BANK OF CANADA
BANQUE DU CANADA
CELEBRATING 75 YEARS
CÉLÉBRONS 75 ANS

Bank of Canada Review

Spring 2010

Special Issue

Uncertainty in Monetary Policy-Making

www.bankofcanada.ca

The *Bank of Canada Review* is published four times a year under the direction of an Editorial Board, which is responsible for the editorial content. The contents of the *Review* may be reproduced or quoted provided that the *Bank of Canada Review*, with its date, is specifically quoted as the source.

Back issues of the *Review* and other publications are available on the Bank's website at **http://www.bankofcanada.ca**.

Subscriptions for print are available, as follows:

Delivery in Canada:	Can$25
Delivery to the United States:	Can$25
Delivery to all other countries, regular mail:	Can$50

Canadian government and public libraries and libraries of Canadian and foreign educational institutions may subscribe at one-half the regular price. Single copies are $7.50.

Remittances in Canadian dollars should be made payable to the Bank of Canada. Canadian orders must include 5 per cent GST, as well as PST, where applicable.

Copies of Bank of Canada documents may be obtained from:

Publications Distribution
Communications Department
Bank of Canada
234 Wellington Street, Ottawa, ON
Canada K1A 0G9
Telephone: 613 782-8248
Toll free in North America: 1 877 782-8248
Email address: publications@bankofcanada.ca

Inquiries related to interest rates or exchange rates should be directed to 613 782-7506.

ISSN 0045-1460 (Print)
ISSN 1483-8303 (Online)
Printed in Canada on recycled paper

Bank of Canada Review

Spring 2010

Merchant Scrip: Grand River Pulp and Lumber Company

David Bergeron, Curator, Currency Museum

Before banks became established in the more remote parts of the country, it was not uncommon for private companies to pay their employees with scrip. Between 1902 and 1911, the Grand River Pulp and Lumber Company of Halifax, issued its own notes in denominations of 5, 10, 25, and 50 cents and $1, $2, and $5 to pay lumbermen working at its mill located at the mouth of the Grand River (now the Churchill River) in Labrador. The area was remote and isolated, and issuing scrip, redeemable either in local goods or in currency back in Halifax, appeared to be a good solution to avoid the risk of shipping coins and notes into the wild.

Alfred Dickie, a prominent Halifax lumberman and politician, established the Grand River Pulp and Lumber Company in 1900. With long-term leases on 500 square kilometres of land around Hamilton Inlet in Labrador, Dickie approached the Newfoundland government to apply for a timber licence. Little did he know that his application would trigger a border dispute between Quebec and Newfoundland. When it learned of Dickie's application, the Quebec government intervened, stating that the land in question belonged to Quebec. According to the Newfoundland legislature, the area that made up Labrador was situated north of the 52nd degree of latitude and east of the 64th degree of longitude, as defined in 1876 by letters patent. Thus, according to Dickie, the area leased by his company was well within Newfoundland's jurisdiction. Newfoundland refused to cancel the licence, maintaining its right to the watershed of all rivers flowing into the Atlantic Ocean.

In 1904, the dispute over the Labrador boundary escalated when the Quebec government requested that Ottawa present the case to the Judicial Committee of the Privy Council in London, England. The Newfoundland government agreed. The process dragged on, with no progress until 1922, when the British Privy Council was asked to decide on the border for Labrador. The panel of judges became mired in the definition of "coast," which was used to describe Labrador in the statutes, orders-in-council, and proclamations. Contrary to Newfoundland's claim, Canada contended that "coast" referred to a strip of land 1.6 km wide along the seashore. In 1927, the Privy Council decided in Newfoundland's favour, and Canada accepted the verdict. Newfoundland's entry into Confederation in 1949 confirmed Labrador's boundary, and in 1971, Quebec cancelled its appeal of the 1927 decision. In 1999, the Newfoundland House of Assembly passed a resolution renaming the province "Newfoundland and Labrador."

As Labrador's only known merchant scrip, the notes of the Grand River Pulp and Lumber Company tell a remarkable story about a part of Canada that is still sparsely inhabited and largely unexplored. The National Currency Collection holds a complete set of scrip from the Grand River Pulp and Lumber Company.

Photography by Gord Carter, Ottawa.

The image of the map of Labrador on the cover is used courtesy of the Centre for Newfoundland Studies, Memorial University Libraries.

Contents

Special Issue
Uncertainty in Monetary Policy-Making

Uncertainty in Monetary Policy-Making

Greg Tkacz, Guest Editor

Uncertainty permeates the monetary policy process across several time dimensions. Uncertainty about the future is well known, since policy-makers require forecasts of key variables when making decisions. Uncertainty about the present exists in the form of model and parameter uncertainty, which can affect the analysis of possible policy actions. Finally, there is also uncertainty about the past, since key economic variables are subject to revision, which can affect the perceived strength of the economy.

This special issue presents four articles that deal with uncertainty in monetary policy-making and how such uncertainty can be potentially minimized. The first two articles relate to uncertainty about the future, the third to uncertainty about the present, and the final one to uncertainty about the past.

Ron Alquist and Elif Arbatli discuss three ways that oil-futures prices can improve our understanding of current conditions and future prospects in the global market for crude oil in "Crude Oil Futures: A Crystal Ball?" First, the response of the oil-futures curve can be used to identify the persistence of oil-price shocks and to obtain an indicator of the rate at which they will diminish. Second, the spread between the current futures price and the spot price of oil can be interpreted as an indicator of the precautionary demand for oil. Third, because oil-futures prices are volatile, forecasts of the future spot price of oil using futures prices should be supplemented with other information to improve their accuracy.

In "Inflation Expectations and the Conduct of Monetary Policy: A Review of Recent Evidence and Experience," Rose Cunningham, Brigitte Desroches, and Eric Santor explore the role of inflation expectations in the conduct of monetary policy. They review the various measures of inflation expectations used by central banks, including surveys and market-based indicators, and consider their advantages and disadvantages. They examine the critical role of inflation expectations in the framework that central banks use to understand, forecast, and control inflation. They also look at their role as an indicator of central bank credibility. The behaviour of inflation expectations over the past two years is analyzed, and policy conclusions are offered.

In "Monetary Policy Rules in an Uncertain Environment," Gino Cateau and Stephen Murchison examine recent research on the influence of various forms of economic uncertainty on the performance of different classes of monetary policy rules: from simple rules to fully optimal monetary policy under commitment. The authors explain why uncertainty matters in the design of monetary policy rules and provide quantitative examples from the recent literature. They also present results for several policy rules in ToTEM, the Bank of Canada's main model for projection and analysis, including rules that respond to price level, rather than to inflation.

Finally, Greg Tkacz shows in "An Uncertain Past: Data Revisions and Monetary Policy in Canada" how many important economic variables are subject to revision. This article explains how, when, and why such revisions occur; how revisions to Canadian gross domestic product (GDP) compare with GDP revisions in some other countries; which GDP components are subject to the largest revisions; and how data revisions can affect policy decisions.

Crude Oil Futures: A Crystal Ball?

Ron Alquist and Elif Arbatli, International Economic Analysis

- *It is common for policy-makers and market analysts to use the prices of crude-oil-futures contracts to interpret developments in the global crude oil market. Based on recent research, this article discusses three ways that oil-futures prices can improve our understanding of current conditions and future prospects in this important international commodity market.*
- *First, the response of the oil-futures curve can be used to identify the persistence of oil-price shocks and to obtain an indicator of the rate at which a given shock will diminish.*
- *Second, the spread between the current futures price and the spot price of oil can be interpreted as an indicator of the precautionary demand for oil.*
- *Third, oil-futures prices can be used to forecast spot prices, but because such forecasts are volatile, they should be supplemented with other information to improve their accuracy.*

The run-up in the price of crude oil since 2002 and its sharp collapse in the autumn of 2008 have renewed interest in understanding the determinants of spot and futures prices in the oil market **(Charts 1** and **2)**. Such shifts highlight the importance of understanding the relationship between the prices of oil-futures contracts and market expectations. Indeed, it is common for policy-makers and market analysts to interpret the price of the crude oil-futures contract traded on the New York Mercantile Exchange, or NYMEX, as a measure of market expectations of the future spot price of oil. In light of this widespread use, it is important to understand the information that can be recovered from the prices of oil futures. Recent studies shed light on the information that these prices provide about developments in the global crude oil market.

Chart 1: Spot and futures prices for crude oil

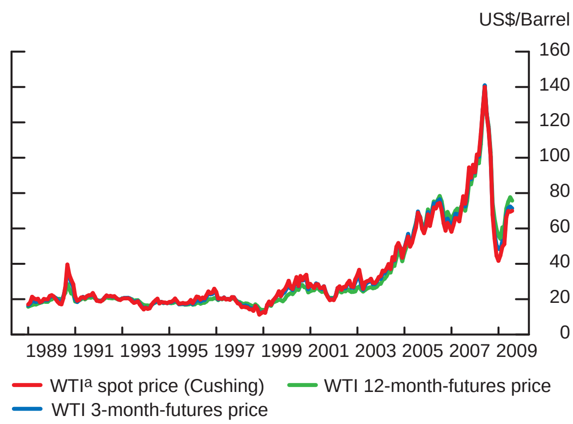

a. West Texas Intermediate
Source: NYMEX

Chart 2: Futures-spot spread

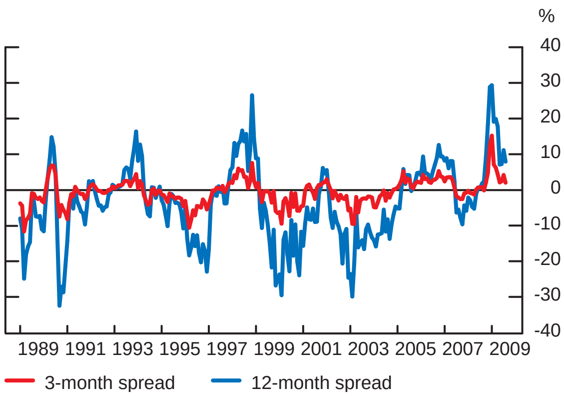

Source: NYMEX

We begin by reviewing the theory of storage as a way to organize thinking about the relationship between spot and futures markets. In this type of model, commodity processors choose how much of the commodity they will use today versus tomorrow and, hence, determine the level of the spot price relative to the futures price. We then assess whether movements in the futures curve capture market expectations of the future path of oil prices, as predicted by the theory of storage under risk neutrality. Finally, we discuss three ways of using the prices of oil futures to understand current developments and future prospects in that market: namely, inferring the persistence of shocks from the response of the futures curve to shocks in the spot price; using the futures-spot spread as an indicator of shifts in expectations about future oil-supply shortfalls; and forecasting the spot price of oil in real time, using futures prices.

Price Determination in the Market for Crude Oil Futures

A standard framework for thinking about the determination of futures prices in the market for crude oil is the theory of storage, which is generally applicable to markets for storable commodities. The spot price is the price at which the commodity is immediately available, and the futures price is the price at which the commodity is available for delivery at a specified future date. Taking the supply of the commodity as given, the framework, in its simplest form, assumes that risk-neutral commodity processors operate in a competitive environment and will optimally choose the quantity of the commodity that they wish to consume today and the quantity that they wish to store.[1] The assumption of risk neutrality ensures that the current futures price equals the expected spot future price, adjusted for the costs and benefits associated with storing oil and having ready access to it.

In this model, the spread between the spot and futures prices adjusts to equate the marginal cost to the marginal benefit of storing a barrel of oil as inventory. The difference between contemporaneous spot prices and futures prices reflects the interest foregone from storing the commodity, the cost of physical storage, and the convenience yield associated with holding inventory. The convenience yield is the benefit of holding a barrel of oil as inventory that accrues to the firm storing oil. It reflects a precautionary motive for holding oil inventory and is assumed to exhibit diminishing marginal returns to storing oil.

Economists appeal to the idea of the convenience yield to explain an apparent puzzle observed in commodity-futures markets. Current futures prices often lie below the current spot price—that is, futures prices are backwardated—at the same time that firms carry over stocks of the commodity from one period to the next.[2] Firms therefore hold stocks at an apparent capital loss. If stocks of a commodity yield benefits to the firm, then it can be rational for a firm to hold inventories even when the futures market is backwardated. That is, the value of having ready access to a stock of oil can justify holding inventory when the futures curve is in backwardation. The West Texas Intermediate oil futures contract—the most liquid, widely traded, and closely monitored energy-futures contract in North America—is frequently in backwardation and yet refiners also hold positive levels of inventory (Litzenberger and Rabinowitz 1995).

A convenience yield associated with holding crude oil as inventory is consistent with the operational requirements of oil refineries. Because of technological constraints, oil refineries have a strong incentive to hold stocks of oil to optimize the production of different types of petroleum products (National Petroleum Council 2004). Stocks of crude oil give a

1 This type of model has a long lineage, beginning with Kaldor (1939), Working (1949), Brennan (1958), and Gustafson (1958). More recent papers include Scheinkman and Schechtman (1983), Williams and Wright (1991), Deaton and Laroque (1992), and Ng and Ruge-Murcia (2000).

2 When futures prices lie above spot prices, the market is said to be in contango. The terms "backwardation" and "contango" originated in the London Stock Exchange during the nineteenth century. "Backwardation" referred to a fee paid by the seller of a security for the right to delay delivery; and "contango" referred to a fee paid by the buyer of a security for the right to delay delivery and payment.

refinery operational flexibility, and the value of this flexibility can be captured by the convenience yield. Considine (1997) finds that the convenience yield net of interest and physical storage costs is about 20 per cent of the spot price on an annual basis.[3]

Futures Prices and Market Expectations

We can use futures prices as a measure of the expected spot price and interpret the term structure of futures prices as the expected time path of oil prices only if futures prices represent the rational expectation of the spot price of oil. The argument for using futures prices to represent market expectations thus relies on the premise that futures prices are unbiased predictors of the future spot price of oil. The available evidence is broadly consistent with that assumption. Although there is some evidence that the futures prices are biased predictors of the spot price, the bias is small, on average.

The argument for using futures prices to represent market expectations relies on the premise that futures prices are unbiased predictors of the future spot price of oil.

Bias and the forecasting efficiency of futures prices

Forecast-efficiency tests are one way to detect if there is bias associated with using futures prices to predict the future spot price. The tests involve regressing the ex post percentage change in the spot price of oil on a constant and the futures-spot spread, the percentage difference between the current futures price and the current spot price.[4] The regression equation is

$$\Delta s_{t+h} = \alpha + \beta \left(f_t^{(h)} - s_t \right) + \varepsilon_{t+h}, \qquad (1)$$

where Δs_{t+h} denotes the ex post change in the log spot price; $f_t^{(h)}$ denotes the log price of a futures contract that matures in h months; s_t denotes the current spot price, and ε_{t+h} is a random error term. If futures prices are unbiased predictors of the future spot price, we expect that $\alpha = 0$ and $\beta = 1$.[5] It is common to interpret failing to reject the null hypothesis that $\beta = 1$ in such regressions as evidence against a time-varying risk premium (see, among others, Chernenko, Schwarz, and Wright 2004).

We estimate forecast-efficiency regressions for 3-, 6-, and 12-month contracts using data over the period January 1989 to August 2009. **Table 1** reports the results from these regressions. The average bias appears to increase monotonically with the maturity of the futures contract, but is significantly different from zero at only the 12-month horizon. We also fail to reject the null hypothesis that $\beta = 1$ at all horizons. These conclusions are very similar to those obtained in other studies that have used different subsamples, such as Chernenko, Schwarz, and Wright (2004), Arbatli (2008), Chinn and Coibion (2009), and Alquist and Kilian (2010). This evidence thus indicates that treating oil-futures prices as the expected future spot price is a good first approximation.

Table 1: Results of forecast-efficiency regressions for oil-futures contracts
January 1989-August 2009

	3-month contract	6-month contract	12-month contract
α (p-value)	0.02 (0.29)	0.04 (0.18)	0.09 (0.05)
β (p-value)	1.51 (0.46)	0.91 (0.85)	0.79 (0.54)
Reject H_0: $\alpha = 0, \beta = 1$	No	No	Yes
Adjusted R-squared	0.03	0.04	0.05
T	246	243	237

Notes: The p-values are based on standard errors that are robust to autocorrelation and heteroskedasticity.

At this juncture it is important to discuss a subtlety surrounding statistical tests of predictability that helps us to understand the relationship between these results and the evidence that futures prices tend to be

3 Considine (1997) derives the convenience yield from a crude oil refinery's dynamic profit-maximization problem, using disaggregated data on the type of petroleum products that refineries typically produce. He finds that significant cost savings are associated with adjusting oil stocks to minimize variable costs. Apart from a difference in sign, the cost savings are equivalent to the convenience yield.

4 Such tests implicitly assume that the goal of market participants is identical to that of the econometrician in that they both pick parameters α and β to minimize the sum of squared errors. If that is not the case, forecast-efficiency tests are biased in favour of the alternative hypothesis (see Elliott, Komunjer, and Timmermann 2005).

5 It is also possible to adjust for the cost-of-carry by including interest rates and the cost of storage. Of the two, only interest rates are directly observable, and including them does not affect the conclusion. The available evidence on the cost of storage from the Energy Information Administration indicates that changes in such costs occur at low frequency and, therefore, cannot account for the size of the high-frequency fluctuations in the futures-spot spread.

less accurate real-time predictors of the future spot price than the no-change forecast.[6] If the price of oil futures and the conditional expectation of the price of oil were equal, then the oil-futures price would be the most accurate predictor according to standard metrics for measuring forecast accuracy (Granger 1969). The forecast-efficiency tests are consistent with this assumption, but they are in-sample tests of predictability that use the full range of data available at a point in time. It is also possible to test for predictability using an out-of-sample test. This type of test employs a subsample of the available data to conduct a real-time forecast that uses data only up to a specific point in time. It is widely recognized among forecasters that there is no necessary connection between detecting significant in-sample predictability and detecting significant out-of-sample predictability, and the two tests can deliver different inferences (for example, Amato and Swanson 2001; Chao, Corradi, and Swanson 2001; and Inoue and Kilian 2006). Put differently, predictability that exists in a population may not be exploitable in real time. This fact explains why there is no logical tension between the forecast-efficiency regressions and the ability of the futures price to predict spot prices out-of-sample.

Other measures of market expectations

Another way to assess whether futures prices for crude oil represent market expectations of future spot prices is to compare the market expectations recovered from futures prices with those provided by market commentaries and professional forecasters. The comparison provides another source of evidence regarding the relationship between futures prices and market expectations. One advantage of such a comparison is that it permits us to link developments in oil markets to movements in futures prices and to understand more clearly the relationship between real-time developments in the crude oil market and futures prices.

Arbatli (2008) compares the market expectations obtained from the futures curve with those from two other sources: commentaries in the *Oil & Gas Journal* and forecasts published by Consensus Economics. The *Oil & Gas Journal* is a major industry journal that contains commentaries on developments affecting the spot and futures markets for oil. This procedure is similar to that used in other studies to identify oil-price shocks associated with exogenous events (for example, Cavallo and Wu 2006). Arbatli identifies episodes with large movements in oil prices, because such episodes are associated with news about underlying supply and demand conditions in the global crude oil market, making the relevant events easier to detect.

She finds that changes predicted by the futures curve, as captured by the slope of the curve, coincide with the predictions suggested by market commentaries. For example, during the Gulf War there was a sharp upward spike in the spot price of oil, whereas the price of long-horizon futures contracts did not move very much. Market commentaries during that episode reveal that oil industry analysts expected the change in the spot price to be transitory. A similar picture emerges from studying the behaviour of oil prices during the Asian financial crisis of 1997–98. During that period, the spot price fell significantly, whereas the price of the long-dated futures contract did not, again suggesting that the market perceived the decline in oil prices to be transitory. Similarly, Arbatli identifies periods during which expectations of more persistent changes in underlying supply and demand conditions are detectable in the prices of oil futures. During such episodes, the entire futures curve shifts up or down. Examples of persistent changes in the price of oil are the collapse in prices in 1986 and the run-up in prices during 2003–06. Both periods were associated with commentaries that emphasized the persistent nature of the price changes.

Since interpreting market commentaries requires forming a subjective judgment about the implications of the statement for the future price of oil, Arbatli uses forecasts from Consensus Economics. **Chart 3** reproduces and extends the data from that paper. It plots the difference between the forecasts for prices 12 and 3 months ahead from Consensus Economics relative to the current spot price and compares that with the difference between prices for 12- and 3-month oil futures relative to the current spot price for the same month. The gap between the 12- and 3-month-ahead forecasts reflects what market participants expect to happen to prices. A positive number indicates that the market expects an increase in prices; a negative value indicates an expected decrease in prices. The chart shows that there is a strong historical correlation between the futures-based forecasts and those obtained from professional forecasters. In one sense, this finding is unsurprising: It may simply demonstrate that professional forecasters use futures prices to inform their forecasts. Furthermore, while the correlation between Consensus forecasts and futures-based forecasts is high, it is not perfect. Evidently, forecasters use futures prices, as

6 The no-change forecast uses the current spot price to forecast the future spot price.

Chart 3: Implied change in the spot price of crude oil

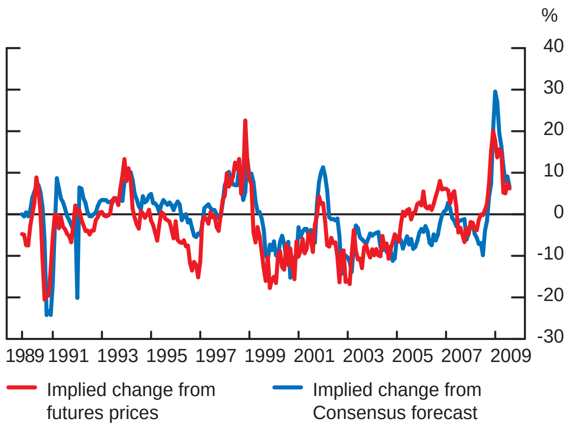

Sources: Consensus Economics and NYMEX

well as other sources of information, to predict the future path of the price of oil.

Forecasters use futures prices, as well as other sources of information, to predict the future path of the price of oil.

In conjunction with the statistical evidence obtained from the forecast-efficiency regressions, the narrative evidence supports the view that futures prices, imperfect as they are, provide a way to measure market expectations. In the next section, we examine in greater detail how to use futures prices to shed light on real-time developments in the global crude oil market.

Interpreting the Behaviour of Crude Oil Futures Prices

The persistence of price shocks and the futures curve

If we assume that the futures curve represents a measure of the expected future path of spot prices, it can be used to capture expectations about the persistence of shocks to the spot price of oil. Bessembinder et al. (1995), for example, estimate the rate at which the price of oil reverts to its mean, using the response of the slope of the futures curve to a change in the spot price. Within their framework, a large response of the slope to changes in spot prices suggests a large expected mean reversion in spot prices. According to estimates presented in the paper, almost half of a spot-price shock is expected to be reversed within eight months. This estimate of mean reversion is consistent with other estimates based on the futures curve (see, for example, Arbatli 2008). In a similar vein, Schwartz and Smith (2000) use the term structure of futures prices to construct a real-time decomposition of the spot price into a long-run and a short-run component. The identification procedure in that paper relies on the assumption that the change in futures prices over different maturities constitutes the impulse response of the spot price to oil-price shocks. Arbatli (2008) uses the same assumption to identify permanent and transitory shocks to oil prices and, hence, to summarize the information about the persistence of shocks embedded in the futures curve.[7]

In conjunction with other models, the permanent-transitory decomposition derived from the futures curve provides information that can guide the conduct of monetary policy. In general, the optimal response of monetary policy to oil-price shocks depends on the persistence of the shock, because of lags in the effect of monetary policy on the economy. If the oil-price shock is expected to be reversed quickly, a more aggressive policy response may be destabilizing and, therefore, inappropriate. In an oil-exporting country like Canada, a persistent increase in the price of oil represents a positive terms-of-trade shock that can generate a large and persistent real appreciation of the exchange rate. Although the appreciation exerts downward pressure on prices through less-expensive imports, the wealth effect of such a persistent change in the price of oil also exerts upward pressure on prices. The permanent-transitory decomposition can suggest the type of shock to feed into a structural macroeconomic model to study the response of the economy and, thus, to design the appropriate policy response.[8] The increasing liquidity in the oil-futures market and the expanding range of actively traded maturities open up the possibility of using long-dated futures contracts to obtain more reliable estimates of the persistence of oil-price shocks.

7 Since both papers include a constant in their specification, they admit the possibility that futures prices are biased predictors of the future spot price.

8 The permanent-transitory decomposition provides an estimate of the long-run price of oil and its behaviour over time. It is important to recognize that the estimated long-run price is not necessarily an estimate of the long-run equilibrium level of the price of oil. The reason is that the market for long-horizon contracts is illiquid, and therefore the longest maturity contract used in both papers is 12 months.

The futures-spot spread and precautionary demand

Alquist and Kilian (2010) propose a model in which the futures-spot spread may be viewed as an indicator of shifts in expectations about future oil-supply shortfalls. In their model, an oil-producing country exports oil to an oil-consuming country that uses the oil to produce a final good to be traded for oil or consumed domestically. Oil importers may insure against uncertainty about oil-supply shocks by holding above-ground oil inventories or by buying oil futures. Oil producers may sell oil futures to protect against endowment uncertainty.

One implication of the model is that increased uncertainty about future oil-supply shortfalls causes the oil-futures spread to fall and raises the current real spot price of oil, as precautionary demand for oil inventories increases. Increased uncertainty about future oil-supply shortfalls thus causes the real price of oil to overshoot and then to decline gradually to a new steady-state value that is higher than the original one.

Alquist and Kilian present three pieces of evidence consistent with the model's predictions. First, the proposed indicator moves as expected during events, such as the Persian Gulf War, that a priori should be associated with large shifts in the precautionary demand for crude oil. They also find evidence of such shifts in the spread associated with the Asian financial crisis, the attacks on September 11, and the 2003 Iraq War. Second, their indicator is highly correlated with an independent estimate of the precautionary demand component of the spot price of oil that is proposed by Kilian (2009). That alternative estimate is based on a structural vector autoregressive model of the global crude oil market that does not rely on data from the market for oil futures. The model decomposes unexpected changes in the real price of oil into shocks attributable to changes in the global supply of crude oil, shocks to global real economic activity, and oil-specific demand shocks that can be interpreted as precautionary demand shocks (see Kilian 2009). Over the period from January 1989 to December 2006, the two measures exhibit a very high correlation. Third, they show that the overshooting pattern in the response of the real price of oil to a precautionary demand shock in the Kilian model is consistent with the predictions of the theoretical model.[9]

This evidence lends credibility to the interpretation of the futures-spot spread as an indicator of fluctuations in the spot price of oil driven by shifts in the precautionary demand for crude oil. Although such shifts in expectations can be difficult to quantify in real time, the paper provides a way to interpret such movements using readily available price data. The availability of such data is especially important in light of the evidence presented in Kilian (2009) that the contribution of oil-supply shocks to changes in the price of crude oil has been smaller than previously thought. He concludes that demand shocks in general and precautionary demand shocks in particular play an economically important role in explaining the variability of oil prices. Since the data on which Kilian's argument is based are not readily available in real time, one can use the futures-spot spread as a real-time indicator of the shifts in expectations associated with precautionary demand shocks.

Using futures prices to forecast the spot price of crude oil

In this section, we survey the evidence on the ability of futures prices to forecast the spot price of oil out-of-sample.[10] The main conclusion is that while futures prices tend to produce forecasts that are correct on average, such forecasts are also highly volatile relative to no-change forecasts. Therefore, futures-based forecasts may be very inaccurate at a given point in time. The variability of futures-based forecasts makes it advisable to use the information contained in oil-futures prices in conjunction with other types of information when arriving at a judgment about the future trajectory of oil prices.

Some early studies found evidence that futures prices were accurate out-of-sample predictors of the future spot price of oil. Ma (1989) reports that futures prices outperform the no-change forecast, as well as other simple time-series models, in out-of-sample forecasting exercises. Kumar (1992) reaches similar conclusions.

9 It is important to point out that the economic environment in the Alquist and Kilian model is risk neutral. Although risk aversion can imply a precautionary motive for holding stocks of crude oil, it is not required. For example, a convenience yield can arise from the convex adjustment costs of firms rather than from the risk aversion of consumers (see Pindyck 1994). Thus, the existence of a convenience yield is equally consistent with risk-averse and risk-neutral preferences.

10 There is a related literature on the use of forward contracts traded in currency markets as indicators of the expected spot price of foreign currency (see Froot and Thaler 1990).

He finds that futures prices provide more accurate forecasts than those obtained from alternative time-series models, including the random-walk model.

In a study that uses data through the end of 2003, Chernenko, Schwarz, and Wright (2004) provide evidence that futures-based forecasts have a marginally lower mean-squared prediction error than the no-change forecast. Three related papers are Chinn, LeBlanc, and Coibion (2005), Wu and McCallum (2005), and Chinn and Coibion (2009). Chinn et al. conclude that futures-based forecasts are unbiased predictors of the spot price of oil and that they perform better than the random-walk forecast according to the mean-squared prediction error. Chinn and Coibion (2009) update the results from their earlier paper, and find that futures prices do not systematically outperform the random-walk forecast although they are superior to forecasts generated by other types of time-series models. Moreover, while Wu and McCallum report that futures prices tend to be less accurate than the no-change forecast, they also observe that spread regressions have a lower mean-squared prediction error than the no-change forecast at short horizons. Similarly, Coppola (2008) obtains improvements in forecast accuracy only at the 1-month horizon, and at longer horizons finds no improvements in forecast accuracy compared with the no-change forecast.

This evidence seems to suggest that the futures price is a useful tool for forecasting the spot price out-of-sample, at least over certain horizons. But in a comprehensive recent study, Alquist and Kilian (2010) consider the price data available from January 1989 through February 2007 and conduct out-of-sample forecasts using data available in real time. They conclude that futures-based forecasts are not more accurate than the no-change forecast for horizons out to 12 months. This finding is robust at all horizons from 1 month to 12 months and for a range of loss functions, including the quadratic and absolute loss functions. In particular, the no-change forecast tends to be more accurate than forecasts based on futures prices, other econometric models, and professional survey forecasts of the price of oil.

The difference between Alquist and Kilian's conclusions and those of prior studies can be traced to the longer sample period. Sensitivity analysis suggests that evidence of accuracy gains, sometimes obtained in shorter samples, tends to vanish when the full sample is examined. The inability of alternative models to forecast more accurately than the random walk may also be attributable to a risk premium, so that adjusting forecasts by the risk premium can improve the model's ability to forecast out-of-sample (Sadorsky 2002; Pagano and Pisani 2009). But the forecast-efficiency regressions reported in Alquist and Kilian, which are qualitatively similar to those reported in this article, do not reveal evidence consistent with the presence of a risk premium.

Alquist and Kilian document why futures-based forecasts are inferior to the no-change forecast. Whereas the bias of futures prices relative to the no-change forecast is small, the variability around the no-change forecast is not. At a point in time, the discrepancy between the futures price and the spot price may be large and may go in either direction. This variability in the deviation of futures prices from spot prices, rather than differences in the mean, drive the larger mean-squared prediction error of futures-based forecasts. Thus, policy-makers and financial analysts who use futures prices to forecast the spot price of oil will tend to be correct on average, but they will also run the risk of obtaining a very inaccurate forecast at a given point in time. This conclusion suggests that it is important not to rely solely on oil-futures prices to predict the future price of oil and instead to use them in conjunction with other pieces of information to arrive at a view of what the price of oil will be.

Policy-makers and financial analysts who use futures prices to forecast the spot price of oil will tend to be correct on average, but they will also run the risk of obtaining a very inaccurate forecast at a given point in time.

Although there is no single rule of thumb that guarantees being able to forecast the price of oil reliably, forecasters can take consolation in the fact that this conclusion is consistent with the views of oil-industry experts. For example, in a 2007 speech to petroleum economists, Peter Davies, chief economist for British Petroleum, noted that "we cannot forecast oil prices with any degree of accuracy over any period whether short or long" (Davies 2007). Thus, even economists with detailed knowledge of the technological and geological constraints related to the extraction of oil find it challenging to produce accurate forecasts.

Concluding Remarks

The findings discussed in this article have immediate policy implications. The decomposition of oil-price shocks into permanent and transitory components can be used to estimate the persistence of oil-price shocks in real time. Such an estimate can be used to simulate the effects of an oil-price shock with particular time-series characteristics. The result of such a policy experiment can guide and inform decisions about the appropriate response to a given type of oil-price shock. Another implication is that one should exercise caution in using futures prices to forecast the future spot price of oil out-of-sample. Such forecasts will be correct on average, but at a given point in time they tend to be very inaccurate.

The findings also suggest some avenues for further study. A natural next step, for example, would be to get a better understanding of the microeconomics of storage in the market for crude oil. Given the available evidence on the significance of the convenience yield in the crude oil market, as well as the importance of precautionary demand shocks as a driver of oil-price shocks at the macroeconomic level, it makes sense to examine the nature and implications of the precautionary motive for holding stocks of crude oil in finer detail. Studying the incentives facing oil refineries for storing oil would shed light on both the details of this important commodity market, as well as on the wider implications of the decision to store oil.

Literature Cited

Alquist, R. and L. Kilian. 2010. "What Do We Learn From the Price of Crude Oil Futures?" *Journal of Applied Econometrics*. Forthcoming.

Amato, J. and N. Swanson. 2001. "The Real-Time Predictive Content of Money for Output." *Journal of Monetary Economics* 48 (1): 3–24.

Arbatli, E. 2008. "Futures Markets, Oil Prices and the Intertemporal Approach to the Current Account." Bank of Canada Working Paper No. 2008-48.

Bessembinder, H., J. Coughenour, P. Seguin, and M. Monroe Smoller. 1995. "Mean Reversion in Equilibrium Asset Prices: Evidence from the Futures Term Structure." *Journal of Finance* 50 (1): 361–75.

Brennan, M. 1958. "The Supply of Storage." *American Economic Review* 48 (1): 50–72.

Cavallo, M. and T. Wu. 2006. "Measuring Oil-Price Shocks Using Market-Based Information." Federal Reserve Bank of San Francisco Working Paper No. 2006-28.

Chao, J., V. Corradi, and N. Swanson. 2001. "An Out-of-Sample Test for Granger Causality." *Macroeconomic Dynamics* 5 (4): 598–620.

Chernenko, S., K. Schwarz, and J. Wright. 2004. "The Information Content of Forward and Futures Prices: Market Expectations and the Price of Risk." Board of Governors of the Federal Reserve System International Finance Discussion Paper No. 808.

Chinn, M., M. LeBlanc, and O. Coibion. 2005. "The Predictive Content of Energy Futures: An Update on Petroleum, Natural Gas, Heating Oil and Gasoline." National Bureau of Economic Research (NBER) Working Paper No. 11033.

Chinn, M. and O. Coibion. 2009. "The Predictive Content of Commodity Futures." La Follette School of Public Affairs Working Paper No. 2009-016, University of Wisconsin, Madison.

Considine, T. 1997. "Inventories under Joint Production: An Empirical Analysis of Petroleum Refining." *Review of Economics and Statistics* 79 (3): 493–502.

Coppola, A. 2008. "Forecasting Oil Price Movements: Exploiting the Information in the Futures Market." *Journal of Futures Markets* 28 (1): 34–56.

Literature Cited (cont'd)

Davies, P. 2007. "What's the Value of an Energy Economist?" Speech to the International Association for Energy Economics, Wellington, New Zealand, November.

Deaton, A. and G. Laroque. 1992. "On the Behaviour of Commodity Prices." *Review of Economic Studies* 59 (1): 1–23.

Elliott, G., I. Komunjer, and A. Timmermann. 2005. "Estimation and Testing of Forecast Rationality under Flexible Loss." *Review of Economic Studies* 72 (4): 1107–25.

Froot, K. and R. Thaler. 1990. "Anomalies: Foreign Exchange." *Journal of Economic Perspectives* 4 (3): 179–92.

Granger, C. 1969. "Prediction with a Generalized Cost of Error Function." *Operational Research Quarterly* 20 (2): 199–207.

Gustafson, R. 1958. "Carryover Levels for Grains." U.S. Department of Agriculture Technical Bulletin No. 1178.

Inoue, A. and L. Kilian. 2006. "On the Selection of Forecasting Models." *Journal of Econometrics* 130 (2): 273–306.

Kaldor, N. 1939. "Speculation and Economic Stability." *The Review of Economic Studies* 7 (1): 1–27.

Kilian, L. 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review* 99 (3): 1053–69.

Kumar, M. 1992. "The Forecasting Accuracy of Crude Oil Futures Prices." IMF Staff Papers 39 (2): 432–61.

Litzenberger, R. and N. Rabinowitz. 1995. "Backwardation in Oil Futures Markets: Theory and Empirical Evidence." *Journal of Finance* 50 (5): 1517–45.

Ma, C. 1989. "Forecasting Efficiency of Energy Futures Prices." *Journal of Futures Markets* 9 (5): 393–419.

National Petroleum Council. 2004. *Observations on Petroleum Product Supply*. U.S. Department of Energy.

Ng, S. and F. Ruge-Murcia. 2000. "Explaining the Persistence of Commodity Prices." *Computational Economics* 16 (1-2): 149–71.

Pagano, P. and M. Pisani. 2009. "Risk-Adjusted Forecasts of Oil Prices." *The B.E. Journal of Macroeconomics* (Topics) 9 (1): 1–26.

Pindyck, R. S. 1994. "Inventories and the Short-Run Dynamics of Commodity Prices." *RAND Journal of Economics 25* (1) 141–59.

Sadorsky, P. 2002. "Time-Varying Risk Premiums in Petroleum Futures Prices." *Energy Economics* 24 (6): 539–56.

Scheinkman, J. and J. Schechtman. 1983. "A Simple Competitive Model with Production and Storage." *Review of Economic Studies* 50 (3): 427–41.

Schwartz, E. and J. Smith. 2000. "Short-Term Variations and Long-Term Dynamics in Commodity Prices." *Management Science* 46 (7): 893–911.

Williams, J. and B. Wright. 1991. *Storage and Commodity Markets*. Cambridge: Cambridge University Press.

Working, H. 1949. "The Theory of Price of Storage." *American Economic Review* 39 (6): 1254–62.

Wu, T. and A. McCallum. 2005. "Do Oil-Futures Prices Help Predict Future Oil Prices?" Federal Reserve Bank of San Francisco *Economic Letter* 2005-38.

Inflation Expectations and the Conduct of Monetary Policy: A Review of Recent Evidence and Experience

Rose Cunningham, Brigitte Desroches, and Eric Santor, International Economic Analysis

- *Inflation expectations play a critical role in the conduct of monetary policy, providing timely and useful information with respect to the central bank's credibility. Inflation expectations are a key determinant of actual inflation and are thus a crucial part of the analysis used by many central banks to generate inflation forecasts.*
- *Inflation expectations in countries with explicit inflation-targeting monetary policy regimes appear to be more forward looking and better anchored. During the 2008–09 period, despite the high volatility of short-term inflation expectations, expectations for longer-term inflation remained well anchored.*
- *As central banks seek to withdraw from the extraordinary measures enacted during the crisis, inflation expectations will be monitored closely.*

Maintaining price stability is the key objective of most central banks, and the recent financial crisis and global recession have produced important upside and downside risks to price stability. On the upside, inflationary pressure could emerge if monetary policy rates are left too low for too long, if extraordinary measures are withdrawn too slowly, or if fiscal budgets are not consolidated in a timely manner. On the downside are deflationary pressures from substantial and prolonged output gaps. Managing these respective risks to price stability is a key concern for central banks, and inflation expectations can provide them with important information as they consider exit strategies from extraordinary measures and the normalization of monetary policy. Moreover, to achieve their goal of price stability on an ongoing basis, it is crucial that central banks manage inflation expectations through policy actions.

This article explores the role of inflation expectations in the conduct of monetary policy. First, we review the various measures of inflation expectations that are used by central banks, including survey- and market-based indicators, and consider their respective advantages and disadvantages. Second, we examine why inflation expectations are so important in the conduct of monetary policy: namely, their crucial role in the framework used by central banks to understand, forecast, and control inflation. We then explore the role of inflation expectations as an indicator of central bank credibility. Simply, if economic agents view the central bank as credible, inflation expectations are more likely to be well anchored, further enhancing the effectiveness of monetary policy. Interestingly, institutional arrangements, such as adopting inflation targeting, appear to enhance credibility. The importance of credibility is highlighted in the presence of inflation shocks: well-anchored inflation expectations can help the central bank look past temporary shocks to inflation, and required adjustments to the central bank's monetary policy instruments are also greatly reduced.

To shed further light on this subject, we analyze the behaviour of inflation, and inflation expectations, through the lens of the past two years—a challenging episode for central banks, as inflation quickly rose and then fell through 2008 and 2009. We note, however, that inflation expectations in most countries remained remarkably well anchored, despite the massive shocks that were affecting the economy, thus demonstrating the credibility of many central banks. In addition, the maintenance of well-anchored inflation expectations assisted the recovery, as the economy avoided a potentially destabilizing deflationary spiral. Thus, the episode provides valuable lessons with respect to the critical importance of credibility and well-anchored inflation expectations in the conduct of monetary policy. From this experience, we offer policy conclusions and note the need to improve measures of inflation expectations. We also highlight the need to better understand how households and firms form inflation expectations, and how those expectations affect price formation.

Measuring Inflation Expectations

Before examining the importance of inflation expectations in the conduct of monetary policy, it is necessary to look at how they are measured in practice. There are two main sources of information on inflation expectations: surveys and markets. Their relative strengths and weaknesses are considered in turn.

Surveys

Surveys of inflation expectations consider three types of respondents: households, businesses, and professionals (the latter are often referred to as market participants or experts). **Table 1** lists the most commonly referenced surveys, together with details on their structure.[1] Surveys typically ask respondents what they expect inflation to be in the next 4 to 8 quarters and in the next 5 to 10 years. Survey frequency varies from monthly to semi-annually, and most are available from the 1990s onwards.[2] Studies of inflation expectations typically focus on the median range as the relevant indicator, since extreme observations may not be particularly informative. Disagreement among respondents to the same survey can be useful at times, however, since it can be interpreted as disagreement in the population or as a proxy for inflation uncertainty (Mankiw, Reis, and Wolfers 2003).[3]

There are two main sources of information on inflation expectations: surveys and markets.

Most surveys are conducted at the national level: for example, in the United States, the Survey of Professional Forecasters, conducted quarterly by the Federal Reserve Bank of Philadelphia. Other examples include surveys by the University of Michigan, the Banco Central do Brasil, and the Bank of Japan. In Canada, the Conference Board of Canada conducts its Survey of Forecasters each quarter. The Conference Board forecasts are on a calendar-year basis, and the survey reports only the mean of respondents' inflation forecasts. The Bank of Canada's quarterly Business Outlook Survey reports on consultations with about 100 firms across Canada in sectors that broadly reflect the composition of the GDP. The survey asks firms their forecasts of annual consumer price index (CPI) inflation over the next two years, and reports the responses grouped into four ranges: below 1 per cent, 1 to 2 per cent, 2 to 3 per cent, and above 3 per cent.[4]

International surveys, such as Blue Chip Economic Indicators, the IFO World Economic Survey, and Consensus Economics' Consensus Forecast, allow for cross-country comparisons. The most widely used is the Consensus Forecast, which surveys a large cross-section of professional forecasters (currently more than 700 worldwide in more than 85 countries, including Canada), asking each one their predictions

1 For further detail on the Michigan survey, the Livingston Survey, and the Survey of Professional Forecasters, see Curtin (1996), Croushore (1997), and Croushore (1993), respectively.

2 The Michigan Survey of Consumer Attitudes and Behavior has been conducted quarterly since 1946, even though for the first 20 years respondents were asked only whether they expected prices to rise, fall, or stay the same.

3 Disagreement about the future path of inflation tends to rise with the inflation rate or when inflation changes sharply. Surveys of consumers usually reveal greater disagreement than surveys of economists, which show a smaller range of estimates across respondents.

4 The question on inflation expectations in its current form has been part of the survey since 2001.

Table 1: Surveys of inflation expectations

Survey	Participants	Start date	Frequency	Organization	Measures of inflation expectations and horizon
United States					
Michigan Survey of Consumer Attitudes	500 to 700 consumers	1978	Monthly	University of Michigan	Expected change in prices 12 months ahead
Survey of Professional Forecasters	34 professionals	1981	Quarterly	Federal Reserve Bank of Philadelphia	GDP deflator, CPI, PCE, and Core PCE: 6 quarters ahead, 5 and 10 years ahead
Livingston Survey	48 professionals	1946	Semi-annually	Federal Reserve Bank of Philadelphia	CPI: current quarter, 2 quarters ahead, 4 quarters ahead
Europe					
Survey of Professional Forecasters	59 professionals	1999	Quarterly	European Central Bank	CPI: point estimates and density forecasts for 1 year ahead, 2 years ahead, and 5 years ahead
European Commission Consumer Survey	39,900 consumers	1985	Monthly	European Commission	Expected changes in consumer prices 12 months ahead
United Kingdom					
Bank of England-GfK/NOP	2,000 consumers	2001	Quarterly	Bank of England and GfK/NOP	Expected change in shop prices 12 months ahead
Citigroup/YouGov	2,000 consumers	2005	Monthly	YouGov/Citigroup	Expected change in consumer prices of goods 12 months ahead
Canada					
Survey of Forecasters	500 firms	1985	Quarterly	Conference Board of Canada	Percentage of firms expecting price increases over the next 6 months and for the next calendar year
Business Outlook Survey	100 firms	1997	Quarterly	Bank of Canada	Expected annual rate of CPI inflation for the next 2 years
Japan					
Bank of Japan Survey	3,000 consumers	1993	Quarterly	Bank of Japan	Qualitative: will prices go up, down, or stay the same? And reference prices for judging change
Other					
IFO World Economic Survey	1,000 professionals in 90 countries	1991	Quarterly	IFO Research Center, w. support from the European Commission	Expected inflation 6 months ahead
Consensus Economics	700 professionals in 85 countries	1989	Monthly	Consensus Economics Inc.	Inflation for the current year, for next year, and for 5 to 10 years
Blue Chip Economic Indicators	50 professionals	1976	Monthly	Aspen Publishers	Inflation 0 to 7 quarters ahead for the United States, 1 to 2 years ahead for other major economies
Banco do Brasil Business Survey	1,000 professionals	2001	Daily	Banco do Brasil	IPCA inflation over the next 12 months

for growth, inflation, unemployment, and short- and long-term interest rates.

Market-based measures

Inflation expectations can also be inferred from asset prices, such as break-even inflation rates (BEIRs).[5] The break-even inflation rate is the difference between the nominal yield on a fixed-rate bond and the real yield on an inflation-linked (or real return) bond of the same term and maturity. Real return bonds, such as U.S. Treasury Inflation-Protected Securities (TIPS) differ from non-indexed debt securities in that their principal is adjusted for changes in a specified price index.[6] Such indexation protects the purchasing power of the principal, which will have the same real value at maturity in terms of the power to buy items in a consumption basket as when the security was originally issued. BEIRs are easily derived for the United States and the United Kingdom, which have the deepest and most liquid markets for both nominal and real return bonds and issue at a wide range of maturity points (and hence the longest time series for the widest range of forecast horizons). Break-even inflation rates can also be calculated for Canada, France, and some other industrialized countries, but the data are much more limited.

Advantages and disadvantages

Both measures of inflation expectations have advantages and disadvantages. Survey measures have three main advantages: (i) the breadth of coverage is large, including market participants, businesses, and households; (ii) some surveys have been conducted for many decades, allowing comparative analysis from previous inflationary (or deflationary) episodes; and (iii) surveys minimize market distortions, because they avoid certain biases, such as liquidity risk, inflation risk, and institutional distortions, that can affect market-based measures.

Surveys also have several shortcomings: (i) they are often conducted only quarterly or semi-annually and may therefore miss recent changes in inflation expectations. There are also lags between the time they are taken and publication of their results; (ii) surveys may be biased, since households may overweight price changes for frequently purchased goods and services, such as gasoline and food;[7] (iii) comparison of survey results across countries is difficult, given differences in survey methodologies; (iv) responses are equally weighted, irrespective of respondents' ability to forecast inflation; and (v) the answers of some survey respondents may be strategic: market participants may have incentives not to reveal private information and thus tend towards consensus forecasts.

Data from market-based measures have many advantages.[8] First, BEIRs and various other measures of inflation expectations derived from asset prices are available daily. Second, market-based measures may reflect agents' expectations more accurately, since market participants "vote" with real money.[9] In addition, depending on the breadth and depth of the relevant markets, market-based measures can reveal inflation expectations across a wider range of forecast horizons than surveys.

There are, however, some concerns with BEIRs. They may suffer from the fact that the liquidity characteristics of the two instruments (nominal and real return bonds) differ considerably: while there are deep and liquid markets for regular, nominal return bonds, this is less true for real return bonds, and their implied yield may therefore be biased because of a variable liquidity premium between the two.[10] Moreover, during times of market stress, a flight to quality might distort nominal yields disproportionately.[11] Institutional factors and self-selection may also distort the information content of BEIRs, since some investors, such as pension funds and insurance firms, may have strong preferences for real return bonds,[12] thus leading to a premium for those bonds. In addition, BEIRs might suffer from mismatched cash flows. While coupon payments on nominal bonds are fixed, those on real return bonds rise with inflation. This means that each bond will react differently to changes in the expected path and variance of the inflation rate, biasing the BEIR. Finally, if the term structure of

5 Inflation-indexed swaps could also be used to infer inflation expectations. An inflation-indexed swap is a derivative instrument where the payments under the contract depend on the value of an inflation index, such as the CPI.
6 In most cases, the index used is the CPI. A notable exception is the United Kingdom, where the Retail Prices Index is used.
7 Thomas (1999) and Mehra (2002) suggest that the bias in survey forecasts may vary across accelerating versus decelerating inflation environments or across the business cycle.
8 For more details, see Christensen, Dion, and Reid (2004).
9 This is important in the current environment, since feedback between expectations of deflation and postponed consumption and investment would begin only if people act on those expectations.
10 In some markets, inflation-indexed swaps are more liquid than real return bonds, which suggests that inflation-indexed swaps may be a more reliable measure of inflation expectations.
11 During the financial crisis in 2008, the flight to safety implied a large premium for nominal bonds, leading to a large distortion in the BEIR.
12 For example, insurance firms may need to hedge liabilities that are indexed to inflation.

inflation expectations is not flat, BEIRS will be biased, and this bias will be greater at shorter horizons.[13]

Are survey and market measures able to give a reliable picture of current inflation expectations? In the Canadian context, Christensen, Dion, and Reid (2004) find that the BEIR in Canada is not a reliable measure of inflation expectations because of the maturity and liquidity characteristics of Real Return Bonds. Simply, Canada's Real Return Bonds have a 30-year maturity and are considerably less liquid than conventional 30-year bonds, which leads to frequent distortions in the measure of expected inflation. For the United States, Ang, Bekaert, and Wei (2007) find that survey data outperform market-based measures, times-series ARIMA models, and regressions using data on real economic activity. Consequently, the most recent evidence suggests that surveys may be a more reliable guide to inflation expectations for the United States and Canada.

Inflation Expectations and Monetary Policy

Measures of inflation expectations play a key role in the conduct of monetary policy since they provide useful signals with respect to the credibility of the central bank and its long-run inflation objective. Inflation expectations are also a crucial part of the analysis used by many central banks to generate inflation forecasts. Inflation expectations are one of the main drivers of current inflation, because expected inflation influences current wage negotiations, price setting, and financial contracting for investment. Because of this link, central banks can affect current and future inflation by better anchoring agents' expectations of long-term inflation.

Inflation expectations and central bank credibility

The analytical framework used by most central banks assumes that economic agents are mainly forward looking and rational, which has strongly influenced the design of monetary policy (Bernanke 2007). In this framework (and in practice), central banks can manage and stabilize inflation expectations, and hence inflation, through various factors, including the choice of policy regime, their actions, and their communications. For instance, an inflation-targeting regime in which the central bank commits to keep inflation at a specific rate or range over a specified period provides a clear, measurable commitment and a performance target. This policy commitment sends a clear signal to the public and to market participants about the priority of monetary policy and thus helps to anchor inflation expectations.[14] But having the correct regime is not enough: delivery is key. Credibility requires policy actions (Mishkin 2007), since these actions demonstrate the central bank's commitment to price stability and its ability to achieve it—making inflation expectations relatively insensitive to incoming data. Lastly, central banks can improve their credibility through clear and effective communication. Clarity about the goals of the central bank, and how it plans to achieve them, can further anchor inflation expectations, and thus inflation.

Central banks can affect current and future inflation by better anchoring agents' expectations of long-term inflation.

A review of the empirical literature

The theoretical basis for the use of inflation expectations is clear: well-anchored inflation expectations can help the central bank achieve its inflation objective. Whether this is true in practice is essentially an empirical question. To this end, there is a large literature on the interaction between inflation, inflation expectations, and the conduct of monetary policy.

Inflation persistence and inflation expectations

Inflation persistence or inertia is of concern to central banks, since it can inhibit the bank's ability to achieve its inflation objective. In particular, high persistence may suggest that economic agents form inflation expectations in a backward-looking (instead of forward-looking) manner. Backward-looking inflation expectations could therefore indicate that a central bank's credibility is low, potentially impairing the efficacy of its actions. Not surprisingly, postwar inflation data suggest that inflation has often tended to be highly correlated with lagged inflation; i.e., there is persistence in observed inflation. While some persistence may be intrinsic to an economy, the level will likely decline if expectations

13 A detailed explanation of this phenomenon can be found in Christensen, Dion, and Reid (2004).

14 Price-level targeting could have a further stabilizing effect on inflation expectations, and this has been an area of active research by the Bank of Canada. See Ambler (2009) for a review of the research.

become more forward looking as a result of more credible monetary policy (Woodford 2006). Simply, the greater the importance of forward-looking expectations, the less persistent inflation should be (Rudd and Whelan 2007; Sims 2008).[15] Consequently, the question arises as to whether the institutional framework for monetary policy, through its effect on inflation expectations, can lower inflation persistence.

Many have argued that a simple way to make agents more forward looking is to introduce an inflation target. Several recent empirical studies test whether the institutional framework affects inflation persistence. Benati (2008) considers several alternative monetary policy regimes in a recent cross-country study, and finds that for inflation-targeting (IT) countries—Canada, Sweden, the United Kingdom, and New Zealand—inflation was persistent prior to the adoption of the IT regime. But since the adoption of IT, lagged inflation is no longer a statistically significant predictor of current inflation: i.e., persistence has declined. Likewise, inflation persistence in the euro area has declined since the adoption of the euro.[16] More recently, Mendes and Murchison (2009–10) examine inflation persistence in Canada and also find a substantial decline in persistence relative to the 1980s. They emphasize the importance of the adoption of the IT target in 1991. Results for the United States are more ambiguous, however, and seem sensitive to which measure of inflation is used. Benati (2008) finds that inflation measured by the GDP deflator and the Personal Consumption Expenditures (PCE) price index shows considerable persistence in the United States, even after 1995, while estimates of CPI inflation show almost no persistence. Benati's study and others suggest that past inflation experience influences current inflation, as well as expectations about future inflation.[17] However, the extent of this influence appears to decline substantially as the monetary policy regime's commitment to price stability strengthens.[18]

Anchoring inflation expectations

As discussed above, uncertainty about the central bank's objective, or its commitment to a target, can affect expectations of long-run inflation. A key argument in favour of inflation targeting is that it leads to better anchoring of inflation expectations. Several studies empirically test this assertion, which is also an implication of the rational-expectations model: if expectations are perfectly anchored, then long-run inflation expectations should not respond to current inflation (especially periods of higher-than-expected current inflation) or to other news about macroeconomic conditions.

Empirical studies on the United States generally find that its inflation expectations have become more stable since the early 1980s, but they remain somewhat sensitive to current shocks. Stock and Watson (2007) find that changes in the trend component of U.S. inflation are highly persistent, but that the variation in trend inflation has declined substantially since 1983. This implies that unexpected changes to inflation are much less likely to persist in the United States than in the past and, thus, that inflation expectations have become better anchored. Likewise, the response of inflation expectations to shocks from the macroeconomy and from monetary policy has declined over the period, as has the volatility of inflation expectations (Clark and Davig 2008). Nevertheless, there remains enough variability in trend U.S. inflation for Bernanke (2007) to conclude that inflation expectations continue to be imperfectly anchored.

15 This literature also includes related work by Altissimo, Mojon, and Zaffaroni (2009); Cechetti et al. (2007); and Kozicki and Tinsley (2003). The appropriate interpretation for the persistence of inflation is the source of an ongoing debate, particularly among Woodford (2006) and Rudd and Whelan (2007). While Rudd and Whelan are skeptical of rational expectations to describe how expectations are formed, they nevertheless emphasize that the role of expectations in the inflation process is "crucial" (p.32).

16 The number and timing of policy regimes is exogenously determined in Benati's model, and he assumes a single regime from 1971 to 1991. Benati's findings of very high inflation persistence in pre-IT periods may reflect a failure to adequately control for changes in trend inflation for Canada. Crawford, Meh, and Terajima (2009) allow for endogenous timing of regime changes and find that prior to inflation targeting Canada's inflation persistence was considerably lower than Benati's estimate.

17 Other evidence for the United States is mixed: Rudd and Whelan (2007) do not find that U.S. inflation is less persistent after the Volcker disinflation. This is consistent with research by Kozicki and Tinsley (2005, 2009), who find that it took a considerable period to build monetary policy credibility following the Great Inflation of the 1970s. They argue that this was partly due to the lack of a clear inflation target. Cogley and Sbordone (2005, 2008), however, find that inflation persistence in the United States is minimal after controlling for shifts in trend inflation.

18 Improvements in central bank credibility appear to be linked primarily to the choice of an inflation-targeting regime, rather than to additional communication or transparency. A few central banks, in the belief that greater transparency would help anchor expectations, have published their policy interest rate path. Andersson and Hofmann (2009) assess whether these forward guidance strategies of the central banks of New Zealand, Norway, and Sweden have helped anchor expectations of long-term inflation. They find that all three countries already had well-anchored inflation expectations and that publishing the interest rate path, on its own, did not improve the degree to which those expectations were anchored.

Market-based measures of U.S. inflation expectations also suggest imperfect anchoring: Gürkaynak, Sack, and Swanson (2005) demonstrate that forward U.S. interest rates at long horizons react significantly to various macroeconomic and monetary policy surprises. Similarly, Potter and Rosenberg (2007) find that shocks to short-run inflation expectations (2 to 5 years) continue to pass through to measures of long-run inflation expectations (9 to 10 years).

Kozicki and Tinsley (2005, 2009) study the Great Inflation of the 1970s and the post-Volcker disinflation period in detail, and conclude that the lack of an explicit monetary policy objective in the United States contributed to unanchored inflation expectations well into the late 1980s. In contrast, they show that after the Bundesbank announced medium-term targets for money growth in the mid-1970s, bond market measures of German inflation expectations soon began to track the Bundesbank's target.[19] Kozicki and Tinsley (2005) find that private sector expectations were slow to adjust to the lower-inflation regime in the United States, even though actual inflation declined quite quickly after 1979. They interpret their results as consistent with an initial lack of credibility regarding the Federal Reserve's long-term commitment to low and stable inflation following the high-inflation episodes in the 1970s and 1980s.

Several cross-country studies indicate that, as with inflation persistence, inflation expectations seem better anchored in countries with inflation-targeting regimes.

Recent international comparisons also provide evidence on the importance of the monetary policy regime for anchoring expectations. Several cross-country studies indicate that, as with inflation persistence, inflation expectations seem better anchored in countries with inflation-targeting regimes. These studies find that, unlike non-targeting countries, inflation expectations in IT countries, such as Canada, the United Kingdom, and Sweden, are not correlated to actual inflation, nor are they as sensitive to macroeconomic news or monetary policy surprises.[20] Overall, the existing research implies that expectations of long-term inflation in the United States are stable but imperfectly anchored, while countries with explicit inflation targets appear to have better-anchored expectations of long-term inflation. Moreover, better-anchored inflation expectations lead to lower inflation persistence.

Other influences on inflation expectations

The importance of credibility and the monetary policy regime for anchoring inflation expectations may ignore other important features of the inflation process, such as relative prices, especially those for food and energy, which are beyond the immediate control of the central bank. Clark and Davig (2008) find that shocks to food prices have a significant and persistent effect on expectations of long-run inflation in the United States. Energy prices, however, were not found to have a significant impact.[21] Galati, Poelhekke, and Zhou (2008) test whether the sharp increases in food and commodity prices that occurred between 2006 and mid-2008 led to a de-anchoring of inflation expectations in the euro area. Employing market data on interest rate swaps and inflation swaps (a more liquid market than inflation-indexed bonds) to measure inflation expectations in the euro area, they find evidence that inflation expectations became more sensitive to inflation news after June 2007, suggesting some drift in the inflation expectations of market participants away from the ECB's target. This evidence suggests that policy-makers must not take well-anchored inflation expectations for granted.

Recent Trends in Inflation Expectations

The 2008–09 period provides an excellent lens through which to examine the importance of inflation expectations for the conduct of monetary policy.

Survey data: History and the crisis

The historical behaviour of inflation expectations has evolved largely in line with the theory presented above: as central bank credibility has improved, inflation

19 At the end of 1974, the Bundesbank began a regime officially described as money targeting; however, Bernanke and Mihov (1997) argue that inflation, rather than money growth, seemed to be the actual target variable.

20 See for example, Levin, Natalucci, and Piger (2004); Gürkaynak, Levin, and Swanson (2006); Gürkaynak et al. (2006); and Beechey, Johannsen, and Levin (2008).

21 The authors point out that energy prices are volatile and that forecasters may, therefore, expect their movements to be transitory and may not place much weight on price changes. Food prices, however, tend to be more persistent and also make up a larger share of the CPI basket, which may lead forecasters to incorporate food-price movements into their expectations more readily.

Table 2: The development of inflation and inflation expectations

		Canada		Euro area		Japan		Norway		Sweden		Switzerland		United Kingdom		United States	
		mean	std. dev	mean	std. dev	mean	std. dev	mean	std. dev	mean	std. dev	mean	std. dev	mean	std. dev	mean	std. dev
Actual Inflation (CPI y/y)	I	1.4	0.7	1.8	0.6	0.5	0.9	2.1	0.6	0.9	1.2	0.7	0.6	2.0	0.5	2.4	0.6
	II	2.4	0.9	2.1	0.4	-0.6	0.4	2.2	1.4	1.6	0.9	0.9	0.5	1.2	0.3	2.5	0.8
	III	2.0	0.7	2.2	0.8	0.2	0.8	2.1	1.2	1.6	1.3	1.2	0.9	2.5	0.8	2.9	1.6
Expected inflation 4 quarters ahead	I	1.8	0.3	-	-	0.4	0.6	2.9	0.4	2.0	1.1	1.0	0.2	2.8	0.5	2.8	0.4
	II	2.0	0.3	1.6	0.2	-0.4	0.3	1.9	0.6	1.9	0.4	1.1	0.4	2.3	0.2	2.3	0.4
	III	1.8	0.6	1.8	0.5	0.1	0.6	2.0	0.7	1.7	0.9	1.0	0.4	2.3	0.9	2.2	1.1
Expected inflation 8 quarters ahead	I	2.0	0.3	-	-	0.8	0.6	2.6	0.5	2.4	0.8	1.2	0.0	3.1	0.6	3.0	0.4
	II	2.0	0.1	1.7	0.1	-0.2	0.5	2.2	0.2	2.1	0.3	1.3	0.2	2.3	0.1	2.4	0.3
	III	1.9	0.1	1.9	0.3	0.6	0.3	2.0	0.3	1.9	0.3	1.2	0.3	2.5	0.3	2.2	0.2
Expected inflation 5 to 10 years ahead	I	2.0	0.3	-	-	1.4	0.5	2.1	0.1	2.4	0.5	1.9	0.2	3.0	0.4	3.0	0.4
	II	2.0	0.1	1.9	0.1	0.8	0.4	2.3	0.2	2.0	0.0	1.6	0.1	2.3	0.1	2.5	0.1
	III	2.0	0.0	1.9	0.0	1.3	0.2	2.3	0.1	2.0	0.0	0.0	0.1	2.6	0.2	2.3	0.2

Note: Period I runs from the second half of 1994 to the first half of 1999; period II runs from the second half of 1999 to the first half of 2004; and period III runs from the second half of 2004 to the first half of 2009.
Source: Consensus Economics

expectations have become better anchored. **Table 2** presents the mean and standard deviation for actual CPI inflation and for inflation expectations 4 quarters ahead, 8 quarters ahead, and 5 to 10 years ahead, for seven advanced economies and the euro area. The data are divided into three periods: period I ranges from the second half of 1994 to the first half of 1999, period II from the second half of 1999 to the first half of 2004, and period III from the second half of 2004 to the first half of 2009. As a general trend, the mean and variance of inflation expectations converged over time to the mean and variance of actual inflation, especially in the 1994–99 and 1999–2004 periods. Moreover, as the forecast horizon becomes longer, expectations are much closer to actual inflation. Importantly, as a potential signal of growing central bank credibility and well–anchored inflation expectations, the standard deviation of the inflation forecasts 5 to 10 years ahead for most countries has fallen sharply over the past 10 years. Canada is an exception, since inflation expectations and actual inflation had already fallen significantly by 1994, and therefore the improvement over the periods considered in the table is smaller than for other countries.

An examination of more recent data provides valuable insights into the importance of well-anchored inflation expectations. Survey data from Consensus Economics for Canada, the United States, the United Kingdom, and the euro area show that since 2007, expectations of short-run inflation have been quite volatile **(Chart 1)**. In particular, the sharp and seemingly persistent rise in energy, food, and commodity prices in 2008 led to higher headline inflation, a feature that was reflected in rising expectations of shorter-term inflation. For example, in mid-2008, expectations for inflation 4 quarters ahead reached more than 3 per cent in the United States and the United Kingdom, over 2.5 per cent in the euro area, and increased in Canada but to slightly less than

Chart 1: Inflation expectations

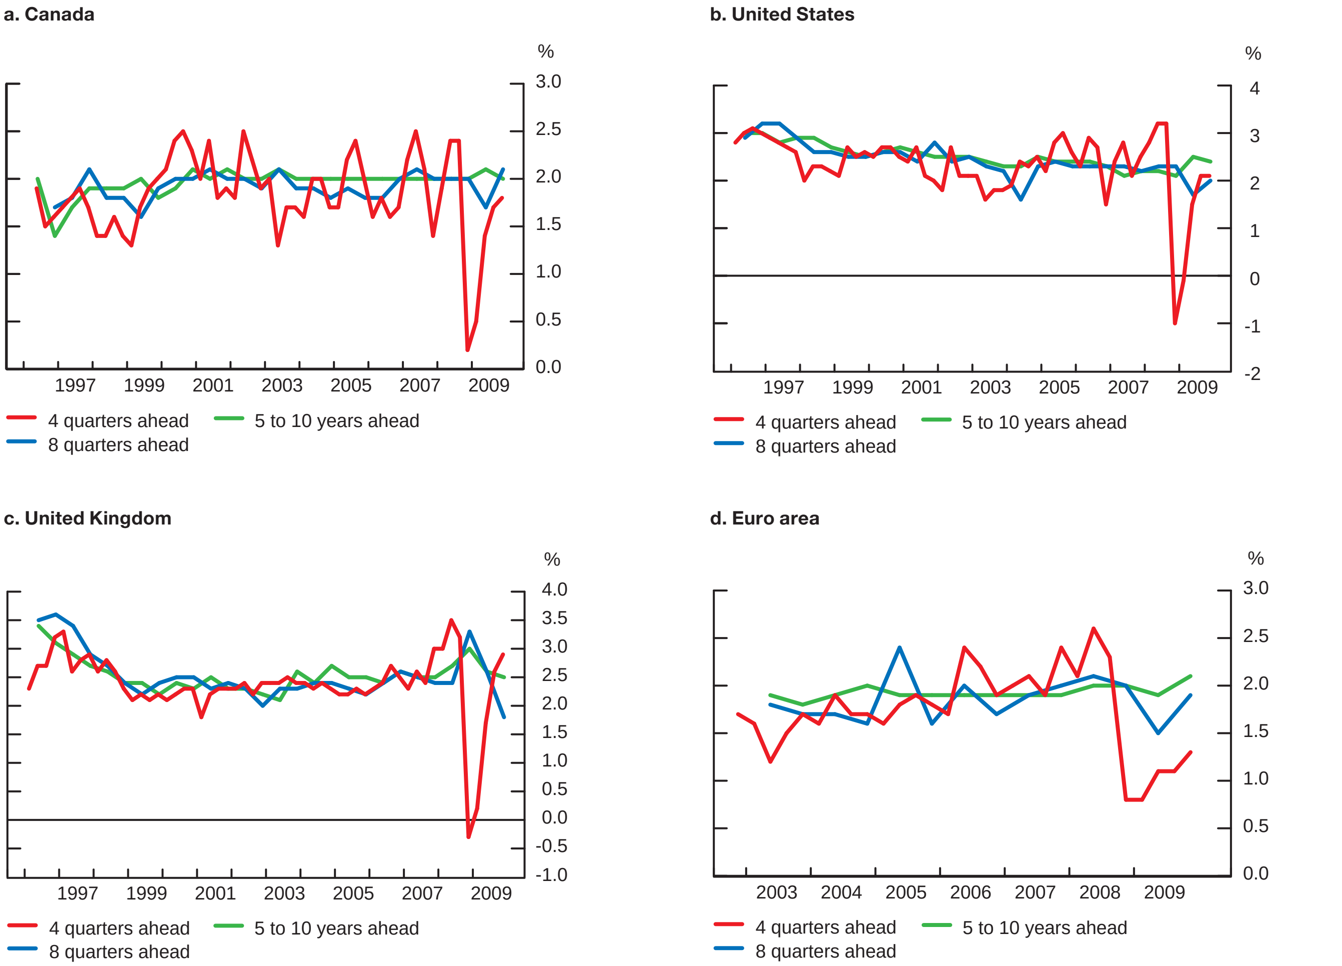

Source: Consensus Economics

2.5 per cent. The increase in inflation expectations in mid-2008 was even more prevalent for emerging-market economies **(Chart 2)**. Following the collapse of Lehman Brothers, the economic and financial turbulence intensified, and expectations for inflation 4 quarters ahead fell sharply, actually turning negative in the United States (December 2008 to March 2009), the United Kingdom (December 2008), and Sweden (March 2009). This was partly due to the collapse in commodity prices and fears of a sharp recession. In Canada, inflation expectations 4 quarters ahead also fell, although to a lesser extent. The sharp decline in expectations of short-run inflation at the height of the credit crisis suggests that market participants in some countries expected deflation in 2009, albeit that expectation was short lived. In fact, inflation expectations began to rise again later in 2009 as economies began to stabilize, although they currently remain lower than the levels prior to the crisis.

Despite the high volatility of short-term inflation expectations, expectations for longer-term inflation remained better anchored. Expectations for inflation 8 quarters ahead declined much less. During the most severe part of the crisis, the Bank of Canada's Business Outlook Survey of inflation expectations 2 years ahead found that over 40 per cent of firms expected inflation below 1 per cent, but by the second quarter of 2009 that share had dropped back to just 11 per cent of respondents **(Chart 3)**. Furthermore, expectations for long-term inflation (5 to 10 years ahead) remained essentially flat in most countries (ranging from 2.0 per cent to 2.5 per cent), despite the observed negative rates of inflation and the length of the recession. In consumer surveys, expected inflation

Chart 2: Inflation expectations in emerging-market countries

Six months ahead

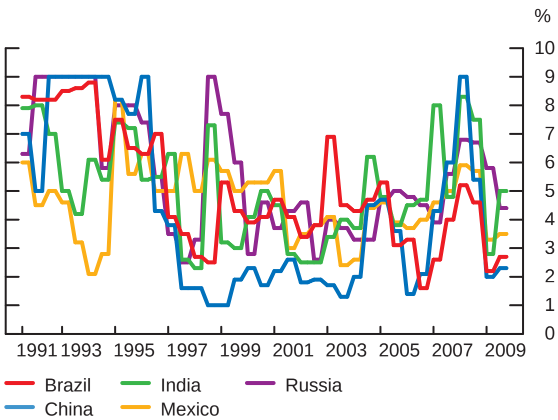

Note: The index is constructed by aggregating responses and assigning a value of 1 to lower inflation, 5 to same inflation, and 9 to higher inflation.
Source: IFO Institute

Chart 3: Bank of Canada Business Outlook Survey

Percentage of firms expecting CPI inflation over the next two years to be:

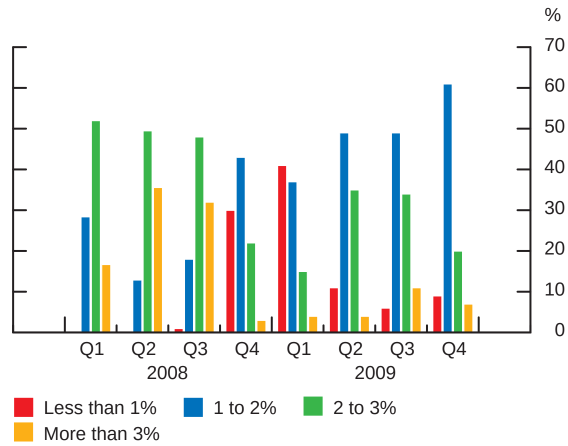

Source: Bank of Canada

Chart 4: Inflation expectations in the United States (Michigan Survey)

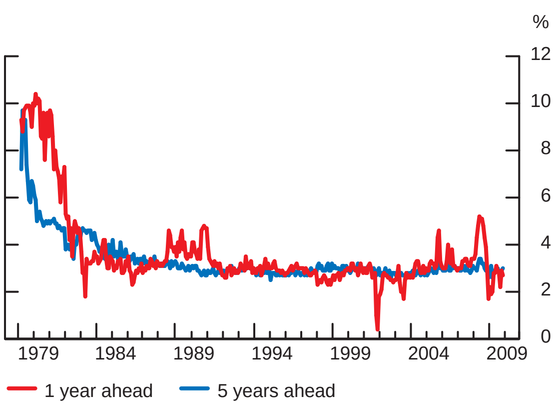

Sources: DataStream/Michigan University

Chart 5: Breakeven inflation rates

Selected maturities

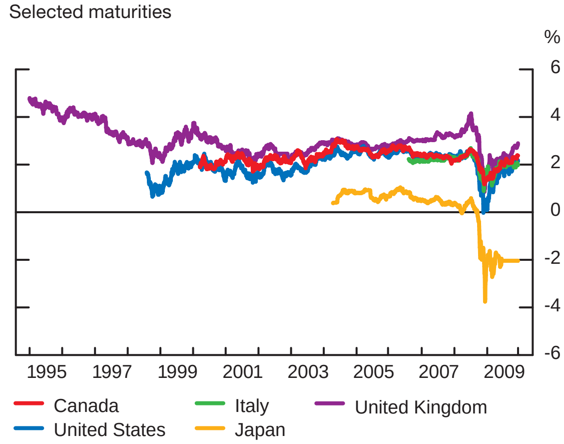

Source: Bloomberg

5 years ahead also remained relatively flat over the crisis period. For example, the Michigan survey indicates that U.S. consumers believed that inflation one year ahead would fall between mid-2008 and late-2008, but they did not expect deflation. In fact, consumers' inflation expectations remained close to the average of 3 per cent reported for the last decade **(Chart 4).**[22] Thus, in both types of surveys (households and professionals), long-term inflation expectations remained well anchored, and the central banks' credibility remained intact.

22 Although there appears to be a persistent upward bias in consumer surveys, one observes that consumers' inflation expectations move roughly in line with the inflation expectations of professional forecasts. For example, the Bank of Japan's consumer survey of expectations of inflation one year ahead for March 2008 was 7.6 per cent, while the Consensus inflation forecast for the same period was 0.4 per cent. While less extreme, the Michigan survey of households also reports inflation expectations that are on average 0.5 percentage points higher than typical expert forecasts (from 1996 to 2008).

Market-based measures

The survey data suggest that inflation expectations, while volatile in the short run, are well anchored for longer horizons. However, since many surveys occur only quarterly or semi-annually, they may not have captured the true volatility of inflation expectations during the crisis. To address this issue, we examine BEIRs for the United States, the United Kingdom, Canada, and Japan.[23] As in the surveys of professionals, expectations for long-term inflation for the United States remained well anchored, except for a brief period at the height of the crisis, when BEIRs fell to close to zero. In Japan, BEIRs fell below zero in 2009 **(Chart 5)**.[24] In contrast, inflation expectations for the inflation-targeting United Kingdom and Canada remained above 1 per cent during the crisis for the BEIR measures considered.

Lessons from recent evidence

The recent financial crisis highlights the usefulness of inflation expectations within the framework for conducting monetary policy. Before the crisis, energy and food prices increased significantly, and expectations for shorter-term inflation rose accordingly. Expectations for long-term inflation remained well anchored, however, since households and firms were able to look through the commodity-price shock. This anchoring of expectations allowed policy-makers to look past the increase in energy prices, avoiding a possible policy mistake (in the absence of well-anchored inflation expectations, central banks may have been forced to raise interest rates just as the crisis was about to intensify, only to reverse them later). Clearly, the gains associated with well-defined inflation objectives and enhanced credibility helped to anchor inflation expectations and thus the inflation outcomes for many central banks.

The experience of the crisis emphasized the importance of well-anchored inflation expectations.

The experience of the crisis itself, from the collapse of Lehman Brothers onwards, again emphasized the importance of well-anchored inflation expectations. In the autumn of 2008, commodity prices fell dramatically, and fears of a severe recession intensified. In fact, expectations for inflation 4 quarters ahead also fell sharply, and even went negative in some jurisdictions (market-based measures revealed a similar pattern). However, longer-term inflation expectations remained well anchored, despite the opening of large and likely persistent output gaps. This clearly indicates that central banks maintained their credibility, despite the massive shocks that were affecting the economy. The maintenance of well-anchored expectations assisted the recovery, since the economy avoided a potentially destabilizing deflationary spiral.

Conclusions and Avenues for Future Research

Inflation expectations play a critical part in the conduct of monetary policy, providing timely and useful information with respect to the central bank's credibility. Inflation expectations form a key part of the information set used by central banks to understand and forecast inflation. Importantly, much of the existing research indicates that central banks that have a clear and credible commitment to low and stable inflation, especially those with inflation-targeting regimes, have been very successful in anchoring inflation expectations over the past two decades.

Inflation expectations will continue to inform policy-making, as central banks seek to withdraw from the extraordinary measures enacted during the crisis and beyond. In normalizing monetary policy rates, inflation expectations will be monitored closely, given the crucial role of credibility in anchoring inflation expectations. Massive fiscal stimulus packages and future pressure on fiscal budgets related to demographic change have led to record fiscal deficits and to high projected ratios of debt to GDP over the coming years in many advanced economies. Some market participants have expressed concerns that debt levels may become unsustainable and will eventually be monetized (although this concern has not yet materialized in measures of inflation expectations). In such an environment, inflation expectations can provide a useful leading indicator of whether fiscal and monetary policy credibility has been maintained.

Further research is required in several areas. First, how households and market participants form inflation expectations is not well understood. Bernanke

23 BEIRs may suffer from liquidity risk: while there are deep and liquid markets for regular, nominal return bonds, this is less true for real return bonds, and their implied yield may therefore be biased.

24 Expectations of negative inflation for Japan are not surprising, given Japan's recent experience with deflation.

(2007) has called for more emphasis on incorporating learning and imperfect information in the modelling of inflation and of inflation expectations. Second, both survey and market measures exhibit biases over time. Accounting for these biases when interpreting measures of inflation expectations requires further consideration. Finally, more cross-country data on inflation expectations are needed, especially on the expectations of firms and business owners. Since business representatives participate directly in setting prices and wages, more insight into the inflation expectations of price-setting firms in a larger set of countries and over different time horizons, would be very helpful to policy-makers.

Literature Cited

Altissimo, F., B. Mojon, and P. Zaffaroni. 2009. "Can Aggregation Explain the Persistence of Inflation?" *Journal of Monetary Economics* 56 (2): 231–41.

Ambler, S. 2009. "Price-Level Targeting and Stabilization Policy: A Review." *Bank of Canada Review* (Spring): 19–29.

Andersson, M. and B. Hofmann. 2009. "Gauging the Effectiveness of Central Bank Forward Guidance." Paper presented at the 24th Annual Congress of the European Economic Association and the 64th European meeting of the Econometric Society (EEA-ESEM), Barcelona, Spain, 23–27 August.

Ang, A., G. Bekaert, and M. Wei. 2007. "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" *Journal of Monetary Economics* 54 (4): 1163–1212.

Beechey, M., B. Johannsen, and A. Levin. 2008. "Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area than in the United States?" Federal Reserve Board Finance and Economics Discussion Series No. 2008-23.

Benati, L. 2008. "Investigating Inflation Persistence across Monetary Regimes." *Quarterly Journal of Economics* 123 (3): 1005–60.

Bernanke, B. 2007. "Inflation Expectations and Inflation Forecasting." Speech to the Monetary Economics Workshop of the NBER Summer Institute, Cambridge, Massachusetts, 10 July.

Bernanke, B. and I. Mihov. 1997. "What Does the Bundesbank Target?" *European Economic Review* 41 (6): 1025–53.

Cecchetti, S., P. Hooper, B. Kasman, K. Schoenholtz, and M. Watson. 2007. "Understanding the Evolving Inflation Process." Paper presented at the U.S. Monetary Policy Forum, Washington D.C., 9 March.

Christensen, I., F. Dion, and C. Reid. 2004. "Real Return Bonds, Inflation Expectations, and the Break-Even Inflation Rate." Bank of Canada Working Paper No. 2004–43.

Clark, T. and T. Davig. 2008. "An Empirical Assessment of the Relationships among Inflation and Short- and Long-Term Expectations." Federal Reserve Bank of Kansas City Research Working Paper No. RWP 08-05.

Cogley, T. and A. Sbordone. 2005. "A Search for a Structural Phillips Curve." Federal Reserve Bank of New York Staff Report No. 203.

———. 2008. "Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve." *American Economic Review* 98 (5): 2101–26.

Crawford, A., C. Meh, and Y. Terajima. 2009. "Price-Level Uncertainty, Price-Level Targeting, and Nominal Debt Contracts." *Bank of Canada Review* (Spring): 31–41.

Croushore, D. 1993. "Introducing: The Survey of Professional Forecasters." *Federal Reserve Bank of Philadelphia Business Review* November/December: 3–15.

———. 1997. "The Livingston Survey: Still Useful after All These Years." *Federal Reserve Bank of Philadelphia Business Review* March/April: 15–27.

Curtin, R. 1996. "Procedure to Estimate Price Expectations." University of Michigan Survey Research Center. Mimeo.

Galati, G., S. Poelhekke, and C. Zhou. 2008. "Did the Anchor of Inflation Expectations in the Euro Area Turn Adrift?" De Nederlandsche Bank Working Paper No. 191/2008.

Gürkaynak, R., A. Levin, A. Marder, and E. Swanson. 2006. "Inflation Targeting and the Anchoring of Inflation Expectations in the Western Hemisphere." Central Bank of Chile Working Paper No. 400.

Gürkaynak, R., A. Levin, and E. Swanson. 2006. "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from Long-Term Bond Yields in the U.S., U.K., and Sweden." Federal Reserve Bank of San Francisco Working Paper No. 2006-09.

Gürkaynak, R., B. Sack, and E. Swanson. 2005. "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models." *American Economic Review* 95 (1): 425–36.

Kozicki, S. and P. Tinsley. 2003. "Alternative Sources of the Lag Dynamics of Inflation?" In *Price Adjustment and Monetary Policy*, 3–47. Proceedings of a conference held by the Bank of Canada, November 2002. Ottawa: Bank of Canada.

———. 2005. "What Do You Expect? Imperfect Policy Credibility and Tests of the Expectations Hypothesis" *Journal of Monetary Economics* 52 (2): 421–47.

———. 2009. "Perhaps the 1970s FOMC Did What It Said It Did." *Journal of Monetary Economics* 56 (5): 842–55.

Levin, A., F. Natalucci, and J. Piger. 2004. "The Macroeconomic Effects of Inflation Targeting." *Federal Reserve Bank of St. Louis Review* 86 (4): 51–80.

Mankiw, N. G., R. Reis, and J. Wolfers. 2003. "Disagreement about Inflation Expectations." National Bureau of Economic Research (NBER) Working Paper No. 9796.

Mehra, Y. P. 2002. "Survey Measures of Expected Inflation: Revisiting the Issues of Predictive Content and Rationality." *Federal Reserve Bank of Richmond Economic Quarterly* 88 (3): 17–36.

Mendes, R. and S. Murchison. 2009–10. "Declining Inflation Persistence in Canada: Causes and Consequences." *Bank of Canada Review* (Winter 2009–2010): 1–14.

Mishkin, F. 2007. "Inflation Dynamics." National Bureau of Economic Research (NBER) Working Paper No. 13147.

Potter, S. and J. Rosenberg. 2007. "Are US Inflation Expectations Anchored, Contained or Unmoored?" Federal Reserve Bank of New York. Mimeo.

Rudd, J. and K. Whelan. 2007. "Modeling Inflation Dynamics: A Critical Review of Recent Research." *Journal of Money, Credit and Banking* 39 (S1): 155–70.

Sims, C. 2008. "Inflation Expectations, Uncertainty, the Phillips Curve, and Monetary Policy." Paper presented at the Federal Reserve Bank of Boston conference "Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective," Chatham, Massachusetts, June 9–11.

Stock, J. and M. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money*, *Credit and Banking* 39 (S1): 3–33.

Thomas, L. B. Jr. 1999. "Survey Measures of Expected U.S. Inflation." *Journal of Economic Perspectives* 13 (4): 125–44.

Woodford, M. 2006. "Interpreting Inflation Persistence: Comments on the Conference on 'Quantitative Evidence on Price Determination'." Available at <http://www.columbia.edu/~mw2230/QEPDcomment.pdf>.

Monetary Policy Rules in an Uncertain Environment

Gino Cateau and Stephen Murchison, Canadian Economic Analysis Department

- *Central banks have increasingly focused on a systematic approach to monetary policy. Simple monetary policy rules help to facilitate the communication of monetary policy to the public and enhance its predictability.*
- *Monetary policy rules have become an integral part of central bank models and are often fine-tuned to maximize economic welfare. However, uncertainty about the "true" model can seriously affect the performance of these rules and should therefore be accounted for when designing robust rules.*
- *Simple policy rules can often provide a good approximation to fully optimal policy under perfect information and are typically more robust to uncertainty.*
- *In ToTEM, an optimized simple rule that responds to a forecast of the price level is more robust to parameter uncertainty than a rule that responds to inflation.*

Monetary policy is most effective when the central bank's objectives, and the means of achieving those objectives, are well understood and regarded as credible by the public. This requires that the central bank communicate clearly what it seeks to achieve, such as inflation control over the medium term, and how its current and future actions can be expected to bring about the desired outcome(s). Since the collection and processing of information is costly for private agents, it is in the central bank's own best interest to respond to economic developments in a predictable fashion that is easy to communicate. Not only does this facilitate a better understanding of current policy actions, but it permits markets to better forecast the central bank's future actions.

Beginning with the seminal work of Taylor (1993), academic researchers and central banks have increasingly focused on the benefits of a systematic approach to the design of monetary policy. Monetary policy rules, or reaction functions, have become an integral part of central bank models and are often fine-tuned to maximize economic welfare. However, such fine tuning is inherently risky when the central bank has an imperfect understanding of how the economy functions.

This article discusses recent research on the influence of various forms of economic uncertainty on the performance of different classes of monetary policy rules: from simple rules to fully optimal monetary policy under commitment. Building on the research discussed in the Summer 2002 issue of the *Bank of Canada Review*, we explain why uncertainty matters for policy-rule design and provide quantitative examples from the recent literature, which has increasingly focused on structural models that feature rational expectations. We also present results for several policy rules in ToTEM, the Bank of Canada's main projection and policy analysis model (Murchison and Rennison 2006), including rules that respond to the price level, rather than to inflation.

The article begins with a brief discussion of the theoretical arguments in favour of commitment to a policy rule and the role played by such rules in the design of real-world monetary policy. It then discuss the four major forms of uncertainty with which central banks must contend when formulating policy and how each type affects the performance of various rules. It concludes with a brief review of strategies for designing so-called robust rules: i.e., rules that perform well across a broad range of economic models.

What Is a Monetary Policy Rule?

For our purposes, a policy rule can be thought of as a mathematical equation that determines the appropriate level for the central bank's policy instrument as a function of one or more economic variables that describe the state of the economy.[1] Given that such rules are specified in terms of the policy instrument, they are often called instrument rules. An essential feature of such a rule is that while the policy interest rate varies through time in response to economic developments, its response to a given shock or state of the economy does not. Therefore, adherence to a rule is synonymous with predictability, and thus private agents in the economy understand how policy will respond now and in the future.

One might question why a central bank would adhere to a single rule, since doing so might constrain it in unfavourable ways. Even if the central banks' objectives do not vary through time, it may wish to maintain a high level of discretion in how it responds to the economy. The simple answer is that no central bank literally sets policy based on a single rule. For various reasons beyond the scope of this article, central banks do exercise a certain degree of judgment or discretion when setting policy. But this does not render the discussion of policy rules academic. What matters is that monetary policy is predictable from the viewpoint of private agents, whose decisions are influenced by current and future policy actions. From this perspective, the central bank's strict commitment to a published rule can be seen as one extreme, whereas choosing policy at each point in time in a purely discretionary fashion can be seen as the opposite extreme.

Adherence to a rule is synonymous with predictability, and thus private agents in the economy understand how policy will respond now and in the future.

Recent empirical research generally supports the idea that monetary policy in many industrialized countries does contain a large systematic component. For instance, much of the interest in the so-called Taylor rule (Taylor 1993) is based on the observation that it predicts the actual behaviour of the federal funds rate in the United States over the period 1987–92 reasonably accurately. Thus, while no central bank literally follows a rule, their actual behaviour may be well approximated by such a rule. This is likely due, at least in part, to the fact that modern central bank projection models feature policy rules and that these models are used to provide policy advice.

So why do central banks behave in a manner broadly consistent with adherence to a rule? One key benefit is predictability. Monetary policy is most effective when households and firms understand both the objectives of monetary policy and how the central bank goes about achieving those objectives. By explicitly or implicitly committing to a certain pattern of behaviour, a central bank can influence private sector expectations of the future path of the policy rate, which, in turn, can help the central bank achieve its objectives. For instance, suppose a central bank has earned a reputation for responding aggressively to inflation whenever it strays from the target. Then, when an unanticipated shock causes inflation to deviate from the target, the deviation will be perceived as short lived. As a result, agents' expectations of future inflation will not respond to the shock, which, in turn, will dampen the current inflation response. In this way, a credible commitment to respond aggressively to shocks that affect inflation, combined with private sector expectations that factor in that commitment, can attenuate the required policy response.

1 This is a somewhat narrow definition. In the economics literature, a rule can either describe how the policy instrument reacts to the state of the economy, or it can prescribe a particular economic outcome, such as the achievement of the central bank's inflation target—hence the label "targeting rules," (Svensson 1999). In the latter case, the behaviour of the policy instrument can be inferred only in the context of a full model that links the policy instrument to the targeting variables included in the rule.

Types of Rules

Since the general consensus among central bankers is that the long-run objective of monetary policy should be price stability, a natural starting point would be to design a rule that ensures long-run price stability. For example, the Bank of Canada aims to maintain the growth rate of the consumer price index (CPI) at the 2 per cent midpoint of a 1 to 3 per cent control range. According to the conventional view of the monetary policy transmission mechanism, inflation tends to decline when interest rates are high, other things being equal, and increase when interest rates are low. Therefore, an appropriate rule would stipulate that the Bank raise the target overnight interest rate[2] when current CPI inflation exceeds 2 per cent and lower it when inflation is below 2 per cent.

Restricting one's focus to the long-run objective of price stability represents an overly narrow view of the role of monetary policy. It is generally acknowledged that monetary policy can focus on, although not necessarily fully achieve, multiple *short-run* objectives. For instance, a central bank may care about stabilizing both inflation around the target and real GDP around potential GDP. To the extent that certain shocks push inflation and the output gap in opposite directions, a short-run trade-off exists, which will be reflected by the inclusion of both inflation and the output gap in the policy rule.

Perhaps the best-known policy rule is the Taylor rule (Taylor 1993), which was estimated using U.S. data and is given by:

$$R_t = 4.0 + 1.5(\pi_t - 2) + 0.5\tilde{y}_t, \tag{1}$$

where R_t is the U.S. federal funds rate, π_t is the rate of price inflation, and $\tilde{y}_t$ is the output gap, all in period t. According to the Taylor rule, when inflation equals 2 per cent and output equals potential output, the federal funds rate should be set equal to 4 per cent—400 basis points (bps). Moreover, that rate should be adjusted by 150 bps up or down for every 1-percentage-point difference between actual inflation and the desired level of 2 per cent, and 50 bps for every 1-per-cent difference between output and potential output. The Taylor rule's greatest virtue may be its simplicity, since the policy rate in a given period can be described in terms of just two economic variables.[3]

The Taylor rule is a special case of a broader class of so-called simple rules. There are important extensions to this basic set-up that include (a) lagged interest rates as an additional argument in the rule, and (b) replacing current inflation by a forecast of future inflation. A lag of the interest rate was initially added because it resulted in a better fit of the data (Clarida, Galí, and Gertler 2000), and it suggests that, in response to a change in economic conditions, central banks adjust the policy rate gradually over several months, rather than all at once, as suggested by the Taylor rule. Woodford (1999) has argued that interest rate smoothing or *inertia* is actually consistent with optimal central bank behaviour when economic agents form their expectations in a forward-looking manner. As the relative weight on the lagged interest rate increases, the future value of the policy rate becomes easier to predict, since it is determined to a greater extent by the current rate.

Responding to a forecast of future, rather than current, inflation is also consistent with optimal behaviour if monetary policy exerts its maximum effect on inflation with a lag and if the central bank is good at forecasting inflation. The policy rule currently used in ToTEM includes a role for both the lagged policy interest rate and a forecast of future inflation, and is described by the equation:

$$R_t = \rho R_{t-1} + (1-\rho)\big[R^* + \varphi_\pi(E_t\pi_{t+k} - \pi^T) + \varphi_y(\tilde{y}_t)\big], \tag{2}$$

where R_t is the target overnight interest rate in period t, R^* is the long-run, neutral rate of interest, $E_t\pi_{t+k}$ is the period t expectation of inflation in period $t+k$, and $\tilde{y}_t$ is the output gap. ρ, φ_π, and φ_y are fixed parameters that determine the degree of interest rate smoothing and the sensitivity of the policy rate to deviations of inflation from target and to the output gap, respectively.[4] Note that k determines the degree to which policy is forward looking and is referred to as the "feedback horizon."

The rules discussed so far summarize the behaviour of monetary policy in terms of just a few economic variables, such as expected inflation and the output

2 The target for the overnight interest rate is the conventional policy instrument in Canada.

3 Potential output was proxied by a simple linear trend of log GDP in Taylor's specification, which is straightforward to calculate.

4 In the current version of ToTEM used for projections, the optimized parameter values are $\rho = 0.95$, $\varphi_\pi = 20$, $\varphi_y = 0.35$, $k = 2$, and $R^* = 4.75$ per cent.

gap. Explaining the movements in the policy rate from one period to the next is, therefore, straightforward. But this simplicity typically comes at the price of reduced performance in terms of economic stabilization. To see why, consider first that the forecast of inflation depicted in equation (2) will depend on every variable in the economic model, and in a fully articulated model, such as ToTEM, the number of economic variables can be considerable. Implicitly, the strength of the central bank's response to each of these variables is governed by a single parameter: φ_π in equation (2).

But suppose that instead of forcing monetary policy to respond to forecast inflation, we allocate a separate response parameter for each variable that influences future inflation, including the exogenous shocks that hit the economy. Such a set-up describes the essential features of fully optimal monetary policy under commitment. Such a rule will better stabilize the economy if the central bank's model is correct and if the data used in the model are well measured. But as we discuss in the next section, such a rule may perform quite poorly if one or both of these assumptions turns out to be false.

Types of Uncertainty Faced by Central Banks

In this section we discuss the four main types of economic uncertainty facing policy-makers and how each affects the performance of different policy rules.

Shock uncertainty

In practice, a monetary policy rule represents one equation in a central bank's model of the economy. At a minimum, the model will also include equations governing the behaviour of the variables that enter the policy rule, such as inflation and the output gap. Taken together, these equations form a self-contained system that can be simulated through time to generate a path for the policy interest rate that is consistent with the outlook for inflation, and vice versa.

Economic models, however sophisticated, are by construction simple caricatures of the true economy (Coletti and Murchison 2002). They are intended to capture those linkages between households, firms, governments, and the central bank believed to be the most important, on average. Nevertheless, the deliberate omission of many idiosyncratic factors means that models will make prediction errors, which are referred to as shocks, and the associated uncertainty is referred to as *shock uncertainty*. To understand the impact of shock uncertainty on the performance of a policy rule, it is helpful to understand how policy rules are parameterized.

Economic models, however sophisticated, are by construction simple caricatures of the true economy.

For central bank models, such as ToTEM, that are used to provide policy advice, the parameters of the policy rule are normally chosen to minimize an assumed loss function,[5] which in ToTEM includes the variance of CPI inflation relative to the 2 per cent inflation-control target, the variance of the economy-wide output gap, and the variance of the change in the target overnight interest rate.[6] The variances of these endogenous variables will depend on the structure and calibration of the economic model, the policy rule, and the variances and covariances of the shocks included in the model, which are normally estimated using historical data. Choosing optimal parameters for the rule involves using the covariance matrix of shocks, in conjunction with the model, to compute variances for the endogenous variables that appear in the loss function. The task then is to choose parameter values in the policy rule that minimize the expected loss.

In general, the optimal parameter values in the rule will depend importantly on which shocks were most important over history, as well as on the covariances among shocks.[7] This is because simple rules must trade off performance for simplicity. As a very simple example, consider an economy with just two shocks: a demand shock that pushes output and inflation in the same direction, and a supply shock that moves them in opposite directions. Also assume that while the central bank seeks to stabilize output and inflation, the policy interest rate responds only to inflation. In this set-up, the optimal response to a demand shock will be larger than the optimal response to a supply shock, since the policy response to a supply shock pushes output away from potential output. Therefore, the optimal response to inflation in the

5 Details of the loss function and of the optimized rule currently used in ToTEM are described in Cayen, Corbett, and Perrier (2006) and Murchison and Rennison(2006).

6 The respective weights in the loss function are 1, 1, and 0.5.

7 Cayen, Corbett, and Perrier (2006) provide examples using ToTEM.

policy rule will depend on the relative importance of demand versus supply shocks in the economy.

This simple example illustrates that the performance of optimal simple rules will depend on the nature of the shocks that hit the economy. If the relative importance of various shocks changes through time, the performance of a simple rule will no longer be optimal. In contrast, since a fully optimal rule responds optimally to each shock, the parameter values of the rule do not depend on the relative importance of the various shocks.[8] Relative to other sources of uncertainty discussed in this article, shock uncertainty is unique in that it renders simple rules less robust than optimal rules.

Data and measurement uncertainty

Much of the data used in economic models, with the notable exceptions of the CPI and the labour force survey in Canada, is subject to periodic revision. As a general rule, recently released data are subject to larger revisions than data that have already been revised several times. When formulating policy, central banks must therefore be aware that the data on which they rely to gauge the current state of the economy contain a potentially important noise component.

In addition to errors associated with data collected by statistical agencies, central banks must often construct data for variables that are not directly measurable. An important example is the trend level of labour productivity. While measures of actual labour productivity are available from Statistics Canada, the underlying trend or permanent component must be estimated, and this is typically done using a statistical filter.[9] Since these filters are often two-sided (i.e., the estimate of the trend in a given period is based on both past and future observations of the data being filtered), their accuracy declines as they approach the end of the sample, since there are fewer future observations on which to condition the estimate.

In designing an optimal monetary policy rule, a central bank would typically respond more cautiously to a variable measured with error. To see why, we refer back to the example in which the estimated level of trend labour productivity is a noisy measure of the true level. Since potential output is constructed using trend labour productivity, the output gap will inherit much of this noise. Now, consider a central bank that uses a policy rule of the form given by equation (1), which can now be written in terms of the true output gap and the noise component, ε_t^y, as

$$R_t = R^* + \varphi_\pi(\pi_t - \pi^T) + \varphi_y(\tilde{y}_t + \varepsilon_t^y). \qquad (3)$$

Equation (3) reveals the nature of the information problem. By choosing to respond positively to the output gap (the variable measured with error), the policy-maker inadvertently reacts to the noise. This introduces undesirable movements in the interest rate, which feed back to the economy and generate unnecessary fluctuations in output and inflation. Cateau, Desgagnés, and Murchison (forthcoming) illustrate this point using an inflation-targeting rule in ToTEM. The results are presented in **Table 1**.

Table 1: Effects of data uncertainty

	σ_π	$\sigma_{\tilde{y}_t}$	$\sigma_{\Delta R}$	Loss
No data uncertainty $R_t = 1.09R_{t-1} + 0.54\pi_t + 0.13\tilde{y}_t$	1.06	1.09	0.55	1
Data uncertainty ignored $R_t = 1.09R_{t-1} + 0.54\pi_t + 0.13\tilde{y}_t$	1.31	1.10	0.56	+12%
Data uncertainty accounted for $R_t = 1.14R_{t-1} + 0.61\pi_t + 0.08\tilde{y}_t$	1.04	1.25	0.51	+6%

The top panel of Table 1 shows an optimized inflation-targeting rule under the assumption that the output gap in ToTEM is perfectly measured; the middle panel evaluates the performance of that rule when the output gap is, in fact, not perfectly measured.[10] Ignoring the measurement errors in the output gap leads to additional volatility in inflation, the output gap, and the change in the interest rate, culminating in a 12 per cent deterioration in the rule's performance.

Of course, a policy-maker who recognizes that the information at his disposal is not accurate need not naively follow a rule that is efficient only in the absence of data uncertainty. Indeed, as is clear from equation (3), by choosing to respond less aggressively

8 For this reason, optimal policy under commitment is said to be *certainty equivalent*.
9 Butler (1996) provides a detailed discussion of the estimations of trend labour productivity and trend labour input that are used in the Bank of Canada's conventional measure of potential output.

10 Cateau, Desgagnés, and Murchison (forthcoming) allow for data uncertainty by computing the discrepancies between the real-time and revised values of the Bank of Canada's conventional estimate of potential output and modelling the resulting measurement errors as an AR(2) process.

to the central bank's measure of the output gap, the influence of the noise can be reduced. The bottom panel of Table 1 presents an optimized rule that accounts for the presence of measurement errors in the output gap. Owing to the difficulty of accurately measuring the output gap, the resulting rule gives it a lower weight[11] but places higher weights on inflation and policy inertia. This leads to a more volatile output gap but allows better control of inflation and of changes in the interest rate. Ultimately, the new rule reduces the influence of output gap mismeasurement relative to the baseline rule by half.

Parameter uncertainty

While economic theory can guide modellers on the nature of certain economic relationships, it rarely provides much guidance on the exact strength of the relationship. For instance, theory says that Canadian exports to the United States will strengthen, other things being equal, following a depreciation of the real Canada/U.S. exchange rate, since Canadian goods become more competitive. But the size of the export response is unknown. It must therefore be estimated using historical data and will be subject to sampling uncertainty, even if the underlying theory is correct. In this sense, policy-makers should regard the parameters of their model as random variables with some underlying distribution, rather than as known, fixed quantities.

Viewed from this perspective, it is natural to ask what differentiates parameter uncertainty from shock uncertainty, since shocks are also modelled as random variables. The crucial difference lies in the fact that a model's parameters enter multiplicatively, meaning that they interact with the model's endogenous variables, whereas shocks are additive. Thus, while the optimal parameter values of a simple policy rule depend on the *relative* variances of the model's shocks, the absolute variances are unimportant.[12] If we think about the model's parameters as random variables, however, absolute variances do matter.

Consider the famous example given by Brainard (1967), in which inflation is linearly related to the policy instrument, and there is an exogenous demand shock, u_t:

$$\pi_t = -\theta R_t + u_t,$$

and the central bank's objective is to minimize the variance of inflation. The optimal policy rule with no parameter uncertainty sets the interest rate in each period to $(1/\theta)u_t$, and inflation is perfectly stabilized at zero each period. However, if the parameter relating the instrument to the target is not known with certainty, the central bank's model will be characterized by:

$$\pi_t = -(\theta - \varepsilon)R_t + u_t = -\theta R_t + u_t + R_t\varepsilon,$$

where ε is a random variable. There are now, in effect, two shocks in the model, and the multiplier on the second one is the nominal interest rate. If the central bank implements the same policy as discussed above, the variance of inflation will be unnecessarily high. The optimal policy rule that accounts for parameter uncertainty in this example is $[\theta/(\theta^2 + \sigma_\varepsilon^2)]u_t$, where σ_ε^2 is the variance of ε. As the degree of parameter uncertainty increases, the optimal response coefficient in the rule declines. This finding is called the "Brainard conservatism principle" (Blinder 1998).

In addition to introducing uncertainty regarding the linkages between observed variables, such as inflation and the policy interest rate, parameter uncertainty also creates uncertainty about the correct level of unobserved, model-defined variables. For instance, in ToTEM, the real marginal cost of production in the consumption-goods sector is the key driver of core CPI inflation (Murchison and Rennison 2006). Since Statistics Canada does not provide a measure of real marginal cost, it is calculated within ToTEM, and its properties reflect both the structure and the parameterization of the model. As a result, parameter uncertainty introduces additional uncertainty about the future evolution of inflation through its influence on marginal cost.

11 This result is in accordance with the literature. Smets (1999) shows that when measurement error in the output gap becomes very large, the efficient Taylor rule parameter on the output gap falls towards zero. Orphanides (2003) shows that once the measurement errors between real-time and ex-post data are properly taken into account, optimized policy reactions are more cautious than otherwise.

12 Slightly more technically, multiplying the covariance matrix of shocks by a scalar will not affect the optimal parameter values of a simple rule, since doing so will not affect the relative variances of the endogenous variables that enter the central bank's loss function.

Table 2: Robustness of optimized inflation- and price-level-forecast rules

Benchmark rule	IF	PLF	FO
No parameter uncertainty			
Performance: $\frac{loss(rule\ j)}{loss\ (IF\ )} - 1$	1	-4.3%	-11.4%
Parameter uncertainty			
Robustness: $\frac{E\ loss(rule\ j \mid Under\ uncertainty)}{loss(rule\ j \mid No\ uncertainty)} - 1$	+80%	+81%	+142%
Overall average performance: $\frac{E\ loss(rule\ j \mid Under\ uncertainty)}{E\ loss\ (IF \mid Under\ uncertainty\ )} - 1$	1	-3.4%	+21%

Finally, any time that a monetary policy rule responds to a forecast of inflation (or of any other variable), the performance of that rule will be influenced by parameter uncertainty, since the forecast will not be as precise. Parameter uncertainty can thus be thought of as introducing noise into the inflation forecast in a manner similar to measurement uncertainty (see equation 3), thereby rendering that variable less reliable as a guide for policy. In the end, whether it is better to respond to current inflation or to a forecast of future inflation, will depend on the benefit of being forward looking, in the absence of parameter uncertainty, relative to the cost of introducing additional noise in the policy rule.[13]

Cateau, Desgagnés, and Murchison (forthcoming) derive optimized inflation-forecast (IF) and price-level-forecast (PLF) rules for ToTEM and compare their performance with fully optimal policy under commitment (FO).[14] They then investigate the robustness of these rules to parameter uncertainty by analyzing how they would fare if the structural parameters that actually characterize the behaviour of private agents differed from those assumed by the policy-maker in deriving the optimized rules **(Table 2)**. These types of comparisons are of particular interest in light of the Bank of Canada's interest in evaluating the potential welfare gains of switching from its current inflation-targeting regime, to a price-level-targeting regime.[15] Furthermore, most of the research to date that explores this issue ignores altogether the issue of uncertainty.

The top panel of Table 2 compares the performance of the optimized inflation-forecast rule, price-level-forecast rule, and optimal policy under full commitment, using ToTEM's baseline calibration. Without parameter uncertainty, fully optimal policy under commitment offers an 11.4 per cent improvement in performance over IF, while PLF would offer a 4.3 per cent improvement.

The authors go on to investigate how parameter uncertainty affects these rankings by evaluating the performance of each benchmark rule in 5000 alternative parameter configurations drawn randomly from the Bayesian posterior distribution of the estimated parameters. The bottom panel of Table 2 contains two important messages. First as recently emphasized by Orphanides and Williams (2008), while fully optimal policy under commitment is the best policy if the parameters are known, it is often the least robust policy under uncertainty. Indeed, relative to the case of no uncertainty, its performance deteriorates 60 percentage points more than the other rules. Second, while IF is slightly more robust than PLF, on average, PLF still performs better than IF under parameter uncertainty. Therefore, while the reduction in loss associated with moving from inflation targeting to

13 The extent of the benefit of setting policy in a forward-looking manner depends on the speed of the monetary policy transmission mechanism. All else being equal, the faster policy actions get transmitted to output and inflation, the less need there is to be forward looking.

14 The optimized inflation-forecast rule for ToTEM is a rule that responds to current inflation, the lagged interest rate, and the output gap. In contrast, the optimized price-level-forecast rule responds to the price-level forecast four quarters ahead, the lagged interest rate, and the output gap. The price-level-forecast rule is an example of a rule that implements price-level targeting, since this rule will eventually return the price level to the desired level following a shock. Optimal policy under full commitment is the policy that is optimally tuned to the model. It is, by design, a very complicated rule that depends on every variable that affects the state of the economy. Optimal policy does not, in general, fully reverse price-level movements following a shock in ToTEM and, therefore, is not fully consistent with a price-level-targeting regime.

15 See Bank of Canada (2006).

price-level targeting in ToTEM is modest, this reduction is robust to parameter uncertainty.[16]

Model uncertainty

So far, we have discussed uncertainty about the underlying shocks that drive business cycles, uncertainty about the data used in a particular model, and uncertainty about the parameter values used in the model. But what about the economic model itself? A model may be misspecified for various reasons: it may be built around an economic paradigm that is further from economic reality than assumed (Engert and Selody 1998); it may ignore economic relationships that are, in fact, relevant; or it may be built under simplifying assumptions that make the model tractable (e.g., linearity) but less realistic. Since a model is ultimately only one view of how the economy works, a policy rule that is tuned to work well in a particular model may perform poorly across alternative but plausible views.

Côté et al. (2002) analyze the performance of various simple rules in 12 models of the Canadian economy. They find that simple outcome-based rules (rules where the policy instrument responds to current and lagged variables) are not particularly robust. In particular, they find that rules with high degrees of inertia often induce substantial volatility in output and inflation and are even unstable in many models.

Since a model is ultimately only one view of how the economy works, a policy rule that is tuned to work well in a particular model may perform poorly across alternative but plausible views.

More recently, Tetlow (2010) evaluates the performance of 8 alternative simple rules in 46 vintages of the Federal Reserve Board FRB/US model used by the Board's staff for forecasting and policy analysis from July 1996 to October 2007. He concludes that model uncertainty is a substantial problem: model properties differ importantly according to vintage and so do the policy rules optimized by vintage. Further, while some rules offer satisfactory performance, many that are promoted as being robust to some specific type of uncertainty perform poorly when confronted with real-time model uncertainty.

Once we acknowledge that any particular model is potentially misspecified, the results above indicate that model uncertainty can seriously affect the performance of policy rules in stabilizing the economy and, hence, should be taken into account when designing effective policy rules. In the next section, we review recent strategies for designing rules that are robust to specific forms of uncertainty, including model uncertainty.

Robust Policy Rules

When designing policy rules, it is important to seek a robust rule—one that yields a satisfactory performance in an uncertain environment. There are two approaches to designing a robust rule. The first involves deriving optimized coefficients that formally account for specific uncertainties. That is, given a specific rule, we determine how strongly the policy instrument should respond to each variable in the rule, taking into account the features about which we are uncertain. The second approach involves determining a functional form for the rule (i.e., what variables the policy instrument responds to) that is less susceptible to yielding a poor performance, given specific uncertainties. These approaches are complementary and are often combined when pursuing a robust simple rule. In this section, we review how they have been or could be applied to design rules robust to each of the uncertainties discussed.

Robustness to data uncertainty

There are two main approaches for designing a rule robust to data uncertainty. The first, alluded to earlier, involves formally taking into account that data are observed with noise and will subsequently be revised. A common strategy for dealing with this problem follows Orphanides (2003) in modelling the measurement errors between real-time and ex-post data and incorporating these equations in the model prior to optimizing the rule. To the extent that future measurement errors may behave like historical errors, this strategy helps the policy-maker design a rule that accounts for likely mismeasurement of the data.

16 This is an important result, since ToTEM's baseline calibration assumes a very small role for dynamic indexation to lagged price and wage inflation. Thus, the distributions for these parameters are positively skewed, and nearly all of the risk is on the upside. In addition, research (see Ambler 2009 for a review) suggests that higher levels of dynamic indexation (or rule-of-thumb behaviour) will cause a deterioration in the relative performance of price-level-forecast rules.

Table 3: Robust inflation- and price-level-forecast rules

$$R_t = \rho R_{t-1} + \varphi_\pi E_t \pi_{t+k} + \varphi_P E_t P_{t+k} + \gamma \tilde{y}_t$$

		Coefficients of rule				Benchmark parameters		Parameter uncertainty	
Rule *j*	ρ	φ_π	φ_P	γ	k	σ_R	$\frac{loss(j)}{loss\,(IF\,)} - 1$	Robustness: $\frac{E\,loss(j)}{loss\,(j\,)} - 1$	Overall: $\frac{E\,loss(j)}{loss\,(IF)} - 1$
No uncertainty									
IF	1.09	0.54	0	0.13	0	1.48	1	+80%	+80%
PLF	0.99	0	0.07	0.17	4	1.84	-4.3%	+81%	+73%
Parameter uncertainty									
IF	1.01	0.46	0	0.14	0	1.56	+1%	+70%	+72%
PLF	1.01	0	0.08	0.21	3	2.04	-4.1%	+68%	+61%

An alternative approach is to design a rule that does not respond to variables measured with error. Taylor's original rule was criticized by Orphanides et al. (2000) and by Orphanides and Williams (2002) for including unobservable variables, such as the natural rate of interest and potential output (or natural rate of unemployment). Given the difficulty of measuring these variables in real time, Orphanides and Williams (2002) propose difference rules in which the short-term nominal interest rate is raised or lowered from its existing level in response to inflation and to changes in economic activity (change in unemployment or growth rate of output). These rules do not require knowledge of the natural rates of interest or unemployment (or potential output) for setting policy and are consequently immune to mismeasurement. Orphanides et al. (2000) and Orphanides and Williams (2002) show that, in the presence of data uncertainty, these difference rules outperform rules that respond to levels of economic activity. But how do such difference rules perform in environments characterized by other forms of uncertainty?

Tetlow (2010) evaluates the performance of the difference rule proposed by Orphanides and Williams (2002) in 46 vintages of the Federal Reserve Board FRB/US model. The experiment provides an ideal laboratory for evaluating the robustness of a rule since it incorporates real-time model and parameter uncertainty in a model used for policy-making. Tetlow observes that the difference rule does lead to robust performance in the sense that a difference rule optimized for a particular vintage maintains good stabilization properties across all other vintages.

Robustness to parameter uncertainty

The most popular approach for deriving a rule robust to parameter uncertainty is the Bayesian approach, which assumes that unknown parameters come from known distributions. That is, even though the precise values of parameters are not known, it is possible to determine the range of values that they can take, together with their associated probabilities. A robust rule is then derived by choosing the coefficients of the rule to minimize the expected loss, given the distribution of parameters. **Table 3** presents the results of Cateau, Desgagnés, and Murchison (forthcoming) who derive robust inflation-forecast and price-level-forecast rules for ToTEM under parameter uncertainty.[17]

The top panel of Table 3 displays the optimized inflation-forecast (IF) and price-level-forecast rule (PLF) with the estimated parameters of ToTEM as benchmark. The bottom panel displays the robust versions of the IF and PLF rule under parameter uncertainty. The results suggest three important messages:

1. PLF is more robust than IF under parameter uncertainty. The last column compares the overall performance of each rule under parameter uncertainty. The robust PLF rule dominates the robust IF rule by 11 percentage points.

17 Cateau, Desgagnés, and Murchison (forthcoming) allow for parameter uncertainty by allowing a set of key parameters to take 5000 possible values drawn randomly from the Bayesian posterior distribution of the estimated parameters. The robust inflation-forecast and price-level-forecast rules minimize expected loss; i.e., the weighted average of the losses across the draws.

2. Robustness to parameter uncertainty in ToTEM leads to more aggressive policy responses. For instance, the robust PLF rule requires more aggressive responses to the lagged interest rate, forecast price level, and output gap. This translates into more aggressive policy responses as shown by an increase in the unconditional standard deviation in the interest rate, σ_R, from 1.84 to 2.04 per cent. The robust IF rule, on the other hand, requires weaker responses to the lagged interest rate and current inflation but stronger responses to the output gap. The stronger response to the output gap dominates, making policy responses slightly more aggressive (the standard deviation of the interest rate increases from 1.48 to 1.56 per cent).[18]
3. While Bayesian robust rules improve policy performance under parameter uncertainty, they do not offer a significant improvement. The second to last column assesses the robustness of the various rules by comparing their average performance under parameter uncertainty with their performance under no uncertainty. Although the robust IF and PLF rules improve performance over the benchmark IF and PLF rules by 10 and 13 percentage points, respectively, they still lead to a high average loss under uncertainty (respectively 70 per cent and 68 per cent higher than the loss that the benchmark IF rule leads to under no uncertainty). Note, however, that this increase in average loss may also reflect that, on average, the alternative parameterizations of the model make inflation and the output gap more difficult to control, relative to the baseline calibration.

The third result illustrates a disadvantage of the Bayesian approach as a tool for deriving robust rules. By design, the Bayesian approach tunes the policy rule to work best across those parameter configurations that are the most probable: i.e., receive the most probability weight. This yields a policy rule that works well in parameter configurations that are most likely to be true but whose performance suffers in the more extreme, but less likely, parameter configurations.

18 Edge, Laubach, and Williams (2010) also find that parameter uncertainty leads to more aggressive policy in a micro-founded model. Uncertainty about the structural parameters in their model leads to uncertainty about the implicit "natural" rates of output and interest. They find that optimal Taylor rules under parameter uncertainty respond less to the output gap and more to price inflation than would be optimal without parameter uncertainty. But the more aggressive response to inflation dominates, making policy more aggressive.

An alternative approach that offers more robustness to extreme parameter configurations is the worst-case approach. For example, Giannoni (2002) proposes a worst-case approach that does not require knowledge of the distribution of the unknown parameters. Instead the policy-maker knows only the bounds for each parameter and seeks robust policy rules that minimize loss in the worst-case parameterization within those bounds. Giannoni (2002) finds that a policy-maker that seeks to mitigate the effect of parameter uncertainty in a standard New Keynesian model would choose Taylor rules that respond more aggressively to both inflation and the output gap.

Both approaches are useful in determining robust versions of a particular choice of rule. Levin et al. (2006) use a micro-founded model to investigate what types of simple rules are effective when the central bank faces parameter uncertainty. They find that the performance of optimal policy is closely matched by a simple operational rule that responds to the lagged interest rate and focuses solely on stabilizing nominal-wage inflation. Furthermore, this simple wage-stabilization rule is robust to uncertainty about the structural parameters and to various assumptions regarding the nature and incidence of the innovations. However, the performance of the rule is sensitive to the specification of wage contracts in the labour market. Indeed, when Taylor contracts rather than Calvo contracts are assumed, rules that respond to price inflation and real economic variables perform better than the wage-inflation rule. Hence, the robustness of wage-inflation rules hinges critically on structure and wage determination in labour markets.

Robustness to model uncertainty

There are two popular approaches to deriving robust rules under model uncertainty. The first allows the policy-maker to consider different candidate models (e.g., those reflecting different paradigms of the monetary policy transmission mechanism) and seeks policy choices that perform well on average (Brock, Durlauf, and West 2007) or on a worst-case basis. Cateau (2007) proposes a decision-making framework where a policy-maker can consider various non-nested models for choosing policy. His framework distinguishes between two types of risk: within-model risk (risk arising because of the stochastic nature of a particular model) and across-model risk (risk arising as a result of contemplating various models). He shows that the policy-maker's aversion to across-model risk determines the extent to which the policy-maker wants to trade off good average performance for robustness: as the degree of aversion to across-model risk increases, the policy-maker wants to achieve more robustness at the expense of good average performance. Cateau shows that when the policy-maker wants to achieve more robustness, the policy-maker chooses less-aggressive Taylor rules that are in line with those estimated in the data.

Levin, Wieland, and Williams (2003) compare the performance of various outcome-based and forecast-based rules with the objective of identifying one rule that would perform well across five distinct models of the U.S. economy. For their model set, they find that a robust rule is a forecast-based rule that responds to a short-horizon forecast of inflation (less than one year), the current output gap, and also involves a high degree of inertia.

The second approach derives policy choices that are robust to misspecification of the policy-maker's baseline model. In this approach, the policy-maker takes into account that his baseline model is only an approximation of some unknown true model and, hence, can potentially be misspecified. In particular, the dynamics of the baseline model may omit important explanatory variables, as in Hansen and Sargent (2008), or parameters affecting the relationship between different variables may be unknown, as in Onatski and Stock (2002). The policy-maker deals with these misspecifications by choosing policy according to the worst-case model in a set of plausible models. Sargent (1999), Onatski and Stock (2002), and Tetlow and von zur Muehlen (2001) find that robust rules are, in fact, more aggressive than those obtained when potential misspecifications are ignored.

Conclusion

Monetary policy is most effective when the central bank's objectives, and the means of achieving those objectives, are well understood and regarded as credible by the public. This requires that the central bank communicate clearly what it seeks to achieve and, further, requires the central bank to respond to economic developments in a predictable and systematic fashion that is easy to communicate.

Since Taylor (1993), academic researchers and central banks have increasingly used simple rules as a guide to setting monetary policy. Simple rules have the advantage of being easier to communicate to the public than more complex policies and, by virtue of their simplicity, offer the promise of making monetary policy more easily understood and predictable. But what simple rule should a central bank use? The various uncertainties that central banks must contend with make the choice and design of a simple rule difficult.

The results surveyed here suggest that uncertainty has a substantial impact on the performance of simple rules. Although simple rules perform better in an uncertain environment than more complex policies, their performance can still deteriorate substantially. It is therefore critical to account for uncertainty in designing rules to ensure that their performance is satisfactory irrespective of the state of the world.

Work with ToTEM suggests that a price-level-forecast rule is more robust to uncertainty than an inflation-forecast rule.

Our work with ToTEM suggests that a price-level-forecast rule is more robust to uncertainty than an inflation-forecast rule. While more research in this area is required, these results suggest that greater insulation from the effects of economic uncertainty may be an additional rationale for considering price-level targeting over inflation targeting. Finally, based on the literature, other rules shown to have good robustness properties, which also warrant further research, include a difference rule, where the change in the interest rate responds to output growth, as well as a wage-inflation rule.

Literature Cited

Ambler, S. 2009. "Price-Level Targeting and Stabilization Policy: A Review." *Bank of Canada Review* (Spring): 19–29.

Bank of Canada. 2006. *Renewal of the Inflation-Control Target*. Ottawa: Bank of Canada.

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

Brainard, W. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review* 57 (2): 411–25.

Brock, W. A., S. Durlauf, and K. West. 2007. "Model Uncertainty and Policy Evaluation: Some Theory and Empirics." *Journal of Econometrics* 136 (2): 629–64.

Butler, L. 1996. "The Bank of Canada's New Quarterly Projection Model Part 4: A Semi-Structural Method to Estimate Potential Output: Combining Economic Theory with a Time-Series Filter." Bank of Canada Technical Report No. 77.

Cateau, G. 2007. "Monetary Policy under Model and Data-Parameter Uncertainty." *Journal of Monetary Economics* 54 (7): 2083–2101.

Cateau, G., H. Desgagnés, and S. Murchison. Forthcoming. "Robustness of Policy Rules." Bank of Canada Working Paper.

Cayen, J.-P., A. Corbett, and P. Perrier. 2006. "An Optimized Monetary Policy Rule for ToTEM." Bank of Canada Working Paper No. 2006-41.

Clarida, R., J. Galí, and M. Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics* 115 (1): 147–80.

Coletti, D. and S. Murchison. 2002. "Models in Policy-Making." *Bank of Canada Review* (Summer): 19–26.

Côté, D., J.-P. Lam, Y. Liu, and P. St- Amant. 2002. "The Role of Simple Rules in the Conduct of Canadian Monetary Policy." *Bank of Canada Review* (Summer): 27–35.

Edge, R., T. Laubach, and J. Williams. 2010. "Welfare-Maximizing Monetary Policy under Parameter Uncertainty." *Journal of Applied Econometrics* 25 (1): 129–43.

Engert, W. and J. Selody. 1998. "Uncertainty and Multiple Paradigms of the Transmission Mechanism." Bank of Canada Working Paper No. 98-7.

Giannoni, M. 2002. "Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model," *Macroeconomic Dynamics* 6 (1): 111–44.

Hansen, L. P. and T. Sargent. 2008. *Robustness*. Princeton, N.J.: Princeton University Press.

Levin, A., A. Onatski, J. Williams, and N. Williams. 2006. "Monetary Policy under Uncertainty in Micro-Founded Macroeconomic Models." In *NBER Macroeconomics Annual 2005*, edited by M. Gertler and K. Rogoff, 229–87. Cambridge, MA: MIT Press.

Levin, A., V. Wieland, and J. Williams. 2003. "The Performance of Forecast-Based Monetary Policy Rules under Model Uncertainty." *American Economic Review* 93 (3): 622–45.

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

Onatski, A. and J. Stock. 2002. "Robust Monetary Policy under Model Uncertainty in a Small Model of the U.S. Economy." *Macroeconomic Dynamics* 6 (1): 85–110.

Orphanides, A. 2003. "Monetary Policy Evaluation with Noisy Information." *Journal of Monetary Economics* 50 (3): 605–31.

Literature Cited (cont'd)

Orphanides, A., R. Porter, D. Reifschneider, R. Tetlow, and F. Finan. 2000. "Errors in the Measurement of the Output Gap and the Design of Monetary Policy." *Journal of Economics and Business* 52 (1/2): 117–41.

Orphanides, A. and J. Williams. 2002. "Robust Monetary Policy Rules with Unknown Natural Rates." *Brookings Papers on Economic Activity* 2: 63–18.

———. 2008. "Learning, Expectations Formation, and the Pitfalls of Optimal Control Monetary Policy." *Journal of Monetary Economics* 55 (Supplement 1): S80–S96.

Sargent, T. 1999. "Comment." In *Monetary Policy Rules*, edited by J.B. Taylor, 144–54. Chicago: University of Chicago Press and NBER.

Smets, F. 1999. "Output Gap Uncertainty: Does It Matter for the Taylor Rule?" In *Monetary Policy under Uncertainty*, edited by B. Hunt and A. Orr, 10–29. Wellington, New Zealand: Reserve Bank of New Zealand.

Svensson, L. 1999. "Inflation Targeting as a Monetary Policy Rule." *Journal of Monetary Economics* 43 (3): 607–54.

Taylor, J. B. 1993. "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* 39: 195–214.

Tetlow, R. 2010. "Real-time Model Uncertainty in the United States: 'Robust' policies put to the test." Board of Governors of the Federal Reserve System Finance and Economics Discussion Series No. 2010-15.

Tetlow, R. and P. von zur Muehlen. 2001. "Robust Monetary Policy with Misspecified Models: Does Model Uncertainty Always Call for Attenuated Policy?" *Journal of Economic Dynamics and Control* 25 (6/7): 911–49.

Woodford, M. 1999. "Optimal Monetary Policy Inertia." National Bureau of Economic Research (NBER) Working Paper No. 7261.

An Uncertain Past: Data Revisions and Monetary Policy in Canada

*Greg Tkacz, Canadian Economic Analysis Department**

- *Policy-makers rely on macroeconomic data released by Statistics Canada, such as consumption and GDP growth, to gauge the current state of the economy. Such variables are necessarily released with a lag, however, and past observations are subject to revision. Such uncertainty complicates the task of forecasters and policy-makers.*
- *In recent years, economists have tried to document the uncertainty inherent in initial data releases by analyzing the nature of the revisions. Analysis of data revisions for Canada is now possible, using newly constructed databases that track the data as they were released.*
- *Revisions to Canadian GDP growth tend to be smaller, on average, than those of some major OECD countries and are also somewhat less volatile.*
- *The growth rates of GDP components tend to be revised more substantially than the growth rate of GDP itself, rendering the analysis and forecasting of components more difficult. The growth of exports and imports tends to be subject to the largest revisions.*
- *Data revisions can affect policy decisions in different ways. We discuss issues that analysts, researchers, and policy-makers may need to confront.*

Economic forecasters must deal with two issues that do not necessarily confront forecasters in other fields: (i) delays in the release of current-period data and (ii) revisions to past data. National Accounts data are released about two months after the end of each quarter. This implies that forecasters trying to predict the *future* path of National Accounts variables are often unsure as to where those variables actually lie *today*. This has led to the development of a specialized area of forecasting, dubbed *nowcasting*, which is described in more detail below. With respect to the second issue, economists are also confronted with possibly non-trivial revisions to past observations of key economic variables. This has implications for the estimation of economic models and for the forecasts produced using them. For example, if the growth of the gross domestic product (GDP) from the previous quarter is revised down by one percentage point, then followers of this variable will likely have to revise their forecasts.

Economic forecasters must therefore confront three forms of uncertainty related to time: uncertainty about the future, the present, and the past. **Chart 1** presents the consequence of these additional layers of uncertainty, using the path of some arbitrary variable x as an example. It is assumed that forecasters are required, at some point in time, t, to produce a forecast about the future path of x. In panel (a), the forecaster is assumed to know the current value of the variable with certainty (this would be the case, for example, of a financial-asset price or a commodity price). A forecast is produced for this variable, depicted by the dashed line. The uncertainty associated with this forecast is arbitrarily depicted by the dotted lines, which provide a confidence interval for the forecast. Typically, but not always, the farther into the future one wishes to forecast, the wider is the confidence interval.

* Thanks to Bob Fay, Sharon Kozicki, Robert Lafrance, John Murray, and Simon van Norden for several valuable comments. All views are those of the author and do not necessarily reflect those of the Bank of Canada.

Chart 1: Different forms of uncertainty as a function of time

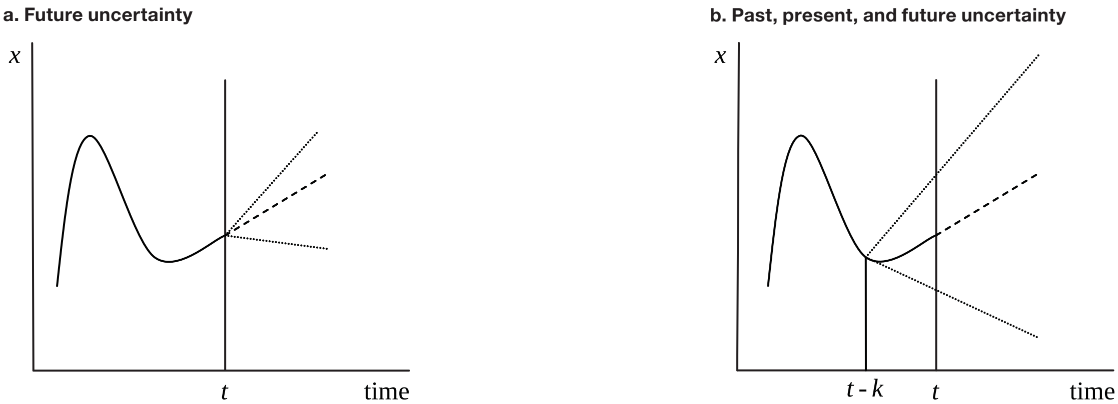

In panel (b), the forecaster is required to forecast a variable that cannot be directly observed at the time the forecast is to be made. This is depicted by the existence of a confidence interval at time t. There is also a possibility that the values for this variable that were observed in the recent past may be revised. Variables such as the National Accounts (GDP, consumption, business investment, etc.) and money and credit aggregates are all subject to revision. The period t-k depicts the time at which variables may no longer be subject to revision, so that observations prior to t-k can be assumed to be measured with certainty.[1] Thus, forecasters of these variables are subject to additional layers of uncertainty that forecasters of precisely measured variables do not confront, which, all other factors being the same, would result in wider confidence intervals around future forecasts.

This article expands on uncertainty about the future and the present, and more thoroughly analyzes uncertainty about the past and how economists have tried to confront it. The challenges posed by data revisions have long been acknowledged by economists but have not been closely scrutinized until recently, owing to lack of databases that incorporate current and past releases of economic variables. The article concludes with a discussion of how policymakers can account for uncertainty about the past in the conduct of monetary policy.

The challenges posed by data revisions have not been closely scrutinized until recently, owing to lack of databases that incorporate current and past releases of economic variables.

Uncertainty as a Function of Time

Uncertainty about the future

Most developments in the field of economic forecasting have tried to address the issue of uncertainty about the future. Recognizing that point forecasts by themselves are of limited value without any associated knowledge about the uncertainty surrounding them, economic forecasters have been trying to better quantify estimates of that uncertainty. In the past several years, methods have been developed to produce and evaluate density forecasts: that is, forecasts of the entire probability distribution of a variable of interest. Density forecasts allow forecast users to easily compute the probability that the variable of interest will lie within a certain range.

As an illustration, Li and Tkacz (2006) demonstrate how density forecasts can be produced for the Canadian inflation rate in the next period. Given that the Bank of Canada wishes to maintain inflation in the centre of a 1 to 3 per cent target band, computing the

1 However, Campbell and Murphy (2006) note that National Accounts can be revised several years after they were first released, with such long-term revisions largely reflecting changes to the methodology used to measure these variables. Revisions to the National Accounts in the near past typically reflect new information received by Statistics Canada, thereby yielding improved estimates of economic activity. See the **Appendix** for details regarding the revision schedule.

probability that inflation would deviate from the target band would be of value. **Table 1** presents the computed probabilities that the next period's inflation rate will lie within various ranges.

Density forecasts reveal that the inflation rate would be within the target band with 97 per cent confidence for the period under study, with the probability of being above the target band being slightly higher than the probability of being below it.

Table 1: Inflation probabilities forecast for different ranges

Inflation range	< 1%	1% to 2%	2% to 3%	> 3%	1% to 3%
Probability	0.007	0.487	0.485	0.021	0.971

Source: Li and Tkacz (2006), Table 4

Uncertainty about the present

Quantities such as GDP, and many other economic variables, are not directly observable and must therefore be estimated by Statistics Canada. The estimates are produced using various surveys and variables covering all sectors of the economy. Because of the time required to compile all this information, data for a given quarter will not be released until about two months after the end of the quarter. For example, data for the first quarter, ending on 31 March, will not be available until about 31 May, which is well into the second quarter. To produce GDP forecasts at any point during the second quarter, forecasters will, at best, have data only up to the previous quarter. This problem of "forecasting" the current value of an economic variable is commonly called "nowcasting." For the purpose of nowcasting, analysts rely on coincident indicators, that is, variables that are correlated with fluctuations in GDP growth but are available on a more timely basis (i.e., with a shorter reporting lag). Nunes (2005) is a recent example of a nowcasting study of GDP growth, but work on identifying coincident indicators of economic activity goes back to Burns and Mitchell (1946) who classified hundreds of economic variables as leading or coincident indicators.

Generally, analysts can monitor developments in variables where the publication lags are shorter, such as employment, housing starts, and manufacturing indexes, in order to gauge economic activity prior to the official release of data on GDP growth. Such monitoring can be used to provide advice to decision-makers prior to the release of National Accounts data.

Uncertainty about the past

This type of uncertainty relates to revisions that occur to variables following their first release. Economists have long recognized that variables get revised (e.g., Stekler 1967), but only in recent years have they made systematic efforts to better understand the revision process. This was mainly because historical data were not maintained. For example, when Statistics Canada releases the latest GDP number, it releases revisions to past GDP figures at the same time. In the process, the new GDP series replaces the old one in the database; so unless researchers systematically saved the old series, they could not analyze the revision process.

At some point, researchers decided to construct their own databases by physically scanning the old series as they originally appeared in the hard copies of statistical agency publications. In the United States, such efforts were spearheaded by the Federal Reserve Bank of Philadelphia[2] and the Federal Reserve Bank of St. Louis, which maintains an extensive real-time database for the United States (dubbed ALFRED, for ArchivaL Federal Reserve Economic Data). Other countries followed with similar databases, which are referred to as "real-time" databases, since they include the data as they were originally reported at each point in time.

Construction of a real-time database for Canada was recently initiated by Campbell and Murphy (2006), and the Organisation for Economic Co-operation and Development (OECD) maintains real-time data for member countries going back a few years. In addition, Keshishbanoosy et al. (2008) document the contents of a real-time database for Canadian money and credit aggregates.

With access to National Accounts data as they were initially published through time, economists can now begin quantifying the uncertainty surrounding the initial estimates of variables of interest, thereby producing confidence intervals around past data as depicted in Chart 1, panel (b). Some effort is also being made to understand whether the revision process can be predicted. If that is the case, it would reduce the uncertainty associated with data revisions.[3] For example, Galbraith and Tkacz (2007) find that debit

2 See the database developed by Croushore and Stark (2001).

3 Of course, some revisions, such as those related to conceptual changes, are necessarily unpredictable. Studies that deal with predicting revisions focus on near-term revisions related to the incorporation of additional information that improves estimates of key National Accounts variables. Proxy variables can potentially be useful for predicting this regular revision process.

card transactions can be useful for predicting revisions to GDP growth up to four quarters in the past.

Revisions to National Accounts

This section presents some updated descriptive statistics of the revision process for data in the Canadian National Accounts, thereby providing some estimates about uncertainty related to the past. The focus is specifically on the annualized quarterly growth rate of GDP or one of its components (consumption expenditures, business investment, government spending, exports, and imports).

The first release of the level of real GDP, or any of its components, is denoted by $x_{1,t}$ for time t, and $x_{2,t}$ denotes the second release of the level of real GDP (or any of its components) for time t.

The initial, or first-release, annualized quarterly growth rate is then calculated as

$$\dot{x}_{1,t} = \ln \frac{x_{1,t}}{x_{2,t-1}} \times 400, \qquad (1)$$

where ln denotes the natural logarithm. Note that the initial annualized quarterly growth rate of GDP for a given period is computed using the first release of the level of GDP for the current period and the second release of the level for the previous period. For example, the annualized growth rate of 0.3 per cent for the third quarter of 2009 is a function of the first estimate for the level of GDP in the third quarter (time t) and the second estimate of the level of GDP for the second quarter of 2009 (time t-1).

Following this logic, the second estimate of the annualized quarterly growth rate for period t is computed as

$$\dot{x}_{2,t} = \ln \frac{x_{2,t}}{x_{3,t-1}} \times 400, \qquad (2)$$

and so forth. If past data were not revised, then the initial and subsequent growth rates would not change; i.e., $\dot{x}_{2,t} - \dot{x}_{1,t} = 0$, and there would consequently be no uncertainty about the past. As new information becomes available, however, the statistical agency will revise its past estimates of GDP and its components, thereby affecting the estimated growth rates. This could be particularly important in instances when economic growth is stagnating and a recession is a possibility.

Revisions to GDP growth: An international comparison

To put revisions to Canadian data into context, revisions to Canadian GDP growth are compared with those reported by a small number of other OECD countries. To ensure that the data are as comparable across countries as possible, all our data are obtained from the OECD. The data were initially published in the OECD's *Main Economic Indicators*, and every issue from 2001 onwards was used to create a real-time database for OECD-member countries and for a select group of other countries.[4] Levels of real GDP are obtained for each country from the first quarter of 2001 to the third quarter of 2009.[5] Once the growth rates are computed as described by equations (1) and (2), the first and last observations are dropped, so that $\dot{x}_{2,t} - \dot{x}_{1,t}$ can be studied.

Release dates for National Accounts differ across countries, and this may influence the size of the reported revisions.

Although the data studied were compiled by a single organization, international comparisons are still complicated by the fact that release dates for National Accounts differ across countries, and this may influence the size of the reported revisions. For example, the first release for GDP growth in the third quarter of 2007 for Canada, Germany, and the United Kingdom appeared in the December 2007 issue of the *Main Economic Indicators*; however, it appeared in the November 2007 issue for the United States, and in the January 2008 issue for Australia. In other words, first-release estimates for GDP growth were available earlier for the United States and later for Australia. If statistical agencies are allowed more time to release their first estimates of National Accounts, they may be able to incorporate additional information and therefore require smaller revisions in the future.

4 The database, located at http://stats.oecd.org/mei, currently contains National Accounts data for 35 countries and the euro area.

5 Note that in 2001, Statistics Canada switched from a Laspeyeres to a Chain Fisher index in computing GDP in order to make Canadian figures more accurate and more directly comparable with those of the United States; see Statistics Canada (2002) for technical details. Conceptual changes may also have been implemented for some countries over the sample in our study, so our cross-country results should be viewed as suggestive rather than conclusive.

Chart 2: Revisions to GDP growth for selected OECD countries

Quarter-over-quarter growth at annualized rates

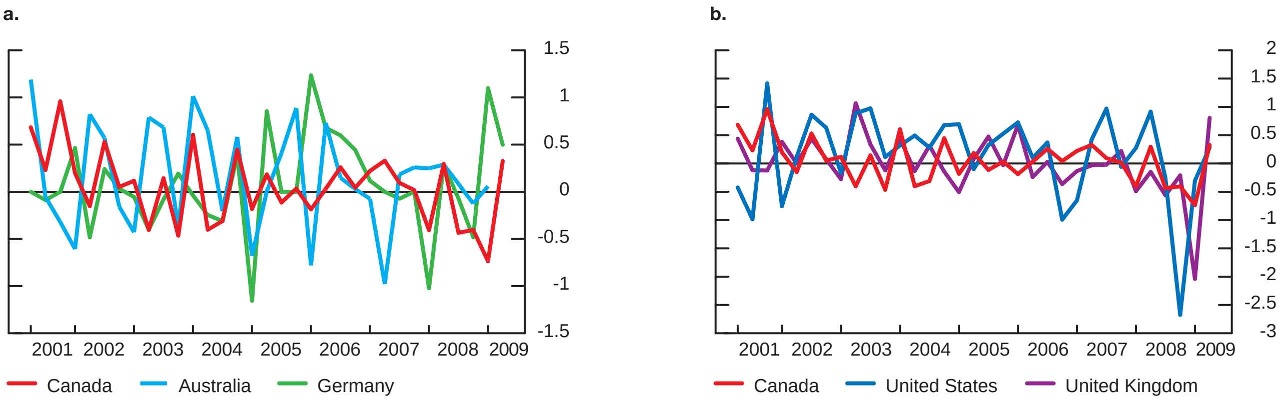

Source: Author's calculations using data from OECD *Main Economic Indicators*

The second release is defined as the revision that accompanies the release of the subsequent quarter's National Accounts data: i.e., the estimate published by the OECD three months later. Thus, the second release of GDP growth for the third quarter of 2007 for Canada, Germany, and the United Kingdom appeared in the March 2008 issue of the *Main Economic Indicators*; in the February 2008 issue, for the United States; and in the April 2008 issue, for Australia.

We now examine the difference between the second and first releases to GDP growth ($\dot{x}_{2,t} - \dot{x}_{1,t}$). The larger this number is in absolute terms, the greater the revisions, and therefore the greater the uncertainty surrounding the initial estimate.

Charts 2a and **2b** present revisions to the growth rate of GDP for Canada, Australia, Germany, the United Kingdom, and the United States.[6] Several features emerge:

- The revisions are not necessarily correlated across countries over time. GDP is computed by each country's statistical agency, and although the definition of GDP and data-collection techniques are very similar across these countries, there are only a few instances where revisions are of the same magnitude or, indeed, in the same direction across a group of countries. This could reflect differences in the business cycle, the differing time constraints imposed on the statistical agency to produce a first estimate of GDP for a given quarter, the resources of the statistical agency, etc.
- Revisions to Canadian GDP are somewhat smaller and less volatile than those of other countries. Although Canadian GDP growth is sometimes revised by more than 0.5 percentage points, this is not unusual for the countries in our sample.

To more accurately depict the revision process in these countries, **Table 2** presents some simple descriptive statistics. The second column is an estimate of the mean revision to GDP growth over the sample period, which provides a measure of bias in the revision process. A mean close to 0.0 indicates that upward and downward revisions tend to offset each other, so the initial growth rate release is said to be *unbiased*. If the mean is above zero, this indicates that GDP growth tends to be, on average, revised upwards in the following quarter; below zero, it would indicate that the growth rate tends to be revised downwards.

Table 2: Revisions to GDP growth for selected OECD countries

Sample: 2001Q2 to 2009Q2

Country	Mean revision	Mean absolute revision	Confidence interval	Largest absolute revision
Canada	0.05	0.30	(-0.70, 0.80)	0.96
Australia	0.15	0.45	(-0.94, 1.24)	1.19
Germany	0.08	0.35	(-0.94, 1.10)	1.23
United Kingdom	0.00	0.35	(-1.05, 1.06)	2.04
United States	0.15	0.60	(-1.39, 1.70)	2.67

Note: Revisions are defined as the differences between the second and first release of the annualized quarterly GDP growth rate for each country. The confidence interval is a simple estimate within which we expect the GDP growth-rate revision to lie 19 quarters out of 20. Data obtained from the OECD *Main Economic Indicators Original Release and Revisions* database.

6 McKenzie (2007) analyzes revisions across a broader set of OECD countries, using different metrics, and over the 1995 to 2007 period.

The mean revision to Canadian GDP growth is 0.05 percentage points, which is trivial. Such a revision is consistent with the revisions of other countries and statistically is not significantly different from 0.0. The largest average revision is for Australia and the United States at 0.15 percentage points. Because our sample is relatively short, however, the associated estimated standard errors are sufficiently large that the average revision for each country is not statistically different from 0.0.

The third column presents the mean absolute revision, which is the average of the absolute value of the revisions. This statistic allows us to gauge the average magnitude of the revisions, regardless of whether the revision is positive or negative. A higher value here indicates that revisions to the GDP growth rate tend to be more substantial; a value of zero would indicate that the initial growth rate is not revised.

We find that the mean absolute revision for Canada is 0.3 percentage points, which is smaller, but not significantly different, than the numbers for other countries in the sample. The United States tends to have the largest revisions, but as mentioned above, this may reflect the fact that its data are released one month sooner than those of the other countries in our sample.

Large revisions to the GDP growth of foreign countries are not simply an issue for policy-makers abroad, but can have implications for policy decisions in Canada. Given that trade is an important component of the Canadian economy, Canadian policy-makers are interested in monitoring economic conditions abroad in order to gauge the potential demand for Canadian goods. If figures for foreign GDP growth are substantially revised, this can potentially complicate policy decisions in Canada. Being aware of revisions to foreign data is therefore important from a Canadian perspective.

Apart from the average size of revisions, analysts and policy-makers are also interested in the volatility of revisions, since less-pronounced revisions lead to less uncertainty about the initial estimate of GDP growth. Using estimates of the standard deviations of the revisions, the fourth column of Table 2 presents confidence intervals for revisions, which correspond to an estimate of the uncertainty around the past growth rate for $k=1$ in Chart 1, panel (b). For Canada, we estimate that the revision to GDP growth will lie in a range of -0.7 to 0.8 percentage points, 19 quarters out of 20. This is narrower than the ranges computed for the other countries. Although not necessarily statistically lower than other countries, it does suggest that Canadian decision-makers may have somewhat more confidence in initial releases of GDP growth than their counterparts in other countries.

Finally, the last column lists the largest (in terms of absolute values) revision for each country. The largest revision for Canada, of almost one percentage point, was recorded for the fourth quarter of 2001, which can be observed in Chart 2a. The United States also had a large positive revision in this quarter (+1.41 percentage points). Since growth for the fourth quarter of 2001 is computed as the percentage change in real GDP from the third quarter to the fourth quarter, we can surmise that the events of 11 September 2001 (which occurred in the third quarter) likely made the task of estimating economic activity especially difficult in both countries.

Revisions to growth rates of Canadian GDP components

Although headline GDP numbers are very important to monetary policy, policy-makers are also interested in the underlying factors that contribute to GDP growth, since some of these components are necessarily more sensitive to interest rate movements and therefore react more strongly to monetary policy actions.

The major components of expenditure-based GDP are

- total household expenditures on goods and services (C);
- business fixed investment (I);
- expenditures by all levels of government (G);
- total exports (X); and
- total imports (IM).

In practice, Statistics Canada obtains growth estimates for each component (which can be further disaggregated) and then aggregates them to obtain an estimate of GDP growth. **Chart 3** presents the same revision series for Canadian GDP growth shown in Chart 2, together with the revisions to the growth rates of each major GDP component. Some observations:

- Revisions to the growth rates of GDP components are more pronounced than the revisions to GDP growth itself. Note that the vertical scale of Chart 3 is wider than that of Chart 2, so revisions to total GDP growth seem almost benign in Chart 3 relative to Chart 2. In contrast, the growth rates of some of the components in Chart 3 are often revised by more than 2 percentage points.

Chart 3: Revisions to estimates of Canadian GDP growth and its components
Quarter-over-quarter growth at annualized rates

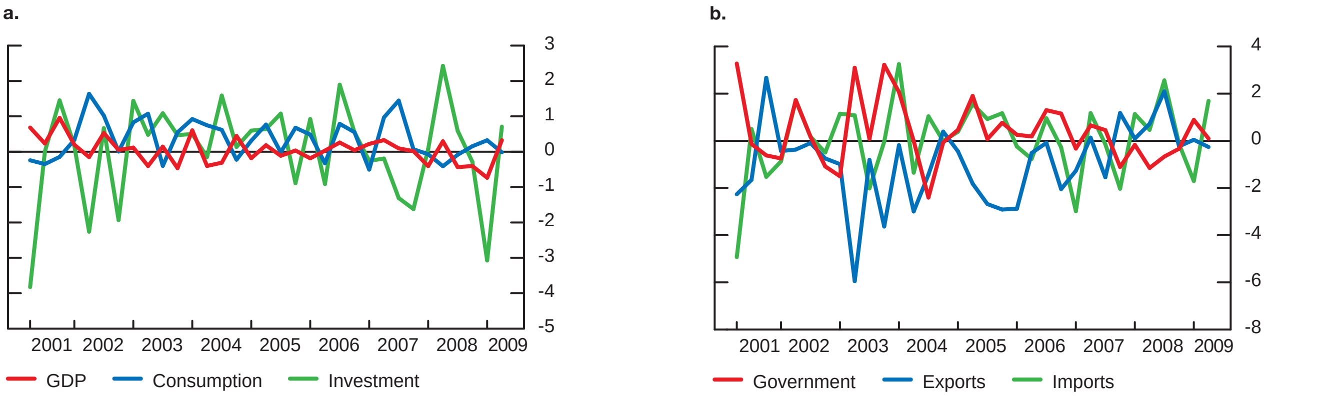

Source: Author's calculations using data from OECD *Main Economic Indicators*

- Revisions to export and import growth are the most pronounced. In particular, export growth was subject to several downward revisions between 2003 and 2006. Revisions to consumption growth appear to be the smallest.
- In the fourth quarter of 2001, export growth was revised upwards by more than 2 per cent, while import growth was revised down by almost 2 per cent. Taken together, net exports ($X - IM$) were raised substantially in this quarter, and combined with an upward revision to business investment this can explain the relatively large upward revision to GDP growth observed for the quarter.

Table 3 presents descriptive statistics for the revisions to GDP growth and its components. The mean revisions deviate more substantially from zero

Table 3: Revisions to Canadian GDP growth and its components
Sample: 2001Q2 to 2009Q2

Series	Mean revision	Mean absolute revision	Confidence interval	Largest absolute revision
GDP	0.05	0.30	(-0.70, 0.80)	0.96
Consumption	0.35	0.51	(-0.78, 1.48)	1.64
Investment	0.02	1.04	(-2.75, 2.80)	-3.83
Government expenditures	0.35	0.98	(-2.30, 3.00)	3.28
Exports	-0.94	1.38	(-4.31, 2.43)	-5.95
Imports	0.04	1.23	(-3.23, 3.31)	-4.93

Note: Revisions are defined as the differences between the second and first release of the annualized quarterly growth rate of GDP and its components. The confidence interval is a simple estimate within which we expect the growth-rate revision for each series to lie 19 quarters out of 20. Data obtained from the OECD's *Main Economic Indicators Original Release and Revisions* database.

for some components. For example, growth of consumption and government expenditure tend to be revised upwards by more than 0.35 percentage points, while export growth is revised downwards by 0.9 percentage points, on average. In terms of mean absolute revisions, consumption growth is revised by over 0.5 percentage points, on average, growth of investment and government expenditures by about 1 percentage point, and export growth by almost 1.4 percentage points. As a result, analysts who are required to monitor and forecast the growth of Canadian trade face a more daunting task than those who focus on other GDP components.

The associated confidence intervals for revisions to the growth rates of the components of GDP are wider, sometimes substantially so, than for GDP growth itself. Among the components, consumption growth is revised between -0.8 and +1.5 percentage points 19 quarters out of 20; growth in investment and government expenditure have roughly similar ranges (-2.7 to 2.8 and -2.3 to 3.0 percentage points, respectively); while export and import growth have the most uncertainty associated with their initial estimates (ranges of -4.3 to 2.4, and -3.2 to 3.3 percentage points, respectively).

Based on the data from 2001 to 2009, we can conclude that greater uncertainty is associated with the first-release growth rates of the components of GDP than with the growth rate of GDP itself. This result can arise because revisions to the components offset each other (for example, higher consumption growth can offset lower business investment), thereby muting the impact of the revisions to the estimate for total GDP, and also because the GDP components are necessarily lower in level than total GDP, so revisions to the levels of the components will result in correspondingly larger changes to the growth rates of the components.

Data Revisions and Monetary Policy: Some Issues and Future Directions

The literature on the consequences of data revisions for economic analysis and forecasting has been growing in the past few years, driven by the availability of real-time databases. With the availability of the OECD's real-time database and the Bank of Canada's real-time money and credit database, researchers and analysts now have access to vintage data that allow them to study some important issues for Canadian policy-makers. Below are some areas for which researchers have recently used real-time data to further our understanding.

The output gap

One of the initial motivations for exploring the impact of data revisions was the analysis of past policy decisions; for example, Runkle (1998) and Croushore and Stark (2003). To conduct a fair assessment, however, one would have to use data that were available to policy-makers at the time decisions were being made. As demonstrated above, GDP growth rates are revised by an average of more than one-half a per cent in some countries, and this after only one quarter. To analyze the policy actions of, say, five years ago, one would have to study the data that policy-makers had at that time.

One of the initial motivations for exploring the impact of data revisions was the analysis of past policy decisions.

A key variable monitored by policy-makers when making policy decisions is the output gap: the difference between the current level of real GDP and the level that would exist if all resources in the economy were fully employed and the inflation rate had no tendency to deviate from the target. The press releases that accompany interest rate announcements by the Bank of Canada on fixed announcement dates often allude to the output gap in statements such as, "Overall, the Canadian economy remained above its production capacity at year-end,"[7] which signals a positive output gap.

Given its importance for policy decisions, researchers have documented the impact of data revisions on the measurement of the output gap. Orphanides (2001) found that, once data revisions are taken into account, estimates of the output gap in the United States may differ by more than two percentage points, a magnitude that is non-trivial from a policy perspective. Kozicki (2004) provides measures of the policy implications of such revisions.

Proper measurement of the output gap requires not only the current level of real GDP, but also an estimate of potential GDP. There are several techniques for estimating the latter,[8] but they tend to provide relatively poor estimates of the output gap in real time.[9]

The output gap also figures prominently in the literature on policy rules, where the policy rate is specified as a simple linear function of the output gap and the deviation of the inflation rate from some target. Taylor (1993) found that policy rates in the United States could be well explained by such a rule in the 1980s; however, if data revisions were taken into account, and policy rules were estimated with the data available to policy-makers at the time decisions were being made, such conclusions might not hold. Côté et al. (2004) is the most comprehensive assessment of policy rules for Canada, and it remains to be determined how their most robust rules would change, given the issue of data revisions.

Finally, the output gap is often used to predict inflation. If it is subject to measurement error, it would be useful to determine how this affects inflation forecasts. Orphanides and van Norden (2005) find that output gaps predict inflation relatively well in the United States, but that the forecast performance diminishes substantially if real-time estimates of the output gap are used instead. An extension of this study using recent Canadian data would be very useful.

The role of money

The various measures of growth in the money supply are not given as much importance in making policy decisions today as was the case 20 years ago. This is because the empirical link between growth in the money supply and future inflation appears to have

7 "Bank of Canada lowers overnight target by 1/2 percentage point to 3 1/2 per cent," Bank of Canada Press Release, 4 March 2008.
8 See St-Amant and van Norden (1997) for a survey.
9 See Orphanides and van Norden (2002) for U.S. evidence and Cayen and van Norden (2005) for Canada.

weakened, partly as a result of innovations in banking products. However, as Keshishbanoosy et al. (2008) show, the money supply itself is subject to revision. It may therefore be worthwhile to further explore the links between money growth, macroeconomic variables, and policy decisions in a real-time context. Garratt et al. (2007) find some evidence that the predictive power of broad money in the United Kingdom did not decrease as much in the 1980s as is popularly perceived when real-time data is used.

Monitoring

From equation (1) we observe that the current growth rate of GDP is a product of this period's GDP level and the revised level of last period's GDP. In other words, success in monitoring this period's growth rate hinges partly on the magnitude of the revision to last period's GDP. Analysts should therefore expend some effort in trying to understand the nature of revisions and, indeed, try to predict them, if possible.

The literature remains divided as to whether past revisions are, in fact, predictable, but preliminary evidence presented by Galbraith and Tkacz (2007) for Canadian data suggests that revisions can be partially anticipated. In future work, analysts can explore other explanatory variables, as well as understanding whether revisions are likely to be more pronounced in some periods than in others. For example, revisions may be larger around the turning points of business cycles, so in such periods of uncertainty analysts may wish to anticipate large revisions and therefore build larger confidence intervals around their estimates of current GDP growth.

Given the large revisions to the components of GDP growth, the payoff for predicting the revisions could be potentially smaller confidence intervals around the monitoring of these variables.

Statistical methodology

Some statistical methods used by economists may not be valid in the presence of data that are subject to revision, and so some empirical findings may have to be reconsidered. New techniques are currently being developed to deal with such issues, but it may be some time before analysts can fully exploit them. For example, Clark and McCracken (2009) propose a new forecast-encompassing test (a test used to select among competing forecasting models) that can be used in the presence of revised data. Up to this point, many tests applied in the field have ignored data revisions, so it is possible that some incorrect inferences may have been made in past studies that have assessed the relative performance of competing forecasting models. In a different context, Jacobs and van Norden (2006) develop a method for constructing optimal forecasts and confidence intervals that are valid in the presence of data revision and use multiple vintages of data.

Development of such new techniques and their application to Canadian data are also important areas for future research.

Conclusion

Data revisions have been recognized as an issue by economists for some time, but research on the impact of data revisions has grown markedly in recent years with the advent of real-time databases. Canadian real-time databases are now available, and Canadian practitioners are expected to use these resources to improve the reliability of their models.

Data revisions can be viewed as uncertainty about the past, which feeds into uncertainty about the future. Revisions to Canadian GDP growth are found to be somewhat lower than those in some other OECD countries. However, revisions to the growth rates of the components of Canadian GDP are appreciably larger, which can lead to greater uncertainty for analysts who must monitor and forecast those components.

Data revisions can affect policy decisions in several ways, notably by yielding more uncertainty around the true values of the variables of interest to policy-makers. Furthermore, they can affect the existence of fundamental relationships between variables and cloud the judgment of analysts. Many outstanding research questions remain to be resolved for policy-makers, but the existence of real-time databases for Canada should help to answer some of these questions in coming years.

Literature Cited

Burns, A. and W. Mitchell. 1946. *Measuring Business Cycles*. New York: National Bureau of Economic Research.

Campbell, B. and S. Murphy. 2006. "The Behaviour of Canadian Data Revisions." Paper presented at the 2nd CIRANO and IPA Workshop on Macroeconomic Forecasting, Analysis and Policy with Data Revision. Montréal, Canada, October 27–28.

Cayen, J.-P. and S. van Norden. 2005. "The Reliability of Canadian Output-Gap Estimates." *North American Journal of Economics and Finance* 16 (3): 373–93.

Clark, T. and M. McCracken. 2009. "Tests of Equal Predictive Ability with Real-Time Data." *Journal of Business & Economic Statistics* 27 (4): 441–54.

Côté, D., J. Kuszczak, J.-P. Lam, Y. Liu, and P. St-Amant. 2004. "The Performance and Robustness of Simple Monetary Policy Rules in Models of the Canadian Economy." *Canadian Journal of Economics* 37 (4): 978–98.

Croushore, D. and T. Stark. 2001. "A Real-Time Data Set for Macroeconomists." *Journal of Econometrics* 105 (1): 111–30.

———. 2003. "A Real-Time Data Set for Macroeconomists: Does the Data Vintage Matter?" *Review of Economics and Statistics* 85 (3): 605–17.

Galbraith, J. W. and G. Tkacz. 2007. "Electronic Transactions as High-Frequency Indicators of Economic Activity." Bank of Canada Working Paper No. 2007-58.

Garratt, A., G. Koop, E. Mise, and S. Vahey. 2007. "Real-Time Prediction with UK Monetary Aggregates in the Presence of Model Uncertainty." Birkbeck Working Paper in Economics and Finance No.0714.

Jacobs, J. and S. van Norden. 2006. "Modeling Data Revisions: Measurement Error and Dynamics of 'True' Values." CCSO Working Paper No. 2006/07.

Keshishbanoosy, R., P. St-Amant, D. Ball, and I. Medovikov. 2008. "A Money and Credit Real-Time Database for Canada." *Bank of Canada Review* (Summer): 55–64.

Kozicki, S. 2004. "How Do Data Revisions Affect the Evaluation and Conduct of Monetary Policy?" *Federal Reserve Bank of Kansas City Economic Review* Q1 89 (1): 5–38.

Li, F. and G. Tkacz. 2006. "A Consistent Bootstrap Test for Conditional Density Functions with Time-Series Data." *Journal of Econometrics* 133 (2): 863–86.

McKenzie, R. 2007. "Relative Size and Predictability of Revisions to GDP, Industrial Production and Retail Trade—A Comparative Analysis across OECD Member Countries." Paper presented at the 3rd CIRANO and IPA Workshop on Macroeconomic Forecasting, Analysis and Policy with Data Revision. Montréal, Canada, October 5–6.

Nunes, L. 2005. "Nowcasting Quarterly GDP Growth in a Monthly Coincident Indicator Model." *Journal of Forecasting* 24 (8): 575–92.

Orphanides, A. 2001. "Monetary Policy Rules Based on Real-Time Data." *American Economic Review* 91 (4): 964–85.

Orphanides, A. and S. van Norden. 2002. "The Unreliability of Output-Gap Estimates in Real Time." *Review of Economics and Statistics* 84 (4): 569–83.

———. 2005. "The Reliability of Inflation Forecasts Based on Output-Gap Estimates in Real Time." *Journal of Money, Credit and Banking* 37 (3): 583–601.

Runkle, D. 1998. "Revisionist History: How Data Revisions Distort Economic Policy Research." *Federal Reserve Bank of Minneapolis Quarterly Review* 22(4): 3–12.

St-Amant, P. and S. van Norden. 1997. "Measurement of the Output Gap: A Discussion of Recent Research at the Bank of Canada." Bank of Canada Technical Report No. 79.

Literature Cited (cont'd)

Statistics Canada. 2002. "Chain Fisher Volume Index —Methodology." Available at <http://www.statcan.gc.ca/concepts/cf/8102792-eng.htm>.

Stekler, H. 1967. "Data Revisions and Economic Forecasting." *Journal of the American Statistical Association* 62 (318): 470–83.

Taylor, J. 1993. "Discretion Versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy* 39: 195–214.

Appendix: Revision Policy for the Canadian National Accounts

Revisions to the National Accounts can be of three different types: (i) regular revisions that take into account new information and/or reconcile data from the myriad surveys and sources that are used as inputs to the construction of the National Accounts; (ii) conceptual revisions arising from changes in definitions within the National Accounts; and (iii) historical revisions that are (infrequently) performed for various reasons.

Regular revisions

In the absence of conceptual changes or major historical revisions, the National Accounts are regularly revised to take into account new information and/or to reconcile data from the sources used to construct them. The revision schedule is as follows:

- Data for preceding quarters of the year are revised when the data for the current quarter are published.
- Revisions extending back four years are made with the publication of first-quarter data for a new year.

These are the revisions that should be of most interest to analysts and policy-makers since they can influence one's perception about the relative strength or weakness of the economy and can therefore influence current decisions.

Source: Statistics Canada <http://www.statcan.ca/english/freepub/13-010-XIE/2003001/revision2003001.htm>

Conceptual revisions

These revisions can arise because of changing perceptions about how certain segments of the economy should be classified or because of fundamental changes in quantifying economic activity. For example, major conceptual changes occurred with the release of the May 2001 National Accounts in which the method for measuring the capitalization of software was changed and the move from a Laspeyres Index to a Chain Fisher Volume Index took place, which made the Canadian and U.S. National Accounts more comparable. This second factor, in particular, renders the comparison of revisions before and after 2001 more challenging. See Statistics Canada (2002) for technical details.

Historical revisions

About once every ten years, Statistics Canada will revise data farther back than the typical four years. Such historical revisions are conducted to "improve estimation methods, eliminate statistical breaks resulting from more limited revisions and introduce conceptual changes into the system." Such revisions would have the greatest impact on users of macroeconomic models, who might find that parameter estimates were affected by such revisions. The latest historical revision occurred in December 1997, and the next is scheduled for 2012/2013. For further details see <http://www.statcan.ca/english/concepts/nateco/ann.htm>.

Bank of Canada Publications

Unless noted otherwise, all publications are available in print and on the Bank's website: <http://www.bankofcanada.ca>.

Monetary Policy Report (quarterly: January, April, July, and October)

Financial System Review (June/December)

Bank of Canada Review (quarterly; see inside cover for subscription information)

Business Outlook Survey (quarterly: January, April, July, and October)*

Senior Loan Officer Survey (quarterly: January, April, July, and October)*

Speeches and Statements by the Governor

Bank of Canada Banking and Financial Statistics (monthly)*

Weekly Financial Statistics (published each Friday)*

Renewal of the Inflation-Control Target: Background Information

Annual Report

A History of the Canadian Dollar
James Powell (available at Can$8 plus GST and PST, where applicable) (2005)

By All Accounts: Outside Perspectives on the Bank of Canada**
A portrait of the Bank from the perspective of outside observers, showing how Canadians have perceived the performance of their central bank over the decades through the eyes of those who monitor its work on the public's behalf. (2010)

Beads to Bytes: Canada's National Currency Collection**
This volume explores the role of money in society through the lens of the National Currency Collection, an extraordinary repository of coins, bank notes, and related paraphernalia from around the world. (2008)

More Than Money: Architecture and Art at the Bank of Canada**
A tour of the head office complex, highlighting the architecture, interior design, and decoration, as well as elements of restoration and preservation. It also features pieces from the Bank's art collection. (2007)

The Art and Design of Canadian Bank Notes**
A journey behind the scenes to explore the demanding world of bank note design. (2006)

The Bank of Canada: An Illustrated History**
To celebrate the Bank's 70th anniversary, this book depicts the history of the Bank from 1935. (2005)

The Transmission of Monetary Policy in Canada
(1996, Can$20 plus GST and PST, where applicable) Available at <http://www.bankofcanada.ca/en/res/other/herm-98.html>.

The Thiessen Lectures (January 2001)
Lectures delivered by Gordon G. Thiessen, Governor of the Bank of Canada 1994 to 2001

A Festschrift in Honour of David Dodge's Contributions to Canadian Public Policy (2009)

Bilingualism at the Bank of Canada (published annually)

Planning an Evolution: The Story of the Canadian Payments Association, 1980–2002
James F. Dingle (June 2003)

Conference Proceedings
Conference volumes published up to and including April 2005 are available on the Bank's website. Print copies can be purchased for Can$15 plus GST and PST, where applicable. Papers and proceedings from Bank of Canada conferences, seminars, and workshops held after April 2005 are now published exclusively on the Bank's website.

Technical Reports, Working Papers, and Discussion Papers*
Technical Reports, Working Papers, and Discussion Papers are published in the original language only, with an abstract in both official languages.

Discussion papers deal with finished work on technical issues related to the functions and policy-making of the Bank. They are of interest to specialists and other central bankers.

For further information, including subscription prices, contact:

Publications Distribution
Communications Department
Bank of Canada
234 Wellington St.
Ottawa, ON
Canada K1A 0G9
Telephone: 613 782-8248
Toll free in North America: 1 877 782-8248
Email address: publications@bankofcanada.ca

* Available only on the Bank's website.

** Each Can$25 plus shipping costs. Sample pages are available on the Bank's website.