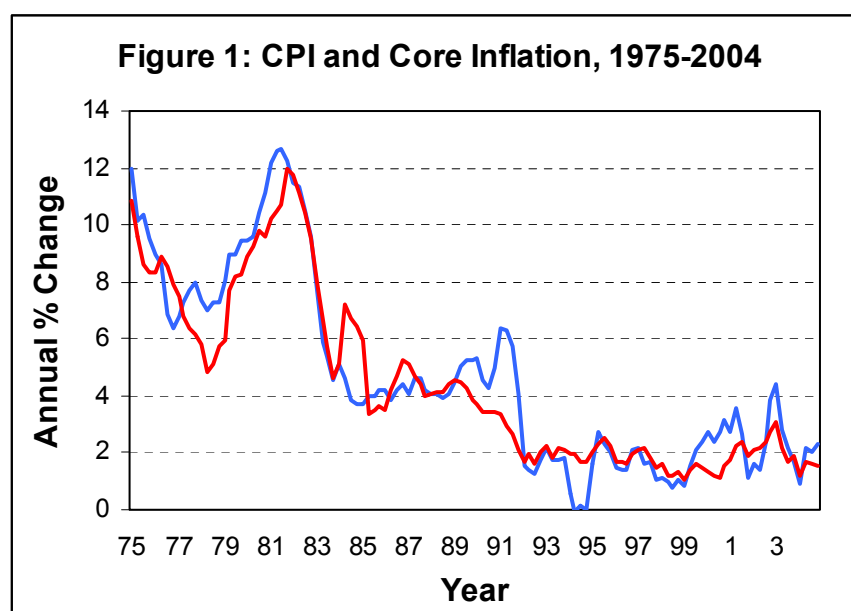


Table 1. Mean and Standard Deviation of Canadian Inflation

	75-88	89-04	75-91	92-04	75-90	95-04
Mean CPI	7.27	2.44	6.89	1.83	6.97	2.02
S.D. CPI	2.90	1.56	2.77	0.90	2.83	0.84
Mean Core	7.01	2.14	6.39	1.82	6.62	1.80
S.D. Core	2.44	0.83	2.60	0.43	2.51	0.47

Figure 1 looks much like similar figures for many countries, whether or not they are inflation targeters. Ball and Sheridan (2003) examine the inflation performance for 20 moderate-

¹ The data are quarterly, at year-on-year rates, and are seasonally adjusted. The CPI measure is CPI all items. The core measure is the CPI excluding the most volatile eight items and including a correction for changes in indirect taxes; before 1984 it is the CPI excluding food and energy.

inflation OECD countries, seven inflation targeters and 13 non-targeters. Since they observe similar declines in inflation in all of their countries, they attribute the decline in average inflation among the IT-countries to “regression to the mean” rather than to their adoption of inflation targeting. In other words, after controlling for the level of inflation in early years, the hypothesis that the adoption of inflation targeting has a significant negative effect on inflation in later years is rejected. For Ball and Sheridan, inflation targeting does not matter.

Hyvonen (2004) questions this interpretation. He shows for the same sample of OECD countries that such regression to the mean does not occur between the 1960s and 1970s, and thus does not appear to be a mechanical property of the data. He also shows that such regression to the mean does occur across the same decades in a sample of inflation within various U.S. metropolitan regions. Since these various U.S. regions obviously have the same monetary policy, his interpretation of the OECD data is that the convergence in inflation apparent in the 1990s may be partly due to the convergence in monetary policies. Finally, and most important, Hyvonen argues that the Ball and Sheridan study suffers from problems of endogeneity. Most of the IT countries in the OECD sample had higher inflation in their pre-IT years than did the non-IT countries. If these higher inflation experiences had anything to do with the countries’ decisions to become inflation targeters, then the IT dummy variable used by Ball and Sheridan cannot be taken to be exogenous.

Rogoff (2003) takes a different tack in arguing that the disinflation observed in many countries over the past 15 years has not been solely the result of better monetary policies. While he is quick to give some credit to greater central-bank independence and improved monetary-policy techniques, he also argues that the forces of globalization have tended to reduce the trend rate of inflation. His basic argument is that globalization leads to greater competition among firms which reduces mark-ups of price above marginal cost and increases the level of equilibrium output toward the level observed in a competitive and frictionless economy. This reduction in the gap between the actual level of output and the “natural” level of output plays a key role in Rogoff’s thinking. In models like Kydland and Prescott (1977) that emphasize the inability of central banks to commit to any particular course of action, this smaller gap reduces the payoff from unexpected inflation and thus reduces the inflationary bias in a world of discretionary monetary policy. Thus, the greater competitive forces from globalization lead to a decline in long-run inflation, independent of any regime change by central bankers.

For several reasons I find it difficult not to give at least some credit to improved monetary policy for the disinflation observed in Canada and elsewhere over the past 15 years. First, since Ball and Sheridan's sample only contains moderate-inflation OECD countries, I am not surprised that their estimated regression coefficient shows no significant difference for IT countries. As I said earlier, with supporting evidence from Collins and Siklos (2004), in the modern era of low inflation, monetary policy looks broadly similar across countries, even when some central banks are formal inflation targeters and others are not. But this is not compelling evidence that IT "doesn't matter" — perhaps it is just an indication that central banks can be effective inflation targeters without using these specific words.

Second, I share Hyvonen's concern with the endogeneity of Ball and Sheridan's IT dummy variable. The higher-than-average inflation experienced in the 1970s and 1980s by specific countries may well have been instrumental in their subsequent adoption of inflation-targeting regimes. In contrast, countries with less severe inflation policies may have been led to undertake less extreme, or at least more vague, policy approaches to fight their inflations. But this possibility leads to another question: is there something special about the IT countries that would have prevented them from achieving their actual disinflations without the explicit adoption of IT regimes, even when other countries got by without such formal arrangements? The openness of those economies and the importance of exchange-rate changes may play some role in the answer, but this is just a hunch.

Finally, I am not convinced by Rogoff's arguments about globalization's effect on the long-run path of inflation. While competitive forces unleashed by greater globalization may tend to accelerate the pace of productivity improvements and thus have temporary downward pressure on inflation, Rogoff's preferred mechanism is quite different and applies more to the long-run trend of inflation. His arguments are that the forces for disinflation come from a reduction in central banks' incentives to produce unexpected inflation. But I share Howitt's (2000) view in finding this game-theoretic approach to monetary policy unenlightening in the modern era of central banking. It is unrealistic to think that today's central bankers intentionally produce unexpected inflation so as to drive the level of output above its natural level. On the contrary, many central banks since the early 1990s have focused on exactly the opposite — being clear about their objectives, transparent about their policy actions, and accountable for their past performance. Thus, while the Kydland and Prescott (1977) class of models was instrumental for

framing and clarifying the rules-versus-discretion debate in monetary policy, I do not think it continues to accurately describe the motivations of increasingly independent central banks.

B. The Volatility of Output Growth

A common view regarding the operation of inflation-targeting regimes is that the success in keeping inflation low and stable may come at a high price — an increase in the volatility of output growth. This is essentially the view behind those opponents of inflation targeting, such as Friedman (2003), who claim that inflation-targeting central banks lack some desirable flexibility in policy that non-inflation-targeting central banks possess. On the other hand, Bernanke and Mishkin (1997), Goodfriend (2003) and others claim that the “constrained discretion” embodied in the inflation-targeting framework is fully consistent with a concern for and a commitment to stable output growth. Indeed, Blanchard (2003) argues that the “divine coincidence” of inflation targeting is that stabilizing inflation is equivalent (in a fairly broad class of models) with stabilizing output around its natural or potential level. This coincidence may provide an important explanation for the widespread support for inflation targeting:

Those who care about inflation volatility like the stated goal of the policy. Those who care about output stabilization see inflation targeting as a commitment by the central bank to stabilize output around its natural level, to get the economy out of recessions, and to slow the economy down in booms. (p.4)

Of course, even if one were to observe a decline in the volatility of output growth following the adoption of inflation targeting, this observation would not establish that monetary policy was the cause of the reduced volatility. Stock and Watson (2003), for example, show that each of the G7 countries has experienced a reduction in the standard deviation of real per capita output growth between the 1960-83 and 1984-02 periods. And since the break-point does not coincide with the adoption of inflation targeting, and since two of the countries are not engaged in inflation targeting, there is obviously some additional force at work.

Blanchard and Simon (2001) have documented the considerable decline in the volatility of U.S. output growth over the past five decades, and have attributed it to many causes, some only loosely related to monetary policy. Kahn, McConnell and Perez-Quiros (2002) attribute the

same observation mostly to changes in inventory behaviour in the durable goods sector. In both studies there is room for improved monetary policy to account for some of the reduced volatility, but the case is not made.

What do the data suggest for Canada? Figure 2 shows the growth rate of real GDP in Canada from 1975 to 2004, and also the output gap (as a percentage of potential output) over the same period. Both output growth and the output gap have become less volatile since the early 1990s.² Table 2 shows the mean and standard deviation of both output measures using the same alternative break-points as in Table 1.

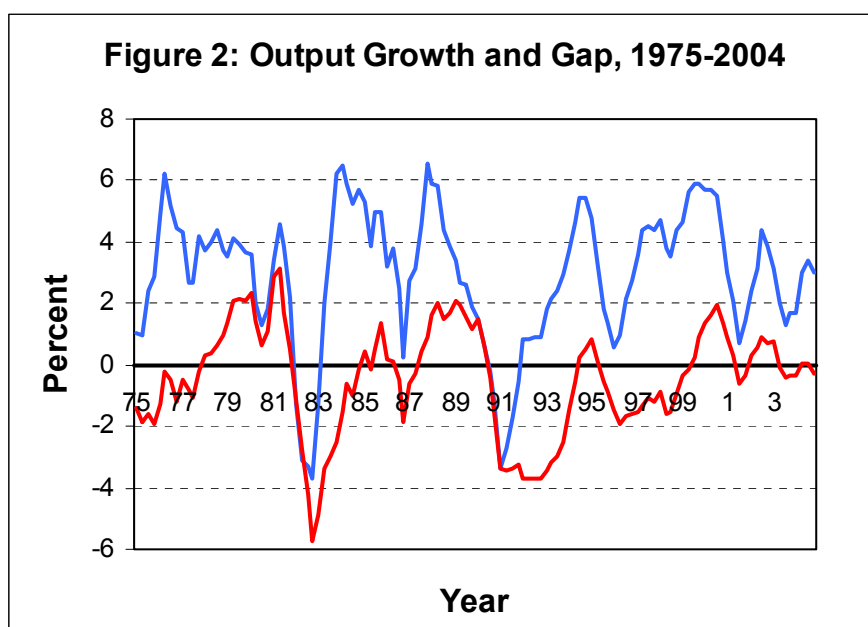


Table 2. Mean and Standard Deviation of Real GDP Growth and the Output Gap

	75-88	89-04	75-91	92-04	75-90	95-04
Mean GDP Growth	3.29	2.65	2.75	3.20	3.06	3.36
S.D. GDP Growth	2.36	2.07	2.59	1.60	2.35	1.54
Mean Output Gap	-0.22	-0.67	-0.28	-0.69	-0.08	-0.20
S.D. Output Gap	1.88	1.64	1.95	1.48	1.85	1.01

Can this reduced output volatility be attributed to Canada's adoption of inflation targeting in 1991? Or has Canada just been lucky by being exposed to a relatively uneventful economic

² The change in the average growth rate depends on how the sample is divided. Interestingly, however, in each case the average output gap is larger in magnitude (more excess capacity) in the latter part of the sample. The data for the output gap are produced by the Bank's QPM projection model and are subject to frequent revision.

environment over the past decade? Cecchetti, Flores-Lagunes, and Krause (2004) address this general issue directly for a sample of 24 countries, including large industrialized countries and small developing ones. They examine the behaviour (using quarterly data) of output volatility and inflation volatility across two periods — 1983-90 and 1991-98. They organize their thoughts by using a downward-sloping efficiency frontier showing combinations of output volatility and inflation volatility that the economy could achieve if policy were optimal, as first outlined by Taylor (1979). In this setting, an improvement in policy results in the actual output and inflation volatilities getting closer to the frontier, whereas changes in the volatility of the economy's underlying shocks result in shifts of the frontier itself. Thus any observed change in the actual volatilities of output and inflation may be due to some combination of a change in policy and a change in the volatility of shocks.

Their approach is to estimate the frontier for each sample country, and to decompose the observed changes in economic performance into the part representing a change in shocks and the part representing a change in policy. They adopt a widespread view that the role of discretionary fiscal policy in economic stabilization is at best limited; thus, improved stabilization policy is interpreted to be improved monetary policy.

For 21 of the 24 countries, their estimates suggest an improvement in policy from the 1980s to the 1990s. Some of these countries experienced reductions in both volatilities, a clear improvement in performance. Others experienced a reduction in one volatility and an increase in the other, but the weighting was such that the measure of overall performance improved. And for several of the countries in their sample, the estimated improvement in policy explains more than 100 percent of the improvement in economic performance — revealing that the volatility of the underlying economic shocks increased from the 1980s to the 1990s but that the improvement in economic policy was large enough to more than outweigh this effect.

Canada was one of these countries with large improvements in policy and economic performance, despite a more volatile economic environment. Moreover, the authors conduct a sensitivity analysis by varying the parameters of the efficiency frontier. For a large range of parameter values, their basic results hold: Canada's monetary policy improved a lot, and more than enough to offset the increase in underlying economic volatility.

Many of the countries that Cecchetti et al (2004) estimate to have improved policies are inflation targeters but some notable ones such as Japan and the United States are not. And while their findings for Canada may make us feel good about the Canadian policy framework, there is nothing in their study pointing specifically to the success of inflation targeting. Furthermore, their identifying assumption that all improvements in policy are taken to be improvements in monetary policy is less tenable in Canada than in many other countries. Though the value of discretionary fiscal stabilization policy probably is quite limited, the improved overall fiscal environment — with major structural advances in the late 1980s and overall advances in the late 1990s — has been of tremendous importance to Canada’s macroeconomic setting.³

Levin, Natalucci, and Piger (2004) use their sample of IT and non-IT countries to examine output and inflation volatility. They first examine the unconditional standard deviation of output growth and find that the volatility in IT countries is no higher than in non-IT countries. They then show that the unconditional standard deviation of core inflation is higher in the IT countries than in the non-IT countries. This appears to contradict their accompanying result (discussed below) that the inflation process in IT countries displays less persistence than in the non-IT countries. After all, if there are stronger forces at work generating mean reversion, we might expect a lower unconditional variance.

They then decompose the volatility of core inflation into two parts, one from the variance of underlying shocks and the second from the propagation of these shocks. The estimates suggest that the IT countries (with the exception of Australia) have most of their inflation variance accounted for by the variance of the underlying shocks, with relatively little due to the propagation mechanism. In contrast, the non-IT countries had a much smaller fraction of their overall inflation variance accounted for by the shocks, and much more due to propagation. Thus, Levin et al (2004) conclude that the IT countries were exposed to more shocks than were the non-IT countries; “good luck” does not seem to be one of the things that can explain the good economic performance in IT countries. This is as true for Canada as for the other IT countries in their sample.⁴

³ Kneebone and Chung (2004) illustrate how changes in the federal government’s overall debt-to-GDP ratio over the past 25 years have hidden very important movements in the underlying structural, cyclical, and rate components.

⁴ In their estimates, the ratio of the variance of inflation to the variance of underlying shocks is 1.23 for Canada and 2.25 for the United States. By this measure, Canada has been much less lucky (in terms of shocks to inflation) than the United States.

C. The Anchoring of Inflationary Expectations

A central part of the appeal of inflation targeting is that the greater transparency of policy objectives, combined with the central bank's acquired credibility, can result in a better anchoring of inflation expectations. If such expectations do become better anchored, then policy gains a degree of freedom in the face of economic shocks. What is the evidence that inflation targeting has such a desirable anchoring effect on expectations?

Johnson (2002) examines the behaviour of forecast-based expected inflation in a sample of IT and non-IT countries from 1984 to 2000. The IT countries include Australia, Canada, New Zealand, Sweden and the United Kingdom; the non-IT countries are France, Germany, Italy, the Netherlands, Japan, and the United States. He regresses his measure of expected inflation on a lagged dependent variable, country and year dummy variables, and IT dummy variables that are sufficiently flexible to allow for a different effect as the IT regime ages. He finds, first, that expected inflation falls after the adoption of inflation targeting and, second, that the magnitude of the reduction increases as the age of the IT regime increases. Given the presence of the lagged dependent variable and the country and year dummy variables, the first result suggests that inflation targeting does indeed help to reduce inflationary expectations. The second result is perhaps more interesting; it suggests that as the IT regime ages, and presumably the central bank's credibility builds, the reductions in inflationary expectations become locked in.

Another interesting result from Johnson's paper is the absence of any direct effect of inflation targeting on the standard deviation of expected inflation or on the absolute forecast errors for inflation. These results may lead one to conclude that, even though inflation targeting reduces inflation, it has no apparent effect on inflation uncertainty. But Johnson notes that if lower average inflation is associated with less inflation uncertainty, as the literature seems to confirm, then inflation targeting may well have a negative effect on uncertainty — just not through a direct announcement effect.

The anchoring of inflationary expectations, however, is less about their level than about their response to economic shocks and policy actions. Levin, Natalucci, and Piger (2004) make this point using a similar sample of IT and non-IT countries over the 1994-2003 period. Consistent with Johnson's (2002) findings, they first show that the unconditional volatility of

their survey-based measure of inflation expectations in the IT countries is no lower than in the non-IT countries.

Their more interesting result is about how expected inflation at various forecast horizons reacts to changes in current actual inflation. Their regression equation is:

$$\Delta E[\Pi_{i,t,q}] = \lambda_i + \beta \Delta \Pi_{i,t} + \varepsilon_{i,t}$$

where $E[\Pi_{i,t,q}]$ is the expectation of inflation in country i formed at time t for q years into the future. In this setting, well-anchored expectations are suggested by an estimate of β close to zero; significantly positive estimates of β suggest that inflationary expectations can be easily dislodged by current inflation shocks. They find that for one-year-ahead forecasts, the estimate of β is not significantly different from zero for either IT or non-IT countries. However, for forecast horizons of 3, 5 and up to 10 years, the estimated values of β are significantly lower in the IT countries than in the non-IT countries. In both the Euro area and in the United States, the estimated values of β are larger than for the average of the non-IT countries.

Finally, the authors show that inflation persistence appears to be lower in the IT countries than in the non-IT countries. Following an inflation shock, inflation more quickly returns to its starting level in the IT countries. From the implied impulse-response functions, they show that nearly half of a one-unit shock to inflation in the IT countries disappears after only one quarter, and 90 percent disappears after four quarters. In contrast, for the non-IT countries it takes four quarters before half of the shock has disappeared and almost three years for 90 percent of the effect to wear off.

The anchoring of inflationary expectations is a large part of the appeal of inflation-targeting regimes. This is one aspect of monetary policy I failed to fully appreciate a few years ago, but it has come to pervade my thinking. We spend much time with our students explaining how monetary policy works — from movements in the policy interest rate to changes in other asset prices, and then to changes in aggregate demand, the output gap, and eventually inflation. But we probably do not emphasize enough to our students that all of these linkages that comprise the monetary transmission mechanism operate best in a setting in which inflation is low and in

which inflation expectations are well anchored. This point may not lend itself easily to simple diagrams, but it must be emphasized, clearly and often.

Well-anchored expectations provide central banks with a greater capacity to pursue stabilization policy without dislodging inflation expectations and setting in motion a dangerous wage-price spiral (Goodfriend, 2003). This is true in the presence of aggregate demand shocks, but it is also true and especially important in the face of aggregate supply shocks, although even with well-anchored expectations supply shocks continue to present central banks with greater challenges, including for communications. As King (2003) says about how the members of the Monetary Policy Committee at the Bank of England view their challenges:

[Having well-anchored inflation expectations] makes monetary policy easier by giving monetary policy a bit more time to respond. We are not worried that an inflationary shock is likely to lead immediately to an upward or downward revision of inflation expectations, feeding through very quickly as it might have done before into inflation expectations, wage bargaining, and then prices. (p.13)

There is considerable debate these days about whether the U.S. Federal Reserve should adopt a formal policy of inflation targeting, and much emphasis has been placed on the flexibility of monetary policy with or without inflation targets. But perhaps of equal importance is the issue of inflation expectations. The greater anchoring of inflation expectations, and the improvement in policy effectiveness that goes along with it, is probably the largest benefit that the United States stands to gain from adopting a formal inflation target. As committed as the Fed may be to maintaining low inflation, nobody knows with any precision what low inflation means in the United States. And this lack of clarity matters for economic performance. Orphanides and Williams (2003) argue that when the public knows the long-run inflation target, rather than having to infer it from outcomes and behaviour, there is an improved tradeoff between the volatility of inflation and the volatility of output growth.

D. A Case for Extending Canada's Inflation Targets

Canada's macroeconomic performance has clearly improved over the past 15 years (Longworth, 2002). There is naturally some debate about the proximate cause of this improved performance,

and whether any credit ought to be given to monetary policy in general, and to inflation targeting in particular. This debate will probably provide job security to macroeconomic researchers for years to come, as the decomposition and identification of the various causal forces will remain contentious and subject to considerable judgement — like most interesting debates in macroeconomics!

Whether inflation targeting gets credit for the improvements in economic outcomes, however, may be missing the relevant point for the decision facing the Bank of Canada. The first question posed in the introduction is whether Canada should extend the inflation targets beyond 2006. But inflation targeting is already with us, and presumably some costs would be incurred in moving to any different regime. Thus the relevant question is not whether IT has been good, but whether there is any reason to move away from IT and toward some other system for monetary policy that offers better performance.

Though it may be difficult, based on the available empirical evidence, to conclude that inflation targeting has been unambiguously good for Canada's economic performance, it is surely even more difficult to conclude that it has been damaging. The most negative comprehensive set of results are probably those of Ball and Sheridan (2003) who conclude only that inflation targeting does not matter. And their results have understandably come into question. To my knowledge, there is no compelling empirical evidence that inflation targeting in general, and specifically in Canada, has actually been bad for economic outcomes.

I have already said that any monetary-policy framework based on a credible commitment to low inflation and a flexible exchange rate is bound to look a lot like inflation targeting. So, if the Bank of Canada were to dispense with inflation targeting, what would fill its place? Any policy without a long-run objective of controlling inflation flies in the face of too much sensible theoretical and empirical research from the past half-century. Any policy that embodies some form of fixed exchange rate must stand up to an almost equally impressive barrage of arguments, especially for an economy as open as Canada's. What options does this leave?

My own reading of the empirical evidence leads me to conclude that inflation targeting has been a success in Canada and elsewhere, and that Canada's improved economic performance over the past 15 years is at least partly due to improved monetary policy. But even for those who

read the evidence differently, there is a compelling case for extending Canada's inflation targets based on the logical coherence of the overall framework.

The logical coherence of inflation targeting is based largely on the recognition of the long-run neutrality (and perhaps super-neutrality) of money.⁵ This is the basis for a policy that states its primary objective in terms of the long-run control of inflation rather than the long-run control of other macroeconomic variables such as output or unemployment. As Bernanke and Mishkin (1997) say:

The rationale for treating inflation as the primary goal of monetary policy is clearly strongest when medium- to long-term horizons are considered, as most economists agree that monetary policy can affect real quantities, such as output and employment, only in the short run. (p. 100)

Many central banks hesitate to speak openly of the neutrality or super-neutrality of money, and rarely even say that in the long run their policies cannot affect real variables like output and employment. Instead, they often focus on the disruptive effects of high inflation on the efficient operation of the price system and then stress that monetary policy can best contribute to the well-functioning of the economy by keeping inflation low and stable. As Goodfriend (2003) says:

A priority for low long-run inflation derives not so much from a belief in its intrinsic value relative to other goals such as full employment and economic growth, but from theory and evidence suggesting that monetary policy encourages employment and growth in the long run mostly by controlling inflation. (p.323)

Whatever the precise language used to defend inflation targeting, and whether or not one openly endorses the proposition that money is either neutral or super-neutral in the long run, the tremendous advantage of inflation targeting is that it places the focus of policy on the long-run control of inflation. It also focuses the public debate on long-run outcomes and on what outcomes are genuinely achievable for the central bank. Like Mishkin (2004b), whose thinking is clearly guided by the idea of long-run neutrality, such a long-run focus embodied in inflation targeting strikes me as very healthy.

⁵ See McCallum (2004) for a recent and insightful discussion of long-run monetary neutrality.

The logical coherence of inflation targeting, while grounded in its long-run focus on inflation, is also apparent in its treatment of short-run stabilization. I disagree with Friedman (2004) that inflation targeting is too inflexible and forces an excessive focus on the stabilization of inflation at the expense of other things that we care about, like the stabilization of real output growth. My sense from listening to policy discussions in the Bank of Canada over the past several months is that, while there is unquestionably a commitment to getting inflation back to target over the medium term, there is an equally clear concern for the path of real output, the output gap, unemployment, and other real variables in the short run. I simply do not see that the long-run focus on inflation comes at the expense of short-run stabilization. Indeed, the mainstream view at the Bank of Canada of how monetary policy works is that changes in the output gap play a crucial role in determining changes in future inflation. Monetary policy is geared to keeping inflation close to its target by keeping output close to potential. This is about as close to Blanchard's (2003) "divine coincidence" as we are likely to get in practice.

If inflation targeting is successful in anchoring inflationary expectations, and I think the evidence on this point is quite strong, monetary policy may be even more effective at stabilization policy than it would be in a non-IT world. The anchoring of inflation expectations means that central banks have greater flexibility in responding to either demand or supply shocks, secure in the knowledge — as long as they constantly affirm their commitment to the inflation target — that inflationary expectations will not be easily dislodged and a wage-price spiral ignited. As Bernanke and Mishkin (1997) say when stating how the "constrained discretion" displayed by inflation-targeting central banks allows them to deal with supply shocks:

Relative to a purely discretionary approach, the inflation-targeting framework should give the central bank a better chance of convincing the public that the consequences of the supply shock are only a one-time rise in the price level, rather than a permanent increase in inflation. (p. 107)

Finally, the logical coherence of inflation targeting means that the objectives of monetary policy can be easily communicated to the public, and the past performance of the central bank can be relatively easily measured. The use of inflation targeting does not make it any easier to explain to the public the details of how monetary policy works, but the long-run focus on inflation implies that, in the face of any new shock to the economy, the public will know that the

relevant question is: what does this imply for future inflation, and thus for the central bank's policy instrument? In an economy where the central bank has not established a clear target for inflation, the public's focus will rarely be this sharp.

To sum up, there is a strong case for extending Canada's inflation targets beyond 2006. I have no clear views on whether the targets should be extended by four years, or six, or eight, or by some other number. One problem with choosing a relatively short extension is that it may unintentionally send a signal that the Bank is considering a change in the near future. On the other hand, a relatively long extension may reduce the pressure on the Bank to re-examine the operation of the regime. One advantage of having to face a potential renewal of the targets is that the Bank is forced to take stock, to re-evaluate, and to think about how things can be improved. It is good to go through this process on a regular basis.

Section 2. Two Possible Amendments to the Inflation Targets

The overall concern for the ongoing costs of (even low) inflation and for uncertainty regarding the long-run path of the price level suggests two possible amendments to Canada's inflation targets — the adoption of price-level targeting and the reduction of the inflation target to one percent. Each policy offers possible benefits and costs, but our knowledge about both policies is still incomplete. Here I merely outline the key issues and raise a few questions.

A. Is It Time to Target the Price Level?

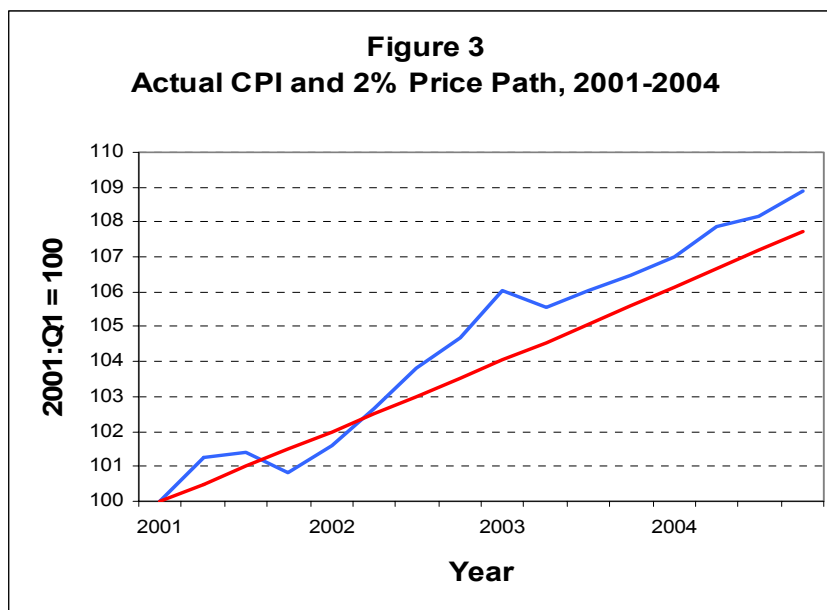
With conventional inflation targeting, a shock that pushes inflation away from target is met by the central bank with a response designed to pull inflation back to target, but the accumulated change in the price level is ignored. In contrast, with price-level targeting, in which the price-level path has a designated "drift" rate, a similar shock is met by a more severe policy response — one that reverses the excessive change in the price level and returns it to its target path.

For long-run certainty regarding the purchasing power of money, it is the long-run behaviour of the price level that matters, not the behaviour of the annual inflation rate (Stuber,

2001). Since the long-run variance of the price level is greater under inflation targeting, price-level targeting is thought by many to be the preferable policy. As Howitt (2000) argues:

Whether or not monetary policy moves all the way to price-level stabilization without drift, I share the opinion that long-term price-level uncertainty is one of the most serious consequences of inflation, because of its ruinous effects on long-term contracting. Inflation targeting with no error correction, and especially with a non-zero inflation target, does little to alleviate this problem. (p.261)

With the 2001 renewal of the inflation targets, the Bank of Canada placed more emphasis on the midpoint of the target band and less on the band itself. This change is evident in recent issues of the *Monetary Policy Report* in which the text almost always describes the near-term return of inflation to the 2-percent target; the 1-3 percent range is rarely mentioned. As Crawford (2001) explains, a systematic commitment to return the inflation rate to the midpoint of the target range leads to greater long-run certainty in the price level. Indeed, if shocks to inflation are drawn from a symmetric distribution (around zero), returning inflation to the target midpoint will, over the long run, produce an average inflation rate very close to 2 percent.⁶ In this case, inflation targeting may yield as much long-run certainty over the price level as would be obtained under a policy of price-level targeting (with 2-percent drift).



⁶ Maintaining such a long-run average also requires a commitment by the Bank to respond symmetrically to positive and negative inflation shocks.

Figure 3 shows the path of the actual headline CPI from the first quarter of 2001, together with a 2-percent target path. The CPI has been above the 2-percent path consistently since the second quarter of 2002, with a gap of just over 1 percent by the end of 2004. A 1-percent deviation from the path after only three years is fairly large, although the Bank may understandably need a longer horizon over which to have fully offsetting positive and negative shocks. But in any event it is clear that long-run uncertainty over the price level continues to exist, especially since there is no guarantee that future inflation shocks will be drawn from a symmetric distribution. While the Bank's commitment to return inflation to 2 percent should be seen as a clear improvement in policy, there is still an argument for the Bank to go even further — to target the price level.

Two main arguments exist against price-level targeting. The first is emphasized by Mishkin (2000) and is based on the possibility of financial crises. His concern is that financial crises can be ignited by deflation, and that such outcomes are more likely to occur in a world in which monetary policy is committed to undoing the excessive price-level effects of past inflation shocks. I have two concerns with this argument. First, in a world of price-level targeting with a 2-percent drift in the target path, a period of deflation would be unlikely to occur, unless the inflation shock is quite large and policy is committed to getting back to the price-level path in a relatively short period of time. Otherwise, the path taken by the price level could easily avoid deflation. Second, the potential for temporary deflation to ignite a financial crisis should be limited in a world in which expectations of (positive, but low) inflation are firmly anchored. After all, if the central bank is able to convince the public of its objective regarding the price level, then any deflation that occurs following a positive inflation shock should be viewed as only a temporary adjustment. Thus, while sustained deflation may indeed present significant threats to the financial system, this is not a likely danger in a regime of credible price-level targeting with a positive drift rate.

The second argument against price-level targeting is often described as the conventional wisdom — that the need for returning the price level to its target path will increase the short-run volatility of both output and inflation as compared to a policy of inflation targeting. To the extent that we care about short-run output and inflation volatility, and perhaps less about long-run price-level uncertainty, there may be a clear policy choice in favour of inflation targeting.

This conventional wisdom has been challenged by Svensson (1999) on the basis that if central banks with an inability for commitment have an inflationary bias, as in Kydland and Prescott (1977), then price-level targeting may involve a “free lunch” as compared with inflation targeting, bringing lower long-run price-level uncertainty without increases in short-run output or inflation volatility. Howitt (2000) argues that the game-theoretic setting in which this free lunch occurs is unrealistic. If anything, central banks these days go to great lengths to do what they said they would do — keep inflation low and relatively stable. Greater independence of central banks has surely had something to do with this, as has the greater transparency to which many of them are committed as integral parts of their inflation-targeting frameworks. And as Howitt argues, if we treat central banks as if they have the power to commit to whatever inflation they announce, the free lunch disappears.

Of more relevance to the debate is the examination of “hybrid” policies that lie somewhere between pure inflation targeting and pure price-level targeting. Barnett and Engineer (2000), Srour (2001), and Cecchetti and Kim (2003) show that if the social loss function includes separate terms for the volatilities of output, inflation, and the price level, then the optimal monetary policy will be a hybrid policy — some amount of price-level shocks will be reversed but less than would be the case under strict price-level targeting. Not surprisingly, the optimal policy depends on the underlying parameter values in the loss function and other aspects of the model, such as the degree to which expectations are forward-looking and the degree of output persistence.

Even if such a hybrid policy was eventually shown to be preferable to pure price-level targeting, it may be a difficult policy to communicate clearly to the public. As Howitt (2000) argues, however, the operation of such a hybrid policy need not look very different from inflation targeting as it currently works. It could be operationalized simply by the central bank committing itself to being a little more aggressive than usual when the price level has been pushed away from its path.

The definitive best choice between inflation targeting and price-level targeting is yet to be determined, and we need to know more before making this choice. Perhaps most important is to determine the appropriate relative weighting on the volatilities of inflation, output and the price level. This set of issues deserves careful thought. Long-run certainty in the price level is too important for either pure or partial price-level targeting to be easily dismissed.

B. Is It Time to Reduce the Inflation Target?

After a decade or more of keeping inflation relatively stable around a 2-percent target, it may be time to consider reducing the target rate of inflation. But the case for such a reduction has yet to be made. It would be unfortunate, however, if complacency won out and the real effort to push the frontiers of our knowledge on this matter was not expended. Given the success of the current system, such complacency might be understandable, but it nonetheless ought to be resisted.

Given the upward bias in the CPI, which Crawford, Fillion, and Lafleche (1998) determine to be approximately 0.5 percent, targeting a zero rate of measured CPI inflation is probably not the place to start. And the obvious attraction of whole numbers suggests that a target of 0.5 percent has some drawbacks. But this leaves an inflation target of 1 percent as a perfectly credible option, and one worthy of consideration. Would there be positive net benefits from reducing the inflation target to 1 percent?

1. Benefits of Lower Inflation

There is little or no compelling evidence that a reduction in the long-run rate of inflation would lead to an increase in the long-run rate of aggregate or per capita output growth, especially a reduction from an already-low rate of inflation (Ragan, 2000). Unfortunately, there is also little evidence that a reduction in inflation would not lead to an increase in growth; the overall relationship between inflation and growth in moderate-inflation countries is difficult to estimate precisely. But new data are generated every quarter, and estimation methods are continually improving, so this issue ought to be an ongoing topic of research. Given the large long-run effects from even very small changes in annual growth rates, it is certainly an important potential source of benefits.

Even in the absence of a desirable growth effect from lower inflation, however, there are likely to be benefits from the more efficient operation of the price system. Though these benefits will often be invisible in much of the economic data that we have at our disposal, they are nonetheless real (Ragan, 1998). The problem is that until we build computable dynamic general-equilibrium models with multiple sectors, relative prices, inflation, and uncertainty, we are

unable to get a sense of the magnitude of these benefits. Such a modelling exercise is not easy, but it should be attempted. It is important to know the magnitude of these efficiency gains.

The upshot is that we don't really know the benefits we stand to gain from reducing inflation further. But there is good reason to think these benefits exist, even if they are not easily observable in published data, and thus it is important that ongoing research effort be directed to measuring their probable size.

2. Transition to Lower Inflation

Few would dispute that the Bank of Canada has acquired a considerable amount of credibility over the past decade. It established a set of inflation targets and has met them consistently for a dozen or so years. One indication of this acquired credibility is that inflation expectations appear to be well-anchored very close to the inflation target.

The degree of central-bank credibility has played a starring role in the ongoing debate regarding the costs of disinflation. For example, the debate in the early 1980s centered on expectations formation and the credibility of the central bank's policy announcement. At one extreme was the argument that a deep and protracted recession would ensue because expectations were largely backward looking, in part because the central bank had little credibility. At the other extreme was the argument that any recession would be small and brief because expectations were largely forward looking and would adjust quickly to the central bank's announced policy shift. As it turned out, the truth was somewhere in the middle — there was a deep but relatively brief recession, and inflation fell very quickly.

If the Bank of Canada has indeed acquired considerable credibility over the past decade, it is difficult not to predict that any clearly announced decision to reduce the inflation target to 1 percent would lead to a quick adjustment of inflation expectations. We should expect this adjustment to be considerably smoother than the adjustments that took place both in the early 1980s and in the early 1990s during the first three years of the inflation-targeting regime. And with a quick adjustment in inflation expectations, there ought to be only very small transitional costs. In other words, the heightened credibility of the Bank should ensure that the sacrifice ratio associated with any reduction in the inflation target is very small.

The problem with this sensible conjecture, as reported by Fortin (2004), is the absence of supporting empirical evidence, at least for Canada. In short, there is little evidence that the nature of wage- and price-setting behaviour has changed in such a way to reduce the sacrifice ratio. Consistent with this lack of evidence is the finding of only a limited role for forward-looking expectations in estimated Phillips curves, although the growing literature on New Keynesian dynamic general-equilibrium models seems to be promising in this regard.

In short, while sensible theory suggests that the transition costs for a small disinflation ought to be small in a setting where the central bank is widely viewed as being very credible, we do not have clear empirical evidence pointing in the same direction.

3. Costs of Lower Inflation

Three main costs of lower inflation have been identified in the debate, and some are explored in other conference papers. The first is that a lower inflation target increases the likelihood that monetary policy will at some point be constrained by the zero lower bound of the nominal policy interest rate. For two reasons, I am not entirely convinced by this potential cost of low inflation.

First, hitting the zero bound seems an unlikely possibility in Canada with an inflation-targeting regime, unless very large shocks occur. Over the past dozen years, with the rate of inflation low and relatively stable, the nominal policy rate has never been below 2 percent. Subtracting 100 basis points to reflect what we (presumably) would have observed in a world with a 1-percent inflation target still leaves a buffer of no less than 1 percentage point. Given that the past decade has witnessed several negative shocks, this buffer may be sufficient.

Second, even in the unlikely event that a large enough shock occurred that the policy rate hit zero, the Bank of Canada has other means of injecting liquidity into the financial system. The Bank could enter the market and directly purchase government securities, thus providing cash directly to financial institutions or other bond-holders, all of whom may admittedly require large increases in bond prices to convince them to sell. I can understand why the Bank, having developed its approach to monetary policy by emphasizing the overnight interest rate and de-emphasising open-market operations, would be hesitant to do this. But in a situation deemed to

be critical, the Bank could clearly announce what it is doing and why, and once again affirm its commitment to low and stable inflation.⁷

A second possible cost from lower inflation is emphasized by Mishkin (2000) and is based on the dangers of the deflation that may occasionally occur in a world in which the inflation target is lowered to 1 percent. For obvious reasons, this is a better argument against a 0-percent inflation target than against a 1-percent target. But as I argued above, I am also not convinced that deflation is a dangerous phenomenon when it happens temporarily within a setting in which a positive inflation rate is being targeted. For deflation to be a problem it must be sustained for a considerable period, which in turn suggests that inflationary expectations would have to lose their anchor.

The third possible cost from low inflation involves the possibility of downward nominal-wage rigidity. Fortin (1996, 2004) has been the most vocal proponent of this view in Canada, based to a large extent on the arguments of Akerlof, Dickens, and Perry (1996, 2000). The first version of this theory is based on an explicit resistance to nominal wage cuts, based on notions of fairness or workers' morale. The second version is based on the idea that inflation is so low that workers and firms tend to ignore it in their calculations. In both versions, a permanent effect on output and unemployment results from very low inflation.

Part of the appeal of this theory is the element of believability in the idea that in a low-inflation environment workers and firms may come to think more in nominal terms than in real terms. In fact, some central bankers appear to believe this as well when they define price stability as a rate of inflation low enough that agents come to ignore it in their planning and purchasing decisions. But the insight of Akerlof, Dickens and Perry (2000) is that even if it is privately optimal for individuals to ignore these computations because inflation is low and the costs of ignoring price changes are small, the aggregate effects of such "rational ignorance" may be significant.

Unfortunately, the empirical evidence is not compelling in either direction, at least not to an agnostic like me, and thus is unlikely to change the views of either true believers or genuine sceptics. Clearer and stronger results will be needed before this debate is resolved.

⁷ See Amirault and O'Reilly (2001) for an extensive review on this topic that concludes that the zero nominal bound is not a significant risk with a 2-percent inflation target, and may not be much riskier with even lower inflation.

There are compelling theoretical reasons to expect efficiency benefits from reducing the inflation target to 1 percent. But we currently have no clear method of estimating the magnitude of these benefits. There are good reasons to think that the transition to a lower inflation target would be relatively painless, but the empirical support is absent. Finally, there are some reasons to expect some costs from a low inflation rate, but only shaky empirical support.

In short, we have a pretty good argument at this point not to reduce the inflation target. But we have an even better argument for trying to fill the gaps in our knowledge, especially by pushing further on the empirical front (O'Reilly, 1998). My own instincts, all things considered, are that the benefits of a 1-percent target would dominate the costs, and that the transition costs of getting there would be small. But we should not be making such decisions on instinct.

Section 3. Further Developments in Communications

The Bank of Canada has made impressive advances in communications over the past 15 years. The obvious list of accomplishments includes: the adoption of explicit inflation targeting in 1991; the bi-annual publication of the *Monetary Policy Report* beginning in 1995; the adoption of Fixed Announcement Dates (FADs) in 2000; the increasing publication of forecasts (for real growth and inflation) in the *MPR*; the modification of the *Bank of Canada Review* to include more articles explaining monetary policy; and the increasing number of public speeches by the Governor and Deputy Governors.⁸

These advances in communications are quite apart from the other significant changes in the Bank's operating procedures, including the elimination of required reserves, the creation of the Large Value Transfer System, the de-linking of the Bank rate from the Treasury Bill market, and the increased emphasis on using the target for the overnight rate as the focal point for monetary policy. Canadian central banking has changed dramatically over the past twenty years. Kuttner (2004) argues that, while there have been many operational changes in the way most

⁸ Mishkin and Posen (1997) offer a detailed discussion of the early years of inflation targeting in Canada.

inflation-targeting central banks do their job, it is probably the communications that have changed the most.

Most of the elements of the Bank's current communications strategy would have been considered unthinkable twenty years ago, and would still have made many people uncomfortable as recently as ten years ago (and perhaps still make some people uncomfortable today). As the benefits of central-bank transparency and accountability began to sink in, these communications strategies were adopted and now they are entrenched in the way the Bank thinks about its responsibilities. Most people inside the Bank have a hard time believing that things used to be done so differently.

Despite these advances, I believe there is more work to be done. My primary recommendation discussed below may be viewed today as unthinkable, undesirable, or infeasible. But it is none of these; it is no more than the natural next step in an overall communications strategy for a central bank that takes transparency seriously.

I argued above that one of the benefits of an inflation-targeting regime is the long-run nature of the policy objective, and the associated long-run focus both inside and outside the central bank. My comments on central-bank communication are driven by a belief that this long-run focus on the objective of monetary policy should be matched by an equally long-run focus in the way the central bank communicates with the public. Just as an excessive focus on short-run economic developments would be inappropriate for the conduct of monetary policy, an undue focus on the short run is inappropriate for communications.

A. The Need for Central-Bank Communications

Imagine a world in which the Bank of Canada has a clear long-run inflation objective, has achieved that objective for several years, and has therefore built up a considerable amount of credibility with the private sector. To this realistic description of the Canadian setting for monetary policy, add three hypothetical conditions:

C1. The Bank has a clear and coherent view of the structural equations for the world and Canadian economies. The public and financial markets have the same view.

C2. The Bank has a clear and coherent view about the transmission mechanism of monetary policy. The public and financial markets share this view.⁹

C3. The Bank has access to a large body of global and domestic economic data with which it conducts its analysis. The public and financial markets have access to the same data.

In this hypothetical world there is plenty of room for uncertainty about future economic events; neither the Bank nor the private sector have a crystal ball. Furthermore, even if nothing unexpected were to occur, there is also plenty of room for policy mistakes — the shared view of the structural equations may be incorrect, the shared view of the transmission mechanism may be wrong, or the (shared) data may be mistaken. Therefore, due to either an uncertain world or a mistaken view about how the world works, monetary policy can make plenty of mistakes.

The key aspect of this hypothetical world is that, whatever possibly incorrect view of the economy guides the Bank's thinking, or whatever possibly incorrect data get employed in the Bank's analysis, the private sector suffers exactly the same problems. In such a world, barring an unusual situation in which the Bank decides to change its policy objective, there is no reason for the Bank to communicate at all to the public. Given the available economic data and the view of how the world economy and monetary policy function, the Bank can conduct its analysis and determine its preferred policy actions. The private sector, with access to the same data and the same views of economic relationships, will reach the same conclusion. There is little or no reason for the Bank to say anything to anybody.

Central bankers often talk about the need for communication to promote transparency and accountability, but these are just means to an end. Presumably the ultimate goal is to have better economic performance, achieved by having a better-understood policy framework and thus fewer surprises for agents outside the central bank. Given the thought experiment above, the fundamental reason for central-bank communication is that one or more of the three conditions above do not hold. Probably none of the three conditions hold exactly, although C3 must come pretty close since the data possessed by the private sector — especially the financial institutions which publish their analyses frequently — is approximately as complete and accurate as that possessed by the Bank of Canada. If C3 is approximately true, then the primary reason for the Bank of Canada to communicate with the private sector is that its own views about the

⁹ Having the same view in C1 and C2 probably requires having similar analytical tools and the ability to use them.

functioning of the economy and about monetary policy differ from the views held by the private sector — C1 and C2 do not hold. Is this reasonable?

My assessment is that there are still significant gaps between the Bank's view and the private sector's view about basic economic relationships and about monetary policy. Driving this assessment, and these are admittedly very casual observations, is the steady stream of media stories that have come across my computer screen over the past eight months, including stories from reputable national newspapers and newsletters from economics departments in commercial banks and other financial firms. Some of these reports are very good, and I read them with interest. But a large fraction of the stories are based on views of economics and monetary policy that are inconsistent with the way these issues are viewed by the majority of economists within the Bank.

Given the violation of C1 and C2, the Bank can use two general types of communication:

1. it can signal to the private sector what it is likely to do in the near future, either through written statements or by publishing policy-rate forecasts; and/or
2. it can explain to the private sector its views of economics and of the functioning of monetary policy, and let the private sector infer the Bank's likely future policy actions.

Within the Bank of Canada, the first type of communication is sometimes called “transparency of the policy path” and the second is “transparency of the response pattern”, although this second label suggests an emphasis more on the workings of policy than on the functioning of the world and Canadian economies. Kohn and Sack (2004), in their study of the impact of the U.S. Federal Reserve's policy statements, focus their attention on the first type of communication. They decompose the Fed's policy statements into two mutually exclusive and exhaustive components — the short-run “policy-inclination” component and the (slightly) longer-run “economic outlook” component.

My classification of central-bank communication comes closest to that used by Issing (2005), who emphasizes the difference between “short-term signalling” and “longer-term consistency”. In what follows, I adopt his term of short-term signalling to describe the first type of communication, but refer to the second type as “longer-term education” to better capture the broader flavour of my arguments.

There is a clear sense in which these two types of communication are substitutes for one another. The perceived need for more active or more precise short-term signalling is one measure of how much work remains to be done to improve longer-term educational communications. If the public and financial markets really understood the way the central bank thought about the functioning of the economy and monetary policy, there would be little need for any short-term signalling.¹⁰ The Bank of Canada should probably not be encouraged to use only a single approach. However, using the Bank's current communications strategy as a reference point, I advocate placing more emphasis on longer-term education and less on short-term signalling.

In the remaining parts of this section I do three things. First, I lay out a case against making further advances in short-term signalling, such as more precise policy-inclination statements or the publication of policy-rate paths. Here I agree with Mishkin (2004a) that this type of communication is likely to do more harm than good. Then I describe three specific speeches or lectures that would be very useful for the Bank's Governing Council to write and deliver to the public. Finally, I address the issue of whether members of the Governing Council should be encouraged to disagree with each other in public; the educational side of me sees some clear advantages in disagreements of a particular type. In the paper's final section, I address some of the possible objections to my recommendation.

B. Less Short-Term Signaling

A large literature has emerged over the past decade about central-bank communications, and much recent interest has focused on whether central banks should make policy-inclination statements that inform the private sector with some precision about the likely future path of policy rates. There is also discussion of the benefits of providing even more precise signals, as would be the case if the central bank published a near-term forecast of its policy-rate path. I view these two options as similar types of short-term signalling — one being just a more precise version of the other.

¹⁰ Of course, one can also argue the reverse — that with sufficient short-term signalling, the private sector may perceive no need for longer-term education. Amato, Morris, and Shin (2002) argue that such an emphasis on short-term signaling is undesirable, as I do later in the paper.

Central banks issuing policy-inclination statements must search for effective “code words” to appropriately signal their intentions (Issing, 2005). It is an interesting game being played, in which the goal of the exercise clearly is to indicate a tendency, but not to indicate a tendency clearly. As Alan Greenspan famously quipped at one hearing: “I know you believe you understand what you think I said, but I am not sure you realise that what you heard is not what I meant.” Quite apart from the game being played — which is hardly consistent with the true spirit of transparency — is that a danger exists if the central bank subsequently comes under pressure to honour the quasi-promises it has made in its statements. If the statement is well-written and accurate, and the economic outlook does not change, then honouring the promise is entirely appropriate. But it is all too easy to make mistakes in the crafting of such policy-inclination statements, precisely because the objective is constructive ambiguity, and being tied to unintended promises in these situations is a problem.

In addition, as Issing (2005), Mishkin (2004a) and others have noted, if the central bank changes its economic outlook for appropriate reasons, it may find it difficult to signal this change to the private sector. This is the problem of explaining the conditional nature of the policy-inclination statements, and it is not a small problem. I recall Doug Purvis many years ago telling me that the two most difficult concepts for students to really understand were comparative advantage and conditional probability. After 16 years of teaching McGill students, I think he is right. Despite the fact that we can use very simple and helpful diagrams in class to illustrate both concepts, and despite the fact that students can successfully complete problem sets forcing them to work through applications of both concepts, students seem to forget the basic insights once they leave the classroom and immerse themselves in the discussions of everyday life. Except for a small fraction of the population, conditional probability is quite unintuitive.

Central banks contemplating publishing their forecast of the policy-rate path also must deal with the problems of explaining the conditional nature of the forecast. But with a published data series — likely a line on a graph — this exercise becomes more difficult, and the danger is that the published series will take on a life of its own. When the Bank of Canada actively spoke about the Monetary Conditions Index and published it with some prominence in the mid 1990s, financial markets quickly came to use it as a formulaic guide for trading. Admittedly, the internal logic of the MCI is more complex to explain than is the target overnight rate, but the problem of explaining the conditional nature of the forecast is the same.

When discussing long-term aid to developing countries, we often hear that it is preferable to teach people how to grow their own crops rather than giving them food directly. In the same spirit, it may be preferable for the Bank of Canada to focus on providing longer-term education rather than more short-term signalling — to explain how the Bank views the determinants of its policy actions rather than providing the actions directly.

Short-term signalling has the potential to discourage financial markets from performing their own analyses of the economy, and encouraging an unhealthy dependence on the Bank's communications. Financial markets may then shift their focus away from thinking about economics and focus instead on interpreting the most recent statements issued by the Bank. An emphasis on short-term signalling which leads to such "laziness" in financial markets produces other dangers, as stressed by Amato, Morris, and Shin (2002). Their concern is that if financial markets come to focus on central-bank pronouncements, and thus spend less time conducting their own economic research, any errors embodied in the central bank's analysis are unlikely to be discovered. The central bank can no longer look to the private sector for an independent check on its own analysis. The result is that mistakes get perpetuated, and perhaps even amplified. Rigorous competition in the marketplace for ideas largely disappears, and with it the benefits from the greater understanding such competition would generate.

C. More Longer-Term Education

My view is that the Bank should not pursue further advances in short-term signalling, whether in the form of more frequent or more precise policy-inclination statements, or published forecasts for the policy-rate path. If anything, a little less emphasis on such signalling might be desirable.

Instead, the Bank should spend more time and effort on longer-term education, on explaining to the private sector how the Bank thinks about economic relationships and how it thinks about the workings of monetary policy. I will make the detailed case in favour of such longer-term education below. For now I lay out three speeches or lectures that would be valuable for the Bank to write and for the public to hear. There are surely other valuable topics that could be added to this list, but these three topics strike me as an obvious starting point. A few months ago there would have been a fourth topic — the role of exchange-rate movements in monetary

policy — but in the past few months the Bank has addressed this topic in a number of speeches that represent exactly the type of educational communication I have in mind (Dodge, 2005).

1. The Bluntness of Monetary Policy

Canadians are used to watching federal and provincial finance ministers announce all sorts of new initiatives. Whatever economists might think about these individual policies, Canadians are left with the clear view that the country's economic policy-makers have many levers at their disposal. When the Governor of the Bank of Canada gives a public speech on exchange rates or productivity or inflation targeting or labour-market flexibility or pension systems, many Canadians will view this policy-maker as just one among several top Canadian officials who have their hands at the policy levers. But few people outside the financial sector (and not everyone inside it) will genuinely understand the fact that the Bank of Canada has only a single policy instrument. This is not news to any economist, but it is very surprising to most people when and if they ever learn it.

The bluntness of monetary policy is especially important in an open economy subject to significant external shocks. In the face of such shocks, the domestic economy is forced to adjust. If the exchange rate moves in response to these shocks, this adjustment will proceed differently than if the exchange rate were fixed. But adjustment will occur in either case, and typically some sectors will be adversely affected by any given shock while others will benefit. The problem for monetary policy is that its bluntness forces it to be geared toward the aggregate of these economic developments; it is helpless to play any systematic role in assisting some sectors at the expense of others. Fiscal policy may be able to help in this way, but not monetary policy.

This point is crucial and is insufficiently understood by the private sector. If this point were as widely understood as it should be, I would not hear comments from very sophisticated business leaders such as “I agree with the Bank's commitment to low inflation, but surely it can do something about the appreciating Canadian dollar”.

2. The Transmission Mechanism and the Importance of Lags

Economists often speak of the long and variable lags in monetary policy. If asked to explain why these lags are long, we respond by describing the various linkages within the monetary transmission mechanism. When asked why they are variable, we go back to the transmission mechanism and explain why the underlying behaviour often seems to respond to changes in the environment that are difficult to capture in a mechanical way. This is usually where “animal spirits” get mentioned in my classes.

The most important implication of the lags in the transmission mechanism is that monetary policy must be forward looking. If central banks wait for inflation to increase before tightening policy, their actions will be too late and risk making the economy less, rather than more, stable. This essential point is under-appreciated in many media analyses of monetary policy. Too often the stories place emphasis on recent data releases, such as the latest monthly unemployment figures or the latest monthly inflation or trade data. The problem is that data releases of this type should have a very limited impact on current policy actions, except to the extent they provide a more accurate estimate of what is likely to occur over the next 1-2 years. Consider one simple example.

Suppose the economy is operating with actual output below but close to potential output but the recent data show no signs of undue inflationary pressures, either in wages or prices. Some observers will argue that the absence of inflationary pressures is a strong argument against any monetary tightening. But any tightening of monetary policy today is designed to fight future inflation, not current inflation. So unless the current absence of wage or price inflation gives you reason not to expect inflation in the future, it is not a powerful part of an argument against tightening.

The public and the financial markets need to place more emphasis on how movements in current data are likely to influence our expectations of the future. And the Bank needs to do its part in explaining in detail with clear and timely examples how this works and why this matters. The obvious advantage of a speech on the importance of lags is that it forces a discussion of the transmission mechanism itself. And this sort of discussion would be very difficult to have without also recognizing the uncertainty inherent to the workings of monetary policy. Thus a coherent discussion of lags will inevitably be a discussion of how monetary policy works.

3. The Role of Oil-Price Shocks

Kuttner (2004) has argued that inflation-targeting countries have not yet been fully tested, as the last decade has not contained the kind of dramatic, negative supply shocks that we experienced in the 1970s. With the world price of oil currently above \$50 (U.S.) per barrel, perhaps it is only a few years before the inflation-targeting regimes are truly stress-tested. How will central banks respond to large supply shocks that increase inflation while reducing growth?

The last sentence is how the press tends to characterize the likely effects of a significant increase in the price of oil, especially for the U.S. economy. But for Canada the story ought to be different, for two reasons: we are a net exporter of oil, and we depend to a large extent on exports of all kinds to a large oil-importing country.

A public speech on the effects of an oil-price shock ought to emphasize the distinct and opposing effects on aggregate demand and aggregate supply, how these effects net out given our position as a net exporter, and how the negative effect on U.S. income generates an additional negative effect on Canadian aggregate demand. In addition, it would be useful to explain how the positive effect on aggregate demand seems to be increasing over time and how greater efficiency in oil use implies that the negative effect on aggregate supply is falling over time. Finally, some attention should be paid to the difficult issue of the uncertain timing of many of these effects.

Such an explanation of how oil-price changes affect the Canadian economy is only possible if the distinction is made between the demand side and the supply side of the economy, a distinction that is rarely evident in current Bank speeches. This takes a little work, but is worth the effort. Once this distinction is made, the various pieces of the complex story can be put together and the public will have a better understanding of how the Bank thinks about oil-price shocks. The overall effects of the shock on output, the output gap, and inflation then leads to a discussion of the Bank's likely policy response.

One further point should be included in this speech: the importance of well-anchored inflationary expectations. Major input-price shocks have always presented central banks with a challenge because validation of the shock risks setting off a wage-price spiral. But to the extent that inflation targeting is successful in anchoring expectations, this policy dilemma is largely avoided: the central bank can validate the shock — arguing credibly that the shock need only

represent a one-time increase in the price level — while affirming its commitment to low inflation.

4. Common Elements

Every speech or lecture of this “educational” variety will be most effective if it shares two elements in common with other such speeches. First, each one should contain some discussion of the overall framework for inflation targeting, including some detailed discussion of the monetary transmission mechanism. The amount of detail can vary depending on the speech, but these elements are essential in order to place the broader message into perspective.

Second, the goal of these speeches is to improve the Bank’s longer-term educational communications, and thus should not be combined with short-term signalling. I therefore recommend not including a section on the short-term economic outlook in these speeches. First, unless there is genuinely new information to announce since the last *Monetary Policy Report (or Update)*, there is little value added in these statements. Second, the press focuses on the short-term signalling statements and tends to disregard broader messages. This occurred in roughly 80 percent of the media stories following the Governor’s Vancouver speech in January 2005 on the role of the exchange rate — a speech that, aside from the section on the economic outlook, was a clear example of the kind of educational communications I am advocating. By having a speech that contains only a longer-term educational component, the reporter who wants any story about monetary policy is forced to think about the speech, figure out the central message, and then explain it to the readers. This is precisely the objective.

D. Should Members of the Governing Council Publicly Disagree?

Senior officials at the U.S. Federal Reserve openly disagree with each other on some issues. Recently, for example, some have come out in favour of the adoption of inflation targeting while others have taken the opposite position. In Canada, no such public disagreement is apparent among the six members of the Governing Council (GC). Would a little disagreement be useful?

One danger with public disagreement among the GC members is that some may come to be known as “hawks” while others come to be known as “doves”. Once such labels get applied, all careful thought is in danger of vanishing as the public comes to place more importance on the messenger than on the economic analysis contained in the message (Issing, 2005). We suffer the same problem within our profession when the labels of “Keynesian” or “Monetarist” are sometimes used, even though the precise meaning is usually difficult to detect from the context. The problem with labels is that they substitute for, rather than complement, effective communications.

Could disagreement among the GC members ever improve the Bank’s overall communications? If done carefully, I think it could be very valuable. In the second speech above I recommend laying out the details of the monetary transmission mechanism, and emphasizing both the length and the variability of the time lags. The Bank faces plenty of uncertainty in this mechanism, including uncertainty about various elasticities, the causes of changes in variables, whether shocks are temporary or permanent, and the confidence in data.

This uncertainty provides an ideal place for some constructive disagreement about economics. Laying out the nature of the uncertainty, and why it matters for monetary policy, underlines the functioning of the overall transmission mechanism. And disagreement over some of the component parts of the transmission mechanism emphasizes to the public an important truth — that much remains to be learned and there are some legitimate debates involved in the process of learning. Some GC members, for example, might argue that the extent of exchange-rate pass-through is likely to remain low; others might argue that the recent lows are likely to be reversed in the near future. Both could draw out the implications for monetary policy. It is important in such cases, however, that GC members recognize the other side of the argument, thus clarifying for the audience where the point of contention lies. If done carefully, the result will be a greater appreciation for that element of the transmission mechanism and thus a greater understanding of monetary policy.

It is not enough, however, to have GC members explain the nature of any disagreement about the details of monetary policy. Individual speakers should reveal to the audience how their own thinking on specific issues or debates has evolved — how they used to think about some aspect of monetary policy, how their mind was changed by new arguments or new data, why they now believe that some past policy action was a mistake, and what they would do to avoid

that mistake in the future. If the GC members are willing to reveal this much to the audience, the public will benefit by seeing the thought process behind monetary policy, and their understanding will increase significantly. As students, we all know that we don't learn simply by listening to one position explained — we need to see an engaged debate between opposing positions.¹¹

Section 4. The Monetary Road Less Traveled

I have come out in favour of extending the inflation targets in Canada. The evidence points toward giving monetary policy some of the credit for the improved macroeconomic performance observed since the early 1990s. Moreover, the logical coherence of inflation targeting is an obvious selling point. I have also outlined two issues worthy of further consideration: the case for price-level targeting and the case for lowering the inflation target.

Even if the Bank extends its current system of inflation targets, however, it has a decision to make regarding its overall communication strategy. It can push ahead on short-term signalling or it can push ahead on longer-term education. And though I have argued against further development in short-term signalling, I have not yet made a compelling case in favour of longer-term education. I conclude the paper by building this case as responses to the two arguments opponents are likely to make: such longer-term education cannot be provided by the Bank, and such longer-term education should not be provided by the Bank.

A. Can the Bank Do It Right?

I come from a profession that believes in the possibility of education, to almost any imaginable audience, as long as enough thought and effort is put into the exercise. And there is no doubt that this exercise requires a great deal of effort — in clear thinking, honest debating, and careful writing. But I have every confidence that if senior members of the Bank believe in the value of providing such longer-term education, they can do a good job.

¹¹ By law in Canada, monetary policy is the responsibility of the Governor, not the other members of the Governing Council. This lack of formal responsibility by the five Deputies is more liberating than constraining — it should give them considerable freedom to be candid about their own views on the workings of monetary policy.

The audiences for these speeches will usually be less familiar with economics and with economic data than is the GC member giving the speech. It is crucial that the Bank keep this in mind, for many senior members of the Bank probably are guilty of believing that the issues they think about on a daily basis also come easily to people outside the Bank. This belief is false, so attention must be paid to careful pedagogy. Proper education of the kind I am advocating requires an accurate portrayal of some of the complexities and uncertainties that confront us as economists and policy-makers. But as Issing (2005) notes, it is crucial to strike the right balance between a clear and simple message and the complexities of the real world. Getting this balance right is not easy, and it requires genuine effort, and several drafts.

What worries me more than the Bank's ability to provide this longer-term education is the lack of interest individual members of the GC might have for the project. I hope I am wrong with respect to the GC, although my general concerns in this direction are well-grounded. Many economists, and far too many of them inside academia, are much more interested in explaining their own higher-level research interests to their professional colleagues than in explaining the importance of more basic economics to non-specialists. This has always puzzled me, for many of these same people will recall with precision the first time they fell in love with economics, typically coincident with the sudden illumination of some previously shadowy concept. I hope all senior Bank people can come to the view that there is genuine pleasure in being the one who flips the switch.

B. Should the Bank Do It At All?

For several reasons the Bank should spend the required time and effort to provide better longer-term educational communications. First, this education should not just be left to those of us in colleges and universities. While I am more than happy to do my fair share of teaching economics, the Bank has a clear interest in doing this as well. In fact, the Bank has a clear advantage over most of us — there is surely no group of people outside the Bank better placed to explain to the world how the Bank thinks about the functioning of the economy and the workings of its monetary policy.

Second, by making the effort to explain the Bank's view in greater detail to the public, it will soon become clear what holes exist in the Bank's thinking. And there are many such holes,

even though their existence may be unpleasant to acknowledge. As long as the exercise is undertaken with honesty, the existence of these holes will force those writing the speech to discuss and debate the points, and to make sure that the argument is made correctly. In some cases it will also lead to a redirection of the Bank's research effort toward filling those knowledge gaps. There is a fundamental truth that all teachers know — we learn much from the process of teaching, and typically much more from teaching those who ask questions forcing us to explain statements and beliefs that we have held unquestioned for many years.

The third reason for the Bank to provide longer-term education to the public is that it will stimulate a better debate. As the Bank lays out in a clear and systematic manner its view of how the world economy operates and of how its monetary policy works, the stage is then set for members of the public, and especially the economics departments in universities and financial institutions, to respond and criticize the Bank's view and, in turn, to offer competing views of how the world works. Disagreements may exist over such details as the length of lags or the size of specific elasticities, or over such fundamental issues as the role of money in the transmission mechanism or the determinants of potential output. Such a debate, if pursued with honesty and vigour, has the potential to improve understanding on both sides. The fostering of such a debate should be an important part of central-bank transparency. As Laidler (2001) states:

Transparency is an asset that depreciates very rapidly, so communication has to be ongoing, and if it is to enhance accountability, it must be a two-way process as well. (p.iv)

The value of an ongoing debate between the Bank and the private sector leads me to one final point. I do not believe that the Bank knows more about economics or the world economy or even about the functioning of monetary policy than does the collection of people outside the Bank. There are plenty of well-trained and thoughtful people in academia and in financial institutions that can challenge the Bank on several issues. But it is obvious that the Bank knows more than anyone else about how it thinks about the world and how it thinks about monetary policy. And this is what matters. The goal is not for the Bank to teach economics to the world; the goal is to carefully explain how the Bank thinks about economics to the world. Once these views are clearly articulated, the subsequent debates and research can lead to better understanding, both inside and outside the Bank.

From my short time at the Bank it is clear that issues of communication still generate discomfort among the senior ranks. But this is a clear improvement, for I'm sure that twenty years ago it terrified them!

Over the past decade or so, the Bank of Canada has been among the leaders in the central-bank community in pushing out the frontiers on communications, despite whatever discomfort occurred along the way. In this sense, the Bank has indeed taken “the road less traveled”, and the payoff has been large. Nobody that I know, even in their most candid moments, has ever suggested that the Bank should return to conducting or communicating its policy in the manner of twenty years ago.

But the journey is not over. If the Bank chooses to place more emphasis on longer-term education, and really takes this project seriously, I can guarantee the road ahead will be challenging. Difficult issues will be broached, gaps in the Bank's knowledge will be exposed, details will be difficult to explain, awkward questions will be asked, and the Bank will be forced to question claims that it once held with confidence. But I have no doubt that this tougher road is the right one to take — for the Bank and for Canada. Along this road lies greater understanding for both the Bank and the public, and also a finer appreciation for the important role the Bank plays in the economic health of our nation.

References

- Akerlof, G., W. Dickens, and G. Perry (1996), "The Macroeconomics of Low Inflation", *Brookings Papers on Economic Activity*, p. 1-59.
- Akerlof, G., W. Dickens, and G. Perry (2000), "Near-Rational Wage and Price Setting and the Long-Run Phillips Curve", *Brookings Papers on Economic Activity*, p. 1-44.
- Amirault, D. and B. O'Reilly (2001), "The Zero Lower Bound on Nominal Interest Rates: How Important Is It?", Bank of Canada Working Paper 2001-06.
- Amato, J., S. Morris and H. Shin (2002), "Communication and Monetary Policy", unpublished mimeo, BIS, Yale University, and LSE.
- Ball, L. and N. Sheridan (2003), "Does Inflation Targeting Matter", NBER Working Paper 9577.
- Barnett, R. and M. Engineer (2000), "When Is Price-Level Targeting A Good Idea?", in *Price Stability and the Long-Run Target for Monetary Policy*, Bank of Canada, p. 101-143.
- Bernanke, B. and F. Mishkin (1997), "Inflation Targeting: A New Framework for Monetary Policy?", *Journal of Economic Perspectives*, Volume 11, p. 97-116.
- Blanchard, O.J. (2003), "Comments on 'Inflation Targeting in Transition Economies: Experience and Prospects' by J. Jonas and F. Mishkin", NBER conference on inflation targeting, January.
- Blanchard, O. J. and J. Simon (2001), "The Long and Large Decline in U.S. Output Volatility", *Brookings Papers on Economic Activity*, p. 135-174.
- Cecchetti, S. and J. Kim (2003), "Inflation Targeting, Price-Path Targeting, and Output Variability", in B. Bernanke and M. Woodford (eds.), *The Inflation-Targeting Debate*, University of Chicago Press, p. 173-195.
- Cecchetti, S., A. Flores-Lagunes, and S. Krause (2004), "Has Monetary Policy Become More Efficient? A Cross-Country Analysis", NBER Working Paper 10973.
- Collins, S. and P. Siklos (2004), "Optimal Monetary Policy Rules and Inflation Targets: Are Australia, Canada, and New Zealand Different from the U.S.?", *Open Economies Review* 15, p. 347-362.
- Crawford, A. (2001), "Predictability of Average Inflation Over Long Time Horizons", *Bank of Canada Review*, Autumn, p. 13-20.
- Crawford, A., J.F. Fillion, and T. Lafleche (1998), "Is the CPI a Suitable Measure for Defining Price Stability", in *Price Stability, Inflation Targets, and Monetary Policy*, Bank of Canada, p. 39-73.
- Crow, J. (1988), "The Work of Monetary Policy", Eric J. Hanson Memorial Lecture, *Bank of Canada Review*, February.
- Dodge, D. (2005), "Monetary Policy and Exchange-Rate Movements", Speech to the Vancouver Board of Trade, Bank of Canada.
- Fortin, P. (1996), "The Great Canadian Slump", *The Canadian Journal of Economics* 19, p. 769-787.

- Fortin, P. (2004), "The Bank of Canada and the Inflation-Unemployment Trade-off", in *Macroeconomics, Monetary Policy, and Financial Stability: A Festschrift in Honour of Charles Freedman*, Bank of Canada, p. 3-24.
- Friedman, B. (2004), "Why the Federal Reserve Should Not Adopt Inflation Targeting", *International Finance* 7, p. 129-136.
- Goodfriend, M. (2003), "Inflation Targeting in the United States?", in B. Bernanke and M. Woodford (eds.), *The Inflation-Targeting Debate*, University of Chicago Press, p. 311-337.
- Howitt, P. (2000), "Comments on 'What Have We Learned About Price Stability?'"', in *Price Stability and the Long-Run Target for Monetary Policy*, Bank of Canada, p. 260-265.
- Hyvonen, M. (2004), "Inflation Convergence Across Countries", Reserve Bank of Australia Discussion Paper 2004-04.
- Issing, O. (2005), "Communication, Transparency, Accountability: Monetary Policy in the Twenty-First Century", *Federal Reserve Bank of St. Louis Review*, p. 65-83.
- Johnson, D. (2002), "The Effect of Inflation Targeting on the Behavior of Expected Inflation: Evidence From an 11-Country Panel", *Journal of Monetary Economics* 49, p. 1521-1538.
- Kahn, J., M. McConnell, and G. Perez-Quiros (2002), "On the Causes of the Increased Stability of the U.S. Economy", *FRBNY Economic Policy Review*, p. 183-202.
- King, M. (2003), "What Has Inflation Targeting Achieved?", in B. Bernanke and M. Woodford (eds.), *The Inflation-Targeting Debate*, University of Chicago Press, p. 11-16.
- Kydland, F. and E. Prescott (1977), "Rules Rather Than Discretion: The Inconsistency of Optimal Plans", *Journal of Political Economy* 85, p. 473-491.
- Kneebone, R. and J. Chung (2004), "Where Did the Debt Come From?", in C. Ragan and W. Watson (eds.), *Is the Debt War Over? Dispatches from Canada's Fiscal Frontline*, Institute for Research on Public Policy.
- Kohn, D. and B. Sack (2004), "Central Bank Talk: Does It Matter and Why?", in *Macroeconomics, Monetary Policy, and Financial Stability*, Bank of Canada, p. 175-206.
- Kuttner, K. (2004), "A Snapshot of Inflation Targeting in its Adolescence", in C. Kent and S. Guttman (eds.), *The Future of Inflation Targeting*, Reserve Bank of Australia, p. 6-43.
- Kydland and Prescott (1977), "Rules Rather than Discretion: The Inconsistency of Optimal Plans", *Journal of Political Economy* 85 (June), p. 473-491.
- Laidler, D. (2001), "Forward", *The Thiessen Lectures*, Bank of Canada, 2001.
- Levin, A., F. Natalucci, and J. Piger (2004), "Explicit Inflation Objectives and Macroeconomic Outcomes", European Central Bank Working Paper 383.
- Longworth, D. (2002), "Inflation and the Macroeconomy: Changes from the 1980s to the 1990s", *Bank of Canada Review*, Spring, p. 3-18.
- McCallum, B. (2004), "Long-Run Monetary Neutrality and Contemporary Policy Analysis", Institute for Monetary and Economic Studies Discussion Paper 2004-E-18, Bank of Japan.
- Mishkin, F. (2000), "Issues in Inflation Targeting", in *Price Stability and the Long-Run Target for Monetary Policy*, Bank of Canada, p. 203-222.

- Mishkin, F. (2004a), "Can Central Bank Transparency Go Too Far?", in C. Kent and S. Guttman (eds.), *The Future of Inflation Targeting*, Reserve Bank of Australia, p. 48-66.
- Mishkin, F. (2004b), "Why the Federal Reserve Should Adopt Inflation Targeting", *International Finance* 7, p. 117-127.
- Mishkin, F. and A. Posen (1997), "Inflation Targeting: Lessons from Four Countries", FRBNY *Economic Policy Review*, p. 9-110.
- O'Reilly, B. (1998), "The Benefits of Low Inflation: Taking Stock", Bank of Canada Technical Report #83.
- Orphanides, A. and J. Williams (2003), "Imperfect Knowledge, Inflation Expectations, and Monetary Policy", in B. Bernanke and M. Woodford (eds.), *The Inflation-Targeting Debate*, University of Chicago Press, p. 202-234.
- Ragan, C. (1998), "On the Believable Benefits of Low Inflation", Bank of Canada Working Paper 98-15.
- Ragan, C. (2000), "Should We Expect Higher Growth From Lower Inflation?", *Canadian Business Economics*, August, p. 19-31.
- Rogoff, K. (2003), "Globalization and Global Disinflation", paper prepared for Federal Reserve Bank of Kansas City conference on "Monetary Policy and Uncertainty: Adapting to a Changing Economy", Jackson Hole, WY.
- Srouf, G. (2001), "Price-Level Versus Inflation Targeting in a Small Open Economy", Bank of Canada Working Paper 2001-24.
- Stock, J. and M. Watson (2003), "Has the Business Cycle Changed? Evidence and Explanations", paper prepared for Federal Reserve Bank of Kansas City conference on "Monetary Policy and Uncertainty: Adapting to a Changing Economy", Jackson Hole, WY.
- Stuber, G. (2001), "Implications of Uncertainty About Long-Run Inflation and the Price Level", Bank of Canada Working Paper 2001-16.
- Svensson, L. (1999), "Price-Level Targeting Versus Inflation Targeting: A Free Lunch?", *Journal of Money, Credit and Banking*, Volume 31, p. 277-295.
- Taylor, J. (1979), "Estimation and Control of a Macroeconomic Model with Rational Expectations", *Econometrica* 47, p. 1267-1286.