Should Forward Guidance Be Backward-Looking?

Rhys Mendes and Stephen Murchison, Canadian Economic Analysis

- Constrained by the zero lower bound, several central banks have employed unconventional tools for lowering long-term interest rates and stimulating demand, including forward guidance with respect to the future level of short-term interest rates.
- In some cases, this guidance has included a threshold that must be met before short-term interest rates would be permitted to rise.
- As shown in model simulation results for Canada, forward guidance that is conditional on achieving a price-level threshold can theoretically boost demand and raise inflation expectations by significantly more than unemployment thresholds. This superior performance is attributable to the fact that the price-level threshold has a backward-looking or history-dependent element, since it depends on past inflation outcomes.
- In practice, history-dependent thresholds might, however, be more challenging for central banks to communicate and may be viewed as less credible by the public.

Faced with persistently weak demand and policy interest rates at or close to the zero lower bound (ZLB), central banks in several countries have been forced to explore unconventional tools for stimulating demand and avoiding deflation. These tools can be broadly classified into two categories: those that work primarily through private agents' expectations of the future level of the policy rate, such as forward guidance, and those that involve direct asset purchases by the central bank. This article explores the design and efficacy of forward guidance implemented through state-contingent thresholds, i.e., thresholds that are dependent on economic conditions.[1] Specifically, we use the Bank of Canada's projection and policy-analysis model, ToTEM (Terms-of-Trade Economic Model[2]), to compare the benefits of thresholds that incorporate history dependence, such as temporary targets for the

1 For assessments of the effectiveness of forward guidance during the crisis and its aftermath, see Bernanke (2012), Swanson and Williams (2013), and Filardo and Hofmann (2014).

2 See Dorich et al. (2013) and Murchison and Rennison (2006) for a description of ToTEM.

price level, with thresholds that are not history dependent, such as the unemployment threshold recently used by the U.S. Federal Reserve and the Bank of England.[3]

In normal times, central banks adjust their policy instrument, typically a very short-term interest rate, to achieve their objective (such as price stability or full employment). For example, the Bank of Canada adjusts its target for the overnight interest rate to maintain projected year-over-year consumer price inflation at the 2 per cent midpoint of the Bank's target range of 1 to 3 per cent.

If there is a sufficiently large negative shock, however, such as that experienced by many countries following the 2007–09 global financial crisis, the ZLB on interest rates may become a binding constraint for central banks. In other words, to achieve the central bank's policy objective, the appropriate level of the policy rate may be negative. Nevertheless, central banks may still be able to reduce *long-term* interest rates by issuing forward guidance that lowers market expectations of the future path of the policy rate. The perceived stimulative effect of guidance may cause expected inflation to rise, further reducing both short- and long-term rates in real, or inflation-adjusted, terms. Forward guidance can also decrease the level of uncertainty regarding the future path of the policy rate, which may boost household and business spending.

The Evolution of Forward Guidance

The challenge for a central bank is to find a simple and easy-to-understand way of conveying the most likely path of the policy rate (or how long it is expected to remain at its current level), and the degree of conditionality of that path on the state of the economy. Carney (2013) points out that one can view threshold-based guidance as the culmination of an evolutionary process in which guidance has become increasingly explicit and state contingent over time. In general, extraordinary guidance can be categorized as belonging to one of three generations: qualitative, time contingent and state contingent.[4] The Bank of Japan pioneered qualitative guidance at the ZLB in 1999 when it indicated that rates would stay at zero until "deflationary concerns" were "dispelled" (Fujiki and Shiratsuka 2002). The Bank of Canada's conditional commitment in April 2009 heralded the second generation of guidance,[5] while the unemployment thresholds used by the Federal Reserve and the Bank of England belong to the third generation.

Extraordinary guidance can be categorized as belonging to one of three generations: qualitative, time contingent and state contingent

All types of guidance at the ZLB aim to provide additional stimulus through one or both of the following channels: an increased period of time at the ZLB and reduced uncertainty about the path of short-term interest rates. Guidance has the potential to shift the perceived distribution of future short-term interest rates, such that the most likely outcome involves interest rates staying at the ZLB for longer than markets would have otherwise anticipated. All else being equal, this lowers the expected path of short-term rates, thereby lowering long-term rates.

3 The price-level threshold is history dependent because the entire history of inflation influences the price level. Thus, reacting to the price level is akin to reacting to past inflation rates. The definition of history dependence is discussed in greater detail later in the article.

4 See Carney (2013) for examples of the different types of guidance.

5 While often described as "time contingent," all second-generation guidance involved some degree of explicit or implicit conditionality. In particular, the Bank of Canada's commitment was conditional on the outlook for inflation.

Reduced uncertainty regarding future monetary policy can also lower the expected path of short rates. One can think of this uncertainty as reflecting a combination of uncertainty about which economic variables influence the central bank's monetary policy decisions (i.e., which variables enter the reaction function), uncertainty about the future evolution of those variables, as well as uncertainty about the weights assigned to each variable by the central bank. At the ZLB, uncertainty about interest rates is asymmetric—short-term interest rates can rise, but they cannot fall. As a result, the mean or expected path of short rates will tend to be greater than the mode or most likely path. Thus, even if agents believe that short rates are most likely to remain at the ZLB for an extended period, the expected path of short rates will typically lie above the ZLB over at least some of that period. This is important because long-term interest rates should depend on the expected (mean) path of short rates. By reducing uncertainty about the conditions under which the policy rate may rise, guidance can reduce the perceived probability of rate increases. This can lower the expected path of short rates for any given modal, or most likely, period of time at the ZLB.

For example, in December 2012, the Federal Reserve committed to keep rates at the lower bound at least until the unemployment rate fell below 6.5 per cent, unless projected inflation rose above 2.5 per cent. These conditions rule out the possibility that rates could rise with unemployment above 6.5 per cent and projected inflation below 2.5 per cent. The shaded region in **Chart 1** shows the coefficient values in a simple Taylor-type reaction function that would be consistent with the thresholds. In essence, in this illustrative example, the Fed told the private sector that the probability of the more hawkish reaction functions outside the shaded region was zero, which eliminated some uncertainty for private agents, permitting a more accurate assessment of the future path of interest rates.

State-contingent thresholds can also reveal important information about the economic variables that enter the central bank's reaction function, which in turn helps markets to better understand the precise nature of the conditionality of forward guidance. Using the Federal Reserve's unemployment

Chart 1: Taylor-type reaction function parameters consistent with thresholds for unemployment and inflation

Reaction function: $i_t = i_{ss} + \phi(\pi_t - \overline{\pi}) - \gamma(u_t - \overline{u})$

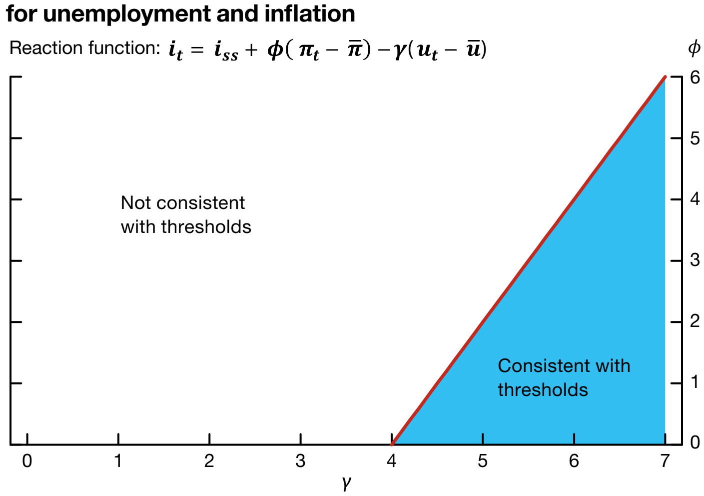

Notes: The Taylor-type reaction function sets the nominal policy rate (i_t) as a function of its steady-state value (i_{ss}), the deviation of inflation from target ($\pi_t - \overline{\pi}$), and the deviation of the unemployment rate from its long-run value ($u_t - \overline{u}$). For illustrative purposes, we set i_{ss} = 4, $\overline{\pi}$ = 2 and $\overline{u}$ = 5.5. The thresholds convey information about the response coefficients, ϕ and γ.

threshold as an example, it is clear that any new data suggesting a longer period of time before the 6.5 per cent threshold is reached would automatically mean the policy rate would remain at the ZLB for longer. This is true not only because it will take longer to reach the threshold, but also because the level of rates will likely remain lower than they would have otherwise during the subsequent period in which rates are increasing. Thus, long-term interest rates will respond to news in a manner consistent with achieving the threshold, while the threshold itself remains unchanged. If the central bank publishes (and regularly updates) a projection for the unemployment rate, the market will have direct insight into how much longer the central bank believes the policy rate will have to remain at the lower bound. Importantly, the private sector can also come to its own assessment of when the threshold will be attained. If the private sector can forecast the unemployment rate more accurately than the central bank, it will also forecast more accurately the length of time that rates remain at the ZLB.

Analyzing Alternative Thresholds

To analyze the properties of alternative types of thresholds, we use the Bank of Canada's main macroeconomic model, ToTEM, to simulate a large and persistent negative demand shock.[6] We assume that the central bank seeks to minimize the squared deviations of inflation from target and output from potential. This determines the optimal date for when interest rates should begin to increase—the "lift-off" date, conditional on returning to inflation targeting afterward.

We focus on the properties of two types of threshold (summarized in **Appendix 1** on page 22):

(i) price-level path: closing the gap between the price level and a threshold path

(ii) unemployment: closing the gap between the unemployment rate and a threshold level subject to an inflation knockout

Although both threshold types can be designed to implement the optimal lift-off date in the absence of additional shocks, there are important differences between them. In particular, they differ in terms of their degree of history dependence, which is a key determinant of their performance in the face of any future shocks that occur between when the threshold is adopted and when it is triggered.

History-dependent monetary policy has received a lot of attention over the past decade, from both academics and policy-makers, and represents an important potential mechanism for influencing private sector expectations. History dependence simply means that, in addition to current and expected future economic conditions, monetary policy responds to "past conditions even when they are no longer relevant to the determination of the current and future evolution of the variables that the [central] bank cares about" (Woodford 2003, 21). Policy will continue to respond to shocks, even after their impact on inflation and/or the output gap has fully dissipated. For example, if a shock initially causes inflation to fall below its targeted level, policy will continue to maintain interest rates below their neutral level until inflation moves above the target. Monetary policy essentially causes inflation to eventually overshoot the target when it is initially below target. This leads

History dependence simply means that, in addition to current and expected future economic conditions, monetary policy responds to past conditions

6 Simulating state-contingent threshold policies at the zero lower bound in a large model such as ToTEM is both technically and computationally demanding. We are very grateful to Nicholas Labelle St-Pierre for providing invaluable assistance in running these simulations.

to higher inflation expectations in the face of larger or more persistent shocks, thus reducing real interest rates and stimulating the economy. The key insight from Woodford (2003) and others has been that, under certain circumstances, the benefits of stabilizing the economy following a shock are greater than the destabilizing effects associated with the eventual inflation overshoot. For an inflation-targeting central bank whose policy rate is already at the ZLB, the temporary adoption of a history-dependent policy would generally result in rates remaining lower for a longer period of time, since it would signal a willingness to accept such an inflation overshoot in the future.

But the performance of a history-dependent policy depends on it being credible and well understood by the public. In reality, history dependence may be less effective if private sector expectations are not fully forward-looking. Estimates for Canada from ToTEM suggest that about half of price-setting firms are not forward-looking, but rather update their prices through time using a simple rule of thumb based on past inflation and the inflation target. Consequently, our simulation results allow for a significant deviation from the ideal fully forward-looking environment.

Price-level and nominal GDP targeting are two important examples of history-dependent monetary policy, since both involve targeting a level rather than a growth rate (such as inflation). Research on the efficacy of price-level targeting as a permanent regime change (including an extensive body of work at the Bank of Canada (2011)), suggests small welfare gains relative to inflation targeting, even when the policy is well understood and credible. In addition, significant uncertainties are associated with how private sector expectations are formed, whether they would adjust to the new regime in the manner predicted by theory, and whether the regime would be considered credible over time.

These considerations may help to explain why no country since Sweden in the 1930s has adopted a price-level or nominal GDP target.[7] However, the cost-benefit analysis may change when considering a temporary, rather than a permanent, switch, particularly at a time when the ZLB is binding. The stabilization benefits of history-dependent forward guidance, if successful, would be much higher at the ZLB precisely because the policy rate cannot be reduced to the level consistent with achieving the central bank's policy objective in a timely manner. For this reason, the hypothetical discussion of various forms of forward guidance that follows focuses on temporary policy measures when constrained at the ZLB.

In our simulations, a large negative demand shock hits the economy in period 0. In the base case, we use a simple rule that sets the policy rate as a function of its own lag, expected inflation and the output gap.[8] This simple rule causes the policy rate to remain at the ZLB for 12 quarters, in the absence of additional shocks. In contrast, if the lift-off date is chosen to minimize the sum of squared deviations of inflation from target and output from potential, then the policy rate remains at the ZLB for 20 quarters. With the policy that implements the optimal lift-off date, the recovery in real GDP growth is faster and stronger than in the base case (**Chart 2**). Similarly, CPI inflation takes almost five years to reach 2 per cent in the base case, compared with just over two years with the optimal lift-off policy (**Chart 3**).

7 For a discussion of the Swedish experience with price-level targeting, see Berg and Jonung (1999).

8 The simple rule exhibits a degree of history dependence because of the weight on lagged inflation. This rule is also used after the lift-off date in the simulations with thresholds.

Chart 2: Real GDP growth after a large negative demand shock

Quarterly simulation, year-over-year growth

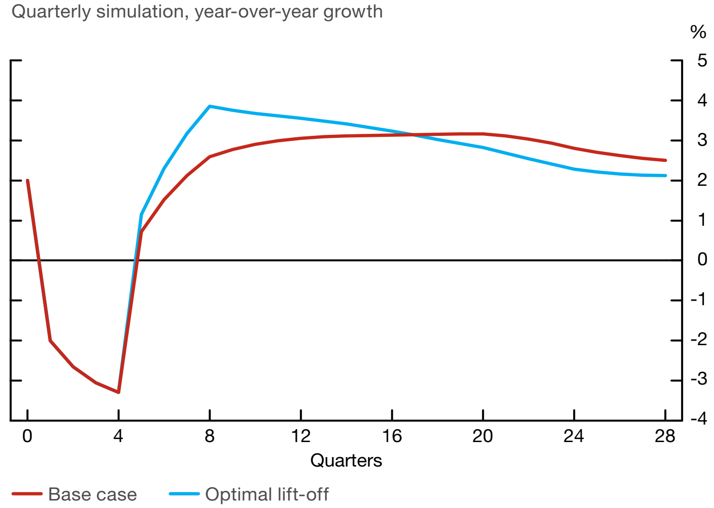

Note: GDP trend growth is assumed to be 2 per cent.

Source: Bank of Canada simulation

Chart 3: CPI inflation after a large negative demand shock

Quarterly simulation, year-over-year growth

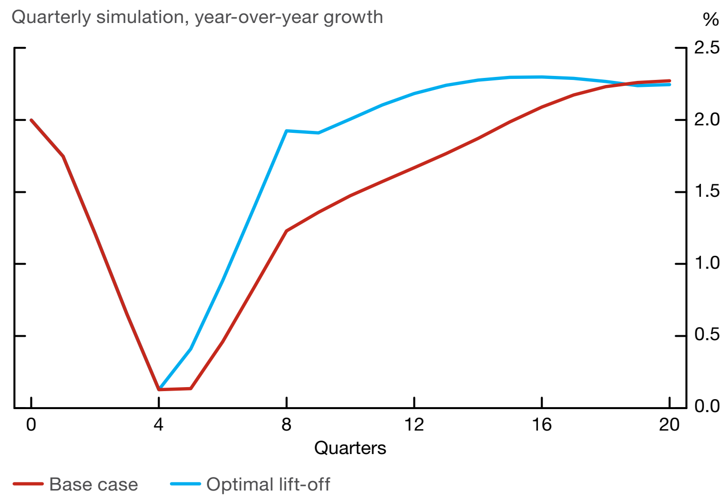

Source: Bank of Canada simulation

We choose the parameters of the price-level and unemployment thresholds such that they are both initially consistent with the optimal lift-off date. Thus, if there are no further shocks after period 0, the two types of threshold imply identical economic outcomes. However, under the more realistic assumption that the economy is buffeted by shocks every period, the choice of threshold variable will influence how the lift-off changes in response to shocks. As we explain in greater detail below, the choice of threshold variable is important because different variables imply different degrees of history dependence in response to shocks that occur after the threshold is established.

Price-level-path threshold

A price-level-path threshold is essentially a temporary price-level target with drift. The central bank would commit to keep the policy rate at the ZLB at least until the gap between the actual level of CPI and an exponentially growing threshold path was closed. In practice, this would be communicated by stating the current gap relative to the threshold along with the growth rate of the threshold. The growth rate of the threshold would typically be the target inflation rate. The choice of the initial gap is less straightforward. While it may seem natural to choose the initial gap as the deviation of the price level from some statistical trend, this will not generally be optimal. For example, if price shocks prevent weak demand from being fully reflected in the price level, it may be optimal to set the threshold above the statistical trend. However, if backward-looking private sector behaviour is important, it may be optimal to make up only part of the shortfall of the price level relative to the pre-shock statistical trend.[9] In our simulations, the latter consideration causes the threshold path to be 1.2 per cent below the steady-state price-level path that prevailed before the initial shock.

In the hypothetical demand shock scenario, the key elements of the price-level-path threshold policy are the following:

- The threshold is initially set 0.6 per cent above the actual level of the CPI (i.e., the initial price-level gap is -0.6 per cent).
- The threshold grows at the target rate of inflation (2 per cent per year).
- The policy rate is maintained at the ZLB at least until the level of the CPI reaches the threshold.

The price-level-path threshold implied by these elements is shown in **Chart 4**. If no shocks were to hit the economy after the introduction of the threshold, the simulated price-level path would be consistent with the initial optimal lift-off date. The advantage of this approach is that it is clearly history dependent, since inflation bygones are not bygones; that is, while the threshold is active, periods of below-target inflation must be subsequently offset by periods of above-target inflation.[10]

To illustrate the implications of history dependence, consider a shock that occurs after the threshold is established and causes the size of the price-level gap to increase. The larger gap will signal the need for greater policy stimulus in the future (relative to that originally announced) in order to make up for the shortfall in inflation. As with a conventional, non-history-dependent inflation target, some additional stimulus will be needed to return inflation to target. However, in order to close the price-level gap, inflation would have to overshoot the target. Achieving this overshooting will require the central bank to provide even more stimulus than is needed to simply return inflation to target.

In order to quantify the implications of history dependence, we allow for the possibility of additional random shocks after the threshold is established.[11] This is intended to capture key features of the real world, including the fact that additional shocks may cause the threshold to be breached sooner or

9 Backward-looking behaviour necessitates a longer period of excess demand (output greater than potential) in order to close a given price-level gap. Since we assume that the squared deviations of output from potential enter the central bank's loss function, greater backward-looking behaviour raises the costs associated with closing a given price-level gap.

10 History dependence is not an all-or-nothing proposition. Since the price level reflects the cumulative effects of all past inflation, the price-level threshold discussed above would make monetary policy contingent upon the average of all past inflation. However, a threshold could be based on, say, a three- or five-year average of inflation—longer than a conventional inflation target but shorter than the infinite window implied by the price-level threshold. Such average inflation thresholds would imply an intermediate degree of history dependence. Another alternative involves thresholds that keep rates at the lower bound at least until the price level (or the level of nominal GDP) reaches some time-invariant threshold level.

11 For each period, we randomly draw a set of shocks from the distribution of historical shocks to the Canadian economy implied by ToTEM.

Chart 4: Price-level-path threshold and the consumer price index

Quarterly simulation

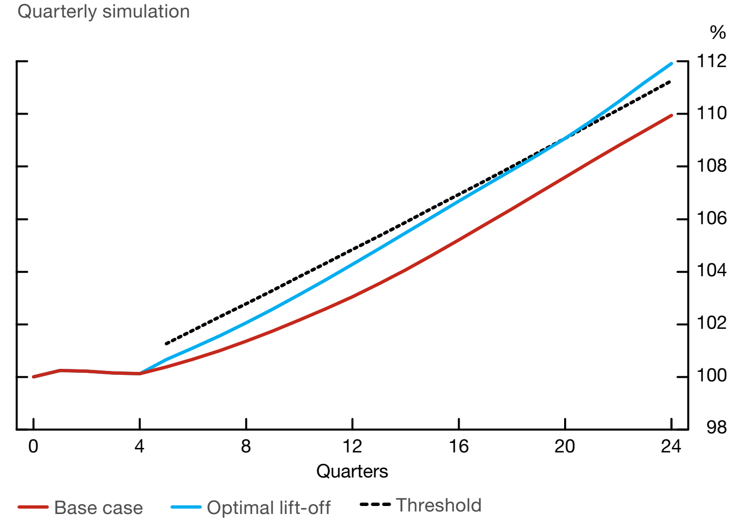

Source: Bank of Canada simulation

Table 1: Average losses following additional shocks (relative to base case)

Threshold	Variance of output gap	Variance of inflation	Loss function
Price-level path	0.80	0.68	0.78
Unemployment	0.92	0.89	0.92

Notes: Values less than 1 reflect variances or losses that are lower than in the base case. The loss function is defined as the sum of the variance of the output gap and the variance of inflation. All variables are relative to their respective base-case values. As a consequence, although the loss function is the sum of the absolute variances, the relative variances in the table do not sum to the relative loss.

later than the initial optimal lift-off date. In this stochastic simulation, the high degree of history dependence causes the price-level-path threshold to yield the lowest expected loss (**Table 1**).

Nevertheless, a price-level-path threshold may entail significant communication challenges. The price level may be an unfamiliar concept to many, since most people are accustomed to thinking in terms of inflation. The fact that the threshold for the price level is a moving target adds an additional layer of complexity. Experimental economics provides one avenue for exploring the importance of these issues. While the Bank has not conducted experimental work on forward guidance, the results reported in Amano, Kryvtsov and Petersen (this issue) indicate that people do, to some extent, understand the difference between inflation targeting and price-level targeting as permanent regimes. This suggests that it may be possible to effectively communicate a temporary price-level-path threshold, though further work would be required to confirm this.

Unemployment threshold

Both the Federal Reserve and the Bank of England have used thresholds for the unemployment rate together with an upper limit for expected inflation. In the context of our hypothetical scenario, the key elements of an unemployment threshold are the following:

- An unemployment threshold of 6.5 per cent.
- The policy rate remains at the ZLB at least until the unemployment rate declines to the threshold level, provided that the rate of inflation one to two years ahead is projected to be no more than 1 percentage point above target.
- If projected inflation rises more than 1 percentage point above target, it is assumed to knock out the threshold policy.

The 6.5 per cent unemployment threshold implements the initial optimal lift-off date in the absence of shocks after the threshold is established. This type of threshold can be breached if either or both of the unemployment and inflation conditions are violated. Our stochastic simulation results suggest that close to three quarters of all breaches would occur with the unemployment rate falling below its threshold level (**Chart 5**). Thus, in most cases, the unemployment threshold succeeds in fostering a recovery without causing inflation expectations to rise above the threshold level, resulting in the unemployment threshold outperforming the base-case rule (**Table 1**). However, this type of threshold is not history dependent. If inflation turns out to be weaker than projected by the central bank at the time the threshold was announced, there is no impetus for the central bank to offset this weaker inflation with higher inflation in the future. The unemployment threshold therefore underperforms the highly history-dependent price-level-path-threshold when faced with shocks.

Nevertheless, an unemployment threshold has several attractive features. This type of policy places an absolute limit on the level of projected inflation that could be obtained without a policy response and is therefore stricter than the price-level-path threshold in this respect. Moreover, it is based on variables—inflation and unemployment—that are familiar to the public, and thus avoids the need to communicate potentially confusing concepts such as a time-varying price-level threshold.

Chart 5: Causes of threshold breaches in the simulation of an unemployment threshold

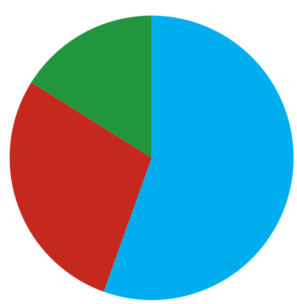

Conclusion

Simulation results presented in this article suggest that there may be material economic benefits to incorporating explicit history dependence when designing forward guidance at the ZLB in general and when employing state-dependent thresholds more specifically. Indeed, the two thresholds considered in this article differ mainly in terms of the degree of history dependence they impose on monetary policy during the period in which the threshold is in effect.

However, there are several ways in which the simulation results may overstate the relative benefits of history-dependent thresholds. In reality, the price-level and unemployment thresholds would differ in terms of their ease of communication. One of the key assumptions underlying the model results is that agents in the economy understand the threshold and what it implies for the path of the policy rate under different economic conditions. One of the appealing aspects of the thresholds used by the Federal Reserve and the Bank of England relative to the price-level-path threshold is that they are stated in terms of a variable most people are familiar with and can relate to—unemployment—and they are easy to understand and remember.

A second way in which our results may exaggerate the benefits of history-dependent thresholds is what is known as the dynamic inconsistency problem. In the model simulations, commitments to future policy actions are fully credible insofar as agents base their own behaviour on a belief that the central bank will keep its promises. In reality, a central bank that normally targets inflation would have an incentive to eventually renege on any promise that involves inflation overshooting the target at some future date. In other words, to minimize its loss function, a central bank would be better off by initially committing to an inflation overshoot in the future in order to realize the near-term benefits of stronger demand and inflation, and then reneging on that commitment once the economy has strengthened and inflation has returned to the level that the central bank normally targets.

Another potential problem with history-dependent thresholds is, in essence, the opposite of the dynamic inconsistency problem. Once it fulfills its commitment to cause inflation to overshoot the target, the central bank may find it difficult (or undesirable) to bring inflation back down to its target if expectations become unanchored and rise with actual inflation.

The benefits of history dependence, and forward guidance more generally, depend importantly on forward-looking private sector expectations and how future actions by the central bank affect economic behaviour today. A related issue is the extent to which private sector expectations regarding the evolution of the economy, even if forward looking, align with those of the central bank. For instance, the central bank may regard its choice of the level of a threshold, such as an unemployment threshold of 6.5 per cent, as implying a different lift-off date than the private sector would. In such an instance, the announcement of this threshold may have a smaller impact on the level of long-term interest rates, not because forward guidance is ineffective or not credible, but simply because the economic forecasts of the central bank and those of the private sector differ.

◂ *The benefits of history dependence, and forward guidance more generally, depend importantly on forward-looking private sector expectations and how future actions by the central bank affect economic behaviour today*

These considerations suggest that the results of model-based simulations, such as those discussed in this article, should be regarded as best-case scenarios. However, two points from the article merit repeating. First, expectations in ToTEM represent a mixture of purely forward-looking, rational expectations and simple rules of thumb, and the relative weight on each has been econometrically estimated using Canadian data. Second, these results omit any benefits associated with reduced uncertainty regarding the future path of the policy rate.

Appendix 1

Summary of Types of Threshold

Type	Threshold	History dependence	Pros	Cons
Price-level path	CPI gap ≥ 0 where gaps are defined relative to a trend that grows at a constant rate	• Fully history dependent • Initial gap chosen to account for any past shortfall in inflation • Gap automatically adjusts to make up for any additional shortfalls if inflation turns out weaker than originally projected	• The most history-dependent option • Limits average rate of inflation	• May require optimistic assumptions regarding initial gap in order to implement optimal lift-off policy • The return to a desired price-level path may be more difficult for the general public to understand than an unemployment threshold
Unemployment	Unemployment ≤ threshold with an upper-limit (knockout) condition on inflation: inflation projection > knockout	• Not history dependent • But automatically makes up shortfalls in real growth (in order to hit unemployment threshold) • Does not make up inflation shortfalls; equivalent to inflation targeting	• Based on familiar variables, therefore simplest to communicate • Places absolute limit on level of inflation	• Not history dependent

Literature Cited

Bank of Canada. 2011. "Renewal of the Inflation-Control Target: Background Information."

Berg, C. and L. Jonung. 1999. "Pioneering Price Level Targeting: The Swedish Experience 1931–1937." *Journal of Monetary Economics* 43 (3): 525–51.

Bernanke, B. S. 2012. "Monetary Policy Since the Onset of the Crisis." Speech to the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, 31 August.

Carney, M. 2013. "Monetary Policy After the Fall." Eric J. Hanson Memorial Lecture, University of Alberta, Edmonton, Alberta, 1 May.

Dorich, J., M. Johnston, R. Mendes, S. Murchison and Y. Zhang. 2013. "ToTEM II: An Updated Version of the Bank of Canada's Quarterly Projection Model." Bank of Canada Technical Report No. 100.

Filardo, A. and B. Hofmann. 2014. "Forward Guidance at the Zero Lower Bound." *BIS Quarterly Review* (March).

Fujiki, H. and S. Shiratsuka. 2002."Policy Duration Effect Under the Zero Interest Rate Policy in 1999–2000: Evidence from Japan's Money Market Data." *Monetary and Economic Studies* 20 (1): 1–31.

Murchison, S. and A. Rennison. 2006. "ToTEM: The Bank of Canada's New Quarterly Projection Model." Bank of Canada Technical Report No. 97.

Swanson, E. T. and J. C. Williams. 2013. "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates." Federal Reserve Bank of San Francisco Working Paper No. 2012-02.

Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.