Extracting Policy Rate Expectations in Canada*

Antonio Diez de los Rios
antonioddr@gmail.com
Bank of Canada

Christopher Reid
chrisreid@bankofcanada.ca
Bank of Canada

July, 2008

Abstract

This paper presents an ordered-probit based approach to extract information about market expectations from asset prices for monetary policy purposes in Canada. We find that the use of three different yield curve slope variables as our set of predictors provides an adequate expectation of future policy rate moves up to one year ahead. Moreover, an analysis of out-of-sample probabilities of target rate moves during 2007 reveals that our model is also able to identify the turns in monetary policy caused by financial turmoil that started during the summer of 2007.

JEL Classification: E43, E47, E52, G12

Keywords: Market expectations, monetary policy, ordered-probit, out-of-sample, yield curve.

* We want to thank Toni Gravelle, Scott Hendry, Grahame Johnson, Wally Speckert, and Greg Tkacz for their comments. We also thank Guy MacKenzie for providing excellent research assistance. We remain responsible for any remaining errors. *Address for correspondence*: Christopher Reid, Financial Markets Department, 234 Wellington Street, Ottawa, Ontario, K1A 0G9, Canada. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada.

1. Introduction

The goal of this paper is to improve the ability to extract information about market expectations from asset prices for monetary policy purposes. In particular, an accurate measure of financial market expectations of changes in the policy rate has several important uses to a central bank. On one hand, the policy overnight rate is a central bank's main source of influence on yields further out the curve and ultimately the real economy through its influence on consumer and corporate borrowing rates. The source of this influence on longer-term rates lies not primarily with the current level of the policy rate but rather in market expectations of future levels of the policy rate and the credibility of the central bank to maintain inflation close to its target. As a result, a good estimate of market participant expectations is an important component in conducting monetary policy. Similarly, central bank communications have over the last several years attempted to provide increasing levels of transparency about the bank's reaction function and outlook for both growth and inflation. As a result, an accurate measure of market policy rate expectations provides a timely and transparent check on which monetary policy decisions would constitute a surprise and which are well anticipated by market participants.

In the United States, monetary policy expectations for future Federal Open Market Committee (FOMC) interest rate outcomes can be gauged through the use of options on federal funds futures (see Carlsson et al, 2005). In particular, one can make use of these options to estimate the (risk-neutral) probability density function associated with different paths for the target federal funds rate over the next several FOMC meetings. However, such an approach cannot be employed in Canada because of the lack of an active market akin to the options on federal funds futures market. In this case, where we only have data on forward interest rates, traditional methods only allow the recovery of the expected mean of future policy rates. Therefore and as noted by Carlsson et al. (2005), the probabilities can only be identified under the very restrictive assumption that the monetary authority will choose between just two targets.

To overcome this problem, we present an alternative econometric approach that allows us to estimate the probability density function for future policy moves in the absence of an equivalent interest rate option market. In particular, our model exploits the present construction of the monetary policy framework to take advantage of the fact that, nowadays, central banks commit themselves to consider discrete changes to the target rate on a series of pre-announced dates, and that these discrete changes are usually twenty-five basis points.[1] In particular, we borrow from the literature on central bank empirical reaction functions (see e.g. Gerlach, 2007) to estimate directly the probability of future policy rate moves using ordered-probit techniques.[2]

As mentioned before, we could have alternatively attempted to derive market expectations of future levels of interest rates from the yield curve using traditional approaches as the expectations hypothesis (see e.g. Soderlind and Svensson, 1997, and Kim and Orphanides, 2007). That is, if the expectations hypothesis holds then the forward interest rate (the short-term rate at which investors agree now to borrow or lend in the future) is the expected future short-term rate. However, it has been shown exhaustively in the literature that the expectations hypothesis does not hold due to the existence of a time-varying risk premium (see Fama and Bliss, 1987, or more recently, Cochrane and Piazzesi, 2005, and Piazzesi and Swanson, 2008). Therefore, if we were to use the expectations hypothesis to derive market expectations we would need to make an adjustment to the forward interest rate to account for the existence of this risk premium. Yet obtaining an accurate estimate of the risk premium that could be used to risk-adjust the forward interest rate is a very delicate task. For example, Kim and Orphanides (2007) note that risk premium estimates are characterized by the lack of robustness to the choice of the sample and regressors.[3] However, we know from Fama (1984) that, while the expectations hypothesis does not hold, the term structure of forward interest rates still

[1] Our treatment of this assumption is explained in greater detail in Sections 2 and 4.2.

[2] Note that in the absence of a monetary policy framework with pre-announced dates, we would need to first estimate a probit model to determine the probability of an intervention by the central bank at each point in time (the so-called events or points) and, second, an ordered probit model to determine the size of the interventions, conditionally on the bank having decided to intervene (the so-called marks) (see Dolado and Maria-Dolores, 2002, and Hamilton and Jorda, 2002).

[3] These drawbacks apply not only to regression-based risk premium estimates but also to those computed using more sophisticated no-arbitrage models.

contains information about future spot rates.[4] This is why this study employs a different approach by abstracting from the estimation of the risk premium and instead asks whether information readily available in financial markets can be used to accurately extract the market expectation of future changes in policy rates.[5]

This study examines the case of policy rate moves of the Bank of Canada that occur after the implementation of a policy of fixed announcement dates (FADs) in 30th October 2000. In particular, this policy commits the Bank to consider changes to the overnight rate on a series of eight pre-announced dates each year (see Johnson, 2003 for additional details). [6] Moreover, when the Bank changes the policy rate it usually does so in twenty-five basis points increments. In fact, the same methodology described in this paper can readily be applied to any central bank with pre-announced policy rate decision dates and which don't benefit from being able to construct risk neutral densities from an active interest rate options market.

Using weekly data from May 2001 to December 2006, we find that the use of three different, non-overlapping yield curve slope variables as our set of predictors provides an adequate expectation of future policy rate moves one, two, four and eight FAD periods ahead.[7] In particular, our set of covariates includes the three-month yield minus the target overnight rate, the one-year yield minus the three-month yield and the difference between the yield on the ten-year Government of Canada bond and the one-year yield.

We find that the estimated probability of future policy rate moves performs well at identifying turns in the monetary policy cycle. This is true not only on an in-sample setting, but also when we perform an out-of-sample exercise. This implies that market

---

[4] Similarly, Clarida and Taylor (1997) find, without modeling the foreign exchange risk premia, that the term structure of interest rate differentials contains information that is useful to predict the exchange rate.

[5] Note that the information contained in the term structure of interest rates comes from two sources. The first one is the forecast component of future interest rates. Second, the risk premium can also contain information about future movements of spot rates as long as it is correlated with the forecast component (as suggested by the rejection of the expectations hypothesis).

[6] In the press release that announced the implementation of the FADs, it was stated that the Bank would retain the option of taking action between fixed dates, although it would exercise this option only in the event of extraordinary circumstances. To date, only one change has been made between FADs: on 17 September 2001, the Bank lowered the overnight rate by 50 basis points following the 11 September 2001 terrorist attacks.

[7] While the FAD process was implemented in 30th October 2000, our data for the three-month yield described below only extends back to May 2001.

participants understand the Central Bank's reaction function and as a result are able to predict future changes in monetary policy reasonably well. Specifically, we estimate our model using data only until the end of 2006 and analyze out-of-sample probabilities of target rate moves during 2007 to find that our model is still able to identify the turns in monetary policy caused by financial turmoil that took place during the summer of 2007.

We also show that our ordered-probit model out-performs other three benchmark models: a "naïve model" that simply assumes that future changes in the overnight rate are the same as its most immediate change, a constant probability model that takes the most frequently observed value in the sample as a constant prediction of the policy outcome, and a model based on the expectations hypothesis of the term structure.

The rest of the paper is organized as follows. Section 2 will provide an overview of the model. Section 3 will provide a summary of the results and an assessment of the model. Section 4 presents some extensions of our methodology, while section 5 will conclude.

2. The model

As noted in the introduction of this paper, we obtain financial market expectations of changes in the policy rate by exploiting the present construction of the monetary policy framework. Thus, we take advantage of the fact that, nowadays, central banks commit themselves to consider discrete changes to the target rate on a series of pre-announced dates. But as noted by Gerlach (2007) and others, it is inappropriate to fit the model using ordinary least squares when the target rate remains unchanged in most months and it only changes by a discrete amount at pre-specified points of time when the Central Bank judges it necessary. For example, note in Table 1 that the target rate was left unchanged in 28 of the 54 interest rate meetings that we consider over the sample period from May 2001 to December 2006. In particular, the rate was raised twelve times and cut fourteen times. Note also that over our time sample, the Bank of Canada never increased the rate by fifty basis points, while it cut it by this quantity three times following the 11 September 2001 terrorist attacks.

Therefore, we borrow from the literature on central bank empirical reaction functions to estimate directly the probability of future policy rates moves using ordered-probit techniques.

However, our goals differ to those previously reported when estimating Central Banks' empirical reaction functions (e.g. see Dolado and Maria-Dolores 2002, and Gerlach 2007). That is, our interest lies in estimating the probability of a *future* target rate move rather than in explaining what is driving a *given* one. Therefore, we construct a new dependent variable that takes the value of one if there is an increase of the target rate on the next FAD date, takes the value of zero if there is no movement, and it will be equal to minus one if the Bank of Canada decides to lower the overnight rate on the next FAD date. That is:

$$y_t = \begin{cases} +1 & \text{if the target rate is increased on the next FAD date} \\ 0 & \text{if the target rate is left unchanged on the next FAD date} \\ -1 & \text{if the target rate is lowered on the next FAD date} \end{cases}$$

We will further assume that this observed dummy variable, y_t, depends on a latent index, y_t^*, according to the following rule:

$$y_t = \begin{cases} +1 & y_t^* \leq \alpha_1 \\ 0 & \alpha_1 < y_t^* \leq \alpha_2 \\ -1 & y_t^* \geq \alpha_2 \end{cases}$$

where y_t^* is taken to be a latent continuous random variable and the α's are the thresholds that the latent variable must cross to change the value of y_t. This underlying index is assumed to depend linearly on a set of covariates, x_t, such that:

$$y_t^* = \beta' x_t + u_t \tag{1}$$

In particular, we choose three different yield curve slope variables as our set of covariates. Finally, assuming that $u_t : i.i.d.N(0,\sigma^2)$ allows us to obtain estimates of the parameter vector (α, β) by maximizing the following likelihood function:

$$l(\alpha,\beta)=\sum_{t\in y_t=-1}\log\Phi(\alpha_1-\beta' x_t)+\sum_{t\in y_t=0}\log[\Phi(\alpha_2-\beta' x_t)-\Phi(\alpha_1-\beta' x_t)] \qquad (2)$$
$$+\sum_{t\in y_t=+1}\log[1-\Phi(\alpha_2-\beta' x_t)]$$

where $\Phi(.)$ is the cumulative Gaussian distribution function.

3. Results

3.1 Data

This section describes our choice of data. We use weekly data from May 2001 to December 2006 on three different slope measures from the Canadian yield curve. While the FAD process was implemented in 30th October 2000, our data for the three-month yield described below only extends back to May 2001. The three slope variables that we use are: a three-month yield minus the target overnight rate, a one-year yield minus the three-month rate and the difference between the yield on the ten-year Government of Canada bond and the one-year yield. These three non-overlapping slope parameters were chosen in order to capture information contained over the entire yield curve. Data for 2007 is left for an out-of-sample diagnostic of the model.

For the first slope measure we use the three month yield from the Canadian Overnight Index Swap (OIS) market.[8] The OIS is a fixed-to-floating interest rate swap that ties the floating leg of the contract to a daily overnight reference rate. OIS yields have a number of advantages relative to other financial instruments such as treasury-bills and banker acceptances. For instance, OIS contracts have a constant maturity, there is no basis risk, they are off-balance sheet, and no cash is transferred between counterparties until final settlement. As a result of these characteristics, OIS yields are thought to have minimal credit and liquidity risk and less trading frictions overall. Therefore, they should provide better information about market expectations than other money market products such as treasury bills and bankers' acceptances.

[8] See Reid (2007) for more information on the OIS market in Canada.

Unfortunately, while the OIS market is very liquid at relatively short-maturities, it is often illiquid further out the yield curve. As a result, for the second slope measure we construct a one-year yield from the BAX market – a futures market for three-month banker acceptances (BAs). The BAX market in many ways resemble that of the eurodollar market (they are both future markets with quarterly maturities, and based on an underlying interest rate product with three months to maturity- the difference being that eurodollar contracts are settled on LIBOR while BAX contracts are settled on the Canadian Deposit Offering Rate).[9] In Canada, this market is very liquid, and has been used in previous studies (Johnson 2003) to examine market participant interest expectations.[10]

For the third slope calculation, we use both the ten-year zero-coupon yields taken from the Bank of Canada's website along with the one-year yield described above.

Additionally, we could also have tried to extract market expectations information from other financial instruments such as stocks, or even exchange rates. We have decided not to do so for two reasons. First, including stock market data introduces a potential identification problem. For example, the central bank may react to higher stock prices by raising interest rates if the monetary authority views the stock price increase as a signal of a strong economy. But, on the other hand, stock prices are also likely to fall if market participants expect a decision to tighten, given the higher discount rate for the expected stream of dividends. Thus, we might arrive at a situation where these two effects compensate each other and, therefore, we would find no relationship between interest rate movements and stock prices (see Rigobon and Sack, 2001, and Lapp et al. 2003). We face a similar identification problem when trying to extract market expectations from exchange rates.

[9] The Canadian Deposit Offering Rate (CDOR) is determined daily from a survey of nine market makers in bankers' acceptances in Canada where the high and low rates are excluded and a simple arithmetic average is calculated from the remaining survey rates.

[10] We do an adjustment relative to Johnson (2003). BAX contracts have quarterly maturities, and as a result the time to maturity declines over the life of the contract. For example consider the third BAX contract, on the first day trading it will have a time of maturity of roughly three quarters of a year – and since it is a future on a three month product this is similar to a one year yield. However, as it approaches the next BAX settlement date (roughly three months forward in time) its time to maturity has likewise declined. Therefore, in order to calculate a constant maturity one-year yield we linearly interpolate between the third and fourth BAX contract and apply the appropriate weights.

Second, we also want to avoid a potential over-fitting of the model, and this is the reason why we only use three yield curve slope variables. In addition, the main objective of this model is its use on a "real-time" basis and these three variables are readily available.

The data is weekly in order to avoid any day-of-week effect, and taken from the first business day of the week.[11] Since FADs are set to occur on the second business day of the week this allows us to capture expectations immediately prior to an interest rate announcement. The timing of the FADs is well known, since they are announced in the fall for the upcoming calendar year.[12]

3.2 One FAD ahead

The first column of Table 2 reports estimates of the model in equations (1) and (2), that is, the model that we use to extract expectations of future policy movements over the next FAD. In particular, we find that the coefficients on all three slope variables are all significant at the one percent confidence level. The coefficient on both the three-month and one-year spreads are, as expected, positive indicating that as these spreads increase the likelihood of an increase in the policy rate at the next FAD also increases (and vice-versa). On the other hand, the coefficient on the ten-minus-one-year slope is both relatively small and negative which may, at first, appear to be counter-intuitive. In principle one would expect a widening spread to signal higher expected future inflation which will lead to an increase of the policy rate. We instead argue that our results are explained by the fact the one-year yield is more sensitive to actual and perceived near-term policy rate changes than the ten-year yield because the latter is mainly driven by the credibility of a central banks' commitment to low inflation (Goodfriend, 1998). As a result, when policy rates are expected to decline, the one-year yield will correspondingly decline more or at least quicker than the ten-year yield when long-run inflation expectations remain anchored. This will result in the widening of the spread between the

---

[11] When Monday was a holiday we used data from Tuesday.

[12] As will be discussed in section 4, when one extends the time frame of the study past one period ahead, this can on occasion require that participants have perfect foresight as to the timing of the next year FAD dates. However as the FADs from year to year remain relatively close in terms of timing this may not be an unreasonable assumption.

two yields (and vice-versa when policy rates are expected to increase). Therefore, the negative sign on the coefficient on the ten-minus-one-year year spread is capturing the relative impact of expected policy rate changes between longer-term and shorter-term yields under a framework where the central bank's commitment to low inflation remains credible.

3.2.1 Estimated Probabilities of Policy Changes

Panel a of Table 3 presents information regarding the models ability to correctly gauge the actual policy rate decision at the next FAD. In particular, we present the predicted outcome from our model, but also how often this prediction is correct in anticipating the monetary policy outcome. Recall that the range of possible outcomes for the next FAD is limited to either no change or +/- 25 basis points.[13] The model then takes the option with the highest probability as the estimated outcome of the FAD decision. We find that the model over-predicts both "no change" and a decrease of 25 basis points, while simultaneously under-predicting an increase of 25 basis points. Yet in our sample of 290 weeks (54 FADs) the model records only 37 misses which implies that we are able to correctly capture the interest rate decision 87 per cent of the time. Taken with the pseudo R^2 of 0.74 (see again table 2), this result suggests that the model is useful in predicting the policy rate decision at the next announcement date. We examine the model over longer horizons, specifically 2, 4, and 8 FADs ahead in section 4.2.

While the ability of the model to predict policy outcomes at the upcoming FAD is re-assuring, it is important to recall that our aim is to provide a reliable gauge of market participant expectations. Given that market participants do not possess perfect foresight, it is not at all surprising that prediction errors occur. Moreover, it is of greater interest for our final goal of extracting market participant policy rate expectations to use our model framework to construct probabilities of the various possible monetary policy actions at

[13] Recall that the range of possible outcomes for the next FAD is limited to either no change or +/- 25 basis points. Since the introduction of the FAD there have been only four instances where the policy rate changed by more than 25 basis points on any one date (and three within the data range we consider). These three instances were in large part a reaction to the events pertaining to the September 2001 terrorist attacks. Given that these changes in the policy rate were obviously a surprise to markets, we partially account for this by truncating each of these large moves in policy rates to be equivalent to a -25 basis point reduction.

the next announcement date. Figure 1 shows the evolution of the probability of the different outcomes over time. We have, for comparison purposes, also included the actual path of the target overnight rate as a solid black line. As can be seen from figure 1, and likely of particular interest to policy makers, the probability is never evenly split between three viable alternatives, but it rather shows that market participants have at the very least a clear bias in one direction. Furthermore, a careful examination reveals that the model also does a remarkable job in predicting turns in monetary policy cycles. This may, in some part, reflect the emphasis the Bank of Canada has placed on improving communications with market participants since the introduction of the FADs.

Furthermore, the model's ability to extract market participant expectations may be even better than that implied by the number of hits and misses depicted in Table 3. Specifically, note that over the time period from the latter half of 2005 until the second half of 2006 the Bank of Canada engaged in a gradual and persistent increase in the overnight rate, and that this episode is well-captured by the model given it implies a very high probability of an increase of the target rate at almost every point of time.[14] However, we can also find during this period that the probability of inaction spiked on the week just after every announcement date. This suggests that, just after every FAD, market participants were expecting no change in policy rates for brief periods of time. While, given the setup of the model, these spikes are recorded as forecast misses, anecdotal evidence suggests that during this period, participants were not immediately pricing in a further increase in the policy rate immediately preceding the last increase, but were rather more cautious and only gradually fully pricing in a rate hike as the next FAD approached (as confirmed by our model).

4. Extensions

The model as outlined in section 2 can easily be modified to examine a number of useful extensions. In particular, we first consider an out-of-sample exercise to examine how well our model performs during a period marked by the financial turmoil that took place

[14] We choose not to smooth expectations through the use of a moving average as used by Gerlach (2007) as we are interested in the evolution of expectations. Furthermore as new information will enter the market at various times and may very well alter expectations, it is unclear that lagged expectations are useful.

during the summer of 2007 if we were to implement it in a "real-time" basis. Second, we examine the performance of the model in predicting policy rate changes beyond one FAD. Finally, we assess the effectiveness of our model in predicting future policy outcomes by comparing our results to those of three simple benchmark models.

4.1 Out-of-sample

In this section, an out-of-sample exercise is conducted to assess the "real-time" performance of the model. Specifically, we estimate our model using data only until the end of 2006. Then, we analyze out-of-sample probabilities of target rate moves during 2007 in order to investigate if our proposed model is able to identify the turns in monetary policy caused by the financial turmoil that took place during the summer of 2007. In particular, we obtain the out-of-sample anticipated probabilities of target rate moves during 2007 by fixing the parameters of the model to the estimates previously obtained using the data from May 2001 to the end 2006 and that we report in Table 2.

The out-of-sample period covers 48 weeks and eight FADs in 2007.[15] Out of these eight FADs, there was a 25 basis point increase in July and a decrease of the same magnitude in December.

Figure 1 displays the out-of-sample probabilities of an interest rate movement in 2007. In particular, note that the market-based estimated probability of an increase in the target rate increased during the weeks previous to the FAD in July. Therefore, these results indicate that the decision in July was anticipated by market participants. Moreover, the probability of an increase remained high for some weeks after this decision implying that markets participants were expecting another increase in the September FAD meeting. This may reflect the effectiveness of Bank communications. Specifically the Bank's press release at the May FAD stated that:

"On balance, the Bank judges that there is an increased risk that future inflation will persist above the 2 per cent inflation target and that some increase in the target for the overnight rate may be required in the near-term to bring inflation back to the target".

---

[15] We drop the last four weeks of 2007 because they reflect the probabilities of a policy rate movement in the first FAD of 2008.

As a result it is not perhaps surprising that market interest rates correctly began to price in a strong probability of an increase in the overnight rate, as can be seen in Figure 1.

The figure also suggests, that once the financial turmoil began in August of 2007, agents started to revise their expectations towards an ease of the monetary policy. In particular, we can see that the probability of an increase rapidly decrease to zero at the same time that the probability of inaction spiked. Our results indicate that markets participants were able to correctly anticipate the policy decisions at the September and October FADs, despite the absence of previous forward looking statements in the central bank's communications. Finally, market participants started to gradually price a reduction in the policy rate that fully anticipated the decision in December. In particular, it was only three weeks before the December FAD meeting that the model suggested that the highest probable outcome was a reduction in the policy rate. We suspect that this was likely related to the continued deterioration of market conditions in the fall of 2007. In this instance it would seem that market participants, even in the absence of a strongly worded statement such as the one in May, were able to correctly determine that the realization of the down side risks outlined in the October Bank of Canada Monetary Policy Report would result in a lowering of the policy rate in December.

Overall, we believe that this model out-of-sample performance is rather good. In particular, in the out-of-sample period, it correctly predicts the forthcoming FAD decision on 40 of the 48 weeks, or 83 per cent of the time, which, given the unusual circumstances that took place in financial markets in 2007, makes the results encouraging.

4.2 More than one FAD ahead

Policy makers already have a number of tools available to them to help them assess the policy rate expectation of market participants at the upcoming interest rate decision. While these tools include surveys, market intelligence and economic commentary, they are only of limited use in assessing market expectations for longer than the upcoming FAD, because they either don't exist or are not available in a timely manner. Therefore, a natural extension to this model is to assess how well it performs at longer horizons.

Specifically, we assess how well the model performs over time periods of two, four and eight FADs (approximately three, six and twelve months) ahead over the same time period May 2001 to the end of 2006.

Note that this change in the forecast horizon induces another change in the definition of our dependent variable, y_t. For example, when we try to predict the value of the policy rate at the next FAD we are confronted with only three choices: the policy rate increases by (at least) 25 basis points, the policy rate does not change, the policy rate decreases by (at least) 25 basis points. However, if our forecast horizon is two FAD periods, we will have more choices. For example, it could be the case that the target rate is increased during two consecutive meetings, or that it does not change over the first meeting but it is lowered during the second one. This is the reason why we re-define our dependent variable when looking two and four FADs ahead as:

$$y_t = \begin{cases} +2 & \text{if the target rate is increased (at least) 50 basis points over the next 2/4 FAD dates} \\ +1 & \text{if the target rate is increased 25 basis points over the next 2/4 FAD dates} \\ 0 & \text{if the target rate is left unchanged over the next two 2/4 dates} \\ -1 & \text{if the target rate is lowered 25 basis points over the next 2/4 FAD dates} \\ \text{-2} & \text{if the target rate is lowered (at least) 50 basis points over the next 2/4 FAD dates} \end{cases}$$

Accordingly, we cover in an exhaustive and meaningful way all possible outcomes that we find when trying to extract market expectations over two and four FADs periods.

Similarly we define our dependent variable as

$$y_t = \begin{cases} +3 & \text{if the target rate is increased (at least) 75 basis points over the next 8 FAD dates} \\ +2 & \text{if the target rate is increased 50 basis points over the next 8 FAD dates} \\ +1 & \text{if the target rate is increased 25 basis points over the next 8 FAD dates} \\ 0 & \text{if the target rate is left unchanged over the next 8 FAD dates} \\ -1 & \text{if the target rate is lowered 25 basis points over the next 8 FAD dates} \\ \text{-2} & \text{if the target rate is lowered 50 basis points over the next 8 FAD dates} \\ \text{-3} & \text{if the target rate is lowered (at least) 75 basis points over the next 4/8 FAD dates} \end{cases}$$

when looking at four or eight FADs ahead.

The estimates of the model at these three new forecast horizons (two, four, and eight FADs ahead) can again be found in table 2. Note that, if compared with the coefficients obtained for the case of one FAD ahead, the coefficients on columns two, three and four have the correct sign and they are all significant at the one per cent level. One exception to this rule is the coefficient on the 3 month spread in the 8 FADs ahead time period which is only significant at the 5 per cent level. Moreover, the pseudo R^2 falls as the forecast horizon examined increases: 0.44, 0.19 and 0.17 respectively. This reveals the difficulty of extracting policy rate expectations over long periods of time.

Overall, the model tends to perform reasonably well in predicting policy outcomes. In particular, we provide two different types of information. First, Table 3 presents information regarding the model's ability to account for specific interest rate changes in the sample. Second, Table 4 collects the corresponding information regarding the model's ability to account for the interest rate movement direction. We provide this second set of results because once we increase the forecast horizon the range of possible outcomes increases too. Therefore, it might be the case that our model does not capture exactly the number of cuts over a given forecast period, but that it captures well the direction of the monetary policy. For example, it can be the case that our model indicates that the most likely outcome over the next four FADs is "two 25 basis points cuts" but there was only "one cut". In that case, we would count this situation as a "miss" when checking the model's ability to capture interest rate changes, but it is clear that such a statement is still useful since it gives us the right direction of the decision.

As can be seen from Tables 3 and 4, the model correctly predicts the outcome over the next two FADs 59.66 per cent of the time, and gets the direction correct 80.35 percent of the time. Similarly the model correctly predicts the specific outcome over the next 4 and 8 FADs 59.66 and 60.69 per cent and gets the direction correct 72.07 and 71.38 per cent of the time, respectively.

While we believe that these results are encouraging, a detailed examination of Tables 3 and 4 reveals certain biases in the predictions. Specifically, we can see in Table 3 that the model is not able to predict any of the small or +/-25 basis points movements when looking to the two longer-time periods (four and eight FADs ahead). In addition, our model is not able to predict any of the 38 times that the policy rate remained the same over the eights FADs period. We in part attribute this bias to the way in which actual policy movements are aggregated in the probit framework relative to the way market participants form their expectations over long-time periods. In particular, note that the actual occurrence of small (or no) movements over the eight and to a lesser extent four FADs forecast horizon is largely the result of policy rate movements in opposite direction that offset each other.[16] That is, if an increase of 25 basis points occurs at one interest rate meeting but interest rates are cut several meetings later within the forecast horizon, this episode will be labeled as a "no change in policy rates". In this case, it is reversals in monetary policy that are largely responsible for the number of small changes (+/- 25 basis points) over the longer FAD forecast horizons. Still, it is hard to think of situations where market participants who expect an increase in near term policy rates might also expect a (partial) reversal of this decision over the coming six to twelve months. This would only occur if market participants felt strongly that the Bank had made a policy mistake. Therefore, although our model will tend to over-estimate the possibility of further interest rates increases in the future, we do not think this is a problem, since the objective of our model is to extract market participant expectations.

We now turn to the results regarding the ability of the model to correctly determine the aggregate direction of monetary policy over the different time horizons in Table 4. That is, we have abstracted from whether the model correctly predicts the exact outcome, but rather are interested in whether on aggregate, it got the direction correct. Despite the fact that the model has problems in capturing the magnitude of the interest rate movements exactly over long-periods of time, it may be the case that capturing the direction of the monetary policy move is more important to policy makers than predicting the actual realization. Note that the model over-predicts no change in policy rate over two FADs ahead horizons. This result likely reflects the fact that market participants will only gradually build in expectations of a policy move as the FAD approaches, as discussed earlier in regards to the one FAD ahead example. This may occur for instance at times of

[16] As can be seen in Figure 1, there are four episodes in which policy reversed direction within a one year time horizon.

uncertainty where the market will look to incoming economic data to validate their expectations of a likely policy move by the Bank.

Similar to our results regarding actual and predicted changes, we find that the model under-predicts "no change" when looking to the four and eight FADs ahead horizon. This again likely reflects monetary policy reversals.

Given that we calculate market participant expectations in terms of probabilities and compare them to actual realizations of policy moves, our results are beholden to the fact that market participants do not have perfect foresight. Given that over long-time periods conditions in the economy often change significantly, we feel that our results are reasonably encouraging.

4.3 Comparison to other models

We finally assess the effectiveness of our model in predicting future policy outcomes by comparing our results to those of three simple benchmark models. We first focus on a "naïve model" that simply assumes that future changes in the overnight rate are the same as its most immediate change as our first benchmark. This model will by definition be unable to detect turns in monetary policy but it may perform reasonably well if a central bank engages in interest rate smoothing (whereby a central bank has a preference for parceling what in aggregate is a large move into a series of smaller uniform changes in the overnight rate). The second benchmark model takes the most frequently observed value in the sample as a constant prediction of the policy outcome. In this instance it allows us to determine the extent that our model improves over the assumption that the Central Bank does not alter policy over the entire time horizon. Finally, we include as a third benchmark the expected future short-term rate that is consistent with the expectations hypothesis: the forward interest rate. We map this forecast into the probability space by linear interpolation. For example, if the target rate is 4% and the forward interest rate is 4.15%, we will assume that there is 60% probability of a 25 basis points increase and a 40% probability that the target rate will remain unchanged. Additional details on the computation of this last benchmark are provided in the appendix.

Table 5 presents the results of this "horse race" among the different benchmark models based on the percentage of correct predictions. For example, we find that forecasts obtained from the expectations hypothesis model are correct 82.69% of the time when looking to the one FAD ahead horizon and 38.62% when looking to the two FADs ahead case. Similarly it predicts the right direction 61.38% of the times at the two FADs ahead horizon.

We find that our model out-performs all the benchmark models in predicting the exact policy outcomes over both the one and two FAD time horizon, and in predicting the direction of the change two FADs ahead. The constant probability model is the least accurate model if we look to the changes in the policy rate one FAD ahead, suggesting that the Bank of Canada has been relatively active in adjusting policy over the time period examined. In addition, the naïve model worsens as the forecast horizon increases: it cannot capture any turns in monetary policy. Finally, note also that the performance of the expectations hypothesis model one FAD ahead is relatively close to that of our model. This result is in line with Longstaff (2000) who finds that the expectations hypothesis might hold when looking to very short-term rates. Yet if we increase the forecast horizon to two FADs, the performance of the expectations hypothesis model significantly worsens. This is probably because term premia becomes more important as we increase the forecast horizon.

Finally, we must acknowledge that we are comparing models on an "in-sample" basis and that, given the complexity of our model, it is perhaps not surprising that our model is better able to fit the data that the expectations hypothesis model. Still, our model does also well in an out-of-sample basis (see section 4.1). Most important, our model can help to identify periods of time where there is higher uncertainty about future actions of the central bank. To show why this is true, we present Figures 2 and 3. Figure 2 presents the evolution of the probability of the different outcomes over time obtained from the order-probit model for the two-FADs horizon, we can see that there are moments in time where this probability tend to be split over three or four outcomes (e.g. during 2006). This signals higher uncertainty in the financial markets if compared to those periods where probability is split among one or two outcomes (e.g. the end of 2001). Figure 3 presents the same probabilities obtained from the expectations hypothesis model and a linear interpolation. Note now that these probabilities are only split between two of the five outcomes. This is due to the linear interpolation that we have to employ to map the forward interest rate into the probability space.

5. Final Remarks

In this paper we have examined the use of an ordered-probit based approach to extract information from asset prices about market expectations of future moves in monetary policy. While we specifically examine the case for Canada, this approach could easily be extended to other central banks with pre-announced policy rate decisions. We find that through the use of information readily available in markets, specifically three different interest rate slope parameters, we are able to adequately extract expectations of future policy moves. Furthermore, the use of the ordered-probit approach allows us to frame our results in terms of probabilities and allows for a greater diversity of outcomes to be considered at once. This may be of particular interest to central banks that do not benefit from being able to construct risk neutral densities from an active interest rate options market.

We also examine a number of useful extensions. Encouragingly, the model perform well out-of-sample in 2007, despite the sudden reversal in the path of policy rates related to the commencement of the recent turmoil in financial markets. In addition, the extension of the model to consider policy movements over longer horizons also produced reasonable results. Given the more limited number of options available to central banks to determine expectations over longer horizons, our model may be of particular practical use to central banks as a complement to existing sources of information.

References

Carlsson, J.B., B.R. Craig and W. Melick (2005): "Recovering Market Expectations of FOMC Rate Changes with Options on Federal Funds Futures," Federal Reserve of Cleveland Working Paper 05-07.

Clarida, R.H., and M.P. Taylor (1997): "The Term Structure of Forward Exchange Premiums and the Forecastability of Spot Exchange Rates: Correcting the Errors," *Review of Economics and Statistics*, 89, 353-361.

Cochrane, J. H and M. Piazzesi (2005): "Bond Risk Premia," *American Economic Review*, 95, 1, 138-160.

Dolado J.J. and R. Maria-Dolores (2002): "Evaluating Changes in the Bank of Spain's Interest Rate Target: An Alternative Approach Using Marked Point Processes," *Oxford Bulletin of Economics and Statistics*, 64, 2, 159-182.

Fama, E. (1984): "The Information in the Term Structure," *Journal of Financial Economics*, 13, 509-528.

Fama, E and R. Bliss (1987): "The Information in Long-Maturity Forward Rates," *American Economic Review*, 77, 4, 680-92.

Gerlach, S. (2007): "Interest Rate Setting by the ECB, 1999-2006: Words and Deeds," *International Journal of Central Banking*, 3, 3, 1-45.

Goodfriend, M. (1998): "Using the Term Structure of Interest Rates for Monetary Policy," *Federal Reserve Bank of Richmond Economic Quarterly*, 84, 13-30.

Hamilton, J.D. and O. Jorda (2002): "A Model of the Federal Funds Rate Target," *Journal of Political Economy*, 110, 5, 1135-1167.

Johnson, G. (2003): "Measuring Interest Rate Expectations in Canada," Bank of Canada Working Paper 2003-26.

Kim D.H. and A. Orphanides (2007): "The Bond Market Premium: What Is It, And How Can We Measure It?", *BIS Quarterly Review*, June, 27-40.

Lapp, J.S, D.K. Pearce and S. Laksanasut (2003): "The Predictability of FOMC Decisions: Evidence from the Volcker and Greenspan Chairmanships," *Southern Economic Journal*, 70, 2, 312-327.

Longstaff, F.A. (2000): "The Term Structure of Very Short-Term Rates: New Evidence for the Expectations Hypothesis," *Journal of Financial Economics*, 58, 397-415.

Piazzesi, M. and E.T. Swanson (2008): "Futures Prices as Risk-Adjusted Forecasts of Monetary Policy," *Journal of Monetary Economics*, 55, 4, 677-691.

Reid, C. (2007): "The Canadian Overnight Market: Recent Evolution and Structural Changes," *Bank of Canada Review*, Spring, 15-29.

Rigobon, R. and B. Sack (2002): "Measuring the Reaction of Monetary Policy to the Stock Market," *Quarterly Journal of Economics*, 118, 2, 639-669.

Soderlind, P and L. Svensson (1997): "New Techniques to Extract Market Expectations from Financial Instruments," *Journal of Monetary Economics*, 40, 2, 383-429.

Appendix

We begin with a brief review of the "expectations hypothesis of the term structure". Denote the price at date t of a default-free pure-discount bond that pays 1 with certainty at date $t+n$ as $b_t^{(n)}$. The continuously-compounded yield on this bond, $r_t^{(n)}$, is defined as

$$b_t^{(n)} \equiv \exp(-nr_t^{(n)}),$$

or

$$r_t^{(n)} = -\frac{1}{n}\log(b_t^{(n)}).$$

We refer to the interest rate, or short-rate, as the yield on the bond with shortest maturity under consideration, $r_t \equiv r_t^{(1)}$. The one-period forward rate, $f_t^{(n\to n+1)}$, implicit in the price of a bond is defined in a similar way,

$$b_t^{(n+1)} = b_t^{(n)}\exp(-f_t^{(n\to n+1)}),$$

or

$$f_t^{(n\to n+1)} \equiv \log(b_t^{(n)} / b_t^{(n+1)}).$$

The expectation hypothesis relates, under risk neutrality, the forward rate to the expectation of a comparably timed future short-rate.

$$E_t r_{t+n} = f_t^{(n\to n+1)}, \tag{A1}$$

which, in turn, can be expressed as

$$E_t r_{t+n} = (n+1)r_t^{(n+1)} - nr_t^{(n)}, \tag{A2}$$

Imagine now that the next FAD is two weeks ahead. In that case, the expectation hypothesis implies that the expected policy rate (that we approximate by its one-week counterpart) has to be equal to the one-week forward rate starting in two weeks. This can be computed from both two and three-week yields. However, a market derived interest rate does not exist for those two dates. Yet, we can make use of a break-even rate calculation in order to approximate the one-week forward rate five and six-weeks ahead (see Johnson, 2004). The break-even calculation relies on the assumption that the overnight rate will only change on FAD dates (no inter-meeting moves). We can then take a market interest rate from that matures after the FAD in the question. Since we know the current overnight rate, the number of weeks to and after the FAD, and the interest rate over the entire time period, we can calculate the implied interest rate after the FAD – known as the break-even rate. The break-even rate is the rate that would make an investor indifferent (no-arbitrage) between investing in the market interest rate over the whole period or receiving the current overnight rate up until the FAD and then the break-even rate over the remaining of the time period.

Retake the example where the FAD is two weeks away, and a BA or OIS one-month rate is available. This interest rate on the 30 day OIS contract is assumed to encompass 15 days at the current overnight rate and 15 days at possibly a different overnight rate (ignoring credit and term premia). Given that we have the market yield of the BA/OIS and the current overnight rate, we can solve for the break-even yield, i.e. the implied no-arbitrage rate embodied in the BA for the period after the FAD.

In the instance where a move is not fully priced in, we can then back-out from the break-even yield the probability of a change in the overnight rate at the next FAD. For example, if the break-even yield was 3.80% and the current overnight rate of 4%, this would imply that the market was assigning roughly an 80% chance of a 25 bp ease in policy at the next FAD. In creating the benchmark model we use the cut-off of greater than 50 per cent or 12.5 basis points to determine whether the benchmark model suggests a policy change or not. In the case of two FADs ahead the same logic is used with a second set of the cut-offs applied. For instance if the break-even yield was 3.55 per cent and the overnight rate was 4 per cent, we would conclude that one rate decrease of 25 basis points was fully factored in, along with an 80 percent chance of second cut. Since the probability of the second cut is above our 50 per cent threshold, it is taken as being equivalent to predicting two rate cuts in the benchmark model.

Table 1

Changes in Target for the Overnight Rate: May 2001 – December 2006

	Small Change ± 25 bps	**Large Change ± 25 bps**	**Subtotal**
Increase	12	0	12
Decrease	11	3	14
Subtotal	23	3	**Total:** 26

Data range from May 2001 to December 2006 (54 observations)

Table 2
Ordered Probit Estimation

	1 FAD Ahead	2 FADs Ahead	4 FADs Ahead	8 FADs Ahead
3-month spread	17.441**	7.111**	2.423**	1.139*
	[9.673]	[9.917]	[3.579]	[2.098]
1-year spread	2.761**	3.780**	4.225**	3.455**
	[3.314]	[7.777]	[7.315]	[8.768]
10-year spread	-0.695**	-0.468	-0.266**	-0.328**
	[4.290]	[5.063]	[2.668]	[4.469]
Number of FADs in sample	53	53	50	46
Pseudo R^2	0.74	0.44	0.19	0.17

Note: Absolute value of z-statistic in square brackets computed using Huber-White standard errors. Data is weekly from May 2001 to December 2006 (290 observations). ** and * denote significance at the 1% and 5% level, respectively

Table 3
Actual and Predicted Policy Rate Changes

Panel a: One FAD ahead

	Actual	Predicted	% Correct
-25 bps	67	68	89.55
No change	134	139	88.06
+25 bps	89	83	84.27
Total	290	290	87.24

Panel b: Two FADs ahead

	Actual	Predicted	% Correct
-50 bps	47	52	68.09
-25 bps	34	16	17.65
No change	96	123	82.80
+25 bps	57	44	31.58
+50 bps	59	55	67.80
Total	290	290	59.66

Panel c: Four FADs ahead

	Actual	Predicted	% Correct
- 50 bps	74	88	83.78
- 25 bps	6	0	0.00
No change	60	58	40.00
+ 25 bps	41	0	0.00
+ 50 bps	109	144	82.57
Total	290	290	60.69

Table 3 (cont.)
Actual and Predicted Policy Rate Changes

Panel d: Eight FADs ahead

	Actual	Predicted	% Correct
- 75 bps	49	96	67.35
- 50 bps	6	0	0.00
- 25 bps	20	0	0.00
No change	38	0	0.00
+ 25 bps	38	0	0.00
+ 50 bps	48	51	12.50
+ 75 bps	91	143	86.81
Total	290	290	40.69

Note: Predictions are assigned to outcome with highest probability as determined by our ordered probit model. Data is weekly from May 2001 to December 2006 (290 observations)

Table 4
Actual and Predicted Policy Rate Changes

Panel a: Two FADs ahead

	Actual	**Predicted**	**% Correct**
Decrease	81	74	83.95
No change	93	108	76.34
Increase	116	108	79.31
Total	290	290	80.35

Panel b: Three FADs ahead

	Actual	**Predicted**	**% Correct**
Decrease	80	96	81.25
No change	60	9	6.67
Increase	150	185	91.33
Total	290	290	72.07

Panel c: Four FADs ahead

	Actual	**Predicted**	**% Correct**
Decrease	75	73	61.33
No change	38	0	0.00
Increase	177	218	90.96
Total	290	290	71.38

Note: Predictions are assigned to outcome with highest probability as determined by our ordered probit model

Table 5
Comparison of Benchmark Models: % of Correct Predictions

	Change		*Direction*
	One FAD Ahead	**Two FADs Ahead**	**Two FADs Ahead**
Naïve	63.10	12.07	32.41
Constant Probability	46.21	32.07	51.72
Expectations Hypothesis	82.69	38.62	61.38
Ordered Probit	87.24	59.66	79.66

Note: Predictions are assigned to outcome with highest probability as determined by the corresponding model