# What Have We Learned About Price Stability?

*Michael Parkin**

## Introduction

This paper summarizes what we have learned about price stability since the C.D. Howe Institute and Bank of Canada conferences of 1996 and 1997.[1] The range of issues that a central bank must consider in its evolving approach to price stability, and that these two conferences covered, is broad. It includes setting targets for the behaviour of the price level, developing techniques for forecasting and influencing the price level, reporting and explaining the bank's actions, and establishing and maintaining credibility.[2]

This seminar and this paper focus on the first of these issues—setting targets. Based on the current knowledge, what kind of targets for the behaviour of the price level would lead to a monetary policy that contributes most to the economic welfare of Canadians?

When our current inflation target ends, the options that we face are: (i) abandon formal targeting; (ii) reaffirm the existing target and formally

1. See Laidler (1997) and Bank of Canada (1998).

2. It also includes deciding whether to surrender monetary independence and become part of a wider currency area, a theme that has recently attracted attention in the Canadian context. While this issue is relevant to the main question addressed in this paper, I have interpreted my brief as examining monetary policy options within a system of flexible exchange rates and monetary policy independence.

* I thank, without implicating, David Laidler for comments on an earlier draft; Tiff Macklem for pointing me towards this topic; Allan Crawford and Tiff Macklem for pulling together a small research library; William Gavin, Jennifer Smith, and Lars Svensson for helping me to understand and interpret their important work; Audra Bowlus for help with some data; Jane McAndrew for library assistance; and Robin Bade for extensive comments.

commit to it for another term—either fixed or open-ended; (iii) move further in the direction of price stability and reaffirm the existing target inflation rate, but formulated as a price-level target; or (iv) move even more aggressively in the direction of price stability and commit to lower, possibly zero inflation, formulated as a price-level target.

My reading is that the current state of knowledge decisively rejects the first option. Targets work. They bring more stable prices, and they do not bring lower or less stable output. Furthermore, our advancing knowledge points clearly in the direction of the fourth option. The benefit-cost calculation has tilted significantly towards favouring a price-*level* target and a low, and possibly zero, inflation rate (appropriately defined).

I reach this conclusion by surveying what we have learned from the recent literature in four areas: (i) the comparative performance of targeters and non-targeters; (ii) the tension between the pursuit of price stability and the attainment of output stability; (iii) the choice between a *price-level* target and an *inflation* target; and (iv) the choice between low positive inflation and zero inflation. I pursue each of these topics in the sections that follow.

In the final section, I return to the ranking of our options.

## 1 Lessons from Comparing Targeters and Non-Targeters

The pursuit of goals for the price level clearly does not require the adoption of formal quantitative targets. Germany and the United States are examples of countries in which the central bank pursues a price-level objective, but without a formal target. But price-level objectives may be formalized as inflation targets or price-level targets. In today's world, no nation has adopted a *price-level* target. But, Canada as well as several other countries, four of which (Australia, New Zealand, Sweden, and the United Kingdom) have much in common with Canada, have adopted formal inflation targets. So, what are the effects of targets?

The most natural way to approach this question is to compare macroeconomic performance with and without targets, which is what two major studies have done.

The first and lengthiest is that of Bernanke, Laubach, Mishkin, and Posen (1999), which studies inflation targeting in Australia, Canada, Israel, New Zealand, Spain, Sweden, and the United Kingdom. In all cases, the basic goals of inflation targeting are the same:

- to clearly state the goals of monetary policy
- to establish a framework of accountability

- to lower the cost of lowering inflation by directly affecting inflation expectations

Targets appear to achieve the first two goals; however, it is not strongly evident in the work of Bernanke et al. that they achieve the third. Estimated vector autoregression (VAR) models of inflation and the output gap and estimates of sacrifice ratios in targeting and non-targeting countries fail to reveal any strong and systematic effects of formal targets.

The other major study, Johnson (1999), has a narrower and sharper focus. It studies five targeters (Australia, Canada, New Zealand, Sweden, and the United Kingdom) and six non-targeters (France, Germany, Italy, Japan, the Netherlands, and the United States) to determine whether inflation targeting changes the expected inflation rate, uncertainty about inflation, and inflation-forecast errors. The focus on inflation expectations is clever, because it gets at the issues that Bernanke et al. wanted to study, but does so in a more controlled manner. If targets significantly lower the expected inflation rate, they help to get the actual inflation rate down with a smaller output gap. If targets lower the degree of inflation uncertainty, they bring welfare gains of a hard-to-quantify type by improving resource allocation. And if targets do not significantly increase forecasting errors, they do not increase the average absolute size of the output gap.

Using direct observations of inflation expectations (of professional forecasters), Johnson finds that targeting significantly lowers the expected inflation rate in all five targeting countries, most effectively in Australia and New Zealand and least in Sweden. He finds no significant effect of targeting on uncertainty about inflation, except in New Zealand. And he finds no significant effect of targeting on absolute forecasting errors. Regardless of whether formal targets are adopted, all countries that lower their inflation rate experience an unexpected disinflation.

Based on Johnson's work, we can say that inflation targeting does appear to contribute to the task of achieving and maintaining lower inflation, and it works by influencing the expected inflation rate.[3]

The macroeconomic performance of the United States and Japan superficially reinforces the lessons learned from the work that I have just reviewed. It is difficult to envision a central bank that takes its inflation targets seriously, getting into either a deflationary hole as Japan has done or into the overheated condition in which the United States now appears to be. These remarks are speculative and are not a substitute for the serious

3. It is not clear whether targets have a pure announcement effect on inflation expectations or whether they work by gradually building credibility. David Laidler has pointed out (in private correspondence) that while both effects are worth having, the latter benefit is the one that matters for a regime in which the extension of existing targets is the issue.

research that is needed to compare macroeconomic performance with and without inflation targets. Further work that studies the performance of these two major non-targeters will have a large payoff.

## 2 Price Stability and Output Stability

There has been a great advance in the clarity with which people now think about the price-stability objective and its place in the broader picture of macroeconomic control. This advance has been gradual, but its pace has quickened during the past few years.

At the heart of the matter is the fundamental idea that *policy is a process, not an event.*[4] This view stems from advances that began around 30 years ago when economists started to envision macroeconomic observations as the outcomes of choices made by rational agents operating through time in a stochastic world. In this new vision, the role of expectations and the effect of policy on expectations assumed a central role in our models and in our understanding of how macroeconomic outcomes are generated and how they might be influenced.

The specific advances of the past few years are usefully described in terms of four misconceptions about inflation targets that have been dispelled. They are:

(i) Inflation targets are for "inflation nutters."[5]

(ii) Inflation targets are alternative and inferior rules to others, like the Taylor interest rate rule, the McCallum money-base rule, or others based on monetary aggregates.

(iii) Inflation targets create a credibility problem because they conflict with central bank discretion.

(iv) Inflation targets are incompatible with the persistent aggregate supply shocks that are a dominant source of aggregate fluctuations.

### 2.1 Are inflation targets for inflation nutters?

Concern over price stability and output stability stems from the same idea. A more stable price level and more stable output enhance economic welfare. During the 1960s, it was presumed that a trade-off existed between price

4. I attribute this description of policy to Neil Wallace, who used these words at a conference on rational expectations held in Maine in 1979.

5. Mervyn King (1996) used this phrase when explaining what Bank of England inflation targeting does and does not mean.

stability and output stability. Higher average output (and employment) could be attained, it was believed, only at the cost of higher average inflation.[6]

With the new vision of the 1970s came the view of the output-inflation trade-off. There is a natural output rate (that grows and fluctuates over time) and a natural unemployment rate (that fluctuates over time), and these "natural rates" are independent of the inflation rate. So there is no trade-off between output and inflation *on the average*. But there is a trade-off between the *variability* of output and inflation.[7] Lower output variability can be attained only at the cost of greater inflation variability. The policy problem is to design control mechanisms that make the output-inflation variability trade-off as favourable as possible, and to select the point on the trade-off that optimally balances the marginal gain in social welfare from lower output variability against the marginal cost in social welfare from greater price-level variability.

In the past, some economists thought that setting a target for the behaviour of the price level was in conflict with this optimal-control view of the policy problem. The idea seemed to be that a price-level target implied placing a large weight on price stability and a low weight on output stability and, in effect, forcing an outcome at a corner of the available output-inflation variability trade-off.

Today, economists of all shades realize that targeting the behaviour of the price level is one of the tools available for influencing the output-inflation variability trade-off, and that adopting formal targets carries no presumption as to the point on the variability trade-off that should be aimed for.

## 2.2 Are inflation targets inferior rules?

An inflation target is not a rule. It is an objective. Achieving that objective is universally agreed to be difficult. And reasonable people agree to differ on the best methods of achieving it.

Among the reasons why hitting an inflation target is difficult is the fact that the central bank's instruments operate with long and uncertain time lags. For this reason, it is impossible to target the inflation rate directly. Instead, the central bank targets its *forecasts* of the inflation rate and of the output gap.

---

6. A few economists still cling to this view, and I will consider the evidence for and against it when I examine the case for a positive inflation target.

7. John Taylor (1979) is usually credited with the most forceful and uncontroversial statement of this, now standard, view.

The central bank may use rules to influence the forecasted macroeconomic outcome. The Taylor interest rate rule or the McCallum money-base rule is a potentially viable method of achieving an inflation target. So are other rules based on monetary aggregates or more complicated rules based on reading econometric models.

## 2.3 Do inflation targets conflict with discretion?

In the older literature, the rules versus discretion dichotomy addressed questions about the politics of central banking and the democratic accountability of the bank. A rule was a legislated constraint on the bank—typically to maintain convertibility of its liabilities into gold. Discretion meant leaving the bank to do what it thought best in the circumstances.

In the modern literature, which begins with Kydland and Prescott (1977), rules mean doing the time-inconsistent but optimal thing, and discretion means optimizing each period (and being time-consistent) but, possibly, achieving a suboptimal outcome.

Inflation targets are an attempt to solve the older problem of discretion, but not the modern version of that problem. Requiring the central bank to pursue an inflation target is equivalent to telling it the nature of the loss function it must seek to minimize. But the adoption of a target still leaves the bank free to optimize each period under discretion (to behave consistently). This combination of an inflation target and freedom for the central bank to pursue the target in whatever way seems best is what Lars Svensson (1999a) has called “constrained discretion.”

Constrained discretion appears to be a good description of how targets actually work. Central banks that have adopted targets have clearly not abandoned discretion. To achieve their inflation targets, these central banks exercise judgment based on reading a wide range of formal and informal models. But the discretion of the central bank is different *with* targets than without them. Stanley Fischer (2000) explains the matter thus:

> The inflation targeting approach helps clarify the rules versus discretion debate. Inflation targeting works by setting out very clear goals of monetary policy and the framework for implementing it. The framework describes who has what responsibilities, and what accountability and how much transparency there should be. But it leaves to the central bank, to the experts, to actually hit the target, using the instruments the framework permits it to use.

> This notion of constrained discretion—of leaving the discretion about how to deploy instruments (instrument independence) to the central bank, while withholding from it the discretion about the goals of policy and the framework in which it is to be made—resolves, in a way that could not have been resolved 30 years ago, what it is that we want central banks to do. It resolves how to combine rules (and what the rules should be about) with the central bank's discretion in the operation of policy.

## 2.4 Are inflation targets incompatible with persistent aggregate supply shocks?

Persistent aggregate supply shocks are an important part of the macroeconomic landscape. They came to prominence with the oil price shocks of the 1970s and 1980s and the productivity slowdown of that period. They came into mainstream macroeconomic models with the development of real-business-cycle theory in the 1980s. And today, they are a routine part of our way of thinking about the supply side of the economy.

A negative supply shock produces stagflation. An inflation target that causes aggregate demand to decrease in the face of a negative supply shock brings an even deeper cut in output than the supply shock alone would have created. On the other side, a positive productivity shock brings rising output and a falling inflation rate. Stimulation to keep inflation on track makes output boom even more, so that targeting the inflation rate in the face of supply shocks exacerbates output fluctuations.

All wrong! With inflation targeting in a broader framework of macroeconomic stabilization, a decrease in aggregate supply that comes from a negative aggregate supply shock is partly accommodated. And because the shock persists, the accommodation persists. The extent to which a supply shock is accommodated depends on the relative weights placed on price stability and output stability in the central bank's preferences.

There is a supply-shock problem, however. Some supply shocks are the response to technological change and are efficient. Others are the response to interventions and distortions and are inefficient. In principle, efficient supply shocks should be ignored and inefficient ones smoothed. In practice, we have no good method for decomposing supply shocks into efficient and inefficient components. So we do not have good ways of deciding the most appropriate way to respond to one.

But this problem is present in all economies. And the adoption of an inflation target does not change the nature of this problem or make it any more difficult to deal with.

To make the ideas that we have just reviewed more concrete, it will be useful to illustrate them in the context of an explicit model.

## 2.5 A formal illustration of inflation targeting

Lars Svensson has developed a simple macroeconomic model that shows how inflation targeting works.[8] The output gap as a proportion of potential GDP, *y*, is generated by the short-run aggregate supply function:

$$y_t = \rho y_{t-1} + \alpha(\pi_t - E_{t-1}\pi_t) + \varepsilon_t. \quad (1)$$

Here, $y_{t-1}$ is the output gap in the previous period, and $\rho(<1)$ is the degree of persistence in the gap. The inflation rate is $\pi_t$, and the rational expectation of the inflation rate conditional on information at the end of the preceding period is $E_{t-1}\pi_t$. The parameter $\alpha$ is the strength of the response of output to an inflation (or price-level) surprise. A supply shock, $\varepsilon$, hits the economy each period. This shock is i.i.d. with a zero mean and variance $\sigma^2$.

The central bank has an inflation target, $\pi^*$, and dislikes deviations of the actual inflation rate from $\pi^*$. But the bank also dislikes output gaps. To formalize this description of the bank's objectives, write the bank's loss function as:

$$L_t = E_t\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t}\left(\lambda y_\tau^2 + (\pi_\tau - \pi^*)^2\right)\right]. \quad (2)$$

In each period, the bank's "loss" is determined by the square of the deviation of the output gap from zero and of the inflation rate from its target. The weight placed on the output gap relative to the weight placed on missing the inflation target is $\lambda$. The bank is concerned not only with the current period, but also with the entire future performance of output and inflation. And it discounts the future with discount factor $\beta(<1)$.

There is no commitment technology, and the central bank uses its discretion each period to minimize *L* in equation (2), subject to the constraint that equation (1) imposes. In the formal story, the central bank controls the inflation rate directly. This assumption is made for convenience.[9]

---

8. The model presented here is that of Svensson (1999b), as simplified by Dittmar, Gavin, and Kydland (1999a).

9. We can add intermediate instruments and equations to describe their links with output and inflation, or we can think of the inflation rate as being the sum of the central bank's forecast of inflation, which it controls, and a zero mean, i.i.d. control error.

The solution to the central bank's problem is a decision rule that allocates the current period supply shock between the inflation rate and the output gap. It is convenient to write this solution as:

$$\pi_t = \pi^* - \frac{b}{1-\alpha b} y_t, \tag{3}$$

where $b$ is chosen optimally by the central bank.

With this decision rule, the inflation rate and the output gap are:

$$\pi_t = \pi^* - \frac{b}{1-\alpha b} \rho y_{t-1} - b\varepsilon_t, \tag{4}$$

$$y_t = \rho y_{t-1} + (1-\alpha b)\varepsilon_t, \tag{5}$$

and the variability of inflation and the output gap, measured by their unconditional variances, are:

$$\sigma_\pi^2 = \frac{b^2}{1-\rho^2}\sigma^2 \tag{6}$$

$$\sigma_y^2 = \frac{(1-\alpha b)^2}{1-\rho^2}\sigma^2. \tag{7}$$

The parameter $b$ depends on all of the parameters and varies in a systematic way with $\lambda$. A central bank that places no weight on output, $(\lambda = 0)$, sets $b$ equal to 0, so that it stabilizes the inflation rate at $\pi^*$, and lets the output gap follow the path $\rho y_{t-1} + \varepsilon_t$ with unconditional variance, $1/(1-\rho^2)\sigma^2$.

A central bank that places a large weight on output, $\lambda = \lambda^*$, sets $b = 1/\alpha$ so that the output gap follows the path $\rho y_{t-1}$ with zero unconditional variance.

The trade-off between output and inflation variability depends on $b$ and on $\alpha, \rho,$ and $\sigma$. These parameter values depend on the length of the interval indexed by $t$. For $t$ = one quarter, reasonable values for the parameters are $\rho = 0.9$ and $\alpha = 0.5$. Figure 1 shows the output-inflation variability trade-off with the variance of the one-period aggregate supply shock normalized at unity.

With these values for $\rho$ and $\alpha$, setting $b$ equal to 2/3 gives inflation and the output gap equal variances at 2.34 times the variance of the one-period supply shock. Setting $b = 2 = 1/\alpha$ stabilizes the output gap at zero and makes the inflation rate highly variable around its target value of $\pi^*$. Setting $b = 0$, an "inflation nutter" central bank stabilizes the inflation rate

**Figure 1**
**The output-inflation variability trade-off**

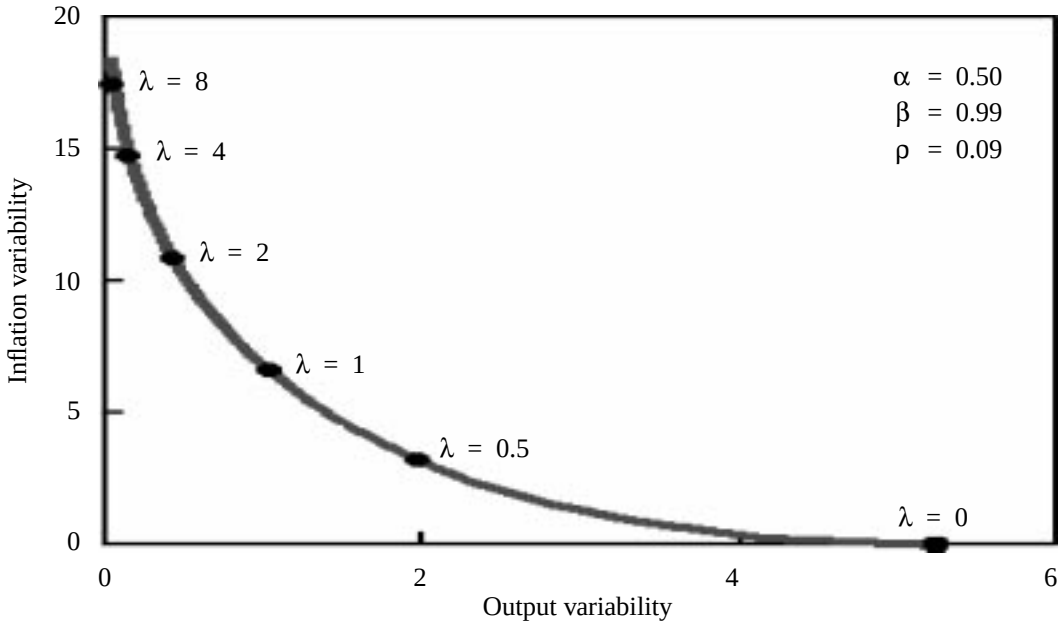


at its target at the expense of output variability of around five times the variance of the one-period supply shock.

We've seen that inflation targeting is a broad framework for pursuing macroeconomic stability that does not preclude placing a large weight on output stability, permits any of the many proposed techniques of monetary control to be used, constrains the central bank to pursue a well-defined objective, yet leaves the bank to use its expertise and discretion to best achieve its goals, and pays due regard to and partly accommodates persistent aggregate supply shocks.

## 3 The Price Level vs. the Inflation Rate

So far, we have thought about price stability in terms of an explicit inflation target. I now turn to the following question: Does it matter whether the price-stability objective is formulated as a target *price-level* path or a target *inflation* rate? This choice does *not* involve different average inflation rates. Rather, it is a choice between a random walk and a deterministic trend for the price level's central target path.

Figure 2 illustrates and contrasts these two possibilities. A price-level target (part "a") specifies a path for the (logarithm of the) price level and a range around that path. An inflation target (part "b") specifies a path for the

**Figure 2**
**Price level vs. inflation target**

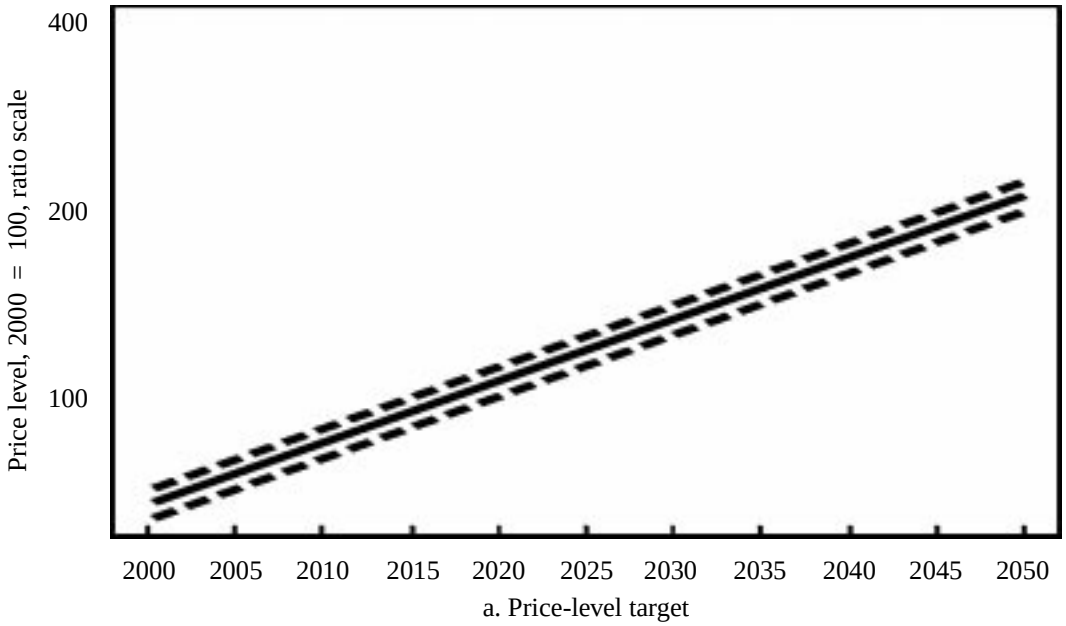


a. Price-level target

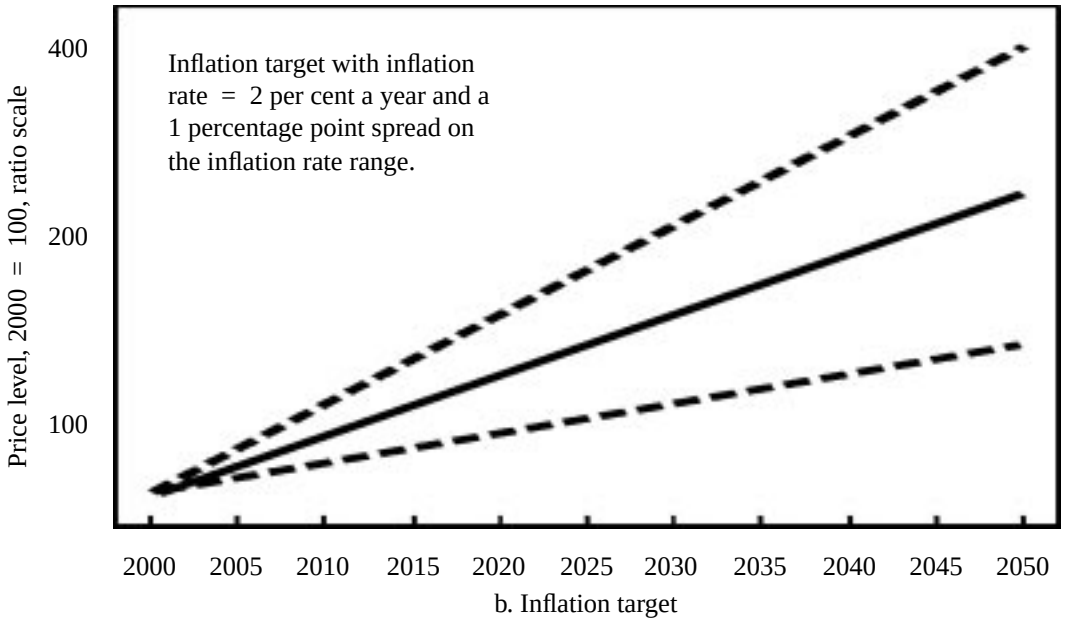


b. Inflation target

inflation rate and a range around that path. It implies a price-level path with a range that is described by an ever-widening cone.

Interestingly, while the choice between these two ways of formulating the price-stability objective was on the agenda at the two conferences that mark the origin for our current discussions, it was not central to them. With an important exception that I will detail below, it was taken for granted that the question had been settled at an earlier Bank of Canada conference in 1993. (See Bank of Canada 1994.)

At the 1993 conference, Scarth (1994, 90) had formulated the question thus: "Can a once-and-for-all devaluation be a 'good' thing? If the answer is 'Yes,' . . . the desirability of . . . price stability is fundamentally threatened." He went on to show that in the deterministic overlapping generations economy that he studied, a once-and-for-all devaluation is welfare-improving. Hence, price stability is undesirable.

Scarth's conclusion, although derived in an entirely different manner, was consistent with the then conventional wisdom, which Svensson summarizes as the view "that the choice between price-level targeting and inflation targeting involves a trade-off between low-frequency price-level variability on one hand and high-frequency inflation and output variability on the other."[10] The intuition behind the conventional wisdom was that "under price-level targeting, higher-than-average inflation must be succeeded by lower-than-average inflation, [hence] higher inflation variability . . . [and] via nominal rigidities, . . . higher output variability." (See Svensson 1999b, 278.)

This view arose from simulations of small macroeconomic models reported by, among others, Lebow, Roberts, and Stockton (1992) and Haldane and Salmon (1995). In all of the cases in which price-level targeting increases output variability, inflation expectations are backward-looking and pay no attention to the monetary policy actually being pursued.

Fillion and Tetlow (1994) challenged the intuition behind this conventional view but did not destroy the nominal-real trade-off view. In their small stochastic model, with a combination of forward-looking rational expectations and backward-looking expectations, simulation results showed that stabilizing the price level *lowers* the variability of inflation but *increases* the variability of output (see Fillion and Tetlow 1994, 153). Laxton, Ricketts, and Rose (1994) obtained a similar result.

---

10. Svensson (1999b, 278). Also consult Svensson for the references to the literature that is the source of what he calls the conventional wisdom.

The new intuition became that with price-level targeting, an increase in the inflation rate would be seen as temporary, so price-level targeting would not lead to as large an increase in the expected inflation rate, and inflation variability would be lowered.[11] But output variability would increase, because the policy response to supply shocks that cause the price level to rise would decrease aggregate demand. To the extent that expectations are backward-looking, this policy response would lower output by even more than the first-round supply-side decrease.

This is the state of knowledge as summarized insightfully by Pierre Duguay (1994) just after the 1993 conference.

Ideas began to change during the mid-1990s with an important paper by Lars Svensson (1999b) that challenged the trade-off view. Svensson showed that with endogenous policy choice and rational expectations, price-level targeting delivers lower price-level and inflation variability than does inflation targeting. But, price-level targeting delivers identical output variability to inflation targeting. So, there is a "free lunch." Dittmar and Gavin (2000) and Vestin (2000) have shown that Svensson's free-lunch result can be derived in models that have either a classical or a Keynesian micro-foundation.

To make the source of the free lunch as clear as possible, we will follow Svensson and work with the same model that we've just examined. We replace the inflation rate and the expected inflation rate with the following price level and the expected price-level identities:

$$p_t \equiv p_{t-1} + \pi_t, \tag{8}$$

$$E_{t-1} \equiv p_{t-1} + E_{t-1}\pi_t. \tag{9}$$

Here, $p$ is the logarithm of the price level. The inflation rate is defined by equation (8) and the expected inflation rate is defined by equation (9).

When these definitions are used, the aggregate supply equation (1) becomes:

$$y_t = \rho y_{t-1} + \alpha(p_t - E_{t-1}p_t) + \varepsilon_t. \tag{10}$$

Although it is written in terms of the price level, this equation is exactly the same as equation (1).

The central bank's target for the price level is $p_t^* = p_{t-1}^* + \pi^*$, and the bank now cares about deviations of the actual price level from this target. So the loss function becomes:

11. Laidler and Robson (1994) gave an early account of a special case of this new intuition.

$$L_t = E_t\left[\sum_{\tau=1}^{\infty}\beta^{\tau-t}\left(\lambda y_\tau^2 + (p_\tau - p_\tau^*)^2\right)\right]. \quad (11)$$

Again, $p$ replaces $\pi$, and $p_\tau^*$ replaces $\pi^*$. But equation (11) is *not* identical to equation (2). In equation (2), the bank is concerned only if the inflation rate misses its target. In equation (11), the bank is concerned if it misses its price-level target, which is the same as being concerned about the *accumulated* inflation miss.

For inflation targeting, the problem is to minimize equation (2), subject to equation (1). For price-level targeting, the problem is to minimize equation (11), subject to equation (10). Notice that $\pi$ enters the first problem in exactly the same way that $p$ enters the second problem. And notice that the $y$ enters both problems identically. So, when we solve the price-level-targeting problem, we get the same solution for $p$ that we get for $\pi$ when we solve the inflation-targeting problem. And we get the same solution for $y$ in both problems.

The *price level* now follows the path:

$$p_t = p_t^* - \frac{b}{1-\alpha b}\rho y_{t-1} - b\varepsilon_t, \quad (12)$$

and its variability is described by:

$$\sigma_p^2 = \frac{b^2}{1-\rho^2}\sigma^2, \quad (13)$$

which is the same as the variance of the inflation rate under an inflation rate target.

The inflation rate now responds to the *change* in the output gap as follows:

$$\pi_t = \pi^* - \frac{b}{1-\alpha b}\rho(y_{t-1} - y_{t-2}) - b(\varepsilon_t - \varepsilon_{t-1}), \quad (14)$$

and the variability of inflation, measured by the unconditional variance of the inflation rate, is:

$$\sigma_\pi^2 = \frac{2b^2}{(1+\rho)}\sigma^2. \quad (15)$$

Comparing the variability of inflation under the two alternative targeting regimes, we can see that price-level targeting delivers a lower inflation variability if $\rho > 0.5$.

**Figure 3**
**The free lunch from a price-level target**

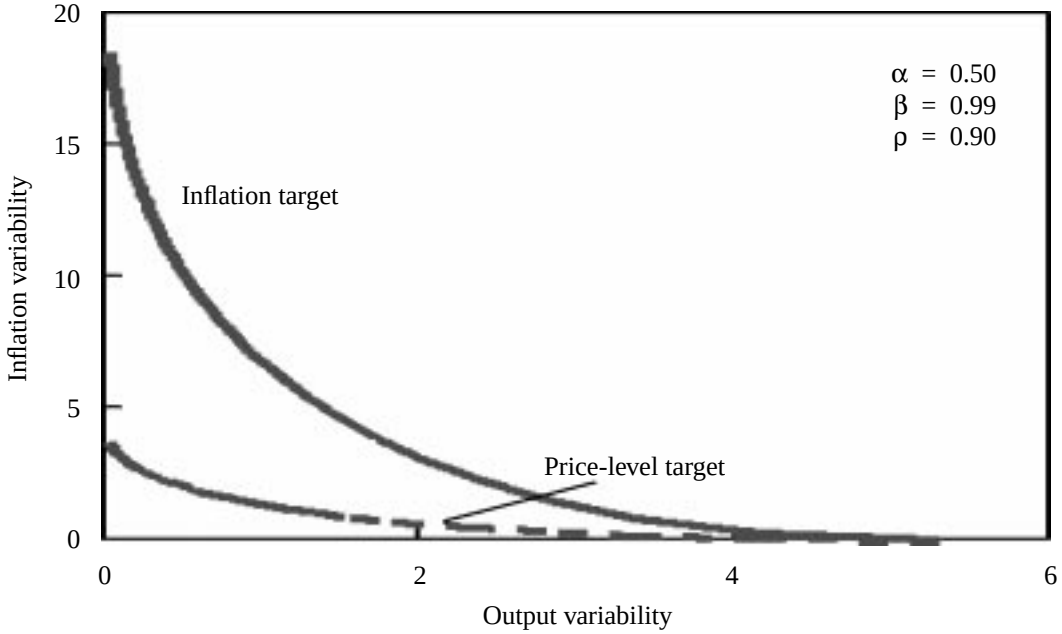


For the likely value of $\rho = 0.9$, the variance of the inflation rate is five times as large when the inflation rate is targeted as it is when the price level is targeted at a given variance of the output gap. Figure 3 shows the two variability trade-offs and the free lunch from a price-level target.[12]

Svensson's work raises two questions. First, what is the source of the free lunch? And second, can we expect the outcome that we generate in this model economy to be relevant to our actual economy?

The source of the gain is easy to see. Under inflation targeting, the decision rule makes the inflation rate respond to the output gap. So the variance of the inflation rate is proportional to the variance of the output gap. Under price-level targeting, the *price level* responds to the output gap, so the inflation rate responds to the *change* in the output gap. With enough persistence in the supply shock ($\rho > 0.5$), the variance of the change in the output gap is smaller than the variance of the output gap.

Does the free-lunch result apply to our actual economy? This type of question is always delicate, and credibility is crucial to the answer. I return to the issue of credibility after reviewing the work that is favourable to a

12. There are some technical issues concerning the existence of equilibrium and the equilibrium concept employed in the work of Svensson, and Dittmar and Gavin. I discuss some of these issues in the Appendix.

price-level target. Four groups of papers make me think that we can rely on the free-lunch result.

First, the result is robust to changes in assumptions. Svensson (1999b) shows that it holds for situations in which the central bank targets the inflation rate or price level indirectly via the control of a monetary aggregate or other instruments. Dittmar and Gavin (2000) show that it arises when the aggregate-supply equation has the new-Keynesian form with current expectations of future inflation rates (price levels) entering the aggregate-supply function via a wage-contract effect. Vestin (2000) shows that it arises if the aggregate-supply function incorporates Calvo (1983) price setting. Interestingly, in these forward-looking, price-setting environments, there is enough persistence in the supply side from the price-setting process for the free-lunch result to hold, even when $\rho$ is less than 0.5.

Second, the free-lunch result holds in the Federal Reserve Board's large-scale, open-economy, macroeconometric model of the U.S. economy. John Williams (1999) studies a class of interest rate rules in which the interest rate responds to targets for the interest rate itself, inflation, and the output gap. Four definitions of the inflation target—one-quarter, one-year, and three-year inflation rates, and the price level—are compared, and the relative weights on inflation and the output gap ($\lambda$ in the above simple models) are varied to generate variability frontiers. In Williams' experiments, the three-year inflation definition wins over the entire range of possible weights for output and inflation. But, over the empirically relevant range, price-level targeting performs almost as well as three-year inflation targeting and better than one-year inflation targeting.

Third, a version of the free-lunch result showed up in the work of Black, Macklem, and Rose (1998), presented at the Bank's 1997 conference. This paper reports the results of simulations with a small stochastic when *both* the inflation rate and the price level are targeted. The rule used is of the form:

$$\text{interest rate} = \ldots + \theta(\pi - \pi^*) + \tau(p - p^*). \qquad (16)$$

With $\tau = 0$, this rule targets only the inflation rate. Points on the inflation variability-output variability frontier are derived for different values of $\tau$ and compared with the zero $\tau$ case. The experiments are run in two stages. First, expectations are constant and the conventional-wisdom result is generated: targeting the price level increases the variability of both the inflation rate and the output gap. Second, expectations adjust to reflect the policy being used and the free-lunch result is generated.

Fourth, the historical evidence from the only instance in which price-level targeting was practised is positive. The Riksbank left the gold standard

and began targeting the consumer price index (CPI) (reset at 100) in September 1931.[13] The experiment ran until April 1937, when it was abandoned for an exchange rate objective. During the 67 months of CPI targeting, the index rose to a peak of 101.5 in 1932, fell to a trough of 98.4 in 1933, and then steadily increased to 104 when the policy was ended. Sweden had one of the least severe recessions through these years, an outcome that has been attributed to its decision to leave the gold standard and avoid a serious fall in the price level (Bernanke 1995).

In his discussion of the Black et al. paper, Gregor Smith (1998) says that the idea that price-level targeting improves performance is speculative, and claims "there is considerable historical experience with price-level targeting" that casts the approach in an unfavourable light. The experience that Smith had in mind is the gold standard, and he provided data from the United Kingdom's infamous 1920s episode to "prove" his point that price-level targeting is highly destabilizing.

I think it is clear that Gregor Smith's example is inappropriate and that the episode in monetary history to which he refers provides an example of what happens when the price level is *not* targeted. It bears repeating that Sweden had to abandon the gold standard to pursue its price-level targeting, and it is widely agreed that its *real* economic performance during the 1930s was among the best precisely because of this policy regime.

I have noted that credibility is the key to making price-level targeting work. If expectations are backward-looking, output can become more variable under a price-level target than under an inflation target. So how crucial is credibility?

Barnett and Engineer (2001) address this question at some length analytically. And Maclean and Pioro (2001) address it empirically in the context of the Bank of Canada's Quarterly Projection Model. Barnett and Engineer show that under discretion, price-level targeting is a good idea when expectations are either forward-looking or indirectly forward-looking through output persistence.[14] Maclean and Pioro perform a large number of simulations and show that, provided some credibility is present, far from

13. See Fisher (1934), Jonung (1979), Black and Gavin (1990), Berg and Jonung (1998), and Dittmar, Gavin, and Kydland (1999b).

14. They also show that, under commitment, price-level targeting is a good idea when the Bank can target current variables, and expectations are directly forward-looking. Otherwise, some version of inflation targeting is optimal where the amount of price-level drift is related to the weight on expectations that are not forward-looking. I am assuming that discretionary outcomes are the only feasible ones under current circumstances and so am ignoring infeasible outcomes that require commitment.

complete credibility is required to achieve the free-lunch result. This empirical finding is broadly in line with the analytical results.

It is worth emphasizing that to achieve the free lunch, output should have an autocorrelation greater than one-half, which it does, and the central bank should use its discretion to transparently pursue its declared price level and output-gap objective so that forward-looking expectations respond to the bank's actual policy.

## 4 Zero vs. a Positive Inflation

Targeting the price level carries no implication for the slope of the target path. Should that slope be positive or should it be zero? Serge Coulombe (1998) cogently and elegantly argued at the Bank's 1997 conference both for a price-level target and for stable prices. In an analysis that focuses on the intertemporal allocation and the role of inflation in distorting intertemporal price signals, he concluded, "A regime with price-level stability is clearly superior from a theoretical standpoint."

The state of knowledge on the benefits of zero inflation has not changed much since the last conference. Everyone agrees that the gains are hard to measure. Some, like Coulombe, believe that those gains are large, nonetheless. Others are skeptical and regard the absence of clearly measured large gains as a sign that they are probably not present. But, even the most skeptical would not claim that inflation is, per se, a good thing. Rather, those who caution against price-level stability assert that the costs of attaining it are too large.

Major advances in our knowledge have come on the cost side of the calculation. There are three reasons why a positive inflation might be preferred to zero inflation:

(i) The measured inflation rate overstates the true inflation rate, and a constant true price level with a rising measured price level is the appropriate target.

(ii) Inflation lubricates the labour market and makes it work more efficiently.

(iii) Nominal interest rates have a zero lower bound, and targeting a stable price level will conflict with the zero lower bound too often and leave monetary policy unable to speed the end of a recession.

### 4.1 Measurement bias

The first reason for aiming at a positive inflation rate is easily disposed of on the basis of what we knew at the time of the last Bank of Canada conference.

Measurement bias in the Canadian CPI is estimated at +0.5 per cent a year (Crawford, Fillion, and Laflèche 1998). This estimate excludes possible downward bias from increases in user fees on government-supplied goods and services that are excluded from the CPI.

So, if measurement bias were the only reason to contemplate a positive target inflation rate, the target rate for the currently measured CPI would be 0.5 per cent a year.

But, it would be a relatively simple matter to direct Statistics Canada to construct a "value of money" index (VMI) that equals the CPI scaled back for the best available estimate of the upward measurement bias. The VMI might also usefully exclude high-frequency components of the CPI (provided these components have a zero mean over a reasonably short period). And it might omit interest-sensitive prices, as the New Zealand CPIX and the U.K. RPIX index do. It might also exclude the first-round effects of indirect taxes to get to the "core," as the targeted Canadian index does. The VMI might also be computed as a longer-term moving average than the CPI. The base value of the index would be 100.

A constant VMI with a range for acceptable deviations would be the price-stability target.

## 4.2 Labour market lubricant

Important work has been done since the last Bank of Canada conference that casts new light on the labour market lubricant view of inflation. And this work must be carefully evaluated.

The labour market lubricant belief is based on the presumption that money-wage rates are downwardly rigid. Consequently, if a decrease in a *real*-wage rate requires a cut in a *money*-wage rate, that real-wage decrease does not occur. The market for the class of labour in question fails to clear, resources are misallocated, and the unemployment rate increases.

Conventional wisdom holds that this problem can be avoided by maintaining inflation at a rate that is sufficiently high to ensure that almost all of the required real-wage rate decreases can be accomplished with a zero change in the nominal-wage rate accompanied by positive changes in other money-wage rates.

The efficient lubricant idea came back into prominence in the mid-1990s with the work of Akerlof, Dickens, and Perry (1996) and Fortin (1996). Each claimed that strong evidence shifted the lubricant idea from belief status to solid fact. Fortin's famous histogram with its huge spike at a zero money-wage change was variously interpreted but was widely held to support the lubricant view.

This issue was a major one in the Bank's conference of 1997 as it is in this seminar. It can be addressed most effectively in two stages. First, what is the evidence from wage data on the existence of downward nominal-wage rigidity? Does the phenomenon exist? Second, if it does exist, what are its consequences? Does it lead to a permanent trade-off between real activity (output and unemployment) and inflation? Good work has been done on both issues that I will review it in the following section.

### *4.2.1 Does downward nominal-wage rigidity exist?*

Crawford and Harrison (1998) provided a careful and comprehensive survey of several Canadian data sets and found two reasons to suspect that money-wage rates are downwardly rigid: (i) the percentage of wage freezes increased during the low-inflation 1992–96 period; and (ii) the percentage of wage cuts is small. Using a hazard model, they estimated an upper bound for the effect of nominal rigidity. That estimate is that the proportion of nominal-wage freezes increases by between 10 per cent and 15 per cent when the inflation rate falls from 6 per cent to 2 per cent.

They also found three facts that point towards flexibility: (i) the wage-change distribution becomes more negatively skewed at low inflation rates; (ii) although the percentage of wage cuts is small, that percentage *increased* during the 1992–96 low-inflation period; and (iii) wage freezes occur at all (observed) inflation rates, not only at low ones.

Crawford and Harrison, and their discussant, Bowlus (1998), were careful to point out the limitations of the data sets available and the need to use multiple ones to provide the best available view of the behaviour of wages.

A recent paper by Jennifer Smith (2000), casts a new light on money-wage rigidity and adds significantly to our understanding of this issue. Smith works with a rich U.K. data set that enables her to control for three potential contributors to the appearance of nominal rigidity where none is present. They are: (i) measurement error; (ii) rounding error; and (iii) long-term contracts.

The data used are from the British Household Panel Study (BHPS), which interviewed around 10,000 individuals each year during 1991 through 1996; interviewees are restricted to employed individuals who remained in the same job from one year to the next—"job stayers." The data include overtime and bonuses, and the wage definition is usual gross weekly pay. Histograms of the raw data[15] contain a large spike at zero, like the Canadian distributions presented in Crawford and Harrison, and Fortin. And the spike

15. Figure 1, p. C181.

is larger, the lower the mean of the wage-change distribution, but is unrelated to the average inflation rate. In the full, five-year sample, 9 per cent of money wages are rigid, but 23 per cent of individuals experience a money-wage decrease.

The inclusion of overtime and bonuses raises the immediate suspicion that they are the sources of the flexibility in the data. Fortunately, people report when they receive these components of pay, so their presence can be controlled for. Surprisingly, perhaps, the sub-sample that has no overtime or bonuses behaves exactly like the full sample.

To estimate the effect of measurement error, Smith (2000) uses the fact, unique to the BHPS, that interviewees are given a chance to check their payslips when reporting their pay. In the "payslip-not-seen" sub-sample, 25 per cent report pay cuts and 10 per cent report rigid pay. In the "payslip-seen" sub-sample, 18 per cent report cuts and 6 per cent report rigid pay. So, assuming (reasonably) that payslip-seen responses are measured without error and payslip-not-seen responses are measured with error, measurement error exaggerates rigidity. This finding is contrary to the *assumption* of Akerlof, Dickens, and Perry, who claim that measurement error in panel data contributes to the appearance of greater flexibility than is truly present.

To estimate the effect of rounding error, Smith examines the raw pay data to determine for each report, whether it is rounded to the nearest £5 up to £1,000. Apparent rigidity due to rounding is not large, but correcting for this bias alone decreases the zero spike to 9 per cent.

To study the effects of long-term contracts, Smith looks at the details of the intervals in months between interviews. Interviews are conducted at intervals that range from 3 to 19 months. She finds that zero wage change is largest at 12.5 per cent for those interviewed 9 months or less after the previous interview. For 10-, 11-, and 12-month intervals, the percentage with zero wage change is roughly constant at 8 per cent. And above 14 months, the zero-change bin falls to 4 per cent. Overall, Smith estimates that interactions between contract length and interview frequency account for about 50 per cent of the apparent rigidity.

Bringing these three sources of explanation for zero reported wage change together, measurement error accounts for 40 per cent of apparent rigidity, rounding for 15 per cent, and long-term contracts for 50 per cent. Smith notes that if the three sources of zero change were independent, there would be less than zero per cent at a zero change!

Because the sources of zero-change reporting are not independent, they actually account for 90 per cent of recorded rigidity. With 10 per cent reporting a zero wage change, just 1 per cent of people who remain in the same job truly receive the same money wage for two successive years.

How confidently can we generalize from Smith's results to other countries and, in particular, to Canada? We would be much better off if we had a Canadian survey with the features of the BHPS that would enable us to check the sources of reported rigidity. But, in areas where the U.K. survey and Canadian and U.S. surveys provide comparable information, there is remarkable similarity across the economies.

First, there is similarity in the raw data of the BHPS and the averages in the PSID data used by Kahn (1997); Lebow, Stockton, and Wascher (1995); and McLaughlin (1994); the PSID and CPS data used by Card and Hyslop (1997); the Canadian sources used by Crawford and Harrison (1998); and the LMAS data reported by Bowlus[16] (1998). Ignoring the large union data of Fortin, zero wage-change ranges from 7 per cent in several PSID studies to 14 per cent in the CPS data, with the Canadian LMAS inside this range at 8 per cent. Negative wage changes range from 10 per cent in the Canadian SLID to 25 per cent in the Canadian LMAS. Table 1 summarizes the raw data.

Second, in the only other study that has proposed corrections to the zero percentage (Card and Hyslop 1997), the corrections are remarkably similar to those suggested by Smith (2000). Using reasonable assumptions about the wage distribution and rounding, they estimate that 4 percentage points of the zero changes result from rounding. This percentage is remarkably close to Smith's 4.5 percentage points. Making further (but sensible) assumptions, Card and Hyslop estimate that long-term contracts account for about one-half of the zero changes. This number is virtually the same as Smith's for the BHPS data. The BHPS data alone permit a correction for reporting error, so we have nothing in the other data sets against which to check Smith's estimate of the scale of this correction, but it is the smallest of her three corrections.

Table 2 summarizes the corrections to the percentage rigid suggested by Card and Hyslop, and Smith. While we cannot say that these corrections apply to the other data sets, it is interesting to note the similarity between the corrections proposed for the PSID and BHPS in these two studies and to note the similar starting points in the raw data across all the data sets.

We cannot conclude with anything approaching certainty that downward nominal-wage rigidity is absent from our labour markets, but we can say that three sources of apparent rigidity have not been properly controlled for in the Canadian (and some U.S.) work. And we can say that when these factors are controlled for in an otherwise similar data set, nominal rigidity effectively disappears.

---

16. The percentage rigid is not reported in Bowlus (1998). Audra Bowlus kindly provided me with the percentage for job stayers shown in the table. She also provided the percentage for job changers, which is 10.3.

**Table 1**
**Summary of data on downward nominal-wage rigidity**

| Source | Data | Decrease (percentage) | Rigid (percentage) |
|---|---|---|---|
| Fortin (1996) | Canada: Labour contracts without COLAs, 1986–88 | Negligible | — |
| Fortin (1996) | Canada: Labour contracts without COLAs, 1992–94 | 6 | 47 |
| Crawford and Harrison (1998) | Canada: Union settlements (>500 workers), 1992–96 | Negligible | 33 |
| Crawford and Harrison (1998) | Canada: SLID, 1993 | 10 | — |
| Crawford and Harrison (1998) | Canada: Sobeco Ernst and Young Survey, 1989–96 | 9–20 | — |
| Bowlus (1998) | Canada: LMAS[a], 1986–87 | 25 | 8 |
| Lebow et al. (1995) | United States: PSID[b], 1971–88 | — | 8 |
| Kahn (1997) | United States: PSID[b], 1971–88 | 18 | 7 |
| McLaughlin (1994) | United States: PSID[b], 1976–86 | 17 | 7 |
| Card and Hyslop (1997) | United States: PSID[b], 1976–79 | — | 7 |
| Card and Hyslop (1997) | United States: PSID[b], 1985–88 | — | 10 |
| Card and Hyslop (1997) | United States: CPS, 1979–93 | 17 | 14 |
| Smith, J. (2000) | United Kingdom: BHPS[c], 1991–96 | 23 | 9 |

— not reported.
[a] Canada, LMAS is for job stayers.
[b] U.S. PSID studies all use job stayers, and wages include bonuses and overtime.
[c] U.K. BHPS is for job stayers and is normal weekly income, including bonuses and overtime.

**Table 2**
**Correcting the zero spike**

| Study and data set | Average others* (percentage) | Card-Hyslop CPS (percentage) | Card-Hyslop PSID (percentage) | Smith BHPS (percentage) |
|---|---|---|---|---|
| Raw data | 7.6 | 14.1 | 8.3 | 9.0 |
| After rounding-error correction | — | 10.1 | 4.3 | 4.5 |
| After contract-length correction | — | — | 3.6 | 5.6 |
| After reporting-error correction | — | — | — | 7.5 |
| After all corrections | — | — | — | 1.0 |

* Bowlus (LMAS) 8.0, McLaughlin (PSID) 7.0, Lebow et al. (PSID) 8.0, Kahn (PSID) 7.5.

If the above conclusion is correct, the second question concerning the effects of downward rigidity on the trade-off between real activity (output and unemployment) and inflation does not arise. But some observers, including Farès and Lemieux (2001), read this literature differently and conclude that the data do show downward nominal-wage rigidity. Because there is a lack of agreement on this matter, the second question must be addressed.

### *4.2.2 Is there a long-run trade-off between inflation and output?*

Audra Bowlus (1998) notes that even if it were demonstrated conclusively that money-wage rates are downwardly rigid, without a structural model of the labour market, this observation on its own would tell us nothing about the effects of inflation on employment, unemployment, and the speed with which the labour market reallocates resources in the face of differential shocks. While less complete than "a structural model of the labour market," recent papers by Hogan and Pichette (1999), Farès and Lemieux (2001), and Beaudry and Doyle (2001) address the central concern that she raises and assess the evidence in favour of the view that we face a permanent trade-off between inflation and output.

Hogan and Pichette argue that the (good) performance of the Akerlof-Dickens-Perry model in dynamic simulations does not constitute evidence in favour of downward nominal rigidity. They argue that static simulations provide a more appropriate test and, that in this setting, the standard model with no downward nominal rigidity outperforms a model with downward rigidity.

After a review of the literature on the extent of downward nominal-wage rigidity, Farès and Lemieux conclude, "that recent studies, mostly

based on U.S. longitudinal micro data, provide compelling evidence that [downward nominal-wage rigidity] is an important labour market phenomenon" (p. 4). They go on to develop a new wage series for 1981–97, based on individual data files from Statistics Canada's Survey of Consumer Finance, and use this new series to estimate real-wage Phillips curves. For Canada as a whole, the significant negative slope of the real-wage Phillips curve that is present before 1992, disappears after 1992. One interpretation of this result might be that the short-run (real-wage) Phillips curve became horizontal when the inflation rate was pushed down to around 2 per cent. The new wage series that Farès and Lemieux use is available on a provincial basis, so they also estimate provincial real-wage Phillips curves in which they control for national factors with year effects. These Phillips curves do not display any slope change as the inflation rate falls.

Beaudry and Doyle estimate Canadian Phillips curves and find that the slope has decreased during the 1990s. They also report evidence of the same phenomenon in the U.S. Phillips curve. They explain this flattening as the consequence of the conduct of monetary policy, which they suggest has gradually incorporated an improved understanding of supply shocks. They show how, in a simple and standard model, such an improved understanding implies a flatter Phillips curve.

More work is always helpful. But, for the present, it appears that, "We should not look to the labour market as the source of major unemployment costs of zero inflation that would justify a positive inflation target" (Smith 2000, C194).

## 4.3 The zero lower bound problem

Summers (1991) resurrected an old idea previously mentioned by Vickery (1954), Phelps (1972), and Okun (1981), that because nominal interest rates cannot fall below zero, there is a trade-off between the inflation rate and macroeconomic stability. The claim is that expansionary monetary policy cannot be used to get the economy out of recession if the nominal interest rate is zero, so in such a situation, recession will last longer. Because the zero bound will be hit more frequently at lower inflation rates, the lower the inflation rate, the longer the economy will spend in recession. The experience of Japan during recent years is often cited as an example of an economy that has hit the zero lower bound.

Before reviewing the recent literature on this topic, it is worth noting that the zero lower bound (not by name but by substance, and sometimes called "the interest rate floor") was a live issue before and during the Great Depression and in much of the subsequent literature that interprets that

episode.[17] It lies at the heart of the Friedman and Schwartz (1963) indictment of the Fed for not injecting money into the banking system. It was repeatedly addressed by Hawtrey (1925, 1932), who insisted that even if interest rates are already very low, there are ways, other than by open market operations, of getting money into circulation, so the inability to drive rates lower does not mean that monetary policy is powerless. The zero lower bound is also a repeated theme in the work of Brunner and Meltzer (1993), who view monetary policy as working through a broad array of rates of return on many assets, including consumer and producer durables. And finally, it is a major theme in several decades of empirical estimates of the demand for money function, not a single one of which has found evidence of a "liquidity trap."

Some of the recent literature has pursued the older themes. Goodfriend (2000) recommends that the central bank impose a cost of carry on its monetary liabilities so that the nominal interest rate can indeed be negative. Goodfriend also suggests open market operations with the non-bank public in long-term bonds, and intervention in the foreign exchange market. Clouse et al. (1999) examine several channels by which a central bank constrained by a zero interest rate can nevertheless influence aggregate demand. These channels include liquidity effects, bank-lending effects, the credibility of the commitment to a prolonged low interest rate, manipulating inflation expectations, and intervention in long-term bond and foreign exchange markets. McCallum (2000) reviews a similar range of options in the context of a dynamic general-equilibrium model.

Other work has operated inside the constraints imposed by interest rate rules. The two most comprehensive studies are Orphanides and Wieland (1998), and Reifschneider and Williams (1999), both of which assess the quantitative consequence of the zero lower bound, by performing stochastic simulations. These two papers reach slightly different conclusions, so we must examine each with a view to identifying the source of the difference, and thus make a judgment about which is correct.

Using a small-scale rational-expectations model of the U.S. economy, Orphanides and Wieland calculate that if the economy is subject to shocks of the magnitude experienced during the 1980s and 1990s, the zero lower bound has no practical consequence at an inflation rate of 2 per cent a year. But, at inflation rates between 0 and 1 per cent, the variability of output increases significantly and the variability of the inflation rate increases

17. David Laidler must be credited with directing me to this literature and insisting on its relevance to the modern discussion of this topic. His own extensive discussion of this matter can be found in Laidler (1999), especially in Part IV, pp. 247–320.

slightly. They find a long-run trade-off between output and inflation, and a zero inflation rate lowers real GDP by 0.1 percentage points.

Two features of the model used by Orphanides and Wieland play a role in reaching their conclusion. First, the model's steady-state real federal funds rate is 1 per cent. This value contrasts with the historical average of 2.5 per cent for 1960 to 1998. The zero lower bound will come into force significantly more frequently and at a given inflation rate with a 1 per cent equilibrium real rate than it would with a 2.5 per cent real rate.

Second, Orphanides and Wieland consider the effects of the zero lower bound only for a central bank that uses the Taylor rule (the interest rate is set at the target real rate, plus the inflation rate, plus 0.5 times the output gap, plus 0.5 times the deviation of the inflation rate from target) or the more aggressive Henderson-McKibbin rule (the same as the Taylor rule, with the weight on the output gap boosted to 2 and that on the deviation of the inflation rate from target boosted to 1). Under either of these rules, when a negative interest rate is called for, the rule is temporarily set aside and the rate is set at zero. Once the rule again calls for a positive interest rate, it resumes. It is evident that this type of rule is insensitive to the zero lower bound in the sense that it does not anticipate it or react to it. We need to know how the economy would respond under an interest rate rule that *does* anticipate and react to the prospect of, or recent experience of, hitting the bound.

Reifschneider and Williams address both of these issues. First, they use the large-scale, rational expectations, open-economy FRB/U.S. model. This model has less inertia than the small-scale one used by Orphanides and Wieland, but it does match the estimated inertia in the U.S. economy. The monetary policy transmission mechanism in this model operates through interest rates and two other channels: a cash-flow variable influences investment, and a liquidity constraint influences consumption expenditure.

A large number (12,000 quarters per set) of stochastic simulations were performed to generate data for the output gap, inflation, and the federal funds rate for target inflation rates of 0, 1, 2, 3, and 4 per cent a year.

Using the Taylor rule, the zero lower bound is hit 1 per cent of the time with a 3 per cent inflation target and 14 per cent of the time with a zero inflation target. The standard deviation of the output gap increases from 2.9 at a 3 per cent inflation target to 3.6 at a zero inflation target. The standard deviation of the inflation rate also increases, but only slightly, from 1.9 to 2.0.

The Henderson-McKibbin rule hits the zero bound more frequently than the Taylor rule—11 per cent of the time at 3 per cent inflation and 31 per cent of the time at zero inflation. But, it delivers a better performance

of output and inflation variability. It also delivers a similar deterioration in the variability of the output gap at lower average inflation rates. Under this rule, the standard deviation of the output gap increases from 1.8 at 3 per cent inflation to 2.4 at zero inflation. The standard deviation of the inflation rate does not change as the inflation rate decreases.

Reifschneider and Williams also consider an augmented Taylor rule that sets the interest rate in a way that takes the zero bound into account. The rule is:

$$i_t = max\left\{ i_t^{Taylor} - \alpha(Z_t, 0) \right\},$$

where $i$ is the funds rate and $Z$ is the accumulated deviation of the funds rate from the Taylor rate that has resulted from past zero-bound constraints.

This rule delivers almost the same standard deviation of the inflation rate (2.2) and output gap (3.0) at a zero inflation rate as it does at a 2 per cent inflation rate. This augmented rule works by immunizing the long-term interest rate from the zero bound. To work, the central bank must possess credibility, but it can deliver the interest rate described by the augmented rule.

Although most of the literature dealing with the zero lower bound has studied the effects of using a Taylor rule (or a similar or augmented rule), we know from the other work of Williams (1999) on alternative policy rules that these rules are inefficient. Reifschneider and Williams also study the zero bound under efficient policy rules.

Efficient rules deliver substantially lower variability of output and inflation than do Taylor-type rules. And the move from a 2 per cent inflation target to a zero inflation target has almost no effect on the position of the efficient variability frontier.

No one appears to have studied the performance of an economy in which policy takes into account anticipated future zero-bound hits. Reifschneider and Williams speculate that such policies would lower the detrimental effects of the zero bound even further.

## 5 The Options We Face

What should we do when the current inflation-target period ends? As I noted in the Introduction, there seem to be four practical possibilities:

(i) Permit the formal target to lapse but keep doing the best job possible to attain low inflation in the broader context of macroeconomic stability.

(ii) Reaffirm the existing inflation target and formally commit to it for another fixed term or for an open-ended term.

(iii) Reaffirm the existing inflation-target rate but formulate it as a target path for the price level.

(iv) Commit to a lower, possibly zero, inflation-target rate formulated as a target path for the price level.

Option (ii) must be declared the winner in the absence of clear and convincing evidence for change. That evidence is not present for option (i), which is a departure from the status quo. But, there is evidence pointing the other way towards either a price-level path or a flat price-level path. There will be a diversity of opinion on the strength of that evidence, and that diversity tells us that more research is needed to sharpen our view of how a constant price-level economy would function.

If my summary of the current state of knowledge and my reasoning are correct, the following targets for the behaviour of the price level would contribute most to the economic welfare of Canadians:

- a formal target for a VMI of 100 on the average
- publication of the Bank's estimate of the trade-off between output and inflation variability that can be achieved with the existing (and presumably best known) techniques of monetary control
- a formal target zone around the trade-off between inflation variability and output variability that monetary policy will aim for

Under this regime, the bank would have a major educational responsibility of two dimensions. The first would be to explain that its objective is a VMI = 100 *on the average* and that departures from 100 will occur because of uncontrollable fluctuations in the economy and as part of the broader objective of achieving the desired balance between price level and output-gap stability.

The second would be to explain the fundamental idea that economists well understand but that few others do, that policy must be judged as a process, not as a series of unrelated events. Each policy event is to be seen against the backdrop of the ongoing process. Each event must be explained and justified with reference to the process.

Policy would be monitored as a process. Bank watchers and commentators would need to learn that second-guessing the bank on whether the interest rate should have been raised (or lowered) on any given occasion is a futile activity. What would matter is whether the bank was operating inside its target distribution for the price level and whether it was using the latitude available inside that distribution to achieve the declared balance between output and price stability.

# Appendix

## Optimal Decision Rules and Equilibrium in the Svensson and Dittmar-Gavin Models

This appendix, which is based on Parkin (2000), probes the free-lunch result and explains why Svensson (1999b), and Dittmar and Gavin (2000), who use the same model, arrive at different decision rules for the central bank. Dittmar and Gavin overlook one part of the margin the central bank faces and so derive a rule that is not optimal. Svensson obtains the optimal decision and equilibrium. But, for reasonable parameter values, the Svensson optimal rule exists only if the central bank places almost no weight on output variability. In contrast, the Dittmar-Gavin rule is available across the whole range of possible relative weights for output variability. This Appendix also shows that the two rules are observationally equivalent if a central bank that uses the Dittmar-Gavin rule appropriately compensates the relative weight placed on output variability. The implication is that in the class of models considered, price-level targeting is superior to inflation targeting, regardless of the relative weight placed on output variability.

## Svensson's Solution for *b*

To derive the central bank's optimal decision rule, rewrite the aggregate supply function as an inflation equation:

$$\pi_t = E_{t-1}\pi_t + \frac{1}{\alpha}(y_t - \rho y_{t-1} - \varepsilon_t). \tag{A1}$$

Use this equation to eliminate the inflation rate from the loss function and obtain:

$$L_t = E_t\left[\sum_{\tau=t}^{\infty}\beta^{\tau-t}\left(\lambda y_\tau^2 + \left(E_{\tau-1}\pi_\tau + \frac{1}{\alpha}\right.\right.\right.$$

$$\left.\left.\left.(y_\tau - \rho y_{\tau-1} - \varepsilon_\tau) - \pi^*\right)^2\right)\right] \tag{A2}$$

The loss in equation (A2) now depends on only the current and future sequence of output gaps, which the bank controls, and on the rational expectation of the current and future inflation rate. It is the treatment of this rational expectation that gives rise to the two decision rules. Svensson treats

this expectation (correctly) as depending on the output gap. Dittmar and Gavin treat it as exogenous. We'll examine the exogenous case later.

To treat the expected inflation rate as endogenous, use the known properties of the decision rule and write it as:

$$\pi_t = \pi^* - a y_{t-1} - b\varepsilon_t, \tag{A3}$$

where $a = b\rho/(1 - \alpha b)$.

The rational expectation of inflation is found using (A3), and is:

$$E_{t-1}\pi_t = \pi^* - a y_{t-1}. \tag{A4}$$

Replace the expected inflation rate in equation (A2) with equation (A4) to obtain:

$$L_t = E_t\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t}\left(\lambda y_\tau^2 + \left(\pi^* - a y_{\tau-1} + \frac{1}{\alpha}\right.\right.\right.$$

$$\left.\left.\left.(y_\tau - \rho y_{\tau-1} - \varepsilon_\tau) - \pi^*\right)^2\right)\right]. \tag{A5}$$

The central bank's problem now is to choose the current output gap, given its current expectations about the sequence of future output gaps, and given the lagged output gap and the current supply shock, to minimize $L$. The first-order condition for this problem is:

$$\alpha\lambda y_t + (1 + \beta(\alpha\rho a - \rho^2))(\pi_t - \pi^*) = 0. \tag{A6}$$

Use the aggregate supply equation with the inflation decision rule to obtain:

$$y_t = \rho y_{t-1} + (1 - \alpha a)\varepsilon_t, \tag{A7}$$

and use equation (A7) in equation (A6) to obtain:

$$\pi_t = \pi^* - \frac{\alpha\lambda\rho}{1 - \beta(\alpha\rho a + \rho^2)} y_{t-1} - \frac{\alpha\lambda(1 - \alpha b)}{1 - \beta(\alpha\rho a + \rho^2)}\varepsilon_t. \tag{A8}$$

Equation (A8), which incorporates the first-order condition, has the same form as equation (A3) and can be solved for the parameters $a$ and $b$. An equivalent solution holds for the price-level targeting case, except that $\pi$ is replaced by $p$, and $\pi^*$ is replaced by $p^*$ at $t$. The solutions for $a$ and $b$ are:

$$a = \frac{(1-\beta\rho^2) - \sqrt{(1-\beta\rho^2)^2 - 4\alpha^2\rho^2\beta\lambda}}{2\alpha\beta\rho}, \quad \text{(A9)}$$

$$b = \frac{a}{\alpha a - \rho}. \quad \text{(A10)}$$

It is the solution for $a$ that restricts the range over which this optimal rule applies. For a real value of $a$ to exist, $\lambda$ must satisfy:

$$\lambda \leq \frac{(1-\beta\rho^2)^2}{4\alpha^2\rho^2\beta}. \quad \text{(A11)}$$

We'll examine how restrictive this existence condition is after looking at the Dittmar-Gavin solution.

## Dittmar-Gavin Solution

Dittmar and Gavin treat the expected inflation rate as exogenous, so their first-order condition for inflation targeting becomes:

$$\alpha\lambda y_t + (1-\beta\rho^2)(\pi_t - \pi^*) = 0. \quad \text{(A12)}$$

For price-level targeting, the equivalent condition is:

$$\alpha\lambda y_t + (1-\beta\rho^2)(p_t - p_t^*) = 0. \quad \text{(A13)}$$

It is a short step from equation (A12) and equation (A13), together with the aggregate supply function, to the solutions for inflation and the output gap. Under inflation targeting, the inflation rate is:

$$\pi_t = \pi^* - \frac{\alpha\lambda\rho}{1-\beta\rho^2} y_{t-1} - \frac{\alpha\lambda}{1-\beta\rho^2+\alpha^2\lambda}\varepsilon_t. \quad \text{(A14)}$$

Under price-level targeting, a suitably adjusted version of equation (A14) describes the price level, and the inflation rate is given by:

$$\pi_t = \pi^* - \frac{\alpha\lambda\rho}{1-\beta\rho^2}(y_t - y_{t-1}) - \frac{\alpha\lambda}{1-\beta\rho^2+\alpha^2\lambda}(\varepsilon_t - \varepsilon_{t-1}), \quad \text{(A15)}$$

and under both regimes, the output gap is:

$$y_t = \rho y_{t-1} + \frac{1-\beta\rho^2}{1-\beta\rho^2+\alpha^2\lambda}\varepsilon_t. \quad \text{(A16)}$$

The variability of the output gap under both regimes is:

$$\sigma_y^2 = \frac{(1-\beta\rho^2)^2}{(1-\rho^2)(1-\beta\rho^2+\alpha^2\lambda)^2}\sigma^2. \quad \text{(A17)}$$

The variability of inflation under inflation targeting is:

$$\sigma_\pi^2 = \frac{\alpha^2\lambda^2}{(1-\rho^2)(1-\beta\rho^2+\alpha^2\lambda)^2}\sigma^2, \quad \text{(A18)}$$

and under price-level targeting is:

$$\sigma_\pi^2 = \frac{2\alpha^2\lambda^2}{(1+\rho)(1-\beta\rho^2+\alpha^2\lambda)^2}\sigma^2. \quad \text{(A19)}$$

## Calibration and the Variability Trade-offs

To find the trade-offs between inflation variability and output variability under the two targeting regimes, we vary $\lambda$ from zero (no weight on the output gap) to a large enough value to make the output variance almost vanish. But, in the case of the Svensson rule, the range of the trade-off is restricted by the existence condition on $\lambda$.

To generate trade-offs and to see how restrictive this existence condition is, we need some parameter values for $\alpha$, $\beta$, and $\rho$. These values depend on the length of the interval indexed as *t* in the model. The longer the interval, the lower are the values of $\beta$ and $\rho$ but the greater is the value of $\alpha$. Dittmar and Gavin suggest that for quarterly data, $\rho = 0.9$ and $\beta = 0.99$. The steady-state slope of the output-inflation trade-off $(1 - \rho)/\alpha = 0.2$, so $\alpha = 0.5$. If we use these numbers, the maximum value of $\lambda$ in Svensson's rule is 0.049.

Extending the length of *t* to a year helps a bit, but not enough. In this case, $\rho = 0.9^4 = 0.656$, $\beta = 0.96$, and $\alpha = 0.5(1 + \rho + \rho^2 + \rho^3) = 1.088$. Using these numbers, the maximum value of $\lambda$ is 0.176.

But, because the available range of $\lambda$ is so small, the Svensson rule can only be used close to the corner of the trade-off, at which the output gap is highly variable—four and one-half to five times the variance of the one-period output shock—and the inflation rate is highly stable, between one-fifth and one-tenth the variance of the one-period output shock.

In contrast, the rule that ignores the endogenous inflation expectation can be used for any value of $\lambda$. By selecting $\lambda \cong 0.4$, inflation and the output gap have approximately equal variances at a bit more than 2 (twice the variance of the one-period supply shock). Selecting $\lambda = 8$ effectively

stabilizes the output gap at zero and makes the inflation rate highly variable around its target value of $\pi^*$. Selecting $\lambda = 0$, "an inflation nutter" central bank stabilizes the inflation rate at its target at the expense of output variability of around five times the variance of the one-period supply shock.

## Observational Equivalence

For given $\lambda$, these two rules are different. But, if the private agents observe the central bank delivering a policy described by the reduced-form decision-rule equations, they have no way of identifying the rule being used. For there exists a $\lambda^*$ in equation (A14) that delivers exactly the same outcome as equation (A8) does. That $\lambda^*$ is given by:

$$\lambda^* = \frac{(1 - \beta\rho^2)}{\alpha\rho} a, \tag{A20}$$

where $a = b\rho/(1 - \alpha b)$, as before.

By using this value for $\lambda$ over the range for which the Svensson equilibrium exists, the Dittmar-Gavin rule traces the same trade-off as that traced by the optimal rule.

Outside the range for which a Svensson solution exists, an appropriate value of $\lambda$ can deliver outcomes with any desired combination of output and inflation variability along the available trade-off.

# References

Akerlof, G.A., W.T. Dickens, and G.L. Perry. 1996. "The Macroeconomics of Low Inflation." *Brookings Papers on Economic Activity* 1: 1–59.

Bank of Canada. 1994. *Economic Behaviour and Policy Choice Under Price Stability.* Proceedings of a conference held by the Bank of Canada, October 1993. Ottawa: Bank of Canada.

Bank of Canada. 1998. *Price Stability, Inflation Targets, and Monetary Policy.* Proceedings of a conference held by the Bank of Canada, May 1997.

Barnett, R. and M. Engineer. 2001. "When Is Price-Level Targeting a Good Idea?" In *Price Stability and the Long-Run Target for Monetary Policy*, 101–43. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Beaudry, P. and M. Doyle. 2001. "What Happened to the Phillips Curve in the 1990s in Canada?" In *Price Stability and the Long-Run Target for Monetary Policy*, 51–82. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Berg, C. and L. Jonung. 1998. "Pioneering Price Level Targeting: The Swedish Experience, 1931–1937." IIES Seminar Paper No. 642, Stockholm.

Bernanke, B.S. 1995. "The Macroeconomics of the Great Depression: A Comparative Approach." *Journal of Money, Credit and Banking* 27 (1): 1–28.

Bernanke, B.S., T. Laubach, F.S. Mishkin, and A.S. Posen. 1999. *Inflation Targeting: Lessons from the International Experience*. Princeton: Princeton University Press.

Black, R., T. Macklem, and D. Rose. 1998. "On Policy Rules for Price Stability." In *Price Stability, Inflation Targets, and Monetary Policy*, 411–61. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Black, S. and W.T. Gavin. 1990. "Price Stability and the Swedish Monetary Experiment." *Economic Commentary.* Federal Reserve Bank of Cleveland. December.

Bowlus, A.J. 1998. "Discussion." In *Price Stability, Inflation Targets, and Monetary Policy*, 219–25. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Brunner, K. and A.H. Meltzer. 1993. *Money and the Economy: Issues in Monetary Analysis*. Cambridge: Cambridge University Press.

Calvo, G.A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

Card, D. and D. Hyslop. 1997. "Does Inflation 'Grease the Wheels of the Labor Market'?" In *Reducing Inflation: Motivation and Strategy*, edited by C.D. Romer and D. Romer, 11–114. Chicago: University of Chicago Press.

Clouse, J., D. Henderson, A. Orphanides, D. Small, and P. Tinsley. 1999. "Monetary Policy When the Nominal Short-Term Interest Rate is Zero." Board of Governors of the Federal Reserve System, Washington, D.C.

Coulombe, S. 1998. "The Intertemporal Nature of Information Conveyed by the Price System." In *Price Stability, Inflation Targets, and Monetary Policy*, 3–28. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Crawford, A. and A. Harrison. 1998. "Testing for Downward Rigidity in Nominal Wage Rates." In *Price Stability, Inflation Targets, and Monetary Policy*, 179–218. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Crawford, A., J.-F. Fillion, and T. Laflèche. 1998. "Is the CPI a Suitable Measure for Defining Price Stability?" In *Price Stability, Inflation Targets, and Monetary Policy*, 39–73. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Dittmar, R. and W.T. Gavin. 2000. "What Do New-Keynesian Phillips Curves Imply for Price-Level Targeting?" *Federal Reserve Bank of St. Louis Review* 82 (2): 21–30.

Dittmar, R., W.T. Gavin, and F.E. Kydland. 1999a. "The Inflation-Output Variability Tradeoff and Price-Level Targets." *Federal Reserve Bank of St. Louis Review* 81 (1): 23–31.

———. 1999b. "Price-Level Uncertainty and Inflation Targeting." *Federal Reserve Bank of St. Louis Review* 81 (4): 23–33.

Duguay, P. 1994. "Some Thoughts on Price Stability versus Zero Inflation." Paper presented at the conference on Central Bank Independence and Accountability, Università Bocconi, Milan, March 1994.

Farès, J. and T. Lemieux. 2001. "Downward Nominal-Wage Rigidity: A Critical Assessment and Some New Evidence for Canada." In *Price Stability and the Long-Run Target for Monetary Policy*, 3–31. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Fillion, J.-F. and R. Tetlow. 1994. "Zero-Inflation or Price-Level Targeting? Some Answers from Stochastic Simulations on a Small Open-Economy Macro Model." In *Economic Behaviour and Policy Choice Under Price Stability*, 129–66. Proceedings of a conference held by the Bank of Canada, October 1993. Ottawa: Bank of Canada.

Fischer, S. 2000. "Opening Remarks." IMF Institute's High-Level Seminar on Implementing Inflation Targets, International Monetary Fund, Washington, D.C., 20–21 March 2000. < URL: http://www.imf.org/external/np/speeches/2000/032000.HTM >.

Fisher, I. 1934. *Stable Money*: *A History of the Movement*. London: Allen & Unwin.

Fortin, P. 1996. "The Great Canadian Slump." *Canadian Journal of Economics* 29 (4): 761–87.

Friedman, M. and A.J. Schwartz. 1963. *A Monetary History of the United States 1867–1960*. Princeton: Princeton University Press.

Goodfriend, M. 2000. "Overcoming the Zero Bound on Interest Rate Policy." *Journal of Money, Credit and Banking* 32 (4) Part 2: 1007–35.

Haldane, A.G. and C.K. Salmon. 1995. "Three Issues on Inflation Targets." In *Targeting Inflation*, edited by A.G. Haldane, 170–201. London: Bank of England.

Hawtrey, R.G. 1925. "Public Expenditure and the Demand for Labour." *Economica* Vol. 5, March, 38–48.

———. 1932. *The Art of Central Banking*. London: Longmans.

Hogan, S. and L. Pichette. 1999. "Downward Nominal Wage Rigidity and the Short-Run Phillips Curve." Research Department, Bank of Canada.

Johnson, D. 1999. "The Effect of Inflation Targeting on the Behaviour of Expected Inflation: Evidence from an 11 Country Panel." Wilfrid Laurier University. Photocopy.

Jonung, L. 1979. "Knut Wicksell's Norm of Price Stabilization and Swedish Monetary Policy in the 1930's." *Journal of Monetary Economics* 5: 459–96.

Kahn, S. 1997. "Evidence of Nominal Wage Stickiness from Microdata." *American Economic Review* 87 (5): 993–1008.

King, M. 1996. "How Should Central Banks Reduce Inflation?—Conceptual Issues." In *Achieving Price Stability*, 53–91. Proceedings of a symposium sponsored by the Federal Reserve Bank of Kansas City, August 1996. Kansas City: Federal Reserve Bank of Kansas City.

Kydland, F.E. and E.C. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85 (3): 473–91.

Laidler, D. (ed.). 1997. *Where We Go from Here: Inflation Targets in Canada's Monetary Policy Regime.* Toronto: C.D. Howe Institute.

Laidler, D. 1999. *Fabricating the Keynesian Revolution*. Cambridge: Cambridge University Press.

Laidler, D. and W.B.P. Robson. 1994. *The Great Canadian Disinflation: The Economics and Politics of Monetary Policy in Canada, 1988–93*. Policy Study 19. Toronto: C.D. Howe Institute.

Laxton, D., N. Ricketts, and D. Rose. 1994. "Uncertainty, Learning and Policy Credibility." In *Economic Behaviour and Policy Choice Under Price Stability*, 173–226. Proceedings of a conference held by the Bank of Canada, October 1993. Ottawa: Bank of Canada.

Lebow, D.E., J.M. Roberts, and D.J. Stockton. 1992. "Economic Performance Under Price Stability." U.S. Board of Governors of the Federal Reserve System Working Paper No. 125.

Lebow, D.E., D.J. Stockton, and W.L. Wascher. 1995. "Inflation, Nominal Wage Rigidity, and the Efficiency of Labor Markets." Federal Reserve Board of Governors Finance and Economics Discussion Series No. 95–45.

Maclean, D. and H. Pioro. 2001. "Price-Level Targeting—The Role of Credibility." In *Price Stability and the Long-Run Target for Monetary Policy*, 153–85. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

McCallum, B.T. 2000. "Theoretical Analysis Regarding a Zero Lower Bound on Nominal Interest Rates." *Journal of Money, Credit and Banking* 32 (4) Part 2: 870–904.

McLaughlin, K.J. 1994. "Rigid Wages?" *Journal of Monetary Economics* 34 (3): 383–414.

Mishkin, F.S. 2001. "Issues in Inflation Targeting." In *Price Stability and the Long-Run Target for Monetary Policy*, 203–22. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Okun, A.M. 1981. *Prices and Quantities: A Macroeconomic Analysis*. Washington: Brookings Institution.

Orphanides, A. and V. Wieland. 1998. "Price Stability and Monetary Policy Effectiveness when Nominal Interest Rates Are Bounded at Zero." Board of Governors of the Federal Reserve System, Washington, D.C.

Parkin, M. 2000. "A Note on Inflation Targeting Versus Price Level Targeting." University of Western Ontario.

Phelps, E.S. 1972. *Inflation Policy and Unemployment Theory: The Cost-Benefit Approach to Monetary Planning*. London: Macmillan.

Reifschneider, D.L. and J.C. Williams. 1999. "Three Lessons for Monetary Policy in a Low Inflation Era." Board of Governors of the Federal Reserve System, Washington, D.C.

Scarth, W. 1994. "Zero Inflation Versus Price Stability." In *Economic Behaviour and Policy Choice Under Price Stability*, 89–119. Proceedings of a conference held by the Bank of Canada, October 1993. Ottawa: Bank of Canada.

Smith, G. 1998. "Discussion." In *Price Stability, Inflation Targets, and Monetary Policy*, 462–8. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Smith, J.C. 2000. "Nominal Wage Rigidity in the United Kingdom." *Economic Journal* 110 (462): C176–95.

Summers, L. 1991. "How Should Long-Term Monetary Policy Be Determined?" *Journal of Money, Credit and Banking* 23 (3), Part 2: 625–31.

Svensson, L.E.O. 1999a. "How Should Monetary Policy Be Conducted in an Era of Price Stability?" Federal Reserve Bank of Kansas City.

———. 1999b. "Price-Level Targeting versus Inflation Targeting: A Free Lunch?" *Journal of Money, Credit and Banking* 31 (3) Part 1: 277–95.

Taylor, J.B. 1979. "Estimation and Control of a Macroeconomic Model with Rational Expectations." *Econometrica* 47 (5): 1267–86.

Vestin, D. 2000. "Price-Level Targeting versus Inflation Targeting in a Forward-Looking Model." IIES, Stockholm University.

Vickery, W.S. 1954. "Stability Through Inflation." In *Post Keynesian Economics*, edited by K.K. Kurihara. New Brunswick, N.J.: Rutgers University Press.

Williams, J.C. 1999. "Simple Rules for Monetary Policy." Board of Governors of the Federal Reserve System, Washington, D.C.

# Discussion

***Peter Howitt***

During the past decade, the economics profession has come a long way in its thinking about monetary policy. When inflation targeting was first adopted in Canada, many economists, myself included, greeted the new policy with skepticism. It seemed rash to stake a central bank's reputation on controlling a variable that was separated from monetary policy by such a long and variable lag. Moreover, the policy ignored the objective of stabilizing fluctuations in real output. As for price-level targeting, hardly anyone thought of it as a practical alternative, because it would have destabilized the economy too much.

As we began to see how inflation targeting actually works, however, many of us started to warm to it. By the end of the 1990s, as Michael Parkin documents in his survey, the profession had developed a solid consensus in favour of inflation targeting. Moreover, a literature had emerged supporting the formerly unthinkable alternative of price-level targeting over inflation targeting, not just as a means of enhancing long-run contracts but as a way to stabilize the economy!

One reason for this change in attitudes is that inflation-targeting central banks have not restricted their attention to the single nominal target they have been assigned but have also taken into account the real consequences of their actions. The targeting regimes in place today are flexible enough to allow monetary policy to aim at stabilizing real output as well as inflation, subject to the usual trade-off. Another reason for the change is that inflation-targeting regimes have provided an unprecedented level of openness, accountability, and transparency in the conduct of monetary policy, all of which are worthy objectives in themselves,

independently of their usefulness in achieving desirable macroeconomic outcomes.

Inflation targeting has also allowed central banks to use discretion where it is most useful, namely in choosing what instruments to use and how. In that respect it has proven to be preferable to monetary targeting, or exchange rate targeting, either of which would tie the central bank's reputation to what might turn out to be a fragile economic relationship. Instead, inflation targeting ties the central bank's reputation to the variable that central banks always have been, and probably always will be, primarily concerned with. When new information refutes the theoretical underpinnings of a particular control technique, there is no need to compromise the bank's reputation and ultimate objectives by introducing face-saving measures. Instead, the bank is free to make use of the new information to improve its performance. The open environment of inflation targeting also gives it a platform for explaining to the public why it is changing tactics, and why this involves no change in its commitment to the regime. In my view, this is the main reason why inflation targeting has worked so well, even though it commits a central bank to controlling something over which it has almost no control from one quarter to the next.

I am in broad agreement with Parkin on these points. I also welcome his suggestion of having a government agency publish a value-of-money index (VMI). 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. Keeping track of a VMI, and also perhaps a slowly growing target value of VMI, with an understanding that monetary policy can be expected to be marginally tighter when the actual VMI exceeds the target than when it falls short, could turn out to be a relatively costless way of dealing with this lingering problem. Publishing the actual VMI in relation to a sensible target path would be a useful first step in that direction.

There are two points, however, on which I think that the recent literature has not taught us as much as Parkin argues. Specifically, (i) I disagree with the approach that Lars Svensson (1999) has been using, and that Parkin has adopted in his survey, for building positive models of an inflation-targeting central bank, and (ii) I think there are serious weaknesses in the recent literature arguing for price-level targeting as opposed to inflation targeting. These two points are related, because one of the main arguments in favour of price-level targeting comes from Svensson's

modelling strategy, in the form of the much-discussed "free-lunch" proposition.

(i) Svensson models an inflation-targeting central bank as having no power of commitment, as in the much-cited Kydland and Prescott (1977) analysis. In my judgment, the time-inconsistency problem at the heart of this analysis is not something that real-world central banks actually confront, except to the extent that they are subject to a great deal of pressure from governments wanting more expansionary policies for short-term political reasons.

More specifically, according to the Kydland-Prescott analysis, a central bank is constantly tempted to engineer surprise movements in inflation, and if it has no power of commitment it will yield to this temptation until inflation has risen high enough to dissuade it from yielding any more. In reality, however, as many of them have testified (see Blinder 1998, for example), central bankers are rarely tempted to produce surprises. Instead, if anything they suffer from a pathological urge to do what they have said they are going to do, because of the capital flight, currency depreciation, and disorderly financial markets that they fear would result if they were to say one thing and do another.

This criticism of the Kydland-Prescott analysis is particularly valid in a regime of inflation targeting, for two reasons. First, the regime is typically endorsed not just by the central bank but also by the government. This and the openness of the regime would make it very difficult for a government to put pressure on a central bank to violate its mandate. Second, the openness of the regime also gives the central bank a clear instrument of commitment. That is, a central bank that was observed always to miss its publicly announced targets would find its reputation severely tarnished. I doubt whether the policy-making committees of any of the world's current inflation-targeting central banks would be willing to pay this price for the dubious benefits of a surprise inflation.

The upshot of this criticism is that, in my view, the Bank of Canada or any of the other inflation-targeting central banks, is best viewed not as if it were subject to the Kydland-Prescott time-inconsistency problem, as in the Svensson analysis, but as if it were seeking to stabilize inflation and real output *with commitment*.

(ii) The main argument in the recent literature supporting price-level targeting is Svensson's free-lunch proposition, a simple version of which runs as follows. Assume that the appropriate social loss function in each period has the form:

$$\lambda y_t^2 + (\pi_t - \pi^*)^2, \lambda > 0, \qquad (1)$$

where $y_t$ denotes the detrended log of real output, $\pi_t$ the rate of inflation, and $\pi^*$ the inflation target. Each period, the central bank chooses $y_t$ and $\pi_t$ to minimize the loss function (1) subject to an expectations-augmented Phillips curve:

$$\pi_t = (1/\alpha)(y_t - y_t^*) + \pi_t^e, \alpha > 0, \quad (2)$$

where $y_t^*$ denotes the detrended log of capacity output, which is assumed to follow a stationary process with some persistence, and $\pi_t^e$ is the public's expectation of inflation conditional on $y_t^*$. The central bank takes $y_t^*$ as given when making its choice. It also takes $\pi_t^e$ as given; that is, it acts *without commitment*.[1]

The expectations-augmented Phillips curve (2) implies that monetary policy will be neutral under rational expectations; real output will always equal capacity output, regardless of the central bank's policy. Thus, the *optimal* monetary policy would be to aim each period at the inflation target; to set $\pi_t \equiv \pi^*$, which would result in a long-run average value of the loss function equal to $\lambda \text{var}(y_t^*)$.

However, the uncommitted central bank will be unable to resist the temptation to attempt to offset the supply shocks that cause capacity output to fluctuate, and the equilibrium inflation rate will be:

$$\pi_t = \pi^* - \lambda\alpha y_t^*.$$

Therefore, persistence in capacity output will translate into persistent deviations of inflation from target value. Since monetary policy is neutral, these deviations of inflation from its target will induce extra variability in inflation without succeeding in reducing the variability of real output; the equilibrium long-run average value of the loss function will be strictly greater than optimal.

It turns out that one way to partially correct for the uncommitted central bank's overeagerness is to instruct it to act like a Zen archer. Don't try to minimize the true loss function (1), but instead aim somewhere else—in this case at the pseudo-loss function:

$$\lambda y_t^2 + (p_{t+1} - p_{t+1}^*)^2, \quad (3)$$

1. The version that Parkin presents assumes that the public and central bank observe only a noisy signal of $y_t^*$. I think the version presented here brings out a little more clearly the logic of the free-lunch argument.

where $p_t$ is the log of the price level and $p_t^* = \pi_t^* t$ is a price-level target. The result of this deliberate misdirection is the price-level-targeting policy function:

$$p_{t+1} = p_{t+1}^* - \lambda\alpha y_t^*.$$

When there is enough persistence in capacity output, the variability of inflation under price-level targeting will be less than under inflation targeting. In effect, the central banker aiming at equation (3) would worry too much about the cumulative price-level error to allow the persistent fluctuations of inflation that are the undoing of the uncommitted inflation targeter.

There are several reasons for my skepticism concerning this proposition. First, it depends on the existence of a stable Phillips curve, and it depends crucially on the dynamic details of that relationship. Although, as Parkin has pointed out, the proposition holds for a New-Keynesian variant of the Phillips curve, it has not been derived from a Phillips curve that exhibits the sort of inflation persistence that Fuhrer and Moore (1995) have shown is an inescapable feature of the data.

More importantly, the free-lunch proposition depends critically on the time-inconsistency framework that I have argued does not fit the typical inflation-targeting central bank. In my view, an inflation-targeting central bank wanting to minimize (1) subject to (2) would have no time-inconsistency problem, and would simply do the optimal thing, namely set $\pi_t = \pi^*$ every period, achieving the optimal outcome.

This is not to argue that we should go back to regarding price-level targeting as not worth considering. On the contrary. I believe that a good monetary policy should attempt to mitigate the long-term price-level uncertainty that would result from a pure inflation-targeting policy. But the design of such a policy needs to take into account two points that the free-lunch argument misses: inflation is persistent and an inflation-targeting central bank can act with commitment.

Consider, for example, the loss function:

$$\lambda_y y_t^2 + \lambda_\pi(\pi_t - \pi^*)^2 + \lambda_p(p_{t+1} - p_{t+1}^*)^2,$$

which takes into account the cost of price-level uncertainty. Suppose further that the Phillips curve takes the form:

$$\pi_t = \pi_{t-1} + (1/\alpha)(y_t - y_t^*),$$

according to which inflation has a momentum of its own, in accordance with the Fuhrer-Moore argument. To simplify the analysis, suppose that there is no persistence in capacity output, so that $y_t^*$ is a serially uncorrelated random variable. A central bank that acted to minimize the long-run average of this loss function subject to this Phillips curve would choose a hybrid policy of the form:

$$\pi_t = \pi^* + \beta(\pi_{t-1} - \pi^*) - \gamma(p_t - p_t^*) - \delta y_t^*,$$

where $\beta$ and $\gamma$ each lie between 0 and 1, and $\delta$ is positive.

The price-level term in this hybrid policy adds an error correction limiting the degree of long-run price-level uncertainty. However, the policy would not attempt to immediately offset any price-level blip the way a pure price-level-targeting policy would do. Instead, with reasonable values of the parameters in the cost function and the Phillips curve, a short-run supply shock that raised inflation would be followed by a gradual reduction in the rate of inflation, with the price-level error being allowed to balloon for several periods before eventually being brought under control. This ballooning is induced by inflation-persistence in the Phillips curve, which would impose a high penalty—in the form of lost output—on a policy that immediately reversed the price-level error.

Such a hybrid policy need not be much different from inflation targeting as it currently works. It could be implemented simply by committing the central bank to being slightly more restrictive than usual when the price level has risen above its long-run target path. Even this small adjustment to current policy could have a major long-run impact, because compared with a pure inflation-targeting policy with no error correction, it would reduce the long-run variance of the price level all the way from infinity to something finite. In my view, this would be a considerable improvement, even though it is far less radical a change than recent proponents of price-level targeting have been advocating.

## References

Blinder, A.S. 1998. *Central Banking in Theory and Practice.* Cambridge, MA: MIT Press.

Fuhrer, J. and G. Moore. 1995. "Inflation Persistence." *Quarterly Journal of Economics* 110 (February): 127–93.

Kydland, F.E. and E.C. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85 (June): 473–91.

Svensson, L.E.O. 1999. "Price-Level Targeting versus Inflation Targeting: A Free Lunch?" *Journal of Money, Credit and Banking* 31 (3) Part 1: 277–95.

# Discussion

*W. Craig Riddell**

Michael Parkin has written a valuable survey of progress in our understanding of price stability, with particular emphasis on recent contributions. The paper covers a range of issues; however, my comments will be more narrowly focused. I first discuss Canada's experience in the 1990s and its relevance to the costs of achieving low levels of inflation. This discussion leads naturally to section 4.2 ("labour market lubricant") of Parkin's paper, and the issue of how the characteristics of wage setting and other labour market institutions influence the costs of achieving and maintaining low rates of price inflation.

## Canada's "Lost Decade"

Several years ago, the Queen of England lamented the previous "annus horribilus." Canadians feel much the same about the 1990s. The decade began with the prolonged recession of 1990–92, the worst Canadian slump since the Great Depression (Fortin 1996). Recovery from this downturn was also painfully slow. Although the 1990s ended with several years of solid growth, the decline in real incomes during the first half of the decade was sufficiently large that only by the end of the decade had the living standards of Canadians returned to their pre-recession (1989) levels. Canadian living standards also fell relative to those in the United States and several other OECD countries.

Two leading explanations for Canada's poor economic performance in the 1990s—in particular for the relatively weak performance during the first half of the decade—have been advanced. For the sake of simplicity, I

* I am grateful to to Paul Beaudry and Thomas Lemieux for helpful discussions.

refer to these as the "Fortin" (see Fortin 1996, 2001) and "Bank of Canada" (see Thiessen 1996 and Freedman and Macklem 1998) positions. Fortin attributes the "great Canadian slump" of 1990–92 and the slow recovery during 1992–95 principally to monetary policy (and the fiscal contraction induced by restrictive monetary policy). The Bank of Canada blames the poor performance on a combination of widespread structural change associated with globalization and technological advances and high interest rates imposed by financial markets because of factors such as high levels of government debt and political instability.

Beginning in the early 1990s, the Bank of Canada sought to reduce inflation from the then-prevailing rate of about 5 per cent. Formal "inflation-reduction targets" were introduced in 1991. They provided for a gradual reduction in inflation, with midpoint ranges of 3 per cent, 2.5 per cent, and 2 per cent over the 1991–95 period. They were subsequently replaced by "inflation-control targets" with a range of 1 to 3 per cent.

The objective of reducing inflation was achieved; the rate of price inflation declined from about 5 per cent in 1990 to 1 to 2 per cent in 1994. However, the cost of doing so was unusually high. Both Debelle (1996) and Bernanke et al. (1999) estimate the sacrifice ratio for the 1990–93 disinflation to have been more than double that experienced by Canada in previous disinflations such as those in the mid-1970s and early 1980s (when the initial inflation rate was much higher). Canada's sacrifice ratio for the 1990–93 disinflation was also higher than that experienced by Australia and New Zealand in reducing inflation to a similar extent during the late 1980s and early 1990s (Debelle 1996 and Bernanke et al. 1999).

Building on work by Akerlof, Dickens, and Perry (1996), Fortin (1996, 2001) argues that the long-run Phillips curve is vertical at moderate to high inflation rates but becomes downward-sloping at low rates of inflation. The principal source of the downward-sloping portion is resistance to nominal-wage cuts—perhaps because of money illusion or notions of fairness in wage setting. In these circumstances, reducing inflation from a moderate to a low level, as Canada did in the early 1990s, has both a temporary cost—the lost output and employment needed to bring about a reduction in inflation—and a permanent cost due to the non-vertical Phillips curve at low rates of inflation.

According to this view, downward wage rigidity has important implications not only for the costs of disinflation but also for the long-run target rate of inflation. In the presence of nominal-wage rigidities, a moderate rate of inflation will contribute to higher output and employment by facilitating the relative wage adjustments needed for the efficient allocation of labour. The question of whether, and to what extent, inflation "greases the wheels" of the labour market by facilitating relative and real-wage changes is thus of

fundamental importance for choosing the long-run target for monetary policy. The greater the degree of nominal-wage rigidity, the higher the output and employment costs of a low inflation target.

An important part of Parkin's survey, and of these comments, assesses the extent to which recent research supports the Akerlof, Dickens, and Perry (1996) and Fortin (1996) emphasis on nominal-wage rigidity and its implications for macroeconomic behaviour with low inflation. As the subsequent discussion will make clear, I am less dismissive of this position than is Parkin. Nominal-wage rigidities do appear to be an empirically important phenomenon in a country like Canada. However, the jury is still out on whether these rigidities have any real economic consequences.

Although there have been useful recent contributions to the debate about the empirical importance of downward wage rigidities, the explanation for Canada's relatively weak economic conditions during the early 1990s has yet to be satisfactorily resolved (Riddell 1999). This is unfortunate because there is an obvious need to better understand the factors that contributed to Canada's poor performance during that period. Furthermore, Canada's weak performance relative to the United States—where the authorities did not attempt to reduce inflation to as great an extent as did authorities in Canada during the early 1990s—provides powerful evidence against the pursuit of low inflation.

As discussed below, recent contributions do not provide substantial additional support for the view that downward wage rigidities are a major factor contributing to high unemployment in periods of low inflation. At the same time, the evidence supporting the Bank of Canada view that the Canadian economy was subject to an unusually large amount of structural adjustment during the early 1990s is less than overwhelming. Riddell (1997) reviews the Canadian experience and finds some evidence of increased structural adjustment during the 1980s and 1990s, particularly the 1990s, but the differences relative to the 1960s and 1970s are not large. Picot and Heisz (2000) report that increased downsizing during the 1990s recovery did contribute to slower employment growth. However, the increased downsizing was concentrated in the public and consumer-services sectors and thus difficult to attribute to globalization and technological innovations. Furthermore, the likelihood of a worker experiencing a permanent layoff did not increase in the first half of the 1990s relative to earlier periods. Slow employment growth mainly reflected reduced hiring rather than an increased incidence of layoffs.

## Aggregate Labour Market Behaviour

Before turning to the research, it is worth examining how much flexibility we see in the aggregate labour market. Figures 1 and 2 show average weekly wages by age group for females and males working full year, full time (FYFT). The data are taken from the Survey of Consumer Finances (SCF), an annual survey carried out in April of each year that gathers retrospective information on annual income and weeks worked during the previous calendar year, thus providing a measure of average weekly earnings. The restriction to FYFT reduces substantially, but does not eliminate, the contribution of hours worked per week to weekly earnings.

These data show that the real wages of different age groups moved very closely together from 1969 to the early 1980s, but that they have subsequently followed markedly different paths. This suggests a moderate amount of real and relative wage flexibility among workers of different ages and labour market experience. In their comparison of the behaviour of average weekly earnings of prime-age males across education groups in Canada and the United States, Kuhn and Robb (1998) also find evidence of considerable downward wage adjustment, especially among the less skilled, in Canada during the 1980s.

Of course, these data are based on a series of cross-sections—they do not result from following a given individual over time. Note, however, that Beaudry and Green (2000), who use the method of "artifical cohorts" to follow representative samples of specific age cohorts through time, also find considerable downward real-wage adjustment with the SCF data, especially among younger cohorts.

To complete this impressionistic picture from the aggregate labour market, it is important to recall that the minimum wage declined dramatically (relative to the average manufacturing wage) in Canada during the 1980s (Benjamin 1996). Recent research by DiNardo, Fortin, and Lemieux (1996) and Fortin and Lemieux (2000) has provided striking evidence for the United States and Canada, respectively, that minimum wages alter the shape of the wage distribution, resulting in considerable mass at the minimum and less mass at lower wage rates. Thus, some of the decline in real wages during the 1980s, as illustrated in Figures 1 and 2, is probably attributable to the decline in the relative minimum wage rather than to downward wage flexibility in an unchanged institutional environment. I will return to the importance of considering institutional changes.

**Figure 1**
**Real annual earnings by age, FYFT females, 1969 = 100**
Women working full year, full time

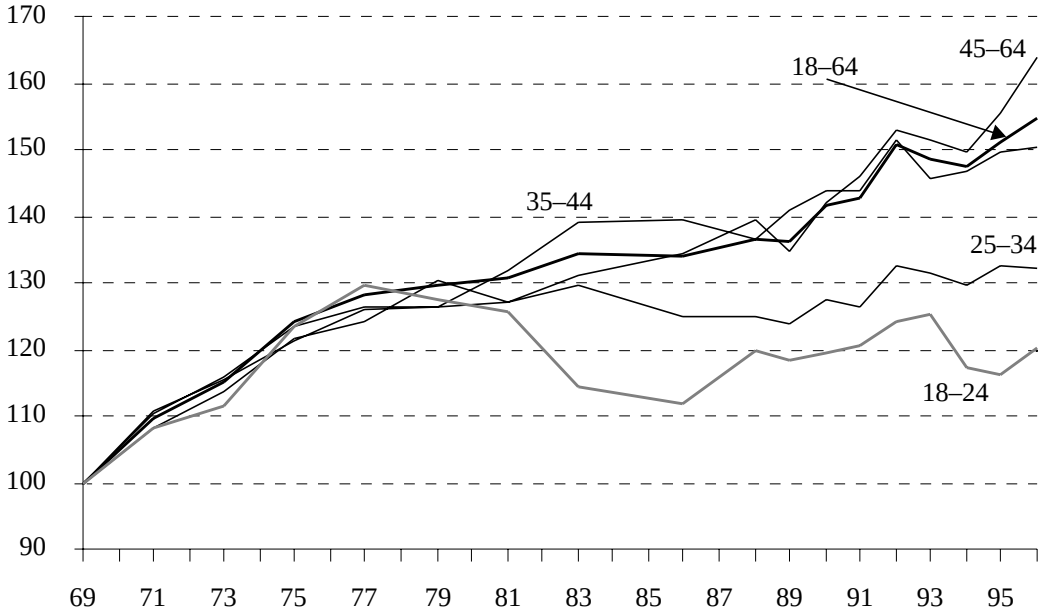


**Figure 2**
**Real annual earnings by age, FYFT males, 1969 = 100**
Men working full year, full time

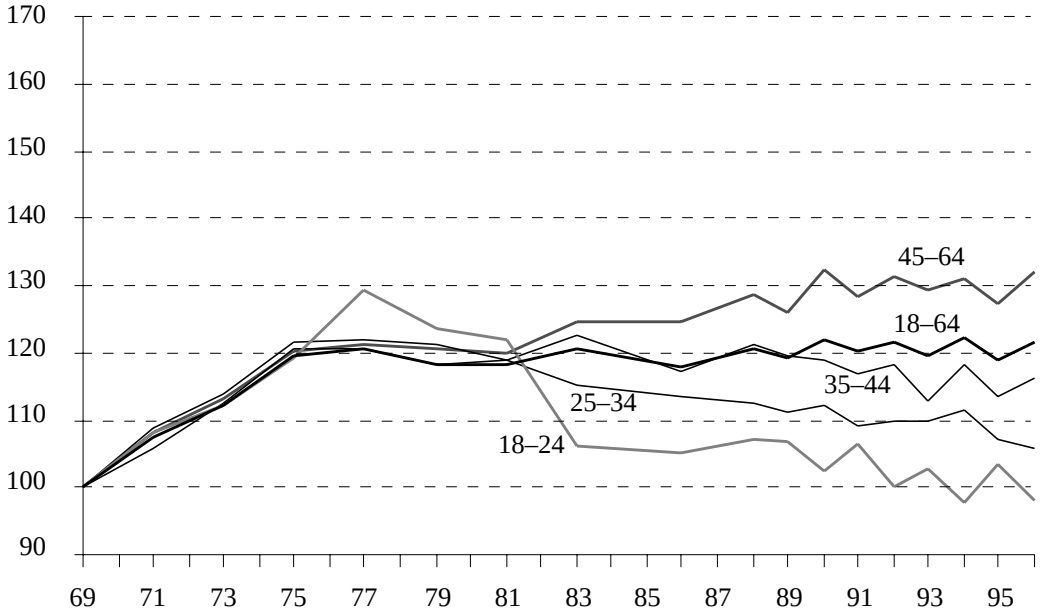

## Empirical Evidence on Downward Wage Rigidity

As Parkin notes, a moderate amount of research had been done on this issue by the time of the Bank of Canada conference of 1997 (Bank of Canada 1998). Contributions included those of McLaughlin (1994), Card and Hyslop (1997), and Kahn (1997) for the United States and Crawford and Harrison (1998) for Canada. Subsequent contributions include Christofides and Stengos (2000) and Farès and Lemieux (2001) for Canada, McLaughlin (1999) for the United States, and Smith (2000) for the United Kingdom.

This research generally proceeds as follows. First, the distribution of nominal-wage changes is characterized. The period over which the change is calculated varies according to the data available, with one year as the most common interval. The wage-change distribution typically has a large spike at zero. The tail of the distribution involving negative wage adjustments also generally has less mass than its counterpart for above-average positive wage adjustments. The next step (not pursued in all studies) is to estimate how much of the spike at zero wage change and the relative absence of negative wage adjustments can be attributed to downward wage rigidity versus other factors. This adjusted measure provides an estimate of the quantitative importance of downward nominal-wage rigidity in the labour market in question.

The importance of downward wage rigidity is likely to depend on a country's institutional arrangements for wage determination. A country in which minimum wages are high relative to the average wage and apply to most of the labour force can be expected to display more downward wage rigidity than a country in which minimum wages are low or in which few workers are covered by the minimum. Similarly, we expect the extent of union representation of the labour force to influence the degree of downward wage flexibility in the economy.

Changes in labour market institutions can therefore be expected to alter the degree of downward wage rigidity in the economy. For the purposes of medium-term macroeconomic policy—such as setting targets for price inflation—the existing institutional arrangements should be taken as given. It is possible that pursuing low inflation or even price stability would eventually alter wage-setting arrangements, but this outcome is highly uncertain. Similarly, and of particular relevance here, the degree of downward wage rigidity in an economy should be characterized for a specific set of wage-setting institutions.

The three countries for which studies of this type have been carried out are Canada, the United States, and the United Kingdom (see Parkin's Table 1). In assessing the relevance of these findings for Canada, it is important to keep in mind that the United States and the United Kingdom

experienced significant changes in labour market institutions during the 1980s and 1990s (the sample period for most of these studies). There is substantial evidence that these changes contributed to increased wage inequality in these countries, and that differences in the rate of institutional change were an important factor influencing the rate of growth of inequality.[1] For example, the United Kingdom, which experienced the most substantial decline in unionization, also experienced the largest increase in earnings inequality among these three countries, whereas Canada, which had the smallest decline in unionization, also had the smallest increase in inequality. Changes in minimum wages in the United States and Canada have also influenced trends in wage inequality in the two countries (Fortin and Lemieux 1997; DiNardo and Lemieux 1997; Green and Paarsch 1997).

Because of the substantial changes in unionization and other labour market institutions that took place during the 1980s and 1990s in the United States and the United Kingdom (particularly the latter), we expect these two countries to display less downward wage rigidity than would have been the case under unchanging wage-setting arrangements. Some of the downward wage changes observed during the period are due to such changes in the institutions affecting wage determination, rather than to the effects of demand and supply shocks in an unchanging institutional environment. Consequently, the U.S. and U.K. experiences are likely to overstate the extent of downward wage flexibility.

Two principal types of data have been used to assess the degree of nominal-wage rigidity. Union contract data, available in Canada for all agreements involving 500 or more employees, have the advantage of precise information on the beginning and end dates of contracts and detailed information on negotiated wage rates during the contract. Their main disadvantages are the limited coverage (only a subset of the union sector, representing less than 10 per cent of the labour force) and their use of the base wage rate, usually the wage associated with an entry-level job and one that is not necessarily representative of the full bargaining unit. The second, and much more extensively exploited, source of information consists of panel data on individuals, available from the Labour Market Activity Study (LMAS) and the Survey of Labour and Income Dynamics (SLID) in Canada, the Panel Study of Income Dynamics (PSID) and matched files of the Current Population Study (CPS) in the United States, and the British Household Panel Survey (BHPS) in the United Kingdom. Several of these data sources are representative of the respective populations (LMAS and

1. See, for example, Fortin and Lemieux (1997) for the United States, DiNardo and Lemieux (1997) for a comparative study of Canada and the United States, and Gosling and Lemieux (2000) for a comparative study of the United States and the United Kingdom.

SLID for Canada, CPS for United States, BHPS for the United Kingdom), have large sample sizes, and follow individuals for one to six years.

While it is straightforward to characterize the distribution of wage changes, the challenge is to estimate what the distribution of wage adjustments would have looked like in the absence of nominal-wage rigidity. The difference between the observed and estimated counterfactual wage-change distributions is the implied estimate of the impact of nominal rigidity on the distribution of wage adjustments.

While economists recognize that the counterfactual is inherently unobservable, classical random-assignment designs provide an ideal solution because they furnish an unbiased estimate of the counterfactual. However, even with non-experimental data, one can frequently obtain estimates of the counterfactual in which considerable confidence can be placed. In this case, however, we are on shakier ground, because we have very little empirical evidence on the shape of the "pure" wage-change distribution—that which would prevail in the absence of nominal rigidities. Consequently, it seems likely that disagreement on the interpretation of available evidence will continue in this area.

Parkin places considerable weight on Smith's new evidence for the United Kingdom. BHPS data are of high quality and have some desirable features for the purpose at hand. Furthermore, some of the findings are striking and perhaps even surprising. For these reasons, it is worth assessing the extent to which her results are relevant for Canadian monetary policy. I discuss this question in four stages: (i) more rapid institutional change in the United Kingdom relative to Canada, (ii) institutional differences between the United Kingdom and Canada, (iii) nature of the counterfactual assumptions, and (iv) the results regarding measurement error in panel data on earnings.

As in previous research, Smith finds a substantial spike at zero in the wage-change distribution—about 9 per cent of employees who remain in the same job from one year to the next experience no wage change. She concludes, however, that about 90 per cent of this spike—8 out of 9 percentage points—can be attributed to "symmetric" causes (long-term contracts, measurement error, and rounding). Thus, only a small fraction of the spike is attributed to asymmetric downward wage rigidity.

Parkin emphasizes the similarity in the raw data on downward wage rigidity in Canada, the United States, and the United Kingdom. However, it is important to consider the prevailing rate of inflation—a key determinant of the location of the distribution of wage changes and thus of the percentages of nominal-wage cuts and rigid wages—in making comparisons

over time within a country or across countries at a point in time.[2] When this is done, there are more downward wage adjustments in the U.K. data (23 per cent of workers experience a wage cut over the sample period) than in the United States at comparable rates of inflation.

During the 1980s and 1990s, the institutional structures of the U.K. labour market changed dramatically (Gosling and Lemieux 2000). Most prominent was the steep decline in unionization, which was especially dramatic after the mid-1980s. In addition, wages councils (who set minimum rates of pay in some low-paying industries) were weakened and finally abolished in 1993. Other changes that affected the determination of workers' pay included declines in public sector employment and increased use of contracting out and competitive tendering of public services. These dramatic changes probably contributed to the substantial amount of downward wage flexibility evident in the United Kingdom in the 1990s.[3]

In addition to the more rapid changes in wage-setting arrangements and structures, there are a number of potentially important differences in Canadian and U.K. institutional and legal arrangements. Canada has a highly legalistic system governing union organizing and collective bargaining, in contrast to the United Kingdom, where custom plays a much larger role. For example, all Canadian jurisdictions use a "Wagner Act" type of model of union formation and recognition. If the union can demonstrate support (through either card-signing procedures or secret-ballot elections) by a majority of the members of the bargaining unit, the union will be certified as the bargaining agent for all of its members. In contrast, in the United Kingdom, union recognition is a voluntary act on the part of employers, and even when a union is present there is nothing to prevent individual bargaining for some employees.[4] Similarly, in the area of union security arrangements, all Canadian jurisdictions require at least the agency shop or Rand formula (not everyone is required to join the union, but all members of the bargaining unit pay union dues) and the union shop (all employees are required to join the union within a certain period after being hired) is common. In the United Kingdom, open-shop arrangements are common, and many employees in the "union sector" are not union members. In the BHPS data, about 40 per cent of the private sector is covered by collective agreements—versus about 20 per cent in Canada—but an astonishing 15 per cent are covered non-members (Hildreth 2000). In

2. For example, Card and Hyslop (1997) find that the fraction of workers with rigid wages is strongly negatively related to the inflation rate.

3. Smith's (2000) study covers the period of 1992–96.

4. Recently (since the end of Smith's sample period), the United Kingdom's Blair government introduced a new labour code with a formal union recognition process similar to that in the United States and Canada.

Canada, the proportion of covered non-members is about 2.5 per cent and in the United States, the comparable figure is under 2 per cent. Thus, almost half the members of the bargaining unit are not union members, and there is no legal requirement for the firm to bargain with the union, circumstances that differ markedly from the typical Canadian bargaining situation.

Another relevant difference is the absence in the United Kingdom of formal wage contracts with specific initiation and termination dates. Although annual negotiations are evidently a common arrangement, this difference makes it difficult to attribute the amount of zero wage changes associated with "long-term contracts."

In the BHPS, the period between interviews varies, and this variation is used by Smith to estimate the counterfactual—in this case, the proportion of zero wage change attributable to long-term contracts. The "percent rigid" shows little change in months 13 and 14 relative to month 12 but a dramatic drop for periods between interviews of more than 14 months. Smith uses the data from interviews spaced more than 14 months apart as the estimate of the proportion of rigidity left after the influence of long-term contracts has been removed, leading her to conclude that more than half of the 9 per cent spike at zero is due to "symmetric" long-term contracts [(9.0-4.0)/ 9.0 = 56 per cent]. However, this estimate is sensitive to the assumption that the observations for interviews spaced 13 and 14 months apart are governed to the same extent as those spaced 12 months apart by long-term contracts.

The general presumption in this literature has been that measurement error overstates the amount of true wage change. Indeed, some authors have argued that most of the downward wage changes observed in the United States can be attributed to measurement error (Akerlof, Dickens, and Perry 1996). Smith's findings regarding measurement error are interesting because they suggest that the reverse may be the case—i.e., that measurement error in wages serves to exaggerate wage rigidity. The separate influence of measurement error is identified because some respondents check their payslips when reporting their earnings, while others do not. The observations for the "payslip seen" subsample are assumed to be measured without error.

Smith's findings are potentially important. However, the assumption that there is no measurement error in the "payslip seen" subsample may be incorrect. The wage concept used in the study is the worker's "usual gross pay (including overtime and bonuses)." She specifically uses usual gross pay rather than latest pay "because latest pay might be distorted by unusual bonuses, overtime payments and so on" (p. C179). However, the payslips consulted by the respondent will contain information on actual pay in the relevant period, not on "usual pay." It seems possible that the "payslip seen" subsample will contain measurement error for the wage concept being used.

Smith's findings are clearly relevant for the U.K. labour market, although it is unclear how much of the observed wage flexibility is attributable to changes in institutional structures and how much to the way wages adjust to demand and supply shocks within a fixed institutional environment. For the reasons discussed above, however, we should be cautious in applying these results to the Canadian setting.

## Real Consequences of Downward Wage Rigidity

The recent contributions to the empirical literature on the distribution of wage changes at the individual level have added to our understanding, although they do not appear to require any significant revisions to the conclusions reached previously by Card and Hyslop (1997) for the United States and by Crawford and Harrison (1998) for Canada. We have made less progress, however, in analysing the implications for market behaviour of wage rigidity at the level of the individual worker. That is, do nominal wage rigidities result in lower output and higher unemployment at the level of an industry, region, or economy?

The "wage rate" that matters for resource allocation is the wage or cost of labour that governs decisions on labour supply and demand. Even if there is considerable nominal-wage rigidity, there are various margins through which employers and employees can adjust to external shocks. They include non-wage compensation, working conditions, and employee effort. The question is whether some of these margins respond to external shocks in ways that alter the true underlying wage rate without necessitating a change in the nominal wage.

A key conclusion of the careful study by Card and Hyslop (1997) was that there is a moderate amount of downward wage rigidity in the U.S. economy. They estimated that downward nominal rigidities in a typical year in the 1980s "held up" real-wage changes of workers by up to 1 per cent per year. However, their analysis of state-level data found little evidence that the rate of wage adjustment across local labour markets is more rapid in a high-inflation environment. This result is consistent with the view that nominal-wage rigidities are quantitatively significant, but that the real effects of such rigidities are reduced by the adjustments of workers and firms on other margins.

Two papers in this volume analyze the prediction that the Phillips curve should become flatter in a low-inflation environment as downward wage rigidities constrain a greater number of employers and employees. Farès and Lemieux (2001) estimate real-wage Phillips curves using SCF data over the period 1981–97. They exploit both the time-series variation for Canada as a whole and the variation in economic conditions across

provinces and over time. The aggregate results provide some weak support for the hypothesis in that the negative relationship estimated for the 1980s is no longer statistically significant during the low-inflation regime of the 1990s. But their analysis of provincial data provides no support for the prediction of a flatter Phillips curve during the 1990s. Overall, their results for Canada tend to reinforce the conclusions of Card and Hyslop (1997) for the United States—there is only weak evidence that downward wage rigidities have negative employment consequences in periods of low inflation.

Beaudry and Doyle (2001) study the behaviour of price inflation versus output gap Phillips curves in Canada and the United States over the past four decades. Their estimated Phillips curves for Canada and the United States are roughly similar. They find that in both countries there has been a substantial reduction in the slope of the Phillips curve during the past two decades (since the late 1970s). In addition, a further substantial decline is observed in the 1990s, although it begins somewhat earlier in the United States (around 1988) than in Canada (around 1992). The much flatter Phillips curve of the 1990s is consistent with the presence of downward wage rigidities constraining firm and worker behaviour. The finding that the slope of the relationship declined to a similar extent in both countries does not fit with expected behaviour, however, given that Canada pursued and achieved much lower rates of inflation.

Beaudry and Doyle attribute the decline in the slope of the Phillips curve to improvements in the process by which central banks gather and react to information about real developments in the economy. They develop a model in which the monetary authorities are imperfectly informed about economic developments and both gather information and set policy optimally given this information. The model implies a reduced-form Phillips-curve relationship, and the slope of this relationship is predicted to become flatter with improvements in the information-gathering function of the central bank. One reason they prefer this interpretation of the observed flattening of the price-output Phillips curve in both countries is that the degree of non-linearity of the Phillips curve did not increase during the more recent lower inflation period, as would be expected if the cause of the flatter slope were downward wage rigidities. The result that the slope of the Phillips curve declined to a similar extent in both Canada and the United States also appears to be consistent with the Beaudry-Doyle interpretation.

The evidence regarding real consequences of downward wage rigidities in Canada is mixed and inconclusive. The Farès-Lemieux study of real wage versus unemployment Phillips curves during the 1980s and 1990s provides no clear evidence that the relationship has become flatter during the low-inflation regime of the 1990s. In contrast, Beaudry and Doyle's study of

price-output Phillips curves over the past four decades finds a substantial decline in the slope of the relationship since the late 1970s, with a further marked decline during the 1990s. This outcome is consistent with downward wage rigidities interfering with relative wage adjustments during the low inflation regime of the 1990s, but it is also consistent with other mechanisms, including the one advanced by Beaudry and Doyle. The finding that the slope of the Phillips-curve relationship has declined to a similar extent in Canada and the United States during the 1990s is more difficult to reconcile with the wage-rigidity explanation.

Although there is not yet convincing evidence of substantial real consequences of nominal wage rigidities, I would not go so far as Parkin, who concludes (quoting Smith, 2000, p. C194) that "We should not look to the labour market as the source of major employment costs of zero inflation that would justify a positive inflation target." The Canadian experience during the 1990s—when we pursued low rather than zero inflation—and the fact that we still do not have a convincing explanation for the relatively weak economic conditions experienced during much of that decade, suggest to me that it would be dangerous to set inflation targets lower than those now in place.

## References

Akerlof, G.A., W.T. Dickens, and G.L. Perry. 1996. "The Macroeconomics of Low Inflation." *Brookings Papers on Economic Activity* 1: 1–59.

Bank of Canada. 1998. *Price Stability, Inflation Targets, and Monetary Policy.* Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Beaudry, P. and M. Doyle. 2001. "What Happened to the Phillips Curve in the 1990s in Canada?" In *Price Stability and the Long-Run Target for Monetary Policy,* 51–81. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Beaudry, P. and D.A. Green. 2000. "Cohort Patterns in Canadian Earnings: Assessing the Role of Skill Premia in Inequality Trends." *Canadian Journal of Economics* 33 (4): 907–36.

Benjamin, D. 1996. "Do Minimum Wages Really Matter?" *Policy Options* (July/August) 37–41.

Bernanke, B.S., T. Laubach, F.S. Mishkin, and A.S. Posen. 1999. *Inflation Targeting: Lessons from the International Experience*. Princeton: Princeton University Press.

Card, D. and D. Hyslop. 1997. "Does Inflation 'Grease the Wheels of the Labor Market'?" In *Reducing Inflation: Motivation and Strategy,* edited by C.D. Romer and D. Romer, 11–14. Chicago: University of Chicago Press.

Christofides, L.N. and T. Stengos. 2000. "The Symmetry of the Wage Change Distribution: Nonparametric Tests Using Contract and Survey Data." University of Guelph. Manuscript.

Crawford, A. and A. Harrison. 1998. "Testing for Downward Rigidity in Nominal Wage Rates." In *Price Stability, Inflation Targets, and Monetary Policy,* 179–218. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Debelle, G. 1996. "The Ends of Three Small Inflations: Australia, New Zealand and Canada." *Canadian Public Policy* 22 (1): 56–78.

DiNardo, J. and T. Lemieux. 1997. "Diverging Male Wage Inequality in the United States and Canada, 1981–88: Do Institutions Explain the Difference?" *Industrial and Labor Relations Review* 50 (4): 629–51.

DiNardo, J., N.M. Fortin, and T. Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach." *Econometrica* 64 (5): 1001–44.

Farès, J. and T. Lemieux. 2001. "Downward Nominal-Wage Rigidity: A Critical Assessment and Some New Evidence for Canada." In *Price Stability and the Long-Run Target for Monetary Policy*, 3–29. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Fortin, N.M. and T. Lemieux. 1997. "Institutional Changes and Rising Wage Inequality: Is There a Linkage?" *Journal of Economic Perspectives* 11: 75–96.

———. 2000. "Income Redistribution in Canada: Minimum Wages Versus Other Policy Instruments." In *Adapting Public Policy to a Labour Market in Transition*, edited by W.C. Riddell and F. St-Hilaire, 211–44. Montreal: Institute for Research on Public Policy.

Fortin, P. 1996. "The Great Canadian Slump." *Canadian Journal of Economics* 29 (4): 761–87.

———. 2001. "Inflation Targeting: The Three Percent Solution." *Policy Matters* 2 (1). Montreal: Institute for Research on Public Policy.

Freedman, C. and T. Macklem. 1998. "A Comment on 'The Great Canadian Slump'." *Canadian Journal of Economics* 31 (3): 646–65.

Gosling, A. and T. Lemieux. 2000. "Labour Market Reforms and Changes in Wage Inequality in the United Kingdom and the United States." University of Essex and University of British Columbia. Photocopy.

Green, D.A. and H.J. Paarsch. 1997. "The Effect of the Minimum Wage on the Distribution of Teenage Wages." University of British Columbia, Department of Economics, Discussion Paper No. 97–02.

Hildreth, A.K.G. 2000. "Union Wage Differentials for Covered Members and Nonmembers in Great Britain." *Journal of Labour Research* 21 (1): 133–47.

Kahn, S. 1997. "Evidence of Nominal Wage Stickiness from Microdata." *American Economic Review* 87 (5): 993–1008.

Kuhn, P. and A.L. Robb. 1998. "Shifting Skill Demand and the Canada-US Unemployment Gap: Evidence from Prime-Age Men." *Canadian Public Policy* 24 (Supplement, February): S170–91.

McLaughlin, K.J. 1994. "Rigid Wages?" *Journal of Monetary Economics* 34 (3): 383–414.

———. 1999. "Are Nominal Wage Changes Skewed Away from Wage Cuts?" *Federal Reserve Bank of St. Louis Review* 81 (3): 117–32.

Picot, G. and A. Heisz. 2000. "The Performance of the 1990s Canadian Labour Market." *Canadian Public Policy* 26 (Supplement, July): S7–26.

Riddell, W.C. 1997. "Structural Change and Adjustment in the Canadian Labour Market." In *Structural Change and Labor Market Flexibility: Experience in Selected OECD Economies*, edited by H. Siebert. Kiel Institute of World Economics.

———. 1999. "Canadian Labour Market Performance in International Perspective." *Canadian Journal of Economics* 32 (5): 1097–134.

Smith, J.C. 2000. "Nominal Wage Rigidity in the United Kingdom." *Economic Journal* 110 (462): C176–95.

Thiessen, G. 1996. "Does Canada Need More Inflation to Grease the Wheels of the Economy?" *Bank of Canada Review* (Winter 1996–1997): 63–9.

# General Discussion*

In his response to Peter Howitt, Michael Parkin noted a dichotomy between "simple," transparent things that monetary policy should aim for and should inform the public of, and the "hard," complex details that underpin the policy and that should, in his view, be subject to central bank discretion. He interpreted much of Howitt's discussion as referring to the complex details of policy. He argued that the Bank should study the nature of the trade-off between the variability of inflation and output and give the public some indication of this range.

Charles Freedman stated that the Bank had provided the public with some of the complex details Michael Parkin referred to. In 1991, for example, a background document to the inflation targets spelled out how monetary policy would react to certain supply shocks that might push inflation away from target. Freedman was uncomfortable with Parkin's dichotomy, suggesting that it could imply that policy rules used inside the Bank might differ from those publicized outside the Bank. He thought that the general public would understand inflation targets much better than price-level targets.

William Robson favoured specifying inflation targets as a range rather than as a single point. If the target was for a specific level of inflation rather than a range, the Bank would constantly be explaining why it had not hit its target. He noted that much of the Bank's research appears to treat the inflation target as if it were for a specific level of inflation rather than a range, and he encouraged the Bank to build into its own models the recognition that the target is a range.

* Prepared by Kim McPhail.

John Murray asked whether the conventional view that the 1990s was a period of excess supply in Canada could be reconciled with the view that nominal wages are flexible rather than rigid downward. Howitt responded that the opposite puzzle existed in the United States: Why was inflation there so stable given such low unemployment? Parkin's view was that there had been very little excess supply in the 1990s. Instead, the NAIRU had shifted up because of churning in the labour market and changes in unemployment insurance early in the decade.

Pierre Duguay said that reaction functions that include price level, inflation rates, and the output gap as factors looked like old-style control rules that incorporated proportional, derivative, and integral factors. Howitt commented that rules that combine price-level and inflation factors are common to most optimal rules.

Jeffrey Fuhrer observed that there was little public support for further declines in inflation rates and questioned whether central banks should pursue such a policy. Parkin suggested that the current period of low inflation could be temporary. Howitt thought that a policy of price-level targeting (as opposed to inflation targeting) could be implemented without much additional cost to the economy.

Robert Lafrance suggested that the period of successful price-level targeting in Sweden during the 1930s was achieved mostly through good luck, not good policy. For example, the Bank of Sweden avoided banking crises by providing liquidity to the banks; before moving to a fixed exchange rate, Sweden devalued, thereby generating a trade surplus for most of the decade. Parkin viewed things differently, arguing that some of these favourable factors in Sweden during the 1930s were the consequence, not the cause, of a successful monetary policy.