BANK OF CANADA
BANQUE DU CANADA

BANK OF CANADA REVIEW

Spring 2014

2014

Articles

Canada, Edward VII, 1908C, gold sovereign, minted at the Ottawa branch of the Royal Mint

The British sovereign has been the standard gold coin of the United Kingdom and its colonies since 1817. Because of the high demand for the coin worldwide in the late nineteenth century, several countries, including Australia, South Africa and India, began to mint sovereigns on behalf of the Crown. At the time, there were no Canadian facilities for refining and coining gold of domestic origin, and gold producers had to ship the unrefined metal to the United States at an enormous cost. This issue was first raised in Parliament in 1890, but legislation for a home-based mint stalled. Finally, in 1901, the Ottawa Mint Act was passed. In 1908, the Ottawa branch of the Royal Mint opened and the first sovereigns minted in Canada rolled off the presses. Only 636 coins were struck, making this sovereign one of the rarest ever produced.

Members of the Editorial Board

The *Bank of Canada Review* is published twice a year. Articles undergo a thorough review process. The views expressed in the articles are those of the authors and do not necessarily reflect the views of the Bank.

For further information, contact:

Public Information
Communications Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9

Telephone: **613 782-8111**; **1 800 303-1282** (toll free in North America)
Email: **info@bankofcanada.ca**
Website: **bankofcanada.ca**

ISSN 1483-8303

The Canadian Dollar as a Reserve Currency

Lukasz Pomorski, Francisco Rivadeneyra and Eric Wolfe, Funds Management and Banking Department

- Over the past five years, central banks and monetary authorities have started adding Canadian-dollar assets to their official foreign reserves portfolios.
- According to survey data from the International Monetary Fund, the Canadian dollar accounted for about 1.8 per cent of reported global foreign reserves in the third quarter of 2013.
- Estimates of the total reserve holdings of Canadian-dollar-denominated securities are around US$200 billion.
- Higher levels of official foreign holdings may lower yields in domestic debt markets and therefore reduce the financing costs for the Government of Canada, but they may also decrease market liquidity.

A clear reflection of Canada's relative economic resilience during the global financial crisis of 2007–09 is the growth in the share of foreign exchange reserves that other countries hold in Canadian-dollar securities, particularly those issued by the Government of Canada.

Foreign exchange reserves are assets held by a central bank (or, less frequently, by finance ministries or monetary authorities) as a precaution for contingencies[1] that would call for foreign exchange market intervention[2] or provision of foreign currency liquidity to domestic financial institutions if access to capital markets were temporarily lost or delayed.

The world's official reserves have quadrupled over the past decade, exceeding US$11 trillion as of December 2013 (**Chart 1**).[3] Canada's foreign exchange reserves were increased in connection with the prudential liquidity plan established in 2011, which expanded the federal government's holdings of liquid assets that could be deployed if necessary.[4] Reserves are typically invested in highly rated government bonds and bills, and may also include

1 For a review of Canada's foreign exchange reserves held in the Exchange Fund Account, see Rivadeneyra et al. (2013).

2 Goldberg, Hull and Stein (2013).

3 This accumulation of reserves outpaced traditional measures of adequacy such as nominal GDP, coverage of short-term debt or broad money aggregates. See IMF (2011) for a detailed discussion.

4 A description of the Government of Canada's prudential liquidity plan is available at http://www.budget.gc.ca/2011/plan/anx2-eng.html.

Chart 1: Total official reserves of advanced and emerging economies, in US$ billions

Quarterly data

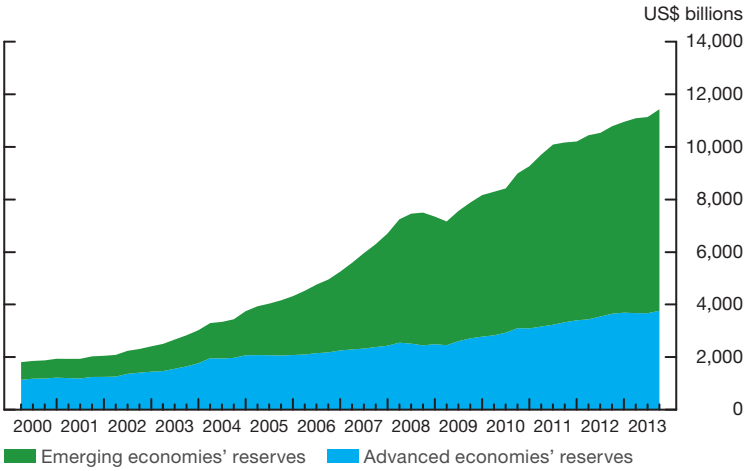

Source: International Monetary Fund COFER

Last observation: 30 September 2013

other securities issued by government agencies or sub-sovereign levels of government. Some reserves managers have also diversified their holdings to include equities.

An important change in the world's official reserves has been the increased diversification in terms of currency composition. Until recently, global reserves were almost entirely invested in five traditional currencies: the U.S. dollar, the euro, the Japanese yen, the British pound and the Swiss franc. According to data from the Currency Composition of Official Foreign Exchange Reserves (COFER) of the International Monetary Fund (IMF), until 2007 the reported allocations to all "other currencies" rarely exceeded 2 per cent of total reserves. By the end of 2013, however, the other currency allocation had more than tripled, to 6.3 per cent, of which the Canadian dollar represents about 1.8 percentage points. This large change in the currency allocation of foreign reserves portfolios included substantial portfolio investment inflows into Canadian fixed-income securities.

An important change in the world's official reserves has been the increased diversification in terms of currency composition, including substantial portfolio investment inflows into Canadian fixed-income securities

The Bank of Canada monitors these developments in reserves management for a number of reasons. In its role as the Government of Canada's fiscal agent, the Bank works with the Department of Finance Canada to provide stable, low-cost funding for the government by ensuring well-functioning markets in government securities.[5] Most foreign reserves investments denominated in Canadian dollars are in government bonds. The Bank monitors the impact of this activity on this market. It also assesses the effect of developments in reserves management on financial stability, particularly since the government debt market and the associated market for repurchase and reverse-repurchase agreements are core funding markets.[6] As part of its continual assessment of possible risks to the function of core funding markets and their stability, the Bank keeps track of the potential

5 The Bank provides policy advice to the federal government on the federal debt distribution framework (outlined in the Debt Management Strategy, available at http://www.budget.gc.ca/2014/docs/plan/anx1-eng.html), and it conducts regular auctions of Government of Canada securities. These securities are then transacted in the secondary markets by foreign reserves managers in other countries. Without well-functioning secondary markets for government debt, it would be difficult to achieve the goal of stable, low-cost funding.

6 Core funding markets are systemically important markets that are necessary for generating liquidity within the financial system (Fontaine, Selody and Wilkins 2009).

impact that the activities of foreign reserves managers may have on these markets. Finally, like many other central banks, the Bank of Canada provides safekeeping services (for example, custody, record keeping and settlement of transactions) to official foreign reserves managers that hold Canadian securities. The growth of reserves held in Canadian dollars has led to increased activity in this function.

This article provides an overview of recent trends in foreign exchange reserves and explores their potential implications for Canadian financial markets. The first part discusses the flows of reserves and resulting holdings in detail, while the second part reviews existing evidence of the potential influence of foreign flows on market functioning.

Demand for Canadian-Dollar-Denominated Assets

Estimating official foreign demand

To estimate the total demand that foreign reserves managers have for Canadian-dollar-denominated assets, we use the IMF's COFER data.[7] As noted, the COFER data had been divided into five traditional reserve currencies and included a catch-all category for the remaining "other currencies." The growing importance of the Canadian dollar and Australian dollar, however, led the IMF to begin reporting on these two currencies separately, recognizing them as de facto reserve currencies. This new breakdown in COFER data was first published in June 2013, presenting Canadian-dollar and Australian-dollar reserve holdings for 2012Q4 and 2013Q1.

According to COFER, the official reported holdings of Canadian-dollar-denominated assets stood at US$112.5 billion (Can$115.9 billion) as of 2013Q3, a share of about 1.8 per cent of reported foreign reserves. The holdings in the Australian dollar were quite similar, at 1.7 per cent. **Box 1** discusses the reserve assets invested in the Australian dollar in more detail.

These figures, while substantial, likely represent only a fraction of total holdings of Canadian-dollar-denominated assets in official foreign reserves, since not all of the polled foreign reserves managers responded to the COFER survey. These figures are based on responses of managers representing about US$6 trillion, or 54 per cent, of the total of US$11 trillion in official foreign reserves worldwide. The currency composition—including the Canadian-dollar holdings—of the remaining 46 per cent is not known.

If the proportion invested in Canadian-dollar-denominated assets for non-respondents is similar to that of responding managers, the total Canadian-dollar-denominated holdings would be US$208 billion (Can$214 billion). For robustness, we use a variety of other approaches to estimate the total holdings, with estimates ranging from US$172 billion (Can$178 billion) to US$219 billion (Can$226 billion), with Can$200 billion roughly at the midpoint of the range.

Chart 2 shows the estimated total foreign official holdings of Can$214 billion as a fraction of total non-resident holdings of Canadian general government securities reported by Statistics Canada. We use the Canadian-dollar equivalent to compare with the Statistics Canada figures of total non-resident holdings. Holdings of foreign reserves managers account for an estimated 43 per cent of the total non-resident holdings. Other official foreign investors, such as sovereign wealth funds, could hold some of the remaining 57 per cent of Canadian-dollar-denominated assets; however, official data that can reliably describe these holdings are limited.

7 COFER is based on a survey of 144 foreign reserves managers, typically national central banks, on the currency breakdown of the assets they hold.

Chart 2: Breakdown of non-resident holdings of Canadian general government securities, as of 2013Q3

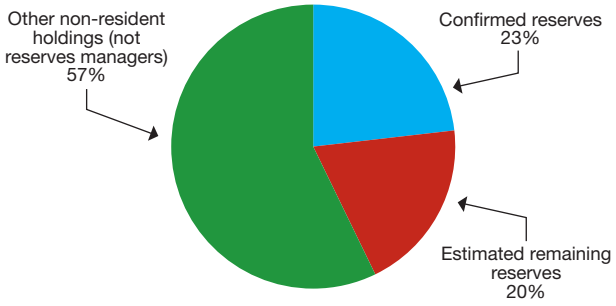

Note: Canadian general government securities include the debt of federal and local governments as well as government enterprises. Other non-resident investors include private institutions and official institutions other than foreign exchange reserves managers, e.g., sovereign wealth funds.

Sources: Statistics Canada and International Monetary Fund COFER Last observation: 30 September 2013

Box 1

Reserve Assets in the Australian Dollar

The growth in reserve assets invested in Canadian markets occurred at the same time as demand increased for Australian securities. As of 2013Q3, the share of the world's official reserves allocated to the Australian currency was 1.7 per cent, according to data from the Currency Composition of Official Foreign Exchange Reserves (COFER) of the International Monetary Fund (IMF), slightly below the share of the reserves invested in the Canadian dollar (1.8 per cent). The analysis presented in this article indicates that Australian markets started receiving reserves inflows at about the same time as Canadian markets. Media reports and our discussions with reserves managers suggest that the drivers of their demand are similar for both Australia and Canada: the desire to diversify reserve holdings to economies with the highest credit rating, the perceived safety of the two markets and the opportunity to earn somewhat higher yields than found in traditional reserve currencies such as the U.S. dollar.

There are also interesting differences between foreign reserves invested in Australian and Canadian securities. For example, countries that have stronger trade links with Canada have a relatively higher weight in the Canadian currency. Bank of Canada data indicate that, across European reserves managers, the average Canadian dollar weight is 4.7 per cent; in Asia, the average weight is only 2.4 per cent.

In contrast, the average European reserves manager assigns a weight of 3.5 per cent to the Australian dollar, and the average Asian manager allocates as much as 8.3 per cent, in line with the importance of economic ties between Australia and Asian countries.

Moreover, foreign investors (both official reserves managers and other non-resident investors) are relatively more important in Australian government debt markets than in Canadian. Between 2004 and 2011, the foreign share in Australian government debt stock almost doubled, from 35 per cent to 68 per cent.[1] Over that same period, the foreign share in Canadian government debt was relatively stable at between 20 per cent and 30 per cent.

Finally, Australia is not only an issuer of, but also an investor in, non-traditional reserve assets. In 2011, the Reserve Bank of Australia added a 5 per cent allocation to the Canadian currency in its benchmark portfolio; in contrast, Canada does not currently invest its reserves in Australian-dollar assets.

1 See the Australian Office of Financial Management at http://www.aofm.gov.au/statistics/non-resident-holdings/.

Assessing demand over time

The Canadian dollar started to attract increased interest from foreign reserves managers around 2009–10

Since COFER provides Canadian-dollar holdings for only the four quarters starting in December 2012, the data do not provide much information on how the demand for Canadian-dollar-denominated assets has evolved over time. However, the data suggest that the Canadian dollar started to attract increased interest from foreign reserves managers around 2009–10. **Chart 3** shows the evolution of COFER's "other currencies," including the Canadian dollar. Until mid-2009, the other currencies were relatively stable, accounting for about 2 per cent of total global reserves. Starting in the second half of 2009, however, they grew substantially, reaching 6.3 per cent of total reserves in 2013.

Anecdotal evidence also suggests that the growth of Canadian-dollar-denominated assets in foreign reserves started during that period. Annual report data and media coverage indicate that several reserves managers, including the central banks of Chile, the Czech Republic, Iceland, Macedonia and Russia, started investing in Canadian assets at that time. **Box 2** presents a case study of the Swiss National Bank to illustrate the evolution of the demand for Canadian-dollar-denominated assets, as well as several possible drivers of flows into the Canadian market.

In addition, according to Statistics Canada data, the overall foreign portfolio investment in Canadian general government bonds and money market instruments more than doubled over the 2007–13 period and, in November 2013, it stood at over Can$493 billion.[8] The increase in the non-resident holdings was in line with the additional issuance of government securities, and therefore the relative importance of non-resident holdings has remained constant: their share of the government debt market has remained at approximately 29 per cent since the 1990s, reflecting a decrease in the relative importance of private foreign investors such as foreign banks or investment-management companies.[9]

Chart 3: The importance of the Canadian dollar and "other currencies" in official foreign reserves

Quarterly data

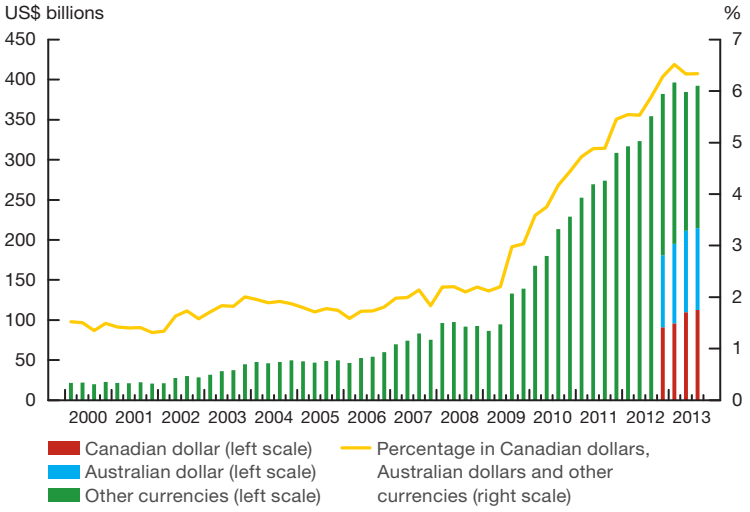

Note: Before December 2012, the Canadian dollar and the Australian dollar were included in "other currencies."
Source: International Monetary Fund COFER Last observation: 30 September 2013

8 Including securities issued by federal, provincial and municipal governments, and by government business enterprises (Statistics Canada, CANSIM Table 376-0146, available at http://www5.statcan.gc.ca/cansim/pick-choisir?lang=eng&p2=33&id=3760146).

9 See, for example, Arslanalp and Tsuda (2012).

Box 2

Canadian-Dollar Reserves of the Swiss National Bank

The Swiss National Bank (SNB), representing the fourth-largest official foreign exchange reserve fund in the world, is an interesting example of a foreign central bank that holds Canadian-dollar-denominated assets.[1]

The SNB decided to add Canadian-dollar assets to its foreign reserves in May 1999, considerably earlier than many other reserves managers. **Chart 2-A** shows the weight of Canadian-dollar assets in Swiss foreign reserves starting in 1999, as well as the overall dollar value of the SNB's Canadian-dollar holdings. Its initial target allocation, 2 per cent, corresponded with roughly $1 billion in Canadian-dollar assets. This target remained unchanged until 2009, although the actual weight and dollar holding in Canadian-dollar-denominated assets varied somewhat with changes in the exchange rates and the overall size of the Swiss foreign exchange reserves. In 2010, the target Canadian-dollar weight doubled to 4 per cent, funded by decreasing allocations to the euro, the U.S. dollar and the British pound. A possible motivation for the change was the desire of Swiss authorities to increase the diversification of their reserves, and perhaps improve their portfolio's resilience against economic uncertainty in Europe and the United States.[2]

Changes in the target weight are only one reason for the growth in the SNB's Canadian-dollar investments. Another important determinant is the overall size of the bank's foreign reserves: constant percentage allocations may still lead to pronounced inflows if the size of the reserves changes, which is what happened after the financial crisis of 2007–09. The heightened global economic uncertainty led to a rapid reserves accumulation worldwide. Until 2008, the SNB foreign currency reserves oscillated around CHF50 billion (Can$58 billion as of 31 December 2008), but in 2009 they doubled to almost CHF100 billion (Can$101 billion as of 31 December 2009). While the target Canadian-dollar weight remained at 2 per cent in 2009, the increase in the size of the reserves was initially funded by currencies other than the Canadian dollar, leading to a temporary drop in the actual Canadian-dollar weight (see **Chart 2-A**). The reserves doubled again in 2010, and once more over the 2011–12 period, this time accompanied by a proportional accumulation of Canadian-dollar-denominated assets.

Chart 2-A: Swiss National Bank Canadian-dollar asset holdings

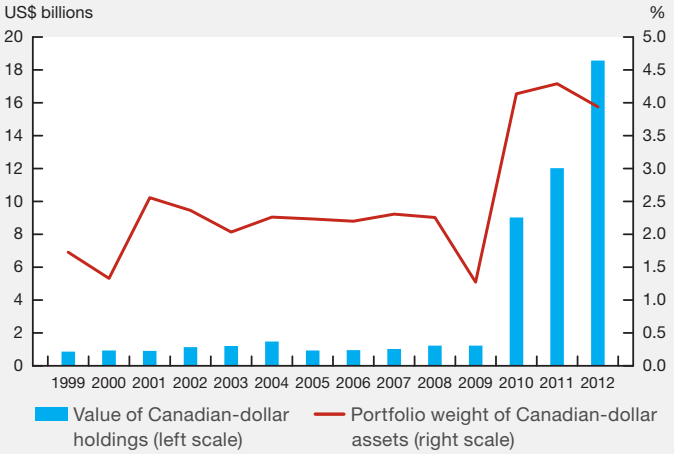

Source: Swiss National Bank annual reports

Last observation: 31 December 2012

The recent growth in Switzerland's foreign reserves illustrates another motive that may underlie flows into and out of Canadian markets: the currency interventions of reserves managers in foreign exchange markets. For the SNB, such interventions were driven by pressures on the Swiss franc–euro exchange rate. Worried by negative developments in the euro area, investors sold the euro and bought the Swiss franc, a traditional "safe haven" currency, leading to a strong appreciation of the currency. The SNB intervened in September 2011 and began to sell the franc, defending the exchange rate of 1.2 euros per franc. The SNB's foreign currency purchases augmented the official foreign reserves and were eventually converted to Canadian dollars, in line with the SNB's target currency allocations, resulting in additional flows into the Canadian markets.

1 The discussion and chart in this box are based on information from the SNB's annual reports, available at http://www.snb.ch.

2 The importance that the Swiss central bank attaches to diversification is further illustrated by its 2012 decision to include the Korean won in its currency allocation, bringing the total number of currencies in its foreign reserves to 10.

Identifying the determinants of demand

To identify the determinants of demand for Canadian-dollar-denominated assets, we consider both the patterns observed in the data and anecdotal evidence from discussions with reserves managers.[10]

Although the details of portfolio strategies and specific holdings are generally not disclosed publicly, most reserves managers share common portfolio-management objectives: preserving capital, maintaining liquidity and, given the required level of risk and liquidity, maximizing expected returns.[11] While there are undoubtedly differences between individual countries, foreign reserves managers are characterized as patient, buy-and-hold investors who may not trade as much as their counterparts in the private sector. Reserves are held for precautionary reasons, as insurance against economic turmoil in the markets. Safety and liquidity are necessary; yield-seeking and speculative behaviour are discouraged.

Foreign reserves managers report that the main reasons for including the Canadian dollar in their portfolios are diversification and safety

Foreign reserves managers report that the main reasons for including the Canadian dollar in their portfolios are diversification and safety. Since reserves are held as insurance, managers want their portfolio to perform well, specifically during times of economic turmoil. Assets that drop in value in times of crisis fit poorly into foreign reserve portfolios, since they are particularly vulnerable precisely at times when reserves managers have the greatest need for funds and liquidity. The recent market stresses in some of the economies of traditional reserve currencies, for example, the euro area, may have led reserves managers to consider non-traditional assets such as the Canadian dollar.

Another perspective on the determinants of demand comes from analyzing the distribution of reserves managers who hold Canadian-dollar-denominated assets. According to COFER data, foreign reserves of advanced economies account for about one-third of the total Canadian-dollar holdings. Emerging economies generally assign a considerably higher weight to the Canadian dollar (2.1 per cent, compared with 1.6 per cent for advanced economies). Given that emerging economies incur higher costs to finance their reserve assets than advanced economies, the historically higher yields that investors earned from Canadian-dollar reserve assets, compared with U.S. government securities or those from core countries in the euro area, will therefore be appealing.

Finally, trade linkages also play a role. Using a small sample of countries for which we have detailed data on Canadian-dollar holdings, we find that countries that have stronger trade links with Canada have a relatively higher weight of the Canadian currency. Papaioannou, Portes and Siourounis (2006) find similar evidence linking the currency composition of a country's reserves to the currencies of its main trading partners and of its own international liabilities.

Possible Impact on Domestic Markets

For the Bank of Canada, understanding the determinants of the flows and the long-term allocation in Canadian-dollar-denominated assets of both private and official foreign investors is crucial for several reasons. As fiscal agent for the federal government responsible for providing advice on the domestic debt program, the Bank needs to understand the long-run demand from the different sectors of the investor base. Potential sudden changes in demand could introduce volatility into the cost of borrowing for

10 In particular, it reflects discussions at the Conference on Foreign Exchange Reserves Management organized by the Bank of Canada in the autumn of 2013.

11 For a discussion of modelling these objectives, see Rivadeneyra et al. (2013). See also IMF (2013).

the federal government. From a financial stability perspective, the Bank monitors the efficiency of financial markets and how their functioning may be affected by these changes in demand.

In this section, we explore the implications of the non-resident flows and the change in the official foreign investor base for average yields of Government of Canada securities, as well as their liquidity and volatility in secondary markets. Inflows from reserves managers may be comparatively stable and, other things being equal, put downward pressure on yields, lowering the funding costs of the government. At the same time, the increasing presence of these large, patient investors might negatively affect the markets' liquidity, making it more difficult for all investors to trade.

A direct effect of increased demand for government securities may be higher bond prices and hence lower yields

A direct effect of increased demand for government securities may be higher bond prices and hence lower yields, since the additional foreign demand is, at least in the short term, independent of the issuance decisions. A large body of literature has explored empirically the effects of foreign inflows into the U.S. Treasury market over the past two decades. Warnock and Warnock (2009) show a negative impact of increased foreign holdings on yields from accumulated net purchases. Bernanke, Reinhart and Sack (2004) measure the short-term effect on yields from Japan's foreign exchange interventions, which typically involve buying U.S. Treasuries, and find a negative but small effect on the 10-year Treasury yield of less than 1 basis point per US$1 billion of purchases. There is additional evidence that, in the context of reserve flows into the U.S. Treasuries market, the effect of foreign holdings on yields tends to be short term. Beltran et al. (2013) estimate that up to two-thirds of this effect can be reversed when domestic investors respond and rebalance their own portfolios, since the downward pressure on yields exerted by non-residents may eventually make the yields less attractive to domestic investors. If domestic investors start selling their holdings, they will push prices downward and yields upward, offsetting some of the impact of foreign investors.

The large-scale asset purchases program that the Federal Reserve conducted in 2009 provides additional evidence of the impact of demand changes. D'Amico and King (2013) find that the cumulative effect of the US$300 billion in purchases of U.S. Treasuries, after controlling for endogeneity, was an average reduction in yields of 30 basis points. In other words, each US$10 billion contributed to a reduction of 1 basis point in the yield curve. Similar to the studies focusing on foreign official flows, D'Amico and King (2013) suggest that an exogenous and permanent reduction in the stock of securities increases prices and consequently reduces yields.

The inflows into the government securities of advanced economies have not been exclusive to Canada and Australia. In fact, in a review of 24 major advanced economies, Arslanalp and Tsuda (2012) find that, between 2004 and 2011, the foreign ownership of government securities increased from 20 per cent to 31 per cent in non-traditional reserve currencies and from 14 per cent to 21 per cent in traditional reserve currencies. In this wider context, Andritzky (2012) finds that a 10 per cent increase in the share of non-resident investors (official and private) in advanced G-20 countries is associated with a decrease in the yields of their respective long-term bonds of between 32 and 43 basis points. Ongoing Bank of Canada research suggests that foreign inflows may have a similar impact on the yields of Government of Canada securities.

A large portion of the rise in foreign ownership of advanced-economy government securities has been due to the increase in the overall size of global reserves. The distinction between the flows from official and private foreign investors is important to understanding the different effects on yields and broadly on financial markets. Although the specific effects are still being debated in the literature, Sierra (2010) finds that official and private flows may have distinct effects. The effect of official foreign net purchases resembles permanent negative supply shocks to the outstanding amount of U.S. Treasury securities available, increasing prices and lowering yields by reducing risk premiums. In contrast, private foreign net purchases appear to correlate positively to an increase in risk premiums.

Similarly, the type of investors has potential effects on the volatility of yields and refinancing risk. Arguably, private foreign investors are relatively more sensitive to external shocks than official reserves managers. Private investors may decide to enter or exit a market quickly and thus represent a relatively less stable source of demand for domestic debt. Growth in official foreign investors may thus be beneficial in lowering the variability of yields. Nevertheless, even official investors may decide to sell their assets at an inopportune time. For example, there is evidence that, during the recent global financial crisis, the procyclical responses of official reserves managers aggravated the funding problems in the U.S. money market and banking sectors (Pihlman and van der Hoorn 2010). Reserves managers may have responded to turbulence in foreign exchange markets by liquidating a portion of their U.S. holdings, which may have depressed prices and possibly worsened market conditions. Indeed, reserves managers are expected to respond to the presence of credit risk in a similar way as private investors. Aizenman and Sun (2009) find varied behaviour from managers regarding the decision to liquidate and use reserves to deal with the crisis.

Foreign reserve flows also have an effect on secondary market liquidity. Bank research indicates that an increased level of foreign reserves managers' holdings of Canadian government securities is an important contributor to changes in the liquidity of that market. One indicator of these changes has been the increased frequency in the use of the securities-lending facility of the Bank of Canada. Securities-lending operations are transactions in which, when specific bond issues are in very high demand in the repo market, the Bank supplies the securities that it holds on its own balance sheet to the market.[12] In 2013, the number of operations the Bank undertook increased substantially to 100, compared with an average of seven in the previous three years.[13] The reasons behind this dramatic increase are complex, but internal analysis indicates that the increase in official foreign holdings is one of the contributing factors. Reserves managers may have less incentive to lend their securities in the repo market, perhaps because they have less appetite than private investors to trade actively, or perhaps they consider potential profits from securities lending insufficient compensation for the provision of liquidity.

There are also indications that the secondary market turnover of securities is negatively related to the level of holdings by foreign reserves managers. The negative impact on liquidity may make it more difficult for investors to trade in the Canadian markets. Lower liquidity may also eventually translate into higher yields, since investors will demand additional compensation for entering illiquid markets.

12 Dreff (2010) discusses the importance of securities lending for market liquidity.

13 The number of securities-lending operations is available on the Bank's website under Banking and Financial Statistics, http://www.bankofcanada.ca/publications/bfs/.

Overall, however, since foreign central bank holdings are not the only factor affecting market liquidity, it is difficult to draw definitive conclusions about the impact of foreign official holdings on market liquidity. It is possible that high foreign central bank holdings coexist or even interact with other factors (for example, the new financial regulation requiring collateral for derivatives trades that are cleared through central counterparties), increasing the demand for government securities.

Since foreign central bank holdings are not the only factor affecting market liquidity, it is difficult to draw definitive conclusions about the impact of foreign official holdings on market liquidity

Conclusion

The potential impact on Canadian financial markets from increased official foreign holdings, particularly the impact on the functioning of core Canadian funding markets, highlights the importance of the Bank of Canada's work to monitor these developments. Given the prospect for continued future growth of reserves, understanding these effects is even more important. Assuming a constant allocation of global reserves to the Canadian dollar of 1.8 per cent, a 1-percentage-point increase in global reserves implies an additional inflow of about $2 billion to Canadian-dollar-denominated assets.[14] This may be a conservative estimate of the official foreign demand for Canadian assets, since most of the growth in reserves is expected to come from emerging economies, which allocate a higher weight to the Canadian currency than advanced countries.

Overall, the holdings of official foreign investors are likely to remain a salient, and may become an even more important, feature of Canadian government debt markets.

14 The average annual growth rate of global reserves since 2007Q3 has been 10 per cent.

Literature Cited

Aizenman, J. and Y. Sun. 2009. "The Financial Crisis and Sizable International Reserves Depletion: From 'Fear of Floating' to the 'Fear of Losing International Reserves'?" National Bureau of Economic Research Working Paper No. 15308.

Andritzky, J. R. 2012. "Government Bonds and Their Investors: What Are the Facts and Do They Matter?" International Monetary Fund Working Paper No. WP/12/158.

Arslanalp, S. and T. Tsuda. 2012. "Tracking Global Demand for Advanced Economy Sovereign Debt." International Monetary Fund Working Paper No. WP/12/284.

Beltran, D. O., M. Kretchmer, J. Marquez and C. P. Thomas. 2013. "Foreign Holdings of U.S. Treasuries and U.S. Treasury Yields." *Journal of International Money and Finance* 32: 1120–43.

Bernanke, B. S., V. R. Reinhart and B. P. Sack. 2004. "Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment." *Brookings Papers on Economic Activity* 35 (2): 1–100.

D'Amico, S. and T. B. King. 2013. "Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply." *Journal of Financial Economics* 108 (2): 425–48.

Dreff, N. 2010. "The Role of Securities Lending in Market Liquidity." Bank of Canada *Financial System Review* (June): 41–45.

Fontaine, J.-S., J. Selody and C. Wilkins. 2009. "Improving the Resilience of Core Funding Markets." Bank of Canada *Financial System Review* (December): 41–46.

Goldberg, L., C. E. Hull and S. Stein. 2013. "Do Industrialized Countries Hold the Right Foreign Exchange Reserves?" Federal Reserve Bank of New York *Current Issues in Economics and Finance* 19 (1): 1–10.

International Monetary Fund (IMF). 2011. "Assessing Reserve Adequacy" (14 February).

———. 2013. "Revised Guidelines for Foreign Exchange Reserve Management" (1 February).

Papaioannou, E., R. Portes and G. Siourounis. 2006. "Optimal Currency Shares in International Reserves: The Impact of the Euro and the Prospects for the Dollar." *Journal of the Japanese and International Economies* 20 (4): 508–47.

Pihlman, J. and H. van der Hoorn. 2010. "Procyclicality in Central Bank Reserve Management: Evidence from the Crisis." International Monetary Fund Working Paper No. WP/10/150.

Rivadeneyra, F., J. Jin, N. Bulusu and L. Pomorski. 2013. "Modelling the Asset-Allocation and Liability Strategy for Canada's Foreign Exchange Reserves." *Bank of Canada Review* (Spring): 29–36.

Sierra, J. 2010. "International Capital Flows and Bond Risk Premia." Bank of Canada Working Paper No. 2010-14.

Warnock, F. E. and V. C. Warnock. 2009. "International Capital Flows and U.S. Interest Rates." *Journal of International Money and Finance* 28 (6): 903–19.

Understanding Platform-Based Digital Currencies

Ben Fung and Hanna Halaburda, Currency Department

- With advances in technology and the rapid spread of the Internet, various digital currencies have emerged. While digital currencies could increase the efficiency of retail payments, they could also raise some important policy issues if they were to become widely used.
- This article focuses on digital currencies issued by Internet platforms such as Facebook and Amazon. Depending on the platform's business model, its digital currency may be equipped with different attributes that affect how users can acquire, transfer or redeem the digital currency. In most cases, platforms restrict the functionality of their digital currencies to enhance the business model and maximize their profits.
- A platform-based digital currency has the potential to become a widely accepted means of payment outside of its platform if it is transferable among its users. None of the platform-based digital currencies with this feature is widely used at this point. This could change, however. Therefore, it is important to closely monitor the evolution of these digital currencies.

Innovations in technology and the widespread use of the Internet have made online commerce, social networks and online gaming a significant part of our lives. The key players in these areas are platforms like Facebook and Amazon.[1] Some of these platforms have issued tokens, such as Facebook Credits or Amazon Coins, that individuals can use to purchase real or virtual goods within the platform. These tokens are often referred to as "digital currency." With millions of users in many countries, Internet platforms have a global reach. Some industry observers have speculated that these currencies could become widely accepted and could even compete with national currencies.

This article distinguishes between digital currencies and the digitization of national currencies such as the Canadian dollar or the U.S. dollar, which involves the electronic transfer of a national currency between two accounts (e.g., using debit or credit cards). Digital currencies, in contrast, have no

1 Platforms are enterprises where the value of using the platform increases with the number of market participants that join. In the case of Facebook, the more friends that are using it, the more attractive the platform becomes. In the Amazon Marketplace, the more sellers selling their products, the more attractive that platform is to buyers, and Amazon's tablet, Kindle Fire, becomes more attractive to users as developers provide more applications for it.

Box 1

Bitcoin: A Digital Currency Without a Central Issuer

Gerald Stuber, Currency Department

In contrast to platform-based digital currencies, Bitcoin is a completely decentralized currency without a central issuer.[1] Based on specialized open-source software, a set amount of bitcoins is given to users in exchange for specific contributions to the operation of the Bitcoin system. Users can transfer bitcoins among themselves or use them to purchase goods and services, provided they can find merchants willing to accept them. Bitcoins can also be bought and sold for national currencies through several unofficial Internet-based "exchanges."

Since bitcoins are purely digital and there is no central institution controlling their use, there is a substantial risk that some users may try to duplicate or counterfeit them. The Bitcoin software solves this problem by using the open community of users to check bitcoin transaction records and validate new transactions.[2] The first user to successfully validate new transactions is rewarded with newly released bitcoins. There is a maximum amount of bitcoins that can be supplied, although this amount will not be reached until 2140.

Although Bitcoin automatic teller machines have been introduced in some major Canadian cities, very few Canadian merchants accept bitcoins as a means of payment. And while the Department of Finance Canada has indicated that the Bitcoin currency is not legal tender in Canada, the Canada Revenue Agency has announced that standard tax rules apply in the use of bitcoins or other digital currencies.

Potential benefits associated with bitcoins include lower transaction costs to online merchants than for conventional payment instruments such as credit cards, since there is no third-party intermediary. Similarly, the costs of international remittances might also be lower than for conventional remittance methods. However, Bitcoin users face a number of challenges, particularly the extreme volatility of the price of bitcoins.[3] As well, it is relatively easy to delete or misplace personal holdings of bitcoins. There have also been a number of security incidents that have compromised either Bitcoin accounts or some other part of the Bitcoin infrastructure (such as Bitcoin exchanges).

In addition, governments may become concerned about a number of legal, security and law-enforcement issues associated with bitcoins. For example, given the private nature of bitcoin transactions, bitcoins could easily be used to facilitate criminal transactions and to evade taxes.[4]

As they do with platform-based digital currencies, central banks are studying and closely monitoring decentralized digital currencies such as Bitcoin. There could be potential risks to overall financial stability if Bitcoin became a significant means of payment and the Bitcoin system remained unstable. As well, Bitcoin users need to be aware of the potential financial risks to which they might be exposed, in light of the ongoing volatility of bitcoin prices and the risk of failure of Bitcoin exchanges.[5] In particular, given that digital currencies such as Bitcoin are not regulated and do not have a centralized issuer, users bear all of the risks themselves and have no legal recourse should they wish to reverse a bitcoin transaction.

1 For more information on Bitcoin, see Brito and Castillo (2013); Elwell, Murphy and Seitzinger (2013); ECB (2012); Nielsen (2013); Šurda (2012); and Velde (2013).

2 The records are public, but each user and each bitcoin are encrypted with unique identities.

3 For example, Bitcoin's price rose to above US$1,200 in early December 2013 and then fell to around US$800 by the middle of that month, after an announcement that the world's largest Bitcoin exchange (BTC China) would no longer accept new customers in China.

4 The federal government recently introduced amendments to the Proceeds of Crime (Money Laundering) and Terrorist Financing Act that would make digital currencies (such as Bitcoin) subject to the application of the Act.

5 Mt. Gox, once the world's largest Bitcoin exchange, recently filed for bankruptcy protection in Japan and the United States.

physical counterpart and do not represent a claim on assets. They are usually not denominated in the national currency and thus provide their own unit of account. The focus of this article is on digital currencies that are centralized and issued by proprietary Internet platforms, rather than digital currencies that are decentralized, for example, Bitcoin (**Box 1**). Platform-based digital currencies are characterized by two main features: (i) the platform maintains control over the design and supply of the currency, and (ii) the platform introduces its currency for objectives other than payment services.

In finance and economics, currency is defined as (i) a unit of account, (ii) a medium of exchange and (iii) a store of value. As will be discussed in the next section, Facebook Credits and other popular platform-based digital

currencies are limited in functionality and may not satisfy this definition. Nevertheless, to facilitate discussion, we refer to Facebook Credits and the other examples in this article as platform-based digital currencies, since this term is used by the popular press.

What drives platforms to introduce their own digital currencies? Why do platforms design their digital currencies in a particular way? Under what circumstances could these digital currencies become more widely used outside the platform? To answer these questions, we conduct an in-depth analysis of Facebook Credits to develop a framework for studying and monitoring developments in platform-based digital currencies more generally. We then use the framework to assess other prominent platform-based digital currencies. Finally, we discuss the relative importance of different elements of the framework for understanding the potential impact of a platform-based digital currency on a national currency.

Central banks and other public authorities are studying and monitoring these platform-based digital currencies for a number of reasons. First, it is important to assess their potential impact on the demand for bank notes. Second, digital currencies represent innovations that could increase the efficiency of retail payments by, for example, providing a cheaper, faster or more convenient alternative to existing payment methods, especially for online transactions.[2] Third, digital currencies could raise a range of public policy and regulatory issues if they were to become a widely used means of payment.[3]

Not All Platform-Based Digital Currencies Are the Same

Facebook Credits

In mid-2009, Facebook, the most popular social networking site in the world, introduced its digital currency—Facebook Credits (FB Credits). With FB Credits, users could purchase premium content for games and applications on Facebook, allowing them to play longer or achieve better results. For example, users could buy fertilizer for virtual plants to increase the "harvest" in their virtual farm. They could earn FB Credits by filling out surveys or testing a beta version of a game, and they could buy FB Credits with national currency. Once acquired, however, FB Credits could not be redeemed for national currency or transferred to another user. They could be "spent" only on the Facebook platform.

Given the large number of Facebook users (over 1 billion) and its international reach, industry observers speculated that FB Credits could become the currency of the web or a global means of payment. However, it was in Facebook's best interest to restrict the functionality of FB Credits (Gans and Halaburda 2013), which made them not viable for use outside the platform. Facebook's main source of revenue is advertising, and the value of advertising space increases with the length of time users spend on the platform. By buying virtual goods with FB Credits, users enhanced their experience and therefore would stay on the platform longer. This had positive spillovers,

2 In Canada, credit cards are the most popular payment instrument for online transactions (Statistics Canada 2012). However, the fees that credit cards charge merchants are very high, particularly for small-value transactions. In addition, concerns about fraud, privacy and identity theft may deter some consumers from using credit cards to make online purchases (The Paypers 2014).

3 For a discussion of issues related to the regulation of digital currencies, including protecting consumer funds and the privacy of data, as well as preventing money laundering and terrorist funding, see ECB (2012). For a discussion of potential policy issues such as the impact on monetary policy and financial stability, see BIS (1996).

since with Facebook, as with many platforms, the value of spending time on the platform (consuming) depends not only on how long an individual is there, but also on the length of time the individual's friends spend there; this is known as "consumption complementarity." When one user acquired FB Credits and therefore logged more hours on the platform, other users stayed on the site longer as well. They were also more likely to acquire FB Credits themselves, further increasing their time on Facebook. As users engaged more with the platform, Facebook's value to advertisers would rise and advertising revenue would increase.

Since Facebook allowed users to acquire FB Credits by earning as well as buying them, it attracted users with less money but more time to earn credits, as well as those with less time but more money to buy credits. Both types of users would likely increase their time on Facebook as they spent their FB Credits on virtual goods. Allowing users to redeem FB Credits for the national currency would have therefore undermined Facebook's objective of enticing users to indulge in longer sessions on the platform. Allowing users to transfer FB Credits among themselves would have also undermined this goal, since users who earned the FB Credits could have sold them to other users (perhaps at a lower price), instead of using the credits to enhance their own Facebook consumption.[4] By limiting the functionality of FB Credits, Facebook sought to maximize the length of time users were on the platform and, hence, its advertising revenue.

Although eliminating the restrictions on FB Credits would have undermined this objective, it could have allowed Facebook to offer a means of payment. And if it became widely accepted, Facebook could have earned both fees and seigniorage revenue.[5] However, the profit resulting from offering payment services would not likely be larger than that coming from advertising. For example, the Bank of Canada's annual seigniorage revenue is in the range of Can$1 billion to Can$2 billion, which is dwarfed by Facebook's 2013 total revenue of almost US$8 billion (Edwards 2014).[6] Introducing a payment service may also be undesirable from a business standpoint, since Facebook would be subject to regulation as a financial institution in many countries, which could limit its use of the personal data collected from users. In the end, Facebook has to weigh the pros and cons of offering payment services.

In mid-2012, Facebook announced plans to phase out FB Credits by September 2013. This decision was driven by conflicts with large developers of Facebook games that had introduced their own digital currencies before FB Credits became available. For example, Zynga found its own currency, zCoins, to be profitable for the same reason that Facebook did with FB Credits—they increased user activity. However, to play the game, users needed to change national currency into FB Credits and then change FB Credits into zCoins. In response to users' complaints, Facebook first tried to convince Zynga to abandon zCoins and adopt FB Credits. When that effort failed, Facebook simplified the process by abandoning FB Credits.

4 Because of the nature of the activities through which users earned FB Credits, the time spent earning them did not contribute to Facebook's advertising revenue.

5 Seigniorage is the revenue earned from the issue of money. In the case of the Bank of Canada, it can be calculated as the difference between the revenue earned on a portfolio of Government of Canada securities—in which the Bank invests the total value of all bank notes in circulation—and the cost of issuing, distributing and replacing those notes. For more details, see www.bankofcanada.ca/wp-content/uploads/2010/11/seigniorage.pdf.

6 Almost all of Facebook's revenue comes from advertising. Although Facebook's user population is larger than the population of Canada, only a small fraction of that user population had ever used FB Credits.

Table 1: Summary of attributes of selected platform-based digital currencies

Platform	Acquirability	Transferability	Redeemability
FB Credits	Both (buy and earn)	No	No
Amazon Coins	Buy	No	No
WoW gold	Earn	Yes	No
Q-coin	Both (buy and earn)	Yes	No
Linden dollars	Both (buy and earn)	Yes	Yes

This analysis of FB Credits allows us to develop a framework for assessing the functionality of platform-based digital currencies, examining not only the size and reach of a platform, but also the reasons why a platform introduces its digital currency. A platform's goal is to increase profits. Depending on the platform's business model, the role of the currency may be different, and thus the platform may equip it with different attributes, specifically whether users can (i) buy or earn the tokens (acquirability), (ii) transfer the tokens between each other (transferability), and/or (iii) redeem the tokens for national currency (redeemability). We call a currency "fully equipped" when it has all three attributes. In most cases, however, a platform prefers to limit the functionality of its currency to fit its business model.

Depending on the platform's business model, the role of the currency may be different, and thus the platform may equip it with different attributes

For a digital currency to be used outside the platform as a means of payment, it must meet two conditions. First, it must be equipped with the attributes just described, which could make it possible for people to adopt it as a currency. Second, individual market participants must decide to adopt the currency. A large existing literature (e.g., Kiyotaki and Wright (1989) and Lagos (2013)) has focused on what induces people to adopt one currency over another. This article focuses on the first condition and investigates whether proprietary Internet platforms have incentives to develop their currencies in a way that allows them to be used outside the platform.

In the remainder of this section, this framework is used to characterize the attributes of other prominent examples of platform-based digital currencies and to assess whether they could be adopted as a currency. **Table 1** summarizes the main attributes found in each of these currencies.

Amazon Coins

Since May 2013, Amazon has been giving away millions of U.S. dollars in Amazon Coins to customers who purchase its newest tablet, the second-generation Kindle Fire. However, Amazon has imposed tight restrictions on the use of its coins. They can be spent only on approved applications (apps) for the Kindle Fire and cannot be used to purchase books or other merchandise from Amazon.com. Moreover, users cannot transfer Amazon Coins to another user or redeem them for dollars.[7] These restrictions are aligned with the role of the coins in Amazon's business model, which is to improve the market position of Amazon's Kindle Fire as a platform for its apps.

Amazon is a relative latecomer to the tablet market. To make the Kindle Fire more attractive to customers, Amazon supports the development of attractive Kindle-specific apps. By offering Amazon Coins to customers (as a gift or for purchase) and limiting the spending of the coins to these apps, Amazon provides incentives to the developers to create popular apps on which the coins will be spent.[8] Allowing Amazon Coins to be a fully equipped currency would be at odds with this objective.

7 In addition to receiving US$50 worth of Amazon Coins when purchasing a Kindle Fire, users can buy coins from Amazon. However, they cannot earn them.

8 The developers (but not the users) redeem Amazon Coins for dollars.

World of Warcraft gold

World of Warcraft (WoW) is the most popular multi-player role-playing game online, with around 8 million gamers worldwide paying subscription fees. Among many activities on the platform, gamers can earn tokens, WoW gold, and use them to buy additional gear for their avatars.[9] Since gamers can buy items from other gamers, WoW gold is transferable among members within the platform. However, users acquire WoW gold only by earning it and cannot buy it with national currency, since earning WoW gold is directly related to activity on the platform. By showing items bought with WoW gold, players are displaying their level of experience and skill to other players—an important element of the game, for example, when choosing a team to go on a quest. Therefore, allowing gamers to buy the items with national currency would break the link between WoW gold and skill level, and would thus undermine the attractiveness of the game.[10] For the same reason, WoW gold is not redeemable for national currency.

Tencent's Q-coin

A popular Chinese social networking site, Tencent, introduced Q-coin to enable users to pay for virtual goods on the site (e.g., to send virtual flowers). Q-coin can be earned or bought, and can also be transferred among members of the platform, although it is not redeemable. Q-coin is an interesting example of a digital currency, because it gained traction outside of its own platform.[11]

While Q-coin was intended for the purchase of virtual goods and services provided by Tencent, it started to be used for peer-to-peer payments. Not only online merchants, but also brick-and-mortar stores started accepting Q-coin (Fowler and Qin 2007).[12] In 2008, the value of Q-coin reportedly reached several billion renminbi (RMB).[13] The Chinese government responded in June 2009 with regulation banning the exchange of a digital currency for real goods and services, in order to "limit its possible impact on the real financial system."[14]

Q-coin had the potential to be used as an alternative to national currency, despite not being redeemable. This is because transferability allowed users to indirectly redeem Q-coin by transferring it among themselves inside the platform and exchanging it for real goods and services, and unofficially for national currency, outside the platform.

Linden dollars

Linden dollars, a platform-based digital currency for the game Second Life, are a fully equipped currency, since economic activity (e.g., setting up and operating shops) is part of the game. Players earn Linden dollars by trading

9 In gaming, an avatar is the graphical representation of a user's character.

10 There is, however, a thriving "black market" outside of the WoW platform, where people buy and sell not only WoW gold, but also other items, including fully equipped avatars from higher levels. Such "impostors" are policed by the WoW community and, when discovered, expelled from the game.

11 Not enough information on Tencent's business model is available to draw firm conclusions on the rationale behind the design of Q-coin.

12 One possible explanation for why Q-coin and other innovative online payment schemes have flourished in China in recent years is that China's less-developed financial system, compared with those of advanced economies, is unable to meet all the payment needs of its consumers. For example, in 2012, the number of credit cards per capita in China was only 0.25, compared with 2.3 in Canada (CPSS 2013). China's inconvertible currency and capital controls may also provide incentives to use digital currencies.

13 See the Government of China news release at http://english.mofcom.gov.cn/aarticle/newsrelease/commonnews/200906/20090606364208.html.

14 Ibid.

with other players for virtual goods; thus, they are transferable. Players bring more Linden dollars into the game by buying them with national currency, and Linden dollars earned in the game can be redeemed for the national currency (at the exchange rate regulated by Linden Labs, the game's developer). Although fully equipped, Linden dollars are not widely used outside the platform. This highlights the point that having all of the desired attributes is not a sufficient condition for a currency to become widely accepted.

Could Platform-Based Digital Currencies Be Widely Used Outside the Platform?

The examples in this article illustrate that platform-based digital currencies may vary considerably in their design. Platforms introduce their own tokens instead of relying on the existing means of payment (e.g., credit cards), because they are able to design the currency's properties to suit their business model. For example, platform-based digital currencies could provide more flexibility in acquiring goods by allowing users to both buy and earn them within the platform (such as with FB Credits), or platforms can impose more restrictions on how these currencies are spent (as seen in the limitation of Amazon Coins to the purchase of Kindle Fire apps). These currencies also facilitate the creation of an economy in a virtual world that is separate from the one in the real world, but enhances the virtual experience (as done in World of Warcraft).

Some of the digital currencies described in this article—FB Credits, Amazon Coins and WoW gold—are too limited in their functionality to become a widely accepted means of payment. Since the respective platforms imposed these limitations to maximize profits, it is not in their best interest to issue fully equipped currencies. At the same time, Q-coin, despite not being redeemable, demonstrated its potential to be widely adopted outside of the platform in the real economy, resulting in a ban by authorities. So, what attributes might contribute to these platform-based digital currencies being widely adopted as a digital alternative to national currency?

According to our assessment, transferability appears to be the attribute that gives a digital currency the potential to become a means of payment outside its platform. Transferability is more important than how users can acquire the tokens (either by buying or earning them), or whether they can redeem them for the national currency. As seen in the examples, with transferability, individuals can buy the tokens even when officially prevented by the platform, as is done with WoW gold.[15] Individuals can also unofficially exchange the digital currency for national currency, as was done with Q-coin.

◀ *Transferability appears to be the attribute that gives a digital currency the potential to become a means of payment outside its platform*

Redeemability is not a necessary attribute for the currency to become widely accepted as a means of payment. All that is needed is the *belief* that the tokens that an individual has acquired will be accepted in the future to pay for the desired goods and services. For example, before 2009, many people had accumulated and spent thousands of Q-coins to buy real goods and services without ever exchanging them for RMB.

Yet, it is important to note that not every currency that is transferable will be used outside the platform. WoW gold is an example, and, more notably, Linden dollars, which, despite being available for over a decade and being fully equipped, have not been widely adopted outside the economy of Second Life. The issue of adoption is beyond the scope of this article and, as noted earlier, has already been discussed extensively in the literature.

15 In World of Warcraft, however, effective community policing prevents such unofficial purchases from becoming widespread.

Conclusion

For the most part, digital currencies issued by proprietary Internet platforms are unlikely to affect existing national payment systems in any significant way. Platforms introduce their currencies to enhance their business model and increase their profits. In most cases, this objective requires limiting the functionality of the platform's currency, which will prevent it from becoming a widely accepted means of payment.[16]

Digital currencies issued by proprietary Internet platforms are unlikely to affect existing national payment systems in any significant way

The online marketplace will continue to evolve, and platforms will develop innovative products and perhaps new digital currencies. Our analysis helps to identify the attribute of digital currencies that needs to be monitored most closely—transferability among users—since it provides the currency with the greatest potential to be adopted as a means of payment outside the platform. However, a digital currency will be adopted by market participants only if it is more attractive to use or better suited to meet their payment needs than existing alternatives.

16 Platforms may also change their business model to adapt to the evolution of the competitive environment. The questions of when and how this could occur are beyond the scope of this article. There is, however, extensive literature on strategic renewal that focuses on these issues (e.g., Mische (2000)).

Literature Cited

Bank for International Settlements (BIS). 1996. "Implications for Central Banks of the Development of Electronic Money" (October).

Brito, J. and A. Castillo. 2013. "Bitcoin: A Primer for Policymakers." Mercatus Center (George Mason University).

Committee on Payment and Settlement Systems (CPSS). 2013. "Statistics on Payment, Clearing and Settlement Systems in the CPSS Countries: Figures for 2012" (December). Bank for International Settlements.

Edwards, J. 2014. "Facebook's Shares Surge on First Ever $1 Billion Mobile Ad Revenue Quarter." *Business Insider* (29 January).

Elwell, C. K., M. M. Murphy and M. V. Seitzinger. 2013. "Bitcoin: Questions, Answers, and Analysis of Legal Issues." Congressional Research Service Report (20 December).

European Central Bank (ECB). 2012. "Virtual Currency Schemes" (October).

Fowler, G. A. and J. Qin. 2007. "QQ: China's New Coin of the Realm? Officials Try to Crack Down as Fake Online Currency Is Traded for Real Money." *Wall Street Journal* (30 March).

Gans, J. S. and H. Halaburda. 2013. "Some Economics of Private Digital Currency." Bank of Canada Working Paper No. 2013-38.

Kiyotaki, N. and R. Wright. 1989. "On Money as a Medium of Exchange." *Journal of Political Economy* 97 (4): 927–54.

Lagos, R. 2013. "Moneyspots: Extraneous Attributes and the Coexistence of Money and Interest-Bearing Nominal Bonds." *Journal of Political Economy* 121 (1): 127–85.

Mische, M. A. 2000. *Strategic Renewal: Becoming a High-Performance Organization*. Upper Saddle River, NJ: Prentice Hall.

Nielsen, M. 2013. "How the Bitcoin Protocol Actually Works." *Data-Driven Intelligence* (6 December). Available at http://www.michaelnielsen.org/ddi/how-the-bitcoin-protocol-actually-works/.

Statistics Canada. 2012. *Canadian Internet Use Survey*.

Šurda, P. 2012. "Economics of Bitcoin: Is Bitcoin an Alternative to Fiat Currencies and Gold?" Master's thesis. Vienna University of Economics and Business.

The Paypers. 2014. *Cross-Border Ecommerce Report Canada*.

Velde, F. R. 2013. "Bitcoin: A Primer." *Chicago Fed Letter*. Number 317 (December). Federal Reserve Bank of Chicago.

The Art and Science of Forecasting the Real Price of Oil

Christiane Baumeister, International Economic Analysis

- Forecasts of the price of crude oil play a significant role in the conduct of monetary policy, especially for commodity-producing countries such as Canada.
- This article explores a range of recently developed forecasting models that can generate, on average, accurate forecasts of the price of oil. Instead of relying on a single model, this article shows that forecast combinations outperform the oil futures curve.
- In addition to accurate forecasts of the price of oil, policy-makers are interested in evaluating the risks associated with the baseline forecast to gauge the implications of alternative oil price paths for the economic outlook. A structural model of the global oil market can be used to develop risk scenarios for oil price forecasts, based on hypothetical assumptions about future demand and supply conditions in the crude oil market.
- Based on this structural model, it can also be shown that changes in demand associated with the global business cycle are the primary determinant of changes in oil prices.

Given the importance of oil prices for the Canadian economy, understanding what drives fluctuations in oil prices and how best to forecast them is critical for monetary policy. Specifically, oil price forecasts play an important role in assessing the future developments of inflation and economic activity in Canada and its trading partners, with implications for Canada's terms of trade.

Until recently, central banks and international organizations tended to rely exclusively on the oil futures curve to forecast the price of oil. Recent research, however, demonstrates that models that include the economic determinants of the price of oil, such as changes in oil inventories, oil production and global real economic activity, may provide more accurate out-of-sample forecasts than oil futures prices (Alquist, Kilian and Vigfusson 2013; Baumeister and Kilian 2014b; Baumeister, Kilian and Zhou 2013). This finding holds even in a real-time forecasting environment, where predictors of the price of oil become available only with a delay and are subsequently revised repeatedly (Baumeister and Kilian 2012).

An important limitation of all these forecasting approaches is that they provide limited insight into what is affecting the oil price forecast. It is, however, possible to derive a structural model of the global oil market from one of the forecasting models that helps policy-makers to interpret oil price forecasts. Such an economic model can also be used to evaluate the risks associated with the baseline forecast—that is, how the future path of the price of oil changes with alternative hypothetical scenarios for the economic environment.

This article begins by describing recent advances in forecasting the real price of oil. It stresses the benefits of combining the forecasts from alternative models that have different strengths and weaknesses, rather than relying on only one forecasting method. It then outlines a framework for constructing forecast scenarios that enhances policy-makers' understanding of the factors affecting oil prices and allows them to formally assess the risks associated with oil price forecasts.

Forecasting Models

The volatility of the real price of oil since 2003 has renewed interest in how best to forecast oil prices (**Chart 1**). This section presents the traditional approach that uses oil futures prices as predictors of the real price of oil, as well as three forecasting methods developed recently at the Bank of Canada. It then compares the relative accuracy of a combination of these forecasting methods with the no-change benchmark forecast.

Chart 1: The real price of oil, 1974 to 2013, in November 2013 U.S. dollars per barrel

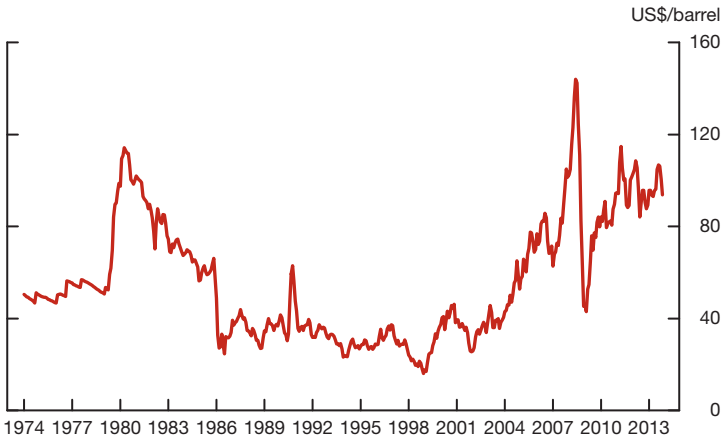

Note: The real price of oil is the nominal refiner acquisition cost of crude oil imports deflated by U.S. CPI.
Sources: U.S. Energy Information Administration, *Monthly Energy Review*; and the Federal Reserve Bank of St. Louis
Last observation: November 2013

Oil futures curve

The traditional approach to constructing out-of-sample forecasts of the real price of oil is to rely on the oil futures curve. Since the oil futures market plays an important role in information aggregation and price discovery, the prices of crude oil futures contracts traded on exchanges such as the New York Mercantile Exchange or the Intercontinental Exchange are commonly perceived to reflect the expectations of market participants about the future

course of oil prices (Alquist and Arbatli 2010). When communicating policy decisions, many central banks have highlighted the importance of oil futures prices for the future evolution of inflation.

When the forecasting performance of oil futures prices is evaluated over a period of 20 years against a simple model that postulates that prices will remain unchanged over the forecast horizon (the no-change forecast), at shorter horizons, there is no significant evidence that the oil futures curve achieves gains in forecast accuracy. Moreover, Alquist, Kilian and Vigfusson (2013) show that, at longer-term horizons that matter for policy-makers, the forecasting performance of the oil futures curve is inferior when compared with the no-change forecast. A possible explanation for this finding is that oil futures prices contain a time-varying risk premium. In fact, Hamilton and Wu (2014) find evidence of considerable changes in risk premiums in oil futures prices after 2005.

When the forecasting performance of oil futures prices is evaluated against a simple model postulating that prices will remain unchanged over the forecast horizon, at shorter horizons, there is no significant evidence that the oil futures curve achieves gains in forecast accuracy

Model of the global oil market

The first of the recently developed alternative approaches uses a model of the global market for crude oil that includes the key determinants of oil prices based on economic theory. Specifically, the current real price of oil is modelled as a function of its own past values and the past values of world oil production, an index of real economic activity that captures fluctuations in the global business cycle and changes in global above-ground inventories of crude oil.

Out-of-sample forecasts generated by this model tend to be more accurate than the no-change forecast at short horizons, even when real-time data constraints are taken into account (Baumeister and Kilian 2012; 2014b).

Spot price of raw industrial materials

The second alternative method is based on the observation that prices of non-oil industrial commodities such as copper and zinc are indicators of shifts in the demand for all industrial commodities, including oil. To the extent that persistent fluctuations in the global business cycle move together with industrial commodity prices, recent cumulative changes in the price indexes of non-oil industrial commodities are expected to have predictive power for the real price of oil.

Based on this insight, Baumeister and Kilian (2012) show that forecasts that extrapolate cumulative changes in the spot price for raw industrial materials adjusted for expected inflation perform well at short horizons relative to the no-change forecast, but become increasingly less accurate at horizons beyond three months. The ability of these forecasts to accurately predict whether the price of oil is increasing or decreasing is consistently high for horizons of up to 12 months.

Refined product spreads

The third promising forecasting approach is based on the idea that the demand for crude oil is driven by the demand for refined petroleum products, such as gasoline, heating oil and diesel. This relationship suggests that spot market prices for petroleum products will ultimately determine the price for crude oil. In fact, many oil industry analysts believe that a widening of the spread between product prices and the price of crude oil signals upward pressures on future oil prices. This insight may be exploited by analyzing whether changes in these price spreads, defined as the extent to which today's price of gasoline or heating oil deviates from today's price of crude oil, have predictive power for future changes in the price of oil.

There are, however, many reasons to expect this predictive relationship to be unstable over time. Given that refined products are produced in approximately fixed proportions, the price of crude oil is likely to be determined by the refined product in highest demand, and this product changes gradually. Another concern is that oil supply shocks, local capacity constraints in refining, changes in environmental regulations or other market turmoil may temporarily undermine the predictive content of these price spreads. To take these possibilities into account, Baumeister, Kilian and Zhou (2013) allow the weights assigned to gasoline price spreads and heating oil price spreads in the forecasting model to evolve smoothly. They find that this model delivers more accurate forecasts relative to the no-change forecast, especially at horizons between one and two years.

Combining forecasts from different models

Given the range of models available for forecasting the real price of oil, the question that arises is which model policy-makers should rely on to obtain the most accurate point forecasts and to correctly predict whether the oil price will go up or down over the projection horizon.

Rather than selecting a single model, it might be beneficial to pool the information contained in the four individual models (Baumeister and Kilian 2013). Combining forecasts from different models is promising for several reasons. First, even the most accurate forecasting models do not work equally well at all times. The global oil market forecasting model, for example, works well during periods when economic fundamentals show persistent variation, as was the case between 2002 and 2011, but not as well at other times. Similarly, there is considerable variation over time in the ability of oil futures prices to forecast the price of oil.

Second, the forecasting performance of individual models differs depending on the forecast horizon. For example, models based on economic fundamentals tend to be more accurate at short horizons, while models based on the spread between the prices of refined products and the price of crude oil tend to be more accurate at longer horizons. Since the policy horizon usually spans a period of two years, no single model provides the most accurate forecasts for the entire projection horizon.

Third, even a forecasting model with a better forecasting track record could be improved upon by incorporating additional information from other models that perform less well, on average.

Pooling information from different models produces more robust forecasts than relying on any one individual model

These arguments suggest that combining forecasts from different models should be more reliable than individual models.[1] To evaluate the forecasting performance of equal-weighted forecast combinations, two criteria are considered. The first is the relative mean-squared prediction error (MSPE), which measures the average squared deviation between the pooled forecasts and the actual realization, relative to the no-change benchmark. An MSPE ratio below 1 indicates that the pooled forecasts are more accurate than the benchmark forecast. Second, the directional accuracy of the pooled forecasts is assessed by the success ratio, which represents the number of times that a method correctly predicts whether the real price of oil is increasing or decreasing. If there were no directional accuracy, the model should be no more successful at predicting the direction of price changes

1 Baumeister and Kilian (2013) conclude that the best way to pool individual forecasts of oil prices is by assigning equal weight to them. This approach is more accurate than basing weights on the recent forecasting performance of each model.

Table 1: Real-time forecast accuracy of pooled forecasts with equal weights

Horizon (in months)	Real U.S. refiner acquisition cost (RAC) for crude oil imports	Real West Texas Intermediate (WTI) price
	Relative mean-squared prediction error (MSPE) ratios	
1	**0.897**	**0.880**
3	**0.874**	**0.873**
6	**0.949**	**0.956**
9	**0.939**	**0.943**
12	**0.892**	**0.902**
15	**0.893**	**0.906**
18	**0.957**	**0.959**
21	1.065	1.064
24	1.029	1.017
	Success ratios	
1	**0.554***	**0.517**
3	**0.609***	**0.592***
6	**0.556**	**0.543**
9	**0.580****	**0.562**
12	**0.609***	**0.605***
15	**0.650***	**0.617***
18	**0.601***	**0.577****
21	**0.550**	**0.550**
24	**0.561**	**0.551**

Notes: Numbers in **bold** indicate improvements relative to the no-change forecast.
* denotes significance at the 5 per cent level and ** at the 10 per cent level, based on the Pesaran and Timmermann (2009) test for the null hypothesis of no directional accuracy. The statistical significance of the MSPE reductions cannot be assessed because none of the currently available tests of equal predictive accuracy applies in this setting.
Source: Baumeister and Kilian (2013)

than a coin toss, with a success probability of 0.5 (or 50 per cent). Thus, success ratios above 0.5 indicate an improvement relative to the no-change forecast.

The forecasts are generated for two alternative measures of the real price of crude oil. The first is the U.S. refiner acquisition cost (RAC) of imported crude oil, which is considered a good proxy for a truly global oil price. The second is the West Texas Intermediate (WTI) spot price, which is the U.S. benchmark. Equal-weighted forecast combinations for the real RAC and WTI yield considerable reductions in MSPE ratios for horizons up to 18 months, ranging from 4 per cent to 13 per cent (**Table 1**).[2] These forecast combinations are also successful at predicting the direction of change for these horizons. For the RAC, improvements in directional accuracy are statistically significant at all but one horizon up to 18 months, and range from 55 per cent to 65 per cent. For the WTI, the highest success ratio is 62 per cent, but the directional accuracy is statistically significant at only four of these horizons.

2 Chart A-1 in the appendix shows the evolution of the recursive root-mean-squared prediction errors for the equal-weighted combination forecast in comparison with the no-change forecast for selected forecast horizons.

A Structural Model for Assessing Risks to Oil Prices

An important limitation of these forecasting approaches is that they do not help policy-makers explore how the forecast would change relative to the baseline forecast under hypothetical assumptions about future economic conditions. For example, it is important to know how much the real price of oil would be affected by civil unrest in the Middle East or by a period of unexpectedly low global demand for crude oil caused by a worldwide recession. Similarly, policy-makers want to understand what drove changes in the real price of oil in the recent past, such as the persistent increase in oil prices between 2003 and 2008, or the rapid decline between 2008 and 2012. To address questions such as these, a structural model of the global oil market is required.

A structural model of the global oil market

Using a dynamic structural model of the global oil market, it is possible to decompose past fluctuations in oil prices into structural driving forces stemming from supply and demand

Kilian and Murphy (2014) propose a dynamic structural model that includes the key determinants of the real price of oil: changes in global oil production, real economic activity worldwide and above-ground crude oil inventories. This econometric model is motivated by a standard stock-flow model with an explicit role for expectations and can be directly derived from the forecasting model of the global oil market by imposing additional economic structure. Within this framework, it is possible to decompose past fluctuations in oil prices into structural driving forces stemming from supply and demand. In particular, the authors distinguish between four types of shocks:

(i) **Flow supply shock**—a classic oil supply shock that captures disruptions to the flow of oil production resulting, for example, from exogenous political events in oil-producing countries, such as war or civil unrest.

(ii) **Flow demand shock**—a shock to the demand for crude oil that is associated with unexpected fluctuations in the global business cycle. A prominent example is the surprisingly persistent demand from emerging-market economies, particularly China.

(iii) **Speculative demand shock**—a shock to the demand for oil inventories arising from shifts in expectations about future demand for and supply of oil that is not otherwise captured in the model. Such shifts could arise, for example, from the anticipation of several factors, including political unrest in oil-exporting countries in the Middle East, peak oil effects or the depletion of oil reserves. This shock thus captures forward-looking behaviour and speculation.

(iv) **Other demand shock**—a residual shock that has no economic interpretation but is designed to capture idiosyncratic oil demand shocks not otherwise accounted for. Examples of such shocks include weather shocks, changes in inventory technology or preferences, or politically motivated releases of the U.S. Strategic Petroleum Reserve.

This model provides a coherent framework both to understand past oil price fluctuations and to assess risks associated with oil price forecasts.

The contribution of each shock to cumulative price changes

Kilian and Lee (2014) use this model to quantify the contribution of each type of shock to cumulative changes in the real U.S.-dollar price of oil during specific historical episodes.

Chart 2: Contribution of each structural shock to the cumulative change in the real price of oil

a. January 2003 to June 2008, in May 2012 U.S. dollars per barrel

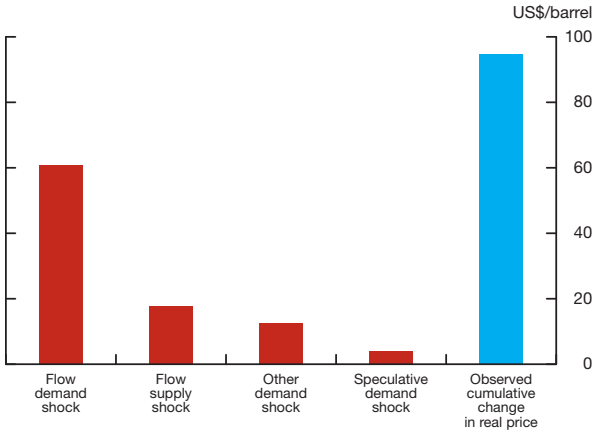

b. June 2008 to May 2012, in May 2012 U.S. dollars per barrel

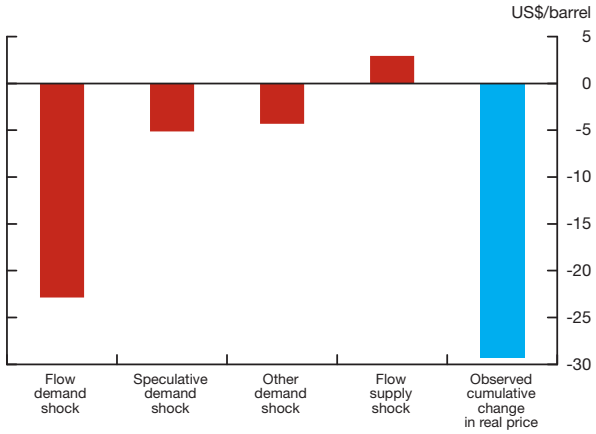

Source: Kilian and Lee (2014)

From 2003 to mid-2008, oil prices experienced an unprecedented surge. This development triggered a debate about whether the run-up in oil prices could be explained by elevated demand from China, or whether it was the result of the financialization of physical oil markets.[3]

Chart 2a provides compelling evidence that an unexpectedly strong world economy was the main cause for the rise in global oil prices. In fact, flow demand shocks associated with shifts in the global demand for oil from emerging Asia and from member countries of the Organisation for Economic Co-operation and Development accounted for US$60 of the observed US$95 increase in the real per-barrel price of oil during that period. While supply-side factors contributed somewhat to the upswing in oil prices, they accounted for less than US$20 of the increase. Speculation by oil consumers, in contrast, was negligible.

A similar picture emerges for the decline in the price of oil between mid-2008 and 2012. **Chart 2b** shows that most of the US$29 decrease in the real per-barrel price of oil can be attributed to a series of unexpected negative flow demand shocks associated with the weak global economy in the wake of the financial crisis. As before, other factors played a limited role. These findings suggest that changes in the demand for crude oil associated with the global business cycle are the primary determinant of changes in oil prices.

From explaining the past to assessing future risks

Over the projection horizon, unpredictable variations in the demand for and supply of crude oil can lead to deviations of the future oil price from its forecasted path. It is therefore useful to assess the sensitivity of the baseline forecast to potential events involving future demand and supply conditions in the crude oil market.

3 Financialization refers to the large increase in investors' participation in commodities as an asset class, as reflected, for instance, in the inflow of investment funds to oil futures markets in the past decade. This trend has led to a debate about the possible influence of financialization on oil price dynamics.

Chart 3: Scenarios for the projected path of the real price of oil, in December 2010 U.S. dollars per barrel

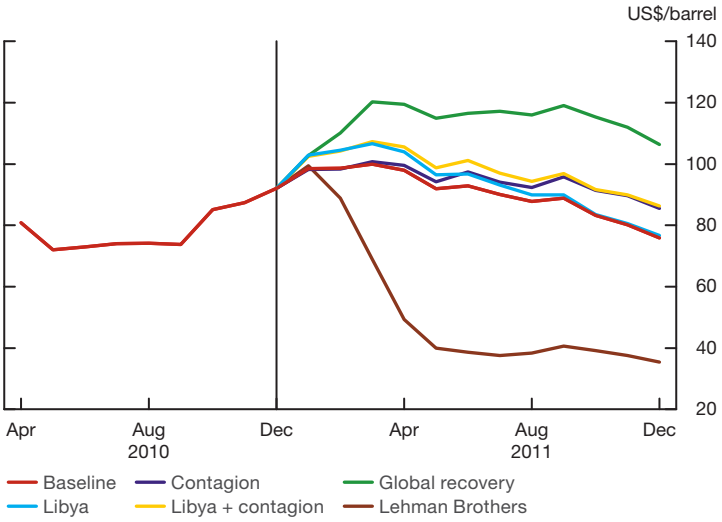

Note: The red line is the real-time out-of-sample forecast for the real U.S. refiner acquisition cost of crude oil in December 2010 U.S. dollars. The vertical line indicates the point in time when the forecast is made.
Source: Baumeister and Kilian (2014a)

Alternative forecast scenarios based on hypothetical sequences of oil demand and supply shocks help policy-makers gauge the possible consequences of unlikely events

To model such departures from the baseline forecast, Baumeister and Kilian (2014a) present alternative forecast scenarios based on the structural model of the oil market (Kilian and Murphy 2014). These scenarios examine the percentage deviation from the baseline forecast if a certain sequence of oil demand or supply shocks were to occur over the projection horizon. They are intended to help policy-makers gauge the possible consequences of unlikely events.

The baseline oil price forecast is generated as of December 2010 (**Chart 3**). To this baseline forecast we add five scenarios taken from Baumeister and Kilian (2014a). The first scenario relates to the supply side of the oil market and is motivated by the political unrest in Libya in early 2011. The authors ask what would have happened to the real price of oil if Libyan production, which accounts for 2.2 per cent of global oil production, were unexpectedly taken off the market. The results from the model show that such a shortfall in Libyan production would raise the price of oil by only 7 per cent after three months. This example illustrates that the observed increase in the price of oil of 21 per cent over that same period (**Chart 1**) cannot be attributed to supply disruptions alone.

Events such as the Arab Spring or the ongoing civil war in Syria can affect the oil price by triggering speculative demand, driven by fears of contagion of political unrest in the Middle East. Such an expectations-driven contagion scenario would increase the real price of oil by 20 per cent after about a year and a half, if the shift in speculative demand were comparable with the sustained speculative frenzy that began in mid-1979 following the Iranian revolution. The third scenario is a combination of the previous two scenarios.

The fourth and fifth scenarios relate to the role of the global business cycle. The global recovery scenario illustrates that an unexpected full recovery of the world economy would raise the real price of oil by an additional 40 per cent after about one year. The prospect of a global collapse shows

that the recurrence of an event such as the financial crisis following the bankruptcy of Lehman Brothers in 2008 would be expected to lower the real price of oil by close to 60 per cent, as global demand drops dramatically.

For expository purposes, it is assumed that all of the scenarios begin in January 2011. Each scenario results in a different projected path for the real per-barrel price of oil, providing the full range of alternative outcomes. The real price of oil may fall as low as US$69 or rise as high as US$120 after one quarter, depending on the scenario. After one year, the range is between US$35 and US$106. Consistent with earlier results, the more extreme movements correspond to scenarios with large shifts in flow demand.

Obviously, policy-makers will not consider all scenarios equally likely; some scenarios will be mutually exclusive, while others might occur in conjunction. Assessing by how much such an alternative path deviates from the baseline, and how sensitive this deviation is to alternative assumptions about the relative likelihood of the underlying scenarios, allows policy-makers to get a better sense of the nature of the upside and downside risks involved. This information can also be used as input into more comprehensive risk scenarios that policy-makers might use to assess potential macroeconomic outcomes.

Conclusion

Combinations of forecasts generated by different models are a useful tool for obtaining more accurate and robust out-of-sample forecasts of the real price of oil. These baseline forecasts can be supplemented by forecast scenarios from a structural model of the global oil market to evaluate upside and downside risks at various horizons. Such an approach is important because central bankers care not only about forecast accuracy but also about the economic interpretation underlying the past, present and future evolution of the real price of oil.

Appendix 1

Evolution of Root-Mean-Squared Prediction Errors

Chart A-1: Recursive root-mean-squared prediction errors for the combination forecast with equal weights and for the no-change forecast for horizons of 1, 12 and 24 months

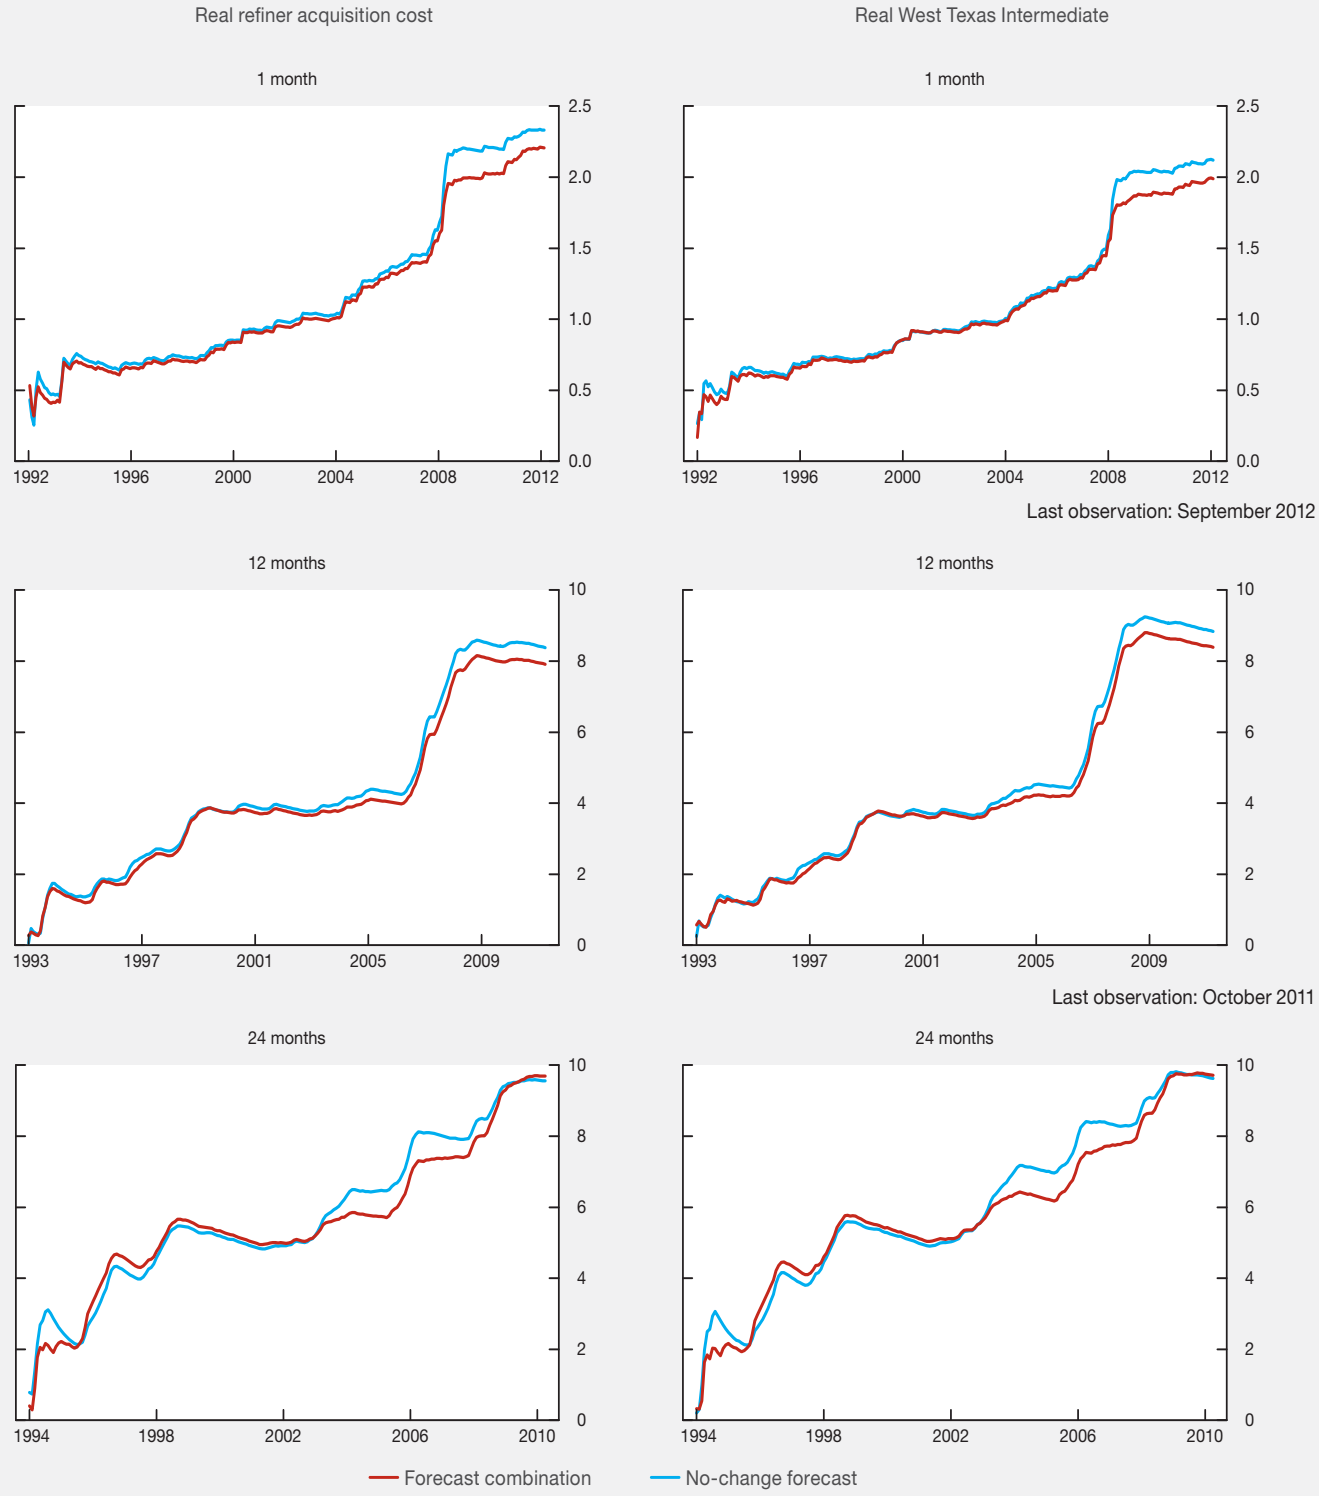

Source: Bank of Canada calculations

Last observation: October 2010

Literature Cited

Alquist, R. and E. Arbatli. 2010. "Crude Oil Futures: A Crystal Ball?" *Bank of Canada Review* (Spring): 3–11.

Alquist, R., L. Kilian and R. J. Vigfusson. 2013. "Forecasting the Price of Oil." In *Handbook of Economic Forecasting*, 2A, edited by G. Elliott and A. Timmermann, 427–507. Amsterdam: North-Holland.

Baumeister, C. and L. Kilian. 2012. "Real-Time Forecasts of the Real Price of Oil." *Journal of Business and Economic Statistics* 30 (2): 326–36.

———. 2013. "Forecasting the Real Price of Oil in a Changing World: A Forecast Combination Approach." Bank of Canada Working Paper No. 2013-28.

———. 2014a. "Real-Time Analysis of Oil Price Risks Using Forecast Scenarios." *IMF Economic Review* (forthcoming).

———. 2014b. "What Central Bankers Need to Know About Forecasting Oil Prices." *International Economic Review* 55 (3): 869–89.

Baumeister, C., L. Kilian and X. Zhou. 2013. "Are Product Spreads Useful for Forecasting? An Empirical Evaluation of the Verleger Hypothesis." Bank of Canada Working Paper No. 2013-25.

Hamilton, J. D. and J. C. Wu. 2014. "Risk Premia in Crude Oil Futures Prices." *Journal of International Money and Finance* 42: 9–37.

Kilian, L. and T. K. Lee. 2014. "Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories." *Journal of International Money and Finance* 42: 71–87.

Kilian, L. and D. P. Murphy. 2014. "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil." *Journal of Applied Econometrics* 29 (3): 454–78.

Pesaran, M. H. and A. Timmermann. 2009. "Testing Dependence Among Serially Correlated Multicategory Variables." *Journal of the American Statistical Association* 104 (485): 325–37.

Measuring Uncertainty in Monetary Policy Using Realized and Implied Volatility

Bo Young Chang and Bruno Feunou, Financial Markets Department

- Measuring the degree of uncertainty in the financial markets around future monetary policy rates and market interest rates is important because interest rate uncertainty affects the real economy through the investment and hiring decisions of firms.
- This article assesses uncertainty surrounding future policy rates set by the Bank of Canada using measures of *realized volatility* computed from the intraday prices of interest rate futures and *implied volatility* computed from the prices of options on interest rate futures.
- According to these measures, interest rate uncertainty decreased following major policy actions that the Bank took in response to the 2007–09 financial crisis.
- Findings also indicate that, on average, uncertainty decreases following the Bank's policy rate announcements.

Central banks monitor various financial variables, such as short-term interest rates and the prices of interest rate derivatives, to gauge market expectations of future monetary policy. The expectations extracted from these variables can then be used to complement other tools adopted by the Bank of Canada to assess the impact of monetary policy. The predicted *level* of future policy rates, captured by these measures of expectations, has a well-known impact on the real economic activity, but *uncertainty* about future monetary policy rates also affects economic activity through the investment and hiring decisions of firms (Bernanke 1983).

The most widely used measure of uncertainty around future policy rates is the dispersion of professional forecasts based on surveys. Although survey-based measures are good indicators of uncertainty, they have several limitations. First, since surveys are conducted infrequently, the uncertainty measures cannot be used to assess the immediate effect of central bank announcements or other macroeconomic events. Second, survey-based measures of uncertainty are based on the opinions of a small number of market participants and may not be reflective of the larger population. Third,

since survey respondents are professionals in the fields of economics and finance, their forecasts are potentially affected by reputational concerns. For example, respondents may be reluctant to deviate too far from consensus for fear of having a wrong forecast when others have the right one, thus harming their reputation. To overcome these limitations, we propose alternative measures based on the prices of interest rate futures and options.

This article describes two measures of the price volatility of interest rate futures that could be used as indicators of policy rate uncertainty: realized volatility, computed from the intraday prices of interest rate futures, and implied volatility, computed from the prices of options on interest rate futures. We construct our volatility measures using futures contracts on the average three-month bankers' acceptance rate (BAX) and options on BAX (OBX).

Using the two measures of policy rate uncertainty, we first analyze the effect of various policy actions taken by the Bank of Canada during and following the 2007–09 financial crisis, and then examine the effect of the Bank's policy rate announcements that took place on fixed announcement dates (FADs) over the 2002–13 period. We examine four crisis-related policy actions that the Bank took between 2007 and 2010: (i) the first introduction of term purchase and resale agreements (PRAs) for liquidity purposes in December 2007; (ii) the second introduction of term PRAs in March 2008; (iii) the inter-meeting cut in the policy rate in October 2008 that was coordinated with other major central banks; and (iv) the conditional commitment to keep the policy rate unchanged that lasted several months in 2009 and 2010.

This article explains how realized volatility and implied volatility can be used to measure policy rate uncertainty in various applications relevant to monetary policy. It first describes briefly how these measures of volatility are calculated and how they can be interpreted. The article then reports the results of empirical applications of the measures of uncertainty.[1]

◀ *Realized volatility and implied volatility can be used to measure policy rate uncertainty in various applications relevant to monetary policy*

Measures of Uncertainty Around Future Policy Rates

The price of interest rate futures reflects expectations of future interest rates, so when uncertainty around future interest rates is high, this price exhibits large variations over time. The most widely used measure of price variation is the standard deviation of daily price changes, referred to as historical volatility. However, since historical volatility is calculated using past daily prices, it is inherently backward-looking. This article uses two alternative measures of price variation:

(i) realized volatility—the volatility of intraday price changes,[2] and

(ii) implied volatility—the volatility of the underlying price process that is implied by the prices of options.

Although historical volatility and realized volatility are both backward-looking, realized volatility is computed with data from a single day and is therefore much more sensitive to the arrival of new information, whereas historical volatility is computed using only past prices.

Implied volatility is a forward-looking measure of future interest rate volatility and thus reflects expectations of future interest rate volatility, while realized volatility is an ex post measure of current interest rate volatility that

1 For more information on how volatility was measured and the empirical tests, see Chang and Feunou (2013).

2 Measured as the square root of the sum of squared changes in high-frequency intraday prices.

estimates the interest rate uncertainty on a given day. Each measure has different applications. Realized volatility can be used to measure the actual change in interest rate uncertainty following a policy action, while implied volatility can be used to assess the effect of a policy announcement on the expected *future* interest rate volatility.

An increase in implied volatility can be due to either an increase in uncertainty or an increase in the risk premium required to compensate for interest rate uncertainty

A large proportion of time variations in both implied and realized volatility is determined by uncertainty around future interest rates. However, other factors also influence these measures. Implied volatility reflects both expected volatility and a volatility risk premium. An increase in implied volatility can therefore be due to either an increase in uncertainty or an increase in the risk premium required to compensate for interest rate uncertainty. Realized volatility is sensitive to noise in high-frequency prices caused by certain trading activities or restrictions, because the measure is computed using intraday prices. Since we cannot separate the effect of these factors from the measure reflecting only uncertainty, our results should be interpreted with these factors in mind.

In Canada, the collateralized overnight rate at which major financial institutions borrow and lend one-day funds among themselves is the main tool used by the Bank to conduct monetary policy. Since the Canadian overnight repo rate average (CORRA) tracks the central bank's policy rate closely, ideal instruments for our study would be futures and options on CORRA. However, futures on CORRA trade with limited liquidity, and options on CORRA have yet to be introduced. Instead, we use futures and options on the three-month bankers' acceptance rate—calculated from bid-side rates of the primary bankers' acceptance market and called the Canadian Dealer Offered Rate (CDOR)—to compute our volatility measures. This approach is consistent with the frequent use of eurodollar futures and options in related studies in the United States.[3]

Bankers' acceptances are tradable short-term corporate obligations that are backed by a line of credit and are therefore guaranteed by the accepting banks. Changes in bankers' acceptance rates are closely related to changes in overnight rates and are known to be good predictors of future policy rates (Johnson 2003). The CDOR is the key reference rate for short-term interbank funding in Canada, similar to the LIBOR (London Interbank Offered Rate) in the United States. Moreover, futures on the three-month bankers' acceptance rate, called BAX, are one of the most liquid instruments in the Canadian money market. Both BAX and options on BAX, which are less liquid, trade on the Montréal Exchange.

Realized volatility of BAX

The realized volatility of BAX is computed as the square root of the sum of squared changes in intraday BAX futures prices, observed at the interval of every five trades, roughly equivalent to every 20 minutes.

To illustrate how realized volatility captures uncertainty around future policy rates, **Chart 1** shows the implied yields[4] from intraday prices of a BAX contract on three consecutive days around 21 April 2009, the day the Bank lowered its policy rate from 0.50 per cent to 0.25 per cent and announced its conditional commitment to keep its policy rate at 0.25 per cent until the end of the second quarter of 2010. Panel b shows the implied yields of a BAX contract with a September 2009 expiry as of 21 April 2009, while panels a

3 See Abken (1995), Rigobon and Sack (2002), Neely (2005) and Bauer (2012).

4 The price of a BAX contract is quoted as 100 minus the annualized yield of a three-month Canadian bankers' acceptance rate. Thus, the implied yield is 100 minus the price of BAX.

Chart 1: Intraday BAX prices

Expiry: 14 September 2009

a. 20 April 2009: realized volatility = 0.69

b. 21 April 2009 fixed announcement date: realized volatility = 1.17

c. 22 April 2009: realized volatility = 0.40

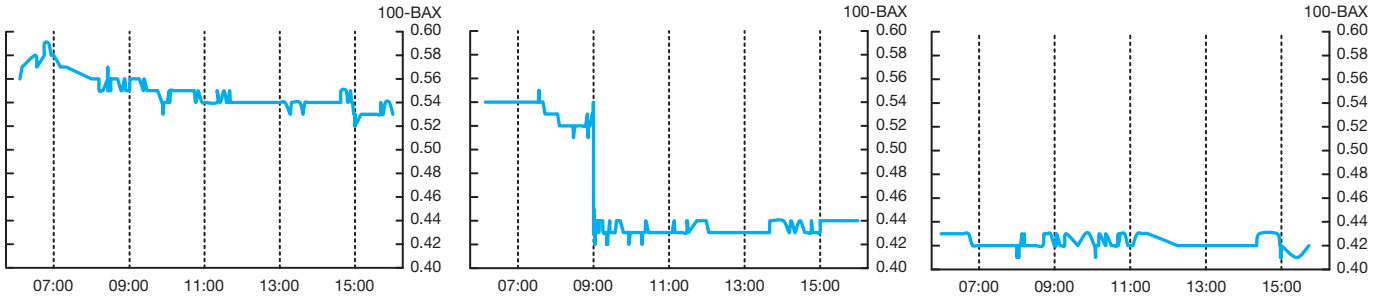

Source: Montréal Exchange

and c show the implied yields of the same BAX contract one day before and one day after 21 April, respectively. There is a large drop in the implied yield at 09:00 on 21 April 2009, the time of the Bank's policy rate announcement. This pattern is typical of FADs, reflecting a shift in expectations immediately following a policy announcement. Any large change (either up or down) in the BAX price leads to distinctly larger-than-average relative volatility on the FADs. We explore this issue in more detail in the section on policy rate announcements.

Implied volatility of BAX

We compute the implied volatility of BAX from OBX option prices using an option valuation formula based on the Vasicek interest rate model (Vasicek 1977).[5] Since the expiry dates of the options that are traded vary, the implied volatility can, in theory, be computed at different maturities and used as a measure of policy rate uncertainty at different horizons. However, the relatively low liquidity of OBX prevents us from computing the implied volatility of BAX for different maturities consistently over time. We therefore compute an average implied volatility for each day, using all options with maturities from one to six months.[6]

Chart 2 shows the daily time series for our measure of implied volatility. Note that the implied volatility cannot be computed between November 2008 and mid-March 2010, since there was no trading in OBX contracts during that period. This temporary halt in trading began in October 2008 at the peak of the financial crisis in the United States, possibly as a result of the increased margin requirements caused by the high volatility of the underlying interest rates and risk premiums at that time. This halt was not a result of actions taken by the exchange or the regulators. Trading in OBX resumed in mid-March 2010, about one month before the removal of the conditional commitment.

5 Since Vasicek's interest rate model allows negative interest rates, the implied volatility that we compute does not take into account the fact that the interest rate cannot fall below zero.

6 Chang and Feunou (2013) provide details on how the implied volatility is calculated.

Chart 2: Implied volatility and realized volatility of BAX

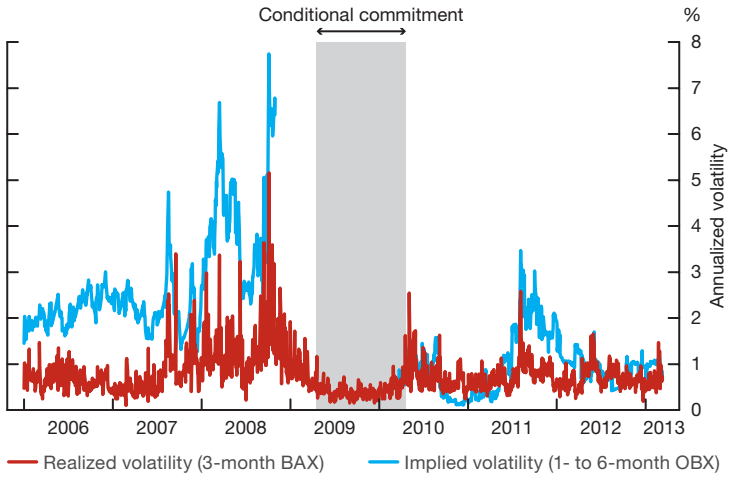

Sources: Montréal Exchange and Bank of Canada

Last observation: 13 March 2013

Impact of Bank of Canada's Policy Actions on Uncertainty

Conditional commitment in 2009 and 2010

The Bank's conditional commitment in 2009 and 2010 reduced a large amount of uncertainty around future policy rates for a specific period of time

The Bank's conditional commitment that lasted several months in 2009 and 2010 had clear implications for policy rate uncertainty, since the commitment reduced a large amount of uncertainty around future policy rates for a specific period of time. On 21 April 2009, the Bank lowered its policy rate from 0.50 per cent to 0.25 per cent and announced that, "Conditional on the outlook for inflation, the target overnight rate can be expected to remain at its current level until the end of the second quarter of 2010 in order to achieve the inflation target."[7] The commitment was eventually removed on 20 April 2010, and the policy rate was subsequently raised back to 0.50 per cent on the following FAD on 1 June, one month earlier than indicated in the initial conditional commitment. The introduction of the conditional commitment coincided with the decision to lower the interest rate to the effective lower bound of 0.25 per cent, at which point interest rates could be moved only upward. In principle, both the reduction of the policy rate to the effective lower bound and the conditional commitment would reduce uncertainty about future policy rates. This section assesses whether uncertainty did in fact decrease during the conditional commitment period.

Chart 2 shows the realized volatility and the implied volatility from January 2006 to March 2013. We choose the three-month maturity for realized volatility because the average maturity of options used in the computation of implied volatility is around three months. Since the asset underlying both BAX futures and OBX options is the three-month CDOR, the time horizon of uncertainty for both measures of volatility reflects uncertainty around the policy rate approximately three to six months ahead.[8]

7 See the Bank of Canada press release at http://www.bankofcanada.ca/2009/04/publications/press-releases/fad-press-release-2009-04-21/.

8 For example, we compute the realized volatility in January 2013 using the BAX contract expiring in March 2013. The final settlement price of this contract is the three-month CDOR on 18 March 2013. Since the three-month CDOR in March reflects expected policy rates between March and June (plus a small spread), the price of the March-expiry BAX in January also reflects these expected policy rates. Therefore, the horizon of uncertainty of the realized volatility computed in January using the prices of BAX expiring in March is between three and six months.

As expected, during the period of the conditional commitment, uncertainty about future policy rates was extremely low, as indicated by the low level of realized volatility. Furthermore, although a measure of implied volatility is not available throughout the conditional commitment period, its level at the time when trading resumed is significantly lower than it was when trading halted in late 2008. This low level is consistent with the decreased uncertainty exhibited by the low level of realized volatility throughout the conditional commitment period.

The timing of the resumption of trading in OBX contracts also provides an interesting insight into the market's expectations regarding the timing of the removal of the conditional commitment. Trading of OBX contracts with a 14 June 2010 expiry date resumed in mid-March 2010, one month before the removal was announced. The fact that options started trading even though they had an expiry date before the end date of the commitment suggests that the market anticipated a possible early removal of the commitment before the actual announcement was made.

The Bank's removal of the conditional commitment resulted in a large increase in the level of both implied and realized volatility

The Bank's removal of the conditional commitment resulted in a large increase in the level of both implied and realized volatility relative to that observed during the conditional commitment period. While realized volatility increased to a level comparable with that observed before the crisis, implied volatility remained much lower compared with both the crisis period and the pre-crisis period. The extremely low level of implied volatility is a result of both little interest rate uncertainty, as reflected in low realized volatility, and a small volatility risk premium, which is measured as the difference between implied volatility and realized volatility. The near-zero volatility risk premium indicates that investors in the post-crisis period did not require a high premium for bearing the risk of interest rate volatility, possibly because any indication of the interest rate moving away from the effective lower bound, and thus increasing interest rate volatility, was linked to an improving rather than deteriorating economic outlook during this period.

Announcement effects of crisis-related policy actions

Announcements of other crisis-related policies led to large decreases in implied volatility

Unlike the conditional commitment, the implication of other crisis-related policies on policy rate uncertainty is not clear a priori. This section examines the impact of these policy actions on uncertainty, as reflected in the change in the implied volatility on the days when policy announcements were made. We examine three crisis-related policy announcements made by the Bank in 2007 and 2008: (i) the first introduction of term PRAs for liquidity purposes in December 2007; (ii) the second introduction of term PRAs in March 2008; and (iii) the intermeeting cut in the policy rate in October 2008 that was coordinated with other major central banks. All of these announcements led to large decreases in implied volatility, ranging between 12 and 51 basis points (**Table 1**).

By far, the largest drop in implied volatility—51 basis points—occurred on 8 October 2008, when, in an intermeeting announcement, the policy rate was cut by 50 basis points, in conjunction with similar moves announced by other central banks. This rate reduction occurred around the height of the crisis, when the implied volatility of BAX had reached its highest level in our sample period, as shown in **Chart 2**. The result is consistent with Bauer (2012), who finds that many of the important announcements made by the Federal Reserve during the financial crisis also led to a larger-than-average drop in the implied volatility of eurodollar futures options. The strong market

Table 1: Implied volatility following crisis-related Bank of Canada policy announcements

Date	Policy announcement	Daily change in implied volatility (basis points)
12 December 2007	First term purchase and resale agreements (PRAs) announced for liquidity purposes (in coordination with similar actions by the Bank of England, the European Central Bank, the Federal Reserve and the Swiss National Bank) (not a FAD[a])	-12
11 March 2008	Second term PRAs announced for liquidity purposes, coordinated with other G-10 central banks (not a FAD)	-25
8 October 2008	Intermeeting cut in target rate, coordinated with other major central banks (not a FAD)	-51
	Average—all days	0
	Average—FADs	-5
	Average—non-FADs	1

a. FAD: fixed announcement date, the Bank of Canada's pre-scheduled policy rate announcement date

reaction indicates that this coordinated action by the central banks was effective at reducing uncertainty around the policy rate and the risk premium associated with the uncertainty at that time.

The Bank announced the term PRA facility on 12 December 2007 as part of a broader initiative that included similar programs announced by other central banks worldwide to alleviate pressures in short-term funding markets.[9] Enenajor, Sebastian and Witmer (2010) find evidence that the term PRA announcements did indeed reduce bank short-term funding costs in Canada. Zorn, Wilkins and Engert (2009) also suggest that the availability of the Bank's extraordinary liquidity facilities such as the term PRAs may have mitigated stresses in the bank funding market in Canada. Since the introduction of a term PRA facility is not directly related to the policy rate, a large decrease in implied volatility observed on the PRA announcement days seems to be related to a lower risk premium rather than lower uncertainty around future policy rates. A lower risk premium is consistent with reduced stress in the short-term bank funding market found in the previous studies.

Policy rate announcements on fixed announcement dates

Many central banks make their policy rate announcements only on pre-scheduled dates to reduce uncertainty about the timing of changes to the policy rate, a practice that is consistent with increased transparency in central bank communications.[10] After each policy rate announcement, central banks are interested in assessing the impact of their decision on the market. Typically, the financial variables they monitor include yields on various fixed-income securities and foreign exchange rates. However, these variables do not indicate whether a particular decision has increased or decreased uncertainty around future policy rates.

To determine whether uncertainty increased or decreased, we look at whether any pattern emerges in the measures of realized and implied volatility around the days of the policy rate announcements. Since the policy rate is fixed until the next FAD, typically six to eight weeks away, we would

9 See the Bank of Canada press release at http://www.bankofcanada.ca/2007/12/bank-canada-temporarily-expands-list-securities/.

10 For further discussion on the impact of fixed announcement dates and forward-looking guidance on uncertainty, see Parent, Munro and Parker (2003) and Fay and Gravelle (2010).

Chart 3: Average level of realized and implied volatility of BAX around policy rate announcement days

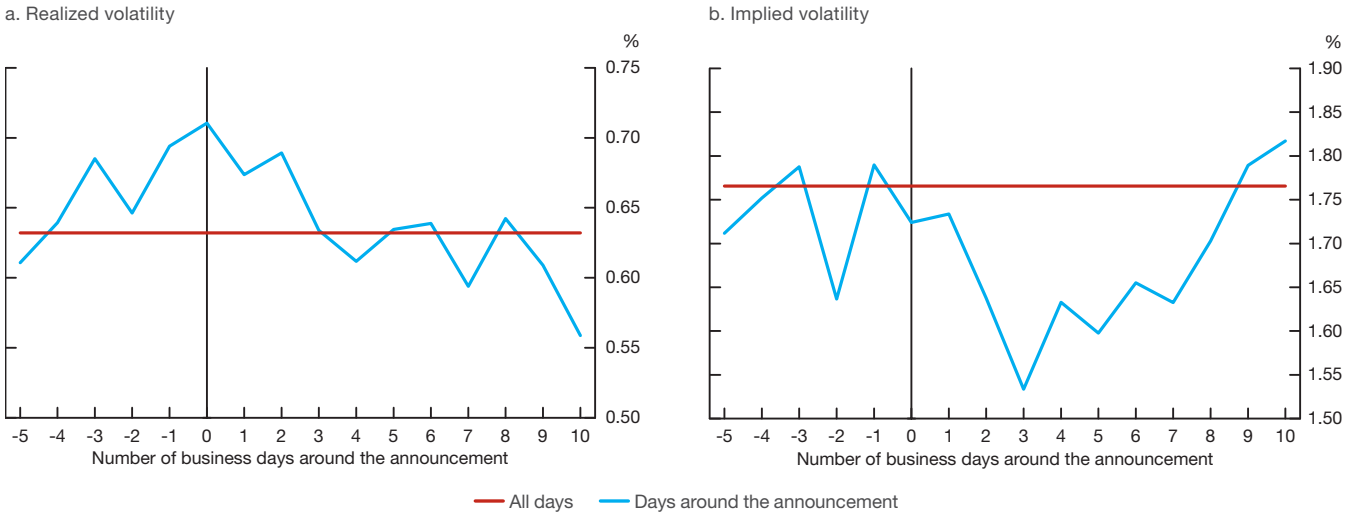

Sources: Montréal Exchange and Bank of Canada

expect any measure of uncertainty over a shorter horizon to drop sharply after each announcement. However, if the horizon extends beyond the following announcement day, as is the case for realized volatility and implied volatility, which have horizons of three to six months, uncertainty can either increase or decrease after a policy rate announcement.

Chart 3 shows the average levels of realized volatility and implied volatility between five business days before and ten business days after policy rate announcements. For this analysis, we compute realized volatility using only trade prices after 09:15 to remove the effect of large changes in BAX prices observed immediately following the policy rate announcements (**Chart 1**). The analysis is based on 90 FADs between January 2002 and March 2013 for realized volatility, and 58 FADs between January 2006 and March 2013 for implied volatility.

The realized volatility is higher than average (with statistical significance) on the FAD and the day before the FAD. It then decreases gradually over the next two weeks. However, this analysis shows that realized volatility is significantly lower than average on only two days within the two-week window, seven and ten business days after the FAD. Thus, based on realized volatility, we find only weak evidence that uncertainty decreases following a policy rate announcement.

The results for implied volatility provide much stronger evidence that uncertainty decreases on average following a policy rate announcement. A regression analysis of the level of implied volatility shows that it is significantly lower than average between two and seven business days after an announcement. A regression analysis of the change in implied volatility shows that statistically significant drops in implied volatility occur on an announcement day and two business days after an announcement. The decrease in implied volatility two business days after an announcement can be explained by the fact that, during a large part of the sample period, the Bank released its *Monetary Policy Report* two days after every other policy rate announcement.[11]

11 Between October 2010 and December 2012, the Bank published its *Monetary Policy Report* one day after its January, April, July and October FAD announcements, and since January 2013, it has published the *Report* concurrently with these FAD announcements.

Our results show that, on average, the Bank's policy rate decisions reduced uncertainty around future policy rates in our sample period. A statistically significant reduction in uncertainty is observed on the days of a policy rate announcement and the release of the *Monetary Policy Report*. The effect of the reduction in uncertainty, however, seems to be temporary, lasting for seven business days, on average. This gradual increase in uncertainty following the initial decrease after a policy rate announcement is reasonable, given that new information and new events tend to increase uncertainty.

Our results show that, on average, the Bank's policy rate decisions reduced uncertainty around future policy rates

Conclusion

Realized volatility computed from the intraday prices of interest rate futures and implied volatility computed from interest rate futures options are useful indicators of uncertainty around future central bank policy interest rates. Based on implied volatility computed from OBX options, we show that, on average, policy rate announcements by the Bank of Canada reduced uncertainty around future policy rates. We also examine the effect of major policy actions that the Bank took in response to the 2007–09 financial crisis. We find that realized volatility was extremely low during the conditional commitment period in 2009 and 2010. Also, the introduction of term PRAs and the intermeeting cut in the policy rate coordinated with other major central banks both resulted in a large drop in implied volatility on the announcement days, indicating that these announcements reduced uncertainty around future policy rates and/or reduced risk premiums on interest rate uncertainty.

Literature Cited

Abken, P. A. 1995. "Using Eurodollar Futures Options: Gauging the Market's View of Interest Rate Movements." Federal Reserve Bank of Atlanta *Economic Review* 80 (2): 10–30.

Bauer, M. D. 2012. "Monetary Policy and Interest Rate Uncertainty." Federal Reserve Bank of San Francisco *Economic Letter* No. 2012-38.

Bernanke, B. S. 1983. "Irreversibility, Uncertainty, and Cyclical Investment." *Quarterly Journal of Economics* 98 (1): 85–106.

Chang, B. Y. and B. Feunou. 2013. "Measuring Uncertainty in Monetary Policy Using Implied Volatility and Realized Volatility." Bank of Canada Working Paper No. 2013-37.

Enenajor, E., A. Sebastian and J. Witmer. 2010. "An Assessment of the Bank of Canada's Term PRA Facility." Bank of Canada Working Paper No. 2010-20.

Fay, C. and T. Gravelle. 2010. "Has the Inclusion of Forward-Looking Statements in Monetary Policy Communications Made the Bank of Canada More Transparent?" Bank of Canada Discussion Paper No. 2010-15.

Johnson, G. 2003. "Measuring Interest Rate Expectations in Canada." *Bank of Canada Review* (Summer): 17–27.

Neely, C. J. 2005. "Using Implied Volatility to Measure Uncertainty About Interest Rates." Federal Reserve Bank of St. Louis *Review* 87 (3): 407–25.

Parent, N., P. Munro and R. Parker. 2003. "An Evaluation of Fixed Announcement Dates." *Bank of Canada Review* (Autumn): 3–11.

Rigobon, R. and B. Sack. 2002. "The Impact of Monetary Policy on Asset Prices." *Journal of Monetary Economics* 51 (8): 1553–75.

Vasicek, O. 1977. "An Equilibrium Characterization of the Term Structure." *Journal of Financial Economics* 5 (2): 177–88.

Zorn, L., C. Wilkins and W. Engert. 2009. "Bank of Canada Liquidity Actions in Response to the Financial Market Turmoil." *Bank of Canada Review* (Autumn): 3–22.

Beyond the Unemployment Rate: Assessing Canadian and U.S. Labour Markets Since the Great Recession

Konrad Zmitrowicz, International Economic Analysis, and Mikael Khan, Canadian Economic Analysis

- Labour market conditions are important for assessing economic well-being and are crucial for informing the conduct of monetary policy. This article uses several measures of labour market activity to provide a broad perspective on the performance of the labour market in Canada and the United States since the Great Recession of 2007–09.
- The article highlights the importance of considering a broad range of information in assessing the state of the labour market and also presents a simple way to summarize much of this information in a single composite labour market indicator (LMI) for both countries.
- The LMI suggests that the unemployment rate in Canada has evolved largely in line with overall labour market conditions since the recession, but may have modestly overstated the extent of recent improvement. This contrasts with the United States, where the unemployment rate appears to have significantly overstated the improvement in broader labour market conditions.

The Great Recession of 2007–09[1] had severe consequences in both Canada and the United States, including significant net job losses, totalling 430,000 in Canada and 8.7 million in the United States. Fortunately, significant progress has been made since the crisis. The Canadian job market has proved to be particularly resilient, recovering the number of jobs it lost during the recession and adding about 600,000 more. As of December 2013, the United States had regained only about 85 per cent of its job losses. Although unemployment rates in both countries are down significantly from the sharp increases seen during the recession, the recovery in labour markets remains

1 The National Bureau of Economic Research dates the recession in the United States as having started in December 2007 and ended in June 2009, while the C.D. Howe Institute dates the recession in Canada as having started in November 2008 and ended in May 2009.

incomplete. For example, an unusually large share of the unemployed have been out of work for six months or more, and many workers who would like to work full time have been able to obtain only part-time employment.

This article attempts to sort through these various signals to create a more comprehensive picture of labour market conditions since the recession. Assessing the health of the labour market is important, since it provides a measure of economic well-being. It is also crucial for the conduct of monetary policy. In Canada, the Bank of Canada's monetary policy goal is defined by its inflation-control target, and inflationary pressures are in part determined by labour market conditions. In the United States, labour market outcomes are an explicit component of the Federal Reserve's dual mandate of achieving maximum employment and price stability. More recently, the Federal Reserve has also tied its unconventional monetary policy programs, such as quantitative easing and forward guidance, to labour market outcomes.

Assessing the health of the labour market is important, since it provides a measure of economic well-being and is crucial for the conduct of monetary policy

The article first presents a broad set of labour market measures, focusing on their behaviour since the Great Recession.[2] It includes measures that capture various facets of the labour market to highlight the importance of considering a broad range of information when assessing its health. The article also presents a simple way to condense much of this information into a single composite labour market indicator (LMI) for both Canada and the United States. The LMI provides a simple benchmark against which to assess whether the unemployment rate—the most widely cited measure of the state of the labour market—is evolving in a manner consistent with broader labour market conditions. In Canada, the post-recession behaviour of the unemployment rate seems to have been largely representative of overall labour market conditions, although it may have modestly overstated the extent of recent improvement. The U.S. unemployment rate, in contrast, appears to have substantially overstated the post-recession improvement in labour market conditions, highlighting the need to consider a broad range of labour market variables.[3]

Measures of the Health of the Labour Market

The unemployment rate

The unemployment rate is the percentage of the labour force that does not have a job and is actively looking for work.[4] In both Canada and the United States, the unemployment rate increased sharply during the recession (**Chart 1**). In Canada, it rose from 5.9 per cent in February 2008 to a peak of 8.7 per cent in August 2009, while in the United States it rose from 4.4 per cent in May 2007 to a peak of 10 per cent in October 2009. Since 2010, unemployment rates have gradually fallen in both Canada and the United States, reaching 7.2 per cent and 6.7 per cent, respectively, in December 2013.

Definitional differences between the Canadian and U.S. unemployment rates make direct comparisons somewhat difficult. In Canada, the official unemployment rate is based on a working-age population that starts at age 15,

2 In cases where the available data are not seasonally adjusted, the authors use the U.S. Census Bureau's X-12-ARIMA approach to seasonally adjust the data.

3 Over time, labour market variables can be affected by country-specific structural and institutional factors. However, such factors are beyond the scope of this article.

4 The labour force is the total number of employed and unemployed. The employed are individuals who have a job or business, while the unemployed are those without work but who are available for work and are actively seeking work.

Chart 1: Unemployment rates

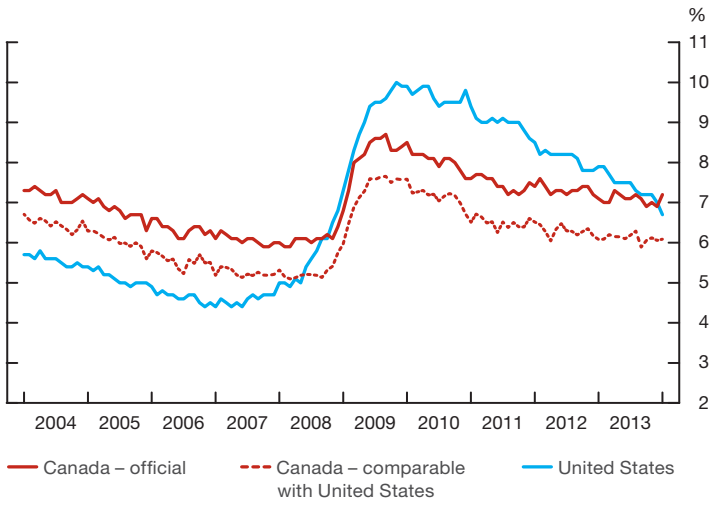

Sources: Statistics Canada and U.S. Bureau of Labor Statistics Last observation: December 2013

whereas in the United States the working age begins at age 16. There are also conceptual differences. For example, individuals who conduct their search for work by merely reading newspaper ads (passive job seekers) are considered unemployed in Canada but are not included in the labour force in the United States.[5] Looking only at the official measures in December 2013, it would appear that the unemployment rate was lower in the United States than in Canada (**Chart 1**). Once adjusted to the U.S. definition, however, the Canadian unemployment rate is in fact lower than in the United States. Note that the gap between the two has been narrowing, since the unemployment rate has fallen at a faster pace in the United States than in Canada.

The unemployment rate does not fully capture the extent of labour underutilization present in the economy...

Although the unemployment rate contains important information about the labour market, it may not be sufficient for gauging overall labour market conditions (Erceg and Levin 2013). The unemployment rate does not fully capture the extent of labour underutilization present in the economy, for example, if workers are discouraged from entering or remaining in the labour force, or if they are working less than they would like to.

...and is best seen within the context of a broad range of indicators

Thus, the unemployment rate is best seen within the context of a broad range of indicators. This article presents seven additional measures of the labour market to provide a broader perspective of underlying labour conditions since the recession, balancing the need to find variables that incorporate important aspects of the labour market with the need for data availability across Canada and United States.

An alternative measure of labour underutilization

Alternative measures of labour underutilization supplement the unemployment rate with broader definitions of joblessness. The most comprehensive measure combines discouraged and marginally attached individuals, as well as involuntary part-time workers, with the unemployed.[6] Because of its construction, the underutilization rate tends to be higher than the official unemployment rate, but the two usually move in tandem. It is therefore instructive to see whether the two measures have followed different trajectories since the recession.

5 See Sorrentino (2000) for further details.

6 Statistics Canada labels this measure R8, and the U.S. Bureau of Labor Statistics labels it U-6.

Chart 2: Unemployment and underutilization rates
Index: December 2007 (start of the recession in the United States) = 100

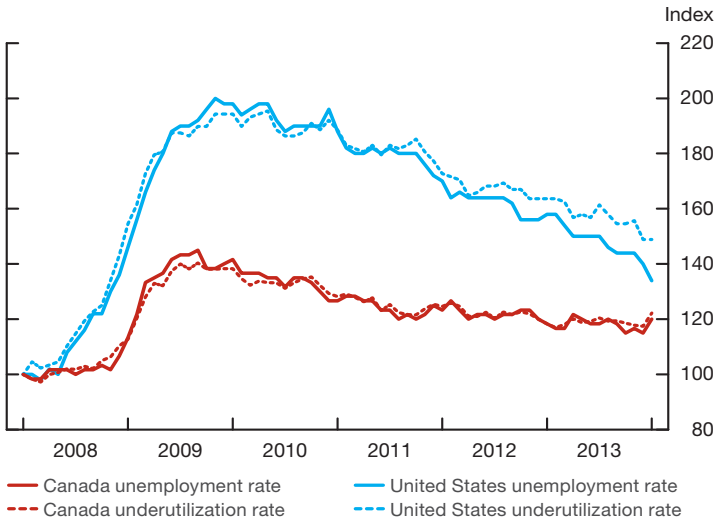

Sources: Statistics Canada and U.S. Bureau of Labor Statistics Last observation: December 2013

By indexing the unemployment and underutilization rates to the same point in time, we can see more easily how the two have recently evolved (**Chart 2**). In Canada, the two measures are virtually indistinguishable over the recession and subsequent recovery. In contrast, the U.S. underutilization rate has not shown the same improvement as the official unemployment rate over the past few years. This divergence suggests that there has been less improvement in labour underutilization in the United States than would be inferred from looking only at the unemployment rate.

Long-term unemployment

A worrying feature of the Great Recession and its aftermath has been the large rise in long-term unemployment in both Canada and the United States. Long-term unemployment, defined as being out of work for at least 27 weeks, has many costs over and above regular spells of unemployment. Research shows that an individual's earnings can suffer permanent damage after a long stretch of unemployment (Jacobson, LaLonde and Sullivan 1993). Most alarming, long-term unemployment can be self-perpetuating, since workers who face extended periods of unemployment, in particular, may find new employment increasingly difficult to obtain. Lower wages and loss of employment opportunities could be the result of the loss of skills (Pissarides 1992) or the stigma employers attach to workers who have been unemployed for a long time (Kroft, Lange and Notowidigdo 2012).

The percentage of unemployed workers who are considered long-term unemployed rose in both countries during the recession and has remained elevated (**Chart 3**). In Canada, that percentage roughly doubled, peaking at just over 20 per cent in June 2011, and has not shown much improvement since. The rise in long-term unemployment has been even more striking in the United States. At its peak in June 2010, almost half of those unemployed had been out of work for 27 weeks or more. In contrast to Canada, however, the share of long-term unemployment in the United States has been on a gradual downward trend since 2011 (although the level of long-term unemployment, at 37 per cent in December 2013, is still much higher than its average of 20 per cent from December 2001 to November 2007). This

Chart 3: Long-term unemployment

Percentage of total unemployed

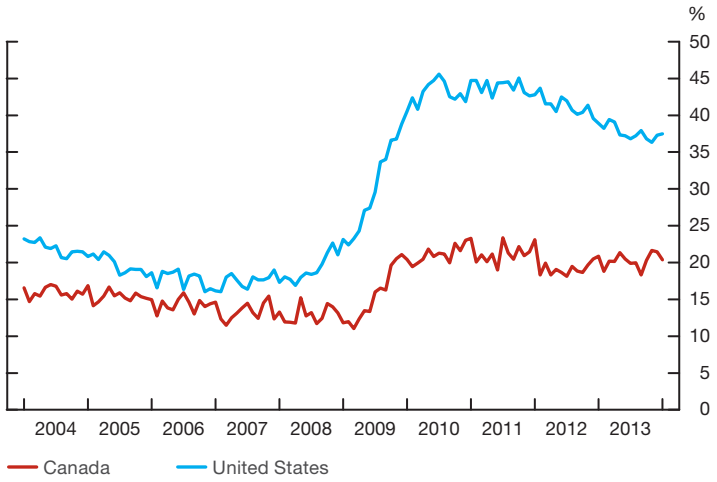

Sources: Statistics Canada and U.S. Bureau of Labor Statistics Last observation: December 2013

decline may reflect improving labour market conditions, but may also be an indication of the long-term unemployed leaving the labour force at a different rate than the unemployed as a whole (Farber and Valletta 2013).

Flows into and out of unemployment

Job-finding and separation rates provide a sense of how employment prospects in the economy are evolving. The job-finding rate measures the proportion of total unemployed workers who flow back into employment, while the separation rate measures the proportion of the total employed workers who enter unemployment. Since these data are not readily available for Canada, we construct them, as explained in **Box 1**.[7]

Box 1

Computing Job-Finding and Separation Rates

Although unemployment flows are published as part of the Job Openings and Labour Turnover Survey (JOLTS) in the United States, no such data exist for Canada. We therefore compute job-finding and separation rates in accordance with the methodology proposed by Shimer (2012). This is a simple method that uses readily available data, as follows:

$$\textit{Job-finding rate} = 1 - \left(\frac{\textit{unemployed}_{t+1} - \textit{short-term unemployed}_{t+1}}{\textit{unemployed}_{t}} \right)$$

$$\textit{Separation rate} = \frac{\textit{short-term unemployed}_{t+1}}{\textit{unemployed}_{t}}$$

The time index t denotes months, and *short-term unemployed* refers to those who have been unemployed for one month or less. We find that the job-finding and separation rates in the United States are correlated with the JOLTS hiring and layoff rates at 0.92 and 0.75, respectively.

7 This method assumes that individuals do not enter or exit the labour force, but rather simply transition between employment and unemployment. Although this is unrealistic, it has been demonstrated that relaxing this assumption does not alter the dynamics of the job-finding and separation rates obtained. See Shimer (2012) for an application to U.S. data and Office of the Parliamentary Budget Officer (2012) for an application to Canadian data.

Chart 4: Unemployment flows

3-month moving average

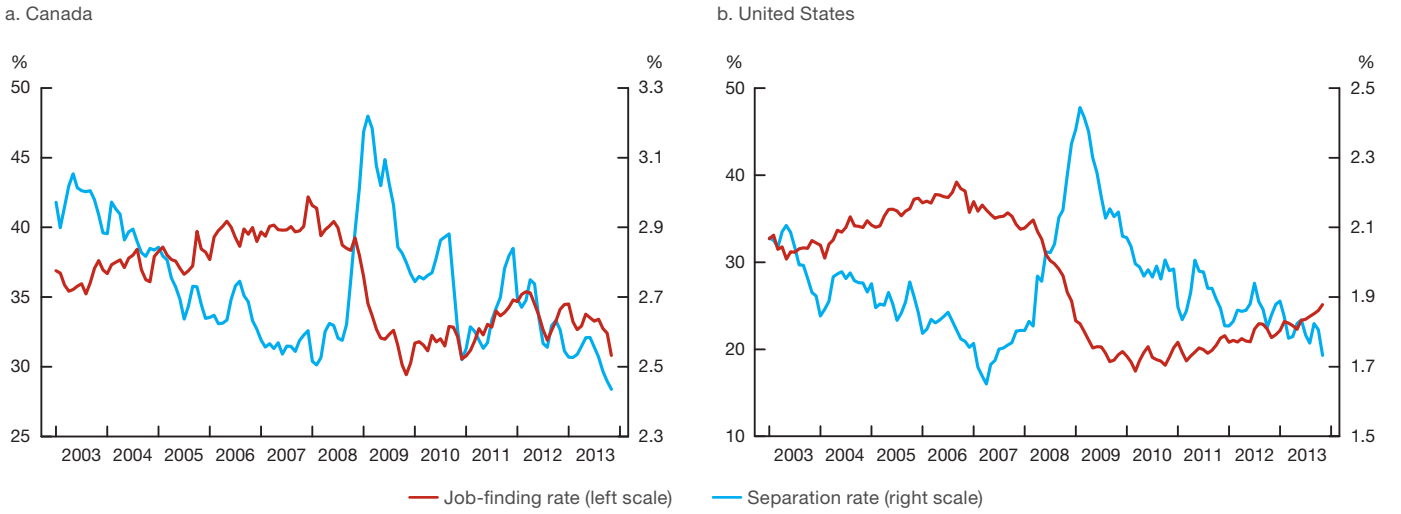

Sources: Statistics Canada, U.S. Bureau of Labor Statistics and Bank of Canada calculations

Last observation: December 2013

In general, job-finding and separation rates have moved in opposite directions (**Chart 4a** and **Chart 4b**). This was particularly true during the recession, when, in both countries, the job-finding rate declined markedly and the separation rate rose sharply. Although the separation rates are back roughly to pre-recession levels, the recovery in job-finding rates has been much more subdued. In Canada, the job-finding rate increased at a relatively robust pace between 2010 and 2012, but has since fallen back to a level only slightly above the low point witnessed during the recession. In the United States, the job-finding rate has trended up since 2010, albeit at a gradual pace. In other words, post-recession employment gains in both Canada and the United States have been driven mainly by a decrease in the number of layoffs rather than by a significant pickup in the pace of hiring.

The labour force participation rate

The labour force participation rate (LFPR) measures the proportion of the working-age population that is either employed or actively looking for work. Slack economic conditions can cause the LFPR to fall, as either previously active job seekers become discouraged and stop searching for work or new job seekers delay their entry into the labour force.

Of course, decisions to exit or enter the labour force are not determined solely by the health of the labour market. For example, labour force participation is also affected by individuals' decisions on how long to stay in school and when to retire. Interpreting movements in the LFPR can therefore be challenging. In fact, the post-recession period has coincided with important demographic shifts on both sides of the border (**Table 1**). Older individuals (defined here as age 55 and over) have represented an increasing share of the working-age population; however, the participation rate of this group is understandably much lower than that of prime-age (25 to 54) or youth (under 25) workers. Indeed, the relative importance of cyclical and demographic factors on the participation rate has been a topic of major debate in the United States since the recession (Bengali, Daly and Valletta 2013; Hotchkiss and Rios-Avila 2013; Erceg and Levin 2013).

Table 1: Share of working-age population
By age group (per cent)

	Canada		United States	
	2007	2013	2007	2013
Under 25 (youth)	16.6	15.5	16.1	15.8
25 to 54 (prime-age)	53.6	50.9	54.2	50.6
55 and over (older)	29.8	33.5	29.7	33.6

Sources: Statistics Canada, U.S. Bureau of Labor Statistics and Bank of Canada calculations

Chart 5: Labour force participation rates
Index: December 2007 (start of recession in the United States) = 100

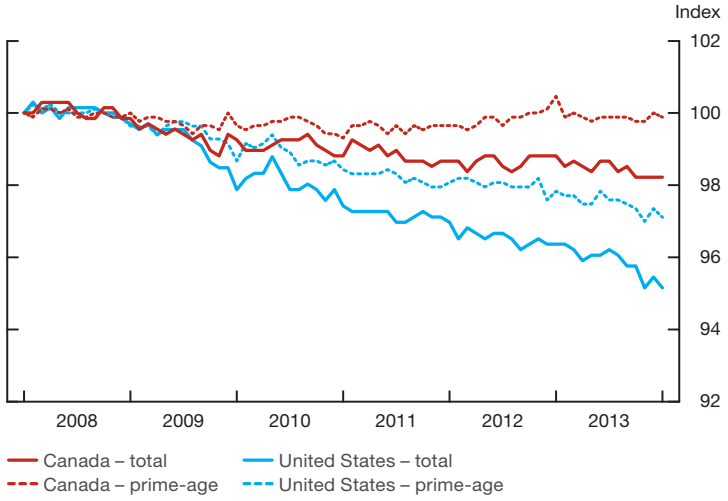

Sources: Statistics Canada, U.S. Bureau of Labor Statistics and Bank of Canada calculations
Last observation: December 2013

The LFPR has trended down in both Canada and the United States since the onset of the recession, with the decline being particularly pronounced south of the border (**Chart 5**). Focusing on the LFPR for workers between the ages of 25 and 54 (prime-age workers) is a simple way to abstract from some of the potential impact of demographic change on the decline in overall LFPR. The prime-age LFPR in Canada fell marginally in the aftermath of the recession and has recently hovered around its pre-recession level, suggesting that much of the recent decline in the aggregate LFPR has been the result of demographic change. The picture is quite different for the United States, where the prime-age LFPR has fallen markedly since the onset of the recession, although not nearly as much as the aggregate LFPR. Thus, demographic factors appear to be only partly responsible for the declining labour force participation in the United States, suggesting worse labour market conditions than in Canada over this period.

Average hours worked

During a recession, firms are likely to cut back on the number of hours that their employees work, since firing and then later rehiring employees can be costly.[8] For the same reason, firms are more likely to increase the hours of

8 Ohanian and Raffo (2012) document the importance of the average number of hours worked in the labour market adjustment process across countries.

Chart 6: Average weekly hours worked

Index: December 2007 (start of recession in the United States) = 100

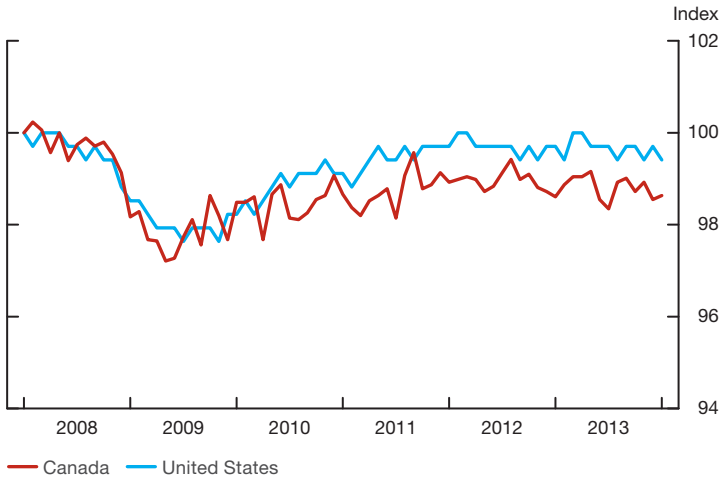

Sources: Statistics Canada, U.S. Bureau of Labor Statistics and Bank of Canada calculations

Last observation: December 2013

their existing workforce rather than immediately hire new employees once a recovery begins. The length of the average workweek can therefore convey important information about momentum in the labour market.

In both Canada and the United States, average hours worked declined sharply during the recent recession and have recovered gradually since (**Chart 6**). In Canada, hours have stabilized at a level somewhat below the pre-recession value, while in the United States the recovery has been slightly more pronounced.

Overall, the recovery in average hours worked indicates an improvement in the intensity with which employed labour is being used. This development bodes well for both economies.

Wage growth

Definitional differences in the measures of wage growth in Canada and the United States make it difficult to compare wage growth across countries, especially since there are a wide variety of measures from which to choose.[9] Nevertheless, the measures of wage growth shown in **Chart 7** have displayed notable similarities over the past decade.

Over the four years before the recession, nominal wage growth trended up in both Canada and the United States. In 2007, it averaged 3.5 per cent in Canada and 4.0 per cent in the United States. Following the start of the U.S. recession in December 2007, wage growth in both countries held near 2007 levels for another year before slowing noticeably in 2009, and by December of that year it had reached 2.5 per cent in both Canada and the United States. These measures of wage growth have recently hovered around the 2 per cent level, averaging 2.1 per cent in Canada and 2.0 per cent in the United States in 2013. This modest wage growth in both countries is consistent with subdued demand for labour, although weak productivity growth could also be a factor.

9 Data on wage growth are from the Labour Force Survey (LFS) in Canada and the Establishment Survey in the United States. This choice is mainly because of the timeliness of the LFS, although the qualitative message does not change if we use Canada's Survey of Employment, Payrolls and Hours (SEPH). Data for Canada represent the total economy, while for the United States the data cover only the private non-farm sectors.

Chart 7: Average hourly wages
Year-over-year percentage change

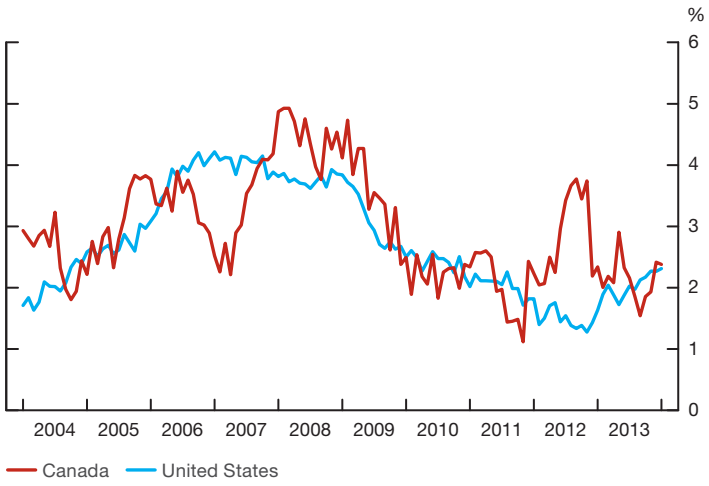

Sources: Statistics Canada and U.S. Bureau of Labor Statistics

Last observation: December 2013

A Broader Measure of Labour Market Activity

To consolidate the information contained in the various labour market measures shown in the preceding section, we construct a labour market indicator (LMI) for both countries using a statistical technique known as principal-component analysis. This technique extracts the common movement across the eight labour variables to create a simple summary measure of labour market activity. The LMI is scaled to be comparable with the unemployment rate,[10] and thus provides a simple benchmark against which to judge whether the unemployment rate is evolving in a manner consistent with broader labour market conditions. **Chart 8a** and **Chart 8b** show the results for Canada and the United States, respectively.

The labour market indicator provides a simple benchmark against which to judge whether the unemployment rate is evolving in a manner consistent with broader labour market conditions

In both instances, the LMIs closely track the unemployment rate, rising rapidly at the onset of the recession and then falling slowly once the recession ends. Over the post-recession period (2010–13), the Canadian LMI declined 0.5 percentage points, while the unemployment rate fell 0.9 percentage points, suggesting that the latter may have modestly overstated the extent of improvement in the labour market. In contrast, there appears to have been a larger and more persistent disconnect between the LMI and the unemployment rate in the United States. From 2010 to 2013, the LMI and the unemployment rate declined 1.1 and 2.3 percentage points, respectively, suggesting that the unemployment rate may have substantially overstated the post-recession improvement in labour market conditions in the United States.[11]

10 Principal-component analysis identifies patterns in data by converting a set of possibly correlated variables into a set of linearly uncorrelated variables called principal components. The first principal component accounts for as much of the variability in the data as possible. The summary labour market indicators calculated for Canada and the United States are set equal to the first principal component, and are then scaled to the unemployment rate by regressing each country's unemployment rate on its LMI and a constant.

11 All variables are in level terms, except for average weekly hours, which are expressed as a year-over-year percentage change. Note that principal-component analysis requires stationary variables. While some of these series fail conventional unit root tests, this could be because of the short time span used or the weakness of these tests. To address these concerns, we remove the trend in our data using the Hodrick-Prescott filter, which is the approach used by Barnes et al. (2007), and find that it does not materially alter the results of our analysis.

Chart 8: Unemployment rate and labour market indicator (LMI)

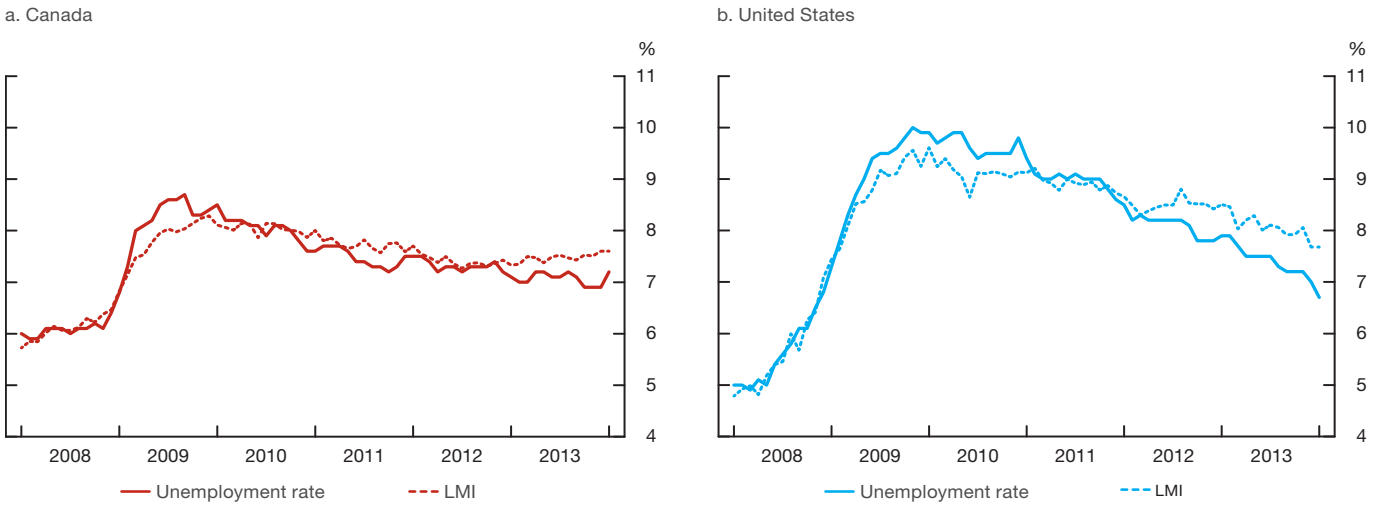

Source: Statistics Canada and Bank of Canada calculations

Last observation: December 2013

Chart 9: Correlations with the labour market indicator

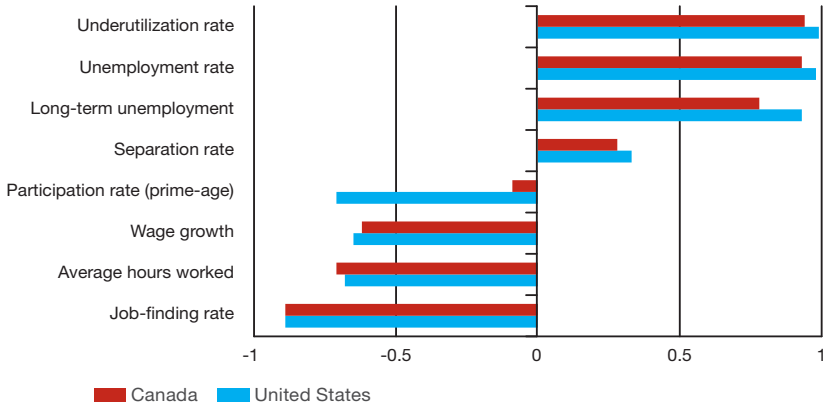

Source: Bank of Canada calculations

The fact that the unemployment rate has fallen more rapidly than the LMI in both countries suggests that other labour market measures have not shown as much improvement as the unemployment rate. In Canada, the modest divergence can likely be attributed to the job-finding rate and the percentage of long-term unemployed. As shown in **Chart 9**,[12] these are among the measures that are most highly correlated with the Canadian LMI and, as indicated in the preceding section, have displayed limited improvement following the recession.

These variables also appear to be partly responsible for the divergence between the LMI and the unemployment rate in the United States, although other factors are also at play. In particular, the U.S. underutilization rate and the prime-age labour force participation rate have been on a less favourable trajectory than the unemployment rate, contributing to relatively less improvement in the LMI.[13]

◂ *The fact that the unemployment rate has fallen more rapidly than the labour market indicator suggests that other labour market measures have not shown as much improvement*

12 Chart 9 shows the correlation between each labour market measure and the LMI for both countries. If a labour measure and the LMI have a positive (negative) correlation, the LMI will generally increase (decrease) as the measure increases.

13 The prime-age participation rate is also much more correlated with the LMI in the United States than in Canada.

Conclusion

This article has discussed the development of labour market conditions since their rapid deterioration during the Great Recession of 2007–09 and has analyzed the recent co-movements in key labour market measures for both Canada and the United States. Given the importance of labour market outcomes to monetary policy decisions, monetary authorities in both Canada and the United States will continue to monitor these developments closely. This article highlights the need to consider a broad range of labour market variables in addition to the unemployment rate. Although the unemployment rate in Canada has evolved largely in line with overall labour market conditions since the recession, the article has shown that it may have modestly overstated the extent of recent improvement. This contrasts with the United States, where the unemployment rate appears to have significantly overstated the improvement in broader labour market conditions.

Literature Cited

Barnes, M., R. Chahrour, G. Olivei and G. Tang. 2007. "A Principal Components Approach to Estimating Labor Market Pressure and Its Implications for Inflation." Federal Reserve Bank of Boston Public Policy Brief No. 07-2.

Bengali, L., M. Daly and R. Valletta. 2013. "Will Labor Force Participation Bounce Back?" Federal Reserve Bank of San Francisco Economic Letter No. 2013-14.

Erceg, C. J. and A. T. Levin. 2013. "Labor Force Participation and Monetary Policy in the Wake of the Great Recession." International Monetary Fund Working Paper No. WP/13/245.

Farber, H. S. and R. G. Valletta. 2013. "Do Extended Unemployment Benefits Lengthen Unemployment Spells? Evidence from Recent Cycles in the U.S. Labor Market." Federal Reserve Bank of San Francisco Working Paper No. 2013-09.

Hotchkiss, J. L. and F. Rios-Avila. 2013. "Identifying Factors Behind the Decline in the U.S. Labor Force Participation Rate." *Business and Economic Research* 3 (1): 257–75.

Jacobson, L. S., R. J. LaLonde and D. G. Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review* 83 (4): 685–709.

Kroft, K., F. Lange and M. J. Notowidigdo. 2012. "Duration Dependence and Labor Market Conditions: Theory and Evidence from a Field Experiment." National Bureau of Economic Research Working Paper No. 18387.

Office of the Parliamentary Budget Officer. 2012. "An Assessment of Canada's Labour Market Performance" (29 October).

Ohanian, L. E. and A. Raffo. 2012. "Aggregate Hours Worked in OECD Countries: New Measurement and Implications for Business Cycles." *Journal of Monetary Economics* 59 (1): 40–56.

Pissarides, C. A. 1992. "Loss of Skill During Unemployment and the Persistence of Employment Shocks." *Quarterly Journal of Economics* 107 (4): 1371–91.

Shimer, R. 2012. "Reassessing the Ins and Outs of Unemployment." *Review of Economic Dynamics* 15 (2): 127–48.

Sorrentino, C. 2000. "International Unemployment Rates: How Comparable Are They?" Bureau of Labor Statistics *Monthly Labor Review* (June): 3–20.