Modeling changes in the Australian cash rate target

Anneke Exterkate
310472

Bachelor Thesis
Econometrics and Operational Research - Major Econometrics
Erasmus University Rotterdam

July 2, 2010

Abstract

This thesis examines the way in which the Australian cash rate target changes over time. The analysis is based on research by Hamilton and Jordà (2002), who proposed to use two models: an autoregressive conditional hazard model to model the durations between two target changes, and an ordered probit model to model the sign and magnitude of a target change. In this way, forecasts can be made of when the central bank will change its rate target, and how large this change will be.

To investigate the adequacy of these models, this thesis uses data on past target changes of the Reserve Bank of Australia to estimate the parameters of the models. Next, one-month-ahead as well as one-year-ahead forecasts are made, and it turns out that the models perform relatively well on predicting one month ahead, but there are some problems with the explanatory variables so that the one-year-ahead forecasts are not that good. It is concluded that the two used models can be used to model the Australian cash rate target, although the models turn out to be somewhat different then the ones that were used by Hamilton and Jordà to forecast the federal funds rate target of the United States.

Contents

1 Introduction

The Reserve Bank of Australia (RBA) determines the level of short-term interest rates in Australia. An important instrument for the RBA to control this rate is their cash rate target. This target is set by the Reserve Bank Board, a committee that meets approximately once per month (eleven times a year, the first Tuesday of each month, except January) to discuss Australian monetary policy. Each time the Board meets, the members decide whether to change the rate target or not. Subsequently, this target is used by the trading desk of the RBA in their daily market operations. Because the cash rate target is set by the RBA, it is an interesting indicator of how monetary policy is translated into practice. Therefore, it is of great interest for traders to get insight in when and how the cash rate target will change.

There are several possibilities to model changes in the cash rate target. One way to do this has been proposed by Hamilton and Jordà (2002). In their study, they modeled the duration between and the magnitude of changes in the federal funds rate target for the Federal Reserve System of the United States. For this purpose, Hamilton and Jordà used an autoregressive conditional hazard model (ACH model) to model the durations between rate changes, and an ordered probit model to predict the size of the change of the rate target.

The main objective of this thesis is to model changes in the Australian cash rate target by applying the models mentioned before, to see whether the models of Hamilton and Jordà can also be applied to rate targets other than those of the Federal Reserve System. Therefore, the research question of this thesis is as follows:

How can changes in the Australian cash rate target be modeled, and is it possible to predict future changes with these models?

Because Hamilton and Jordà showed that their proposed models are capable to make a good prediction of future changes in the American federal funds rate target, it seems reasonable to expect that those models will work for the Australian rate target, too. But the Australian economy differs from the economy of the United States, so that this cannot be said on beforehand. For example, changes in the American economy might have a large impact on the Australian economy, while the reverse impact will not be that large. It cannot be said on beforehand whether or not the explanatory variables that turned out to be useful for the US federal funds rate target, are also useful for the Australian cash rate target.

This thesis is organized as follows. In Section 2, the used ACH model and ordered probit model are clarified. Section 3 gives a description of the data, and Section 4 shows the results of the application of both models. To find out whether the estimated models are useful in practice, Section 5 evaluates the forecast quality of the models. Finally, Section 6 concludes by answering the research question stated above and by discussing some limitations of the research.

2 Models

In this section, the two used models are described. Section 2.1 describes the autoregressive conditional hazard model and Section 2.2 explains the ordered probit model.

2.1 Autoregressive conditional hazard model

The autoregressive conditional hazard (ACH) model as proposed by Hamilton and Jordà (2002) is an extension of the autoregressive conditional duration (ACD) model designed by Engle and Russell (1998). To model the duration between events, in this case the time between two consecutive target changes, let u_n denote the length of time between the n-th and the $(n+1)$-th target change. Engle and Russell specified the expectation of u_n, given past observations u_{n-1}, u_{n-2}, ..., as ψ_n, with the following equation in the case of an ACD(r, m) model:

$$\psi_n = \sum_{j=1}^{m} \alpha_j u_{n-j} + \sum_{j=1}^{r} \beta_j \psi_{n-j} \tag{1}$$

where r is the number of lagged expectations of durations and m is the number of past durations that is taken into account. This equation does not include a constant term, because it was decided always to include a constant term in z_{t-1}, which is defined below.

Because the Reserve Bank of Australia (RBA) can change the cash rate target only once a month, define $N(t)$ as the number of target changes from the first month that was observed until t, where t is the time in months. For example, if the first target change takes place in the fifth month that was observed and the second target change takes place three months later, so in month $t = 8$, $N(t)$ is defined as follows:

$$N(t) = \begin{cases} 0 & \text{for } t = 1, 2, 3, 4 \\ 1 & \text{for } t = 5, 6, 7 \\ 2 & \text{for } t = 8, 9, ... \end{cases} \tag{2}$$

Now equation (1) can be rewritten using (2), as ψ_n only changes each time a target change occurred. The conditional expectation of the duration until a next target change becomes

$$\psi_{N(t)} = \sum_{j=1}^{m} \alpha_j u_{N(t)-j} + \sum_{j=1}^{r} \beta_j \psi_{N(t)-j}. \tag{3}$$

Next, define Υ_{t-1} as a vector that contains past information, that is, a vector that consists of past target changes up to month $t-1$ and explanatory variables in month $t-1$. The possible explanatory variables will be discussed in Section 3.2. Define the conditional probability of a change in the target at time t, given this Υ_{t-1}, as

$$h_t = \mathrm{P}[N(t) \neq N(t-1) | \Upsilon_{t-1}] \tag{4}$$

which represents the probability that a target change occurred in month t, given past information. Hamilton and Jordà (2002) assumed that this probability can be expressed as

$$h_t = \frac{1}{\psi_{N(t-1)} + \delta' z_{t-1}} \tag{5}$$

where $\delta' z_{t-1}$ is a linear combination of known information (Υ_{t-1}). To ensure that this probability is restricted to $0 \leq h_t \leq 1$, (5) is replaced with

$$h_t = \frac{1}{\lambda[\psi_{N(t-1)} + \delta' z_{t-1}]} \tag{6}$$

where

$$\lambda(v) = \begin{cases} 1.0001 & \text{if } v \leq 1 \\ 1.0001 + \frac{0.2(v-1)^2}{0.01+(v-1)^2} & \text{if } 1 < v < 1.1 \\ 0.0001 + v & \text{if } v \geq 1.1. \end{cases}$$

This function of $\lambda(v)$ is chosen, first of all to ensure that for $v \leq 1$ the probability in (6) cannot be larger than 1. For values of v equal to or larger than 1.1, $\lambda(v)$ is the same as the original denominator in (5), only 0.0001 is added because this was also done for $v \leq 1$; $\lambda(v)$ starts at 1.0001 and then becomes larger as v increases. For values of v between 1 and 1.1, this function for $\lambda(v)$ is chosen to ensure a smooth transition between 1.0001 and 1.1001.

Let x_t be a dummy variable that is equal to 1 if the target changes during month t, and 0 otherwise. Then the probability of observing x_t given past information Υ_{t-1} is

$$\mathrm{P}[x_t|\Upsilon_{t-1}; \theta_1] = h_t^{x_t}(1-h_t)^{1-x_t}, \tag{7}$$

with θ_1 the vector of parameters $\theta_1 = (\delta', \alpha', \beta')'$. It follows from (7) that the log likelihood conditional on past information, contained in Υ_{t-1}, is equal to

$$L_1(\theta_1|\Upsilon_{t-1}) = \sum_{t=1}^{T}[x_t \log(h_t) + (1-x_t)\log(1-h_t)]. \tag{8}$$

To obtain maximum likelihood estimates, the function $L_1(\theta_1|\Upsilon_{t-1})$ has to be maximized with respect to θ_1, under the following restrictions:

$$\begin{array}{rl} \alpha_j \geq 0 & \text{for } j \in 1, ..., m \\ \beta_j \geq 0 & \text{for } j \in 1, ..., r \\ 0 \leq \beta_1 + ... + \beta_r \leq 1 & \\ \sum_{j=1}^{m} \alpha_j + \sum_{j=1}^{r} \beta_j < 1 & \end{array}$$

where the latter is required to obtain stationarity (see also Hamilton, J.D. and Jordà, Ò. (2002), pp. 1138).

2.2 Ordered probit model

The ordered probit model is used to describe the size and magnitude of a target change, given that there is a target change in a particular month. So, given that the target changes, the question is whether it will go up or down, and by how many percentage points.

Let w_{t-1} denote a vector of variables, observed in month $t-1$. Because factors that are observed in the previous month may influence the size of the target change in this month, w_{t-1} may have an influence on the target change in month t and therefore the variables in w_{t-1} can be used as explanatory variables. To incorporate all of these variables into one new variable, we define an unobserved latent variable y_t^*:

$$y_t^* = \pi' w_{t-1} + \varepsilon_t, \tag{9}$$

where π is a vector of parameters and where $\varepsilon_t | w_{t-1} \sim$ NID(0,1).

Now suppose that there are k possible sizes of the target change: s_1, s_2, ..., s_k, where $s_1 < s_2 < ... < s_k$. In practice, in the past years the target changes were of many different sizes. To ensure that each group contains enough observations, the target changes are classified into k groups at first. Conditional on the fact that there is a target change in month t, the observed discrete target change y_t is related to the latent variable in (9) as follows:

$$y_t = \begin{cases} s_1 & \text{if } y_t^* \in (-\infty, c_1] \\ s_2 & \text{if } y_t^* \in (c_1, c_2] \\ \vdots & \\ s_k & \text{if } y_t^* \in (c_{k-1}, \infty), \end{cases} \tag{10}$$

where $c_1 < c_2 < ... < c_{k-1}$ are threshold parameters. By combining (9) and (10), the following expression can be derived for the probability that the target changes by s_j, given that the target changes:

$$\text{P}[y_t = s_j | w_{t-1}, x_t = 1] = \text{P}[c_{j-1} < \pi' w_{t-1} + \varepsilon_t \leq c_j] \tag{11}$$

for $j = 1, 2, ..., k$, with $c_0 = -\infty$ and $c_k = \infty$. Because ε_t is standard normal distributed, this can be rewritten as

$$\text{P}[y_t = s_j | w_{t-1}, x_t = 1] = \begin{cases} \Phi(c_1 - \pi' w_{t-1}) & \text{for } j = 1 \\ \Phi(c_j - \pi' w_{t-1}) - \Phi(c_{j-1} - \pi' w_{t-1}) & \text{for } j = 2, 3, ..., k-1 \\ 1 - \Phi(c_{k-1} - \pi' w_{t-1}) & \text{for } j = k \end{cases} \tag{12}$$

where Φ denotes the cumulative standard normal distribution function. The logarithm of this probability (conditional on w_{t-1} and $x_t = 1$), written as $l(y_t | w_{t-1}; \theta_2)$, is equal to

$$l(y_t|w_{t-1};\theta_2) = \begin{cases} \log[\Phi(c_1 - \pi' w_{t-1})] & \text{if } y_t = s_1 \\ \log[\Phi(c_j - \pi' w_{t-1}) - \Phi(c_{j-1} - \pi' w_{t-1})] & \text{if } y_t = s_2, s_3, ..., s_{k-1} \\ \log[1 - \Phi(c_{k-1} - \pi' w_{t-1})] & \text{if } y_t = s_k. \end{cases} \quad (13)$$

Here, $\theta_2 = (\pi', c_1, c_2, ..., c_{k-1})'$, and the conditional log likelihood of this ordered probit model can be written as

$$L_2(\theta_2|\Upsilon_{t-1}) = \sum_{t;x_t=1} l(y_t|w_{t-1};\theta_2). \quad (14)$$

The sum of logarithms of probabilities in this function only includes months where there was actually a target change ($x_t = 1$). Because $x_t = 0$ in all other months, (14) can be written as

$$L_2(\theta_2|\Upsilon_{t-1}) = \sum_{t=1}^{T} x_t l(y_t|w_{t-1};\theta_2). \quad (15)$$

The log likelihood function $L_2(\theta_2|\Upsilon_{t-1})$ has to be maximized with respect to θ_2, with the restriction that $c_j > c_{j-1}$ for $j = 1, 2, ..., k-1$.

2.3 Combination of models

It has to be mentioned that the previous models and their parameters can be modeled apart from each other (so $L_1(\theta_1|\Upsilon_{t-1})$ and $L_2(\theta_2|\Upsilon_{t-1})$ can be maximized separately), as long as the two log likelihood functions have no parameters in common. This can be done separately, because we have to model the joint probability distribution of x_t and y_t, conditional on explanatory variables in the past, which are captured in Υ_{t-1}. This joint probability distribution can be written as

$$\mathrm{P}(x_t, y_t|\Upsilon_{t-1}) = \mathrm{P}(x_t|\Upsilon_{t-1};\theta_1)\mathrm{P}(y_t|x_t, \Upsilon_{t-1};\theta_2). \quad (16)$$

Because we want to maximize the log likelihood, the objective function becomes

$$\begin{aligned} L &= \sum_{t=1}^{T} \log[\mathrm{P}(x_t, y_t|\Upsilon_{t-1})] \\ &= \sum_{t=1}^{T} \log[\mathrm{P}(x_t|\Upsilon_{t-1};\theta_1)] + \sum_{t=1}^{T} \log[\mathrm{P}(y_t|x_t, \Upsilon_{t-1};\theta_2)] \\ &= L_1(\theta_1|\Upsilon_{t-1}) + L_2(\theta_2|\Upsilon_{t-1}). \end{aligned} \quad (17)$$

This sum of $L_1(\theta_1|\Upsilon_{t-1})$ and $L_2(\theta_2|\Upsilon_{t-1})$ can be optimized either together or separately; both methods will lead to the same results.

3 Descriptive data analysis

This section describes the data. Section 3.1 shows the observed target changes and their associated durations, and Section 3.2 discusses the candidate explanatory variables.

3.1 Observed cash rate target changes

A plot of the time series of the Australian cash rate target is given in Figure 1, where some kind of cyclical pattern is visible.

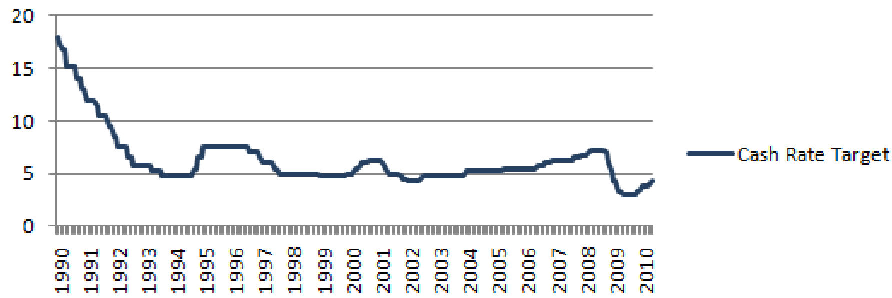

Figure 1: Time series of the Australian cash rate target (1990-2010)

Corresponding to this figure, Table 1 gives an overview of all observed changes in the Australian cash rate target, from January 1990 until the last target change up to now, which occurred on the fifth of May 2010[1]. As can be seen in this table, as well as in Figure 1, the cash rate target was much higher in the early nineties than it is nowadays. This is also the case for the target rates of other countries. Especially in the last two years, 2009 and 2010, many rate targets declined very quickly due to the international financial crisis. Further, most of the times, a negative target change is followed by another negative target change, and positive changes tend to be followed by positive changes.

The table shows further that most target changes occurred in the first week of a month. This is due to the fact that the Reserve Bank Board, which decides whether to change the target or not, meets the first Tuesday of each month, except for January. For the same reason, as can be seen in Table 1, there was never more than one target change in a month. Although the models that were introduced in Section 2 can be applied to daily data, it is not useful in this case to use data with a daily or weekly frequency. Therefore, monthly data will be used to model the target changes.

A frequency table of all durations between 1990 and 2010 can be found in Figure 3 in the Appendix. As can be seen in this figure, in many cases the time between two target changes was only one or two months. Only a few durations were over a year.

[1]These dates and changes are taken from the web site of the Reserve Bank of Australia: http://www.rba.gov.au/statistics/cash-rate.html

Date of change	Target value	Target change	Duration in days	Duration in months	Date of change	Target value	Target change	Duration in days	Duration in months
1990, January 23	17.00 to 17.50	-0.50 to -1.00			2001, February 7	5.75	-0.50	189	6
1990, February 15	16.50 to 17.00	-0.50	23	1	2001, March 7	5.50	-0.25	28	1
1990, April 4	15.00 to 15.50	-1.00 to -1.50	48	2	2001, April 4	5.00	-0.50	28	1
1990, August 2	14.00	-1.00	120	4	2001, September 5	4.75	-0.25	154	5
1990, October 15	13.00	-1.00	74	2	2001, October 3	4.50	-0.25	28	1
1990, December 18	12.00	-1.00	64	2	2001, December 5	4.25	-0.25	63	2
1991, April 4	11.50	-0.50	107	4	2002, May 8	4.50	0.25	154	5
1991, May 16	10.50	-1.00	42	1	2002, June 5	4.75	0.25	28	1
1991, September 3	9.50	-1.00	110	4	2003, November 5	5.00	0.25	518	17
1991, November 6	8.50	-1.00	64	2	2003, December 3	5.25	0.25	28	1
1992, January 8	7.50	-1.00	63	2	2005, March 2	5.50	0.25	455	15
1992, May 6	6.50	-1.00	119	4	2006, May 3	5.75	0.25	427	14
1992, July 8	5.75	-0.75	63	2	2006, August 2	6.00	0.25	91	3
1993, March 23	5.25	-0.50	258	9	2006, November 8	6.25	0.25	98	3
1993, July 30	4.75	-0.50	129	4	2007, August 8	6.50	0.25	273	9
1994, August 17	5.50	0.75	383	13	2007, November 7	6.75	0.25	91	3
1994, October 24	6.50	1.00	68	2	2008, February 6	7.00	0.25	91	3
1994, December 14	7.50	1.00	51	2	2008, March 5	7.25	0.25	28	1
1996, July 31	7.00	-0.50	595	20	2008, September 3	7.00	-0.25	182	6
1996, November 6	6.50	-0.50	98	3	2008, October 8	6.00	-1.00	35	1
1996, December 11	6.00	-0.50	35	1	2008, November 5	5.25	-0.75	28	1
1997, May 23	5.50	-0.50	163	5	2008, December 3	4.25	-1.00	28	1
1997, July 30	5.00	-0.50	68	2	2009, February 4	3.25	-1.00	63	2
1998, December 2	4.75	-0.25	490	16	2009, April 8	3.00	-0.25	63	2
1999, November 3	5.00	0.25	336	11	2009, October 7	3.25	0.25	182	6
2000, February 2	5.50	0.50	91	3	2009, November 4	3.50	0.25	28	1
2000, April 5	5.75	0.25	63	2	2009, December 2	3.75	0.25	28	1
2000, May 3	6.00	0.25	28	1	2010, March 3	4.00	0.25	91	3
2000, August 2	6.25	0.25	91	3	2010, April 7	4.25	0.25	35	1
					2010, May 5	4.50	0.25	28	1

Table 1: Calendar of changes in the Australian cash rate target

3.2 Candidate explanatory variables

Many factors may influence the decision whether or not to change the cash rate target. For example, if the federal funds rate of the United States increases by one or two percent, this might be an indicator for the RBA to increase their own target as well. Further, large changes in the Australian GDP or unemployment rate can have explanatory power for the level of the target. Therefore, the afore mentioned factors, as well as many other possible explanatory variables, are collected. An overview of all variables can be found in Table 2. Whether these variables can truly explain target changes or not is discussed in Section 4.

Variable	Measurement scale
Australian economic measurements	
Gross Domestic Product growth	Quarterly growth in percentages
Consumer Price Index	Quarterly percentage change in the index
Unemployment	Monthly rate in percentages
Budget deficit(+)/surplus(-)	Annualized percentage change
Yield on government bonds, 5 years	Monthly yield in percentages
Consumer sentiment index	Monthly percentage change in the index
Foreign influences	
US $/AUS $ exchange rate	Monthly percentage change in the rate
US federal funds effective rate	Monthly rate change in percentage points
Japanese call rate (collateralized overnight)	Monthly rate change in percentage points
Trading desk variables	
Last target change	Last change that occurred
Reserve Bank Board meeting dates	Dummy: 1 if it is January, 0 elsewhere

Table 2: Candidate explanatory variables for the ACH- and ordered probit model

Some of these variables were originally in percentage points, while others were rates, and some of them where indexes. To ensure that all variables are captured in the models in the same measurement scales, all variables are transformed into percentages.

Next to this, data on Australian GDP and CPI are only available at a quarterly frequency, while the durations between target changes are measured monthly. It was decided to leave these explanatory variables in quarterly frequency and use each data point three months, because the most recent available information is expected to influence the target. For example, the data on GDP and CPI from the last quarter of a year, which becomes available at the end of December, is used to model possible target changes in the first three months of the next year.

Further, it is decided to leave some variables in percentages, while other variables are included as percentage changes. The unemployment is included as the monthly rate of unemployed persons compared to the Australian labour force. It is decided to leave this variable in monthly rate instead of monthly change, because in this way the unemployment rate is expected to have the largest influence on the cash rate target and thus on the inflation. This is also known as the Phillips curve; according to this curve, a high unemployment rate indicates low inflation and a low unemployment rate is related to high inflation. See Figure 4 in the Appendix for an example of a short-run Phillips curve.

The yield on government bonds is also included as the monthly yield in percentages instead of the monthly change, because it is expected that the level of this yield influences the cash rate target, instead of the changes in the yield. In contrast to this, the US federal funds effective rate and the Japanese call rate are transformed into monthly changes in the rate. This is because we want to model whether or not the target will change and if it changes by how many percentage points; it is expected that not the level of foreign interest rates has an influence on the Australian target, but that changes in foreign rates will have an influence. For example, if the federal funds rate was 5 percent last month, this gives us no information on whether or not the Australian target will change this month, because we do not know whether the federal funds rate had increased or decreased to that 5 percent. However, if we take the change in federal funds rate into account, a change of -0.75 percent might indicate the Australian rate target to decrease as well.

Furthermore, the available data on the budget deficit/surplus represent the change in the Australian budget for each month, compared to the previous month. But, in some months of the year the Australian government has to do large expenditures, while in other months of the year the earnings are very large, for example due to revenues obtained from Australian tax payers. This causes the budget to fluctuate from month to month. To get rid of this effect, the data were annualized and then the monthly change in the budget with respect to the previous month was measured.

Moreover, some influences from outside Australia are included in the table. Because the United States are very influential for Australia in terms of trade, the exchange rate of the US dollar as compared to the Australian dollar may affect the Australian target. For the same reason, the federal funds effective rate of the United States is included as possible explanatory variable. As the most important trading partner of Australia is Japan, the Japanese call rate is also included.

There is also a dummy included for January, because this is the only month in the year in which the board does not meet. Therefore, when they meet again in February, they have the information of the past two months at their disposal (instead of one month), which may influence the decision whether or not to change the cash rate target.

Lastly, Hamilton and Jordà used the 6-month Treasury bill spread, relative to the US federal funds rate. In their models it turned out that this variable has a significant influence on target changes. Unfortunately, this variable cannot be included in the models for the Australian target. That is, the Treasury bill yield for Australia is not available between May 2002 and March 2009 and therefore is useless for this research.

For some of the before mentioned variables, a certain influence on the Australian cash rate target can be expected. For instance, a growth of the Australian GDP is expected to have a positive influence on the cash rate target; according to macroeconomic theory interest rates are procyclical, which means that whenever GDP shows a large growth, interest rates tend to increase (see also Burda, M.C. and Wyplosz, C. (2005), pp. 187). This is due to the fact that an increase in GDP increases the demand for money, which causes the interest rates to be higher.

A same influence can be expected of the CPI; large inflation causes higher interest rates, so the coefficient for the CPI is expected to be positive. As already said before, according to the Phillips curve, high inflation implies a low unemployment rate. For this reason, the unemployment rate is expected to have an opposite effect on the rate target.

Further, it is expected that the Australian cash rate target follows the same cyclical pattern as the interest rates for the United States and Japan. Therefore, changes in the foreign interest rates that are included in the models are expected to have a positive effect on the Australian rate.

See Table 9 in the Appendix for the correlation coefficients between all employed variables. Five variables show a significant correlation with changes in the Australian target value, namely GDP growth (positive correlation with target changes), the government yield (negative), the exchange rate between the US dollar and the Australian dollar (positive), the US federal funds effective rate (positive), and the last target change (positive). Other variables, such as changes in the CPI index and the Australian budget deficit/surplus, do not seem to have much in common with changes in the Australian target value.

Most of the signs of the correlation coefficients are in line with the expectations. For example, the GDP growth shows a significant positive correlation with changes in the target value, which means that if the GDP growth increases, the target value is also likely to increase in the next month. The same holds for the exchange rate; if the exchange rate becomes larger, the target value will also increase. This is also explainable: an increasing exchange rate between the US dollar and the Australian dollar means that the Australian dollar becomes more valuable. This is the case when the demand for the Australian dollar increases, which also explains a higher target value.

The correlation between the US federal funds effective rate and the Australian cash rate target is also positive. This implies that an increase (a decrease) in the US federal funds rate is followed by an increase (a decrease) in the Australian target value.

In view of these correlation coefficients, it may be expected that some of the proposed variables have an influence on changes in the cash rate target. See also the scatter diagrams (Figures 5 until 8 in the Appendix) for some important Australian economic measurements against the value of the cash rate target; from this we can also conclude that there is some correlation between the different variables.

To investigate possible effects of the variables on the cash rate target in another way, the behavior of these variables in the past can be examined. Is there for example a difference between the GDP growth during months where the target was changed and during months where the target was not changed? For this purpose, Table 3 shows the average values of all variables as described in Table 2 in three groups: the average values for months before the months in which a negative target change occurred, before the months in which the target remained the same, and before months in which the board decided to increase the target. To compare, a fourth row with the overall averages is added.

Target change?	GDP growth	CPI	Un-emp.	Budget	Govern. yield	Cons. sent.	US$ /AUS$	US fed. funds	Japan call	Last change
Negative	0.46	0.59	7.44	-17.62	7.87	-0.36	-0.18	-1.33	-0.02	-0.58
No	0.82	0.68	7.08	-4.56	6.73	0.48	-0.02	0.15	-0.03	-0.07
Positive	0.88	0.85	5.88	11.82	6.32	-0.20	0.02	1.60	0.01	0.22
All	0.78	0.69	7.00	-4.65	6.84	0.29	-0.03	0.11	-0.02	-0.11

Table 3: Average values for month $t-1$ compared to target changes in month t

From this table, it can be seen that for all variables there is indeed a difference present in the averages for the three groups, although for some variables this difference is larger than for others. For example, the GDP growth is on average the smallest in the months where the target was changed downwards (yet this growth is still positive, as there is an average GDP growth of approximately 0.78 percent over the whole time series). A similar thing can be seen for the unemployment rate: the average unemployment rate is much higher in months where the target change was negative. So, a larger unemployment rate might indicate a decrease in the rate target, which is explained before using the Phillips curve.

A notable number in this table is on the change in the Australian budget; in months where the target change was negative, the budget surplus became much larger. And, as can also be seen in Table 3, a decrease in the United States federal funds rate might cause the Australian target to be lower, which was also expected.

Generally, the average percentage changes for most variables are different for the three groups, where it seems that a negative target change shows the largest differences in variables compared to the other two groups. So, again, it can be expected that some of those variables may have an influence on the rate target decisions.

4 Results from maximum likelihood estimation

This section discusses the results that are obtained with the models from Section 2. In this case, the two models do not have any parameters in common. Because there is also no significant correlation between the parameters of both models, it is allowed to optimize the two models separately, as well as optimizing them together. After comparing these two ways of optimization it turned out that both methods lead to approximately the same results (not exactly the same results due to some insignificant correlation between the parameters), so in the next two paragraphs the results of both models are addressed separately. First, Section 4.1 describes the results of the ACH model and after this, in Section 4.2 the results of the ordered probit model are shown.

4.1 Modeling the target change decision

In this section, the model of Section 2.1 is applied to model the choice whether or not to change the cash rate target. First, we consider an ACH(1,1) model. This model takes one lagged duration and one lagged expectation of a duration into account, so that the model (3) becomes

$$\psi_{N(t)} = \alpha u_{N(t)-1} + \beta \psi_{N(t)-1}. \tag{18}$$

To start with, all variables in Table 2 are included in z_{t-1}. It turned out that, as expected, some of the explanatory variables were not significant. By using a top-down procedure, where the least significant variable is removed one at a time, we ended up with the model shown in Table 4. Apparently, because both α and β are estimated to be zero, the durations between past target changes are not useful to estimate the next duration. This is an important difference between the US federal funds rate target as modeled by Hamilton and Jordà and the Australian cash rate target: in their article (Hamilton, J.D. and Jordà, Ò (2002)) they found that α and β are significantly different from zero, which means that past durations between target changes have an influence on future target changes. Apparently, the moments on which the Reserve Bank Board of Australia decides to change the target are not influenced by durations in the past. If we check for autocorrelation in the durations, we find that there is no significant autocorrelation between them. This can be seen in Figure 9 in the Appendix, and when performing an AR(1) model of the durations on their first lag, the coefficient turns out to be insignificant (p-value = 0.88). Therefore, it is not surprising that both α and β are estimated to be zero; there is just no significant relation between durations and past durations.

Parameter	Variable $(t-1)$	Estimate	Std. error	P-value
δ_1	January dummy	-2.2297	0.5481	0.0001
δ_2	GDP growth	2.1662	0.6668	0.0020
δ_3	CPI change	0.9791	0.3705	0.0121
δ_4	Unemployment rate	0.4687	0.0700	0.0000
δ_5	US federal funds rate	4.7522	1.5802	0.0043
δ_6	US$/AUS$ exchange rate	0.2059	0.0718	0.0066
δ_7	Japanese call rate	-6.3341	2.9335	0.0388

Table 4: Parameter estimates for ACH model, monthly data 1990-2010 (240 observations)

The coefficients in Table 4 seem reasonable. For example, the coefficient of -2.2297 for the dummy variable for January causes the probability of changing the target, as can be found in (6), to be higher in February. On the other hand, a large GDP growth or CPI change causes the probability of changing the target to be lower, because their coefficients are positive. Further, the US federal funds rate and the Japanese call rate have opposite effects: where an increase in the federal funds rate causes the probability of changing the target to be lower, an increase in the Japanese rate causes this probability to be higher.

To provide some further intuition for the estimates in Table 4, we consider two random chosen dates: one month where the target changed and one month where the target did not change. We consider January 2001, where the target did not change. Here, because α is estimated to be zero, the value of the probability in equation (6) is equal to

$$\frac{1}{\delta' z_{Dec'00}},$$

where the coefficients for δ can be found in Table 4 and $z_{Dec'00}$ consists of the seven explanatory variables in Table 4 for December 2000. This probability equals

$$\frac{1}{\lambda[-2.23\cdot 0+2.17\cdot 0.20+0.98\cdot 0.31+0.47\cdot 6.26+4.75\cdot -0.11+0.21\cdot 4.75-6.33\cdot 0]}=0.24,$$

so that it is more likely that the target does not change in January 2001. We also calculated this probability for the next month, February 2001, where the target did actually change. Here, using the explanatory variables for the month January 2001, the probability which follows from the estimated model is equal to

$$\frac{1}{\lambda[-2.23\cdot 1+2.17\cdot 0.20+0.98\cdot 0.31+0.47\cdot 6.37+4.75\cdot -0.42+0.21\cdot 1.57-6.33\cdot 0.01]}=0.99.$$

So, the target was extremely likely to change in this month. At first sight, it looks like this model is able to estimate the decision to change the target or not very well.

Another way to evaluate the adequacy of the model in-sample, is by using a hit rate table. The total percentage of months in which the probability is modeled well can be found by adding the two percentages in Table 5 of where the model estimation gives a probability that is corresponding to what is observed; this total percentage of good estimations is equal to 80.4%.

		Estimated with model		Total
		$h_t < 0.5$	$h_t > 0.5$	
Observed	$x_t = 0$	183 (76.3%)	3 (1.3%)	186 (77.5%)
	$x_t = 1$	44 (18.3%)	10 (4.2%)	54 (22.5%)
Total		227 (94.6%)	13 (5.4%)	240 (100%)

Table 5: Hit rate table for ACH model

Unfortunately, in many cases the model fails to give a probability of larger than 0.5 for months in which a target change was observed. Only few real target changes are predicted by the model. But, for over 80 percent of all months the prediction corresponds to the observation in that month. See also Figure 10 in the Appendix for a plot of the probabilities of a target change according to the model, along with the observed target changes. In this figure it can be seen that in many cases where there was a target change but h_t was estimated to be lower than 0.5, h_t was actually somewhat higher than in months where the target did not change. The problem that h_t is lower than 0.5 in many cases while it had to be larger than 0.5, is partly explained because over the whole time period, there were much more months in which the target did not change than months in which the target did actually change. Because of this, h_t has the tendency not to become higher than 0.5 too easy.

4.2 Modeling the size of the target change

To estimate the parameters in the ordered probit model, as already mentioned in Section 2.2, the target changes have to be divided into a few groups. In this case, they are classified in four groups; given that a target change occurs, the change can be either negative or positive, and given the direction of the change, we make a distinction between small and large changes. Because of the frequencies of different target changes in the data (which can be found in the last column of Table 6), the target changes are classified as can be found in Table 6. The last column in this table gives an indication of how the target is expected to change, based on the most frequent target changes in the data. This expected change is respectively equal to s_1, s_2, s_3, s_4 as stated in (10).

Size of change (%)	Defined as	Frequency in data	Expected change: s_j
$-1.5 < y_t < -0.5$	Large decrease	15	-1 (s_1)
$-0.5 \leq y_t < 0$	Small decrease	18	-0.25 (s_2)
$0 < y_t \leq 0.25$	Small increase	19	0.25 (s_3)
$0.25 < y_t < 1.5$	Large increase	4	1 (s_4)

Table 6: Classification of target changes in four categories

Next to the parameters in π (pertaining to the explanatory variables in the model), there are three threshold parameters that have to be estimated (c_1, c_2 and c_3; the three thresholds between the four groups). Again, we start with the full model and we eliminate the least significant coefficient one at a time. The resulting model is shown in Table 7.

Parameter	Variable ($t-1$)	Estimate	Std. error	P-value
π_1	GDP growth	1.1963	0.4333	0.0088
π_2	US federal funds rate	1.2915	0.7623	0.0950
π_3	US$/AUS$ exchange rate	0.1189	0.0584	0.0504
π_4	Last target change	1.8697	0.4192	0.0000
c_1		-1.1902	0.4323	0.0090
c_2		0.8197	0.4323	0.0662
c_3		2.9451	0.5501	0.0000

Table 7: Parameter estimates for the ordered probit model, monthly data 1990-2010 (240 obs.)

As expected, the last target change up to month $t-1$ has a positive significant effect on the direction of the next target change; most of the time positive target changes are followed by other positive target changes and negative changes are followed by negative changes.

Further, it turned out that the GDP growth has a significant positive influence on the direction of the target change. This is not completely unexpected, because as the GDP grows, it can be expected that the target value will also become larger, as already mentioned in Section 3.2. Another unsurprising factor is the US federal funds rate; as expected, this rate has a positive influence on the Australian rate. So, a positive change in the US rate implies a positive change in the Australian rate and a negative change in the American rate causes a negative change in the Australian cash rate target.

After testing the three threshold variables, it turned out that they were significantly different from each other, which means that the group division as in Table 6 does not have to be changed; no groups have to be merged.

For this ordered probit model, we can also make a hit rate table. We consider the four different groups of target changes; the hit rates can be found in Table 8.

	Estimated with model				Total
Observed	$-\infty < y_t < -0.5$	$-0.5 \leq y_t < 0$	$0 < y_t \leq 0.25$	$0.25 < y_t < \infty$	
$-\infty < y_t < -0.5$	12 (21.4%)	3 (5.4%)	0 (0%)	0 (0%)	15 (26.8%)
$-0.5 \leq y_t < 0$	3 (5.4%)	15 (26.8%)	0 (0%)	0 (0%)	18 (32.2%)
$0 < y_t \leq 0.25$	0 (0%)	1 (1.8%)	18 (32.1%)	0 (0%)	19 (33.9%)
$0.25 < y_t < \infty$	0 (0%)	0 (0%)	2 (3.6%)	2 (3.6%)	4 (7.2%)
Total	15 (26.8%)	19 (33.9%)	20 (35.7%)	2 (3.6%)	56 (100%)

Table 8: Hit rate table for ordered probit model

As can be seen in this table, almost all target changes are estimated correctly; the hit rate is equal to the sum of percentages on the diagonal, which is 83.9%. Only a few target changes were assigned to another group, but never more than one group from the real target change (that is, a large decrease is never estimated as being an increase).

5 Forecast evaluation

This section will pay attention to forecasts; are the used models able to forecast future target changes and their size and magnitude? For this purpose, one-month-ahead forecasts are made as well as one-year-ahead forecasts. This forecasts are compared to the real target changes, and after this we can conclude whether the used models are adequate enough or not.

5.1 One-month-ahead forecasts

In order to make out-of-sample forecasts for one month ahead, the model has to be re-estimated a few times using different sub-samples of the available data. We have chosen to use an expanding window for this, to make one-month-ahead forecasts for the years 2000 until 2010. For example, to make forecasts for the year 2005, the model is re-estimated for the sub-sample 1990-2004 and then the parameter estimates of this model are used to make one-month-ahead forecasts for each month in 2005. So, the parameters are not monthly updated but only yearly, because it is expected that taking one more month into account when estimating the parameters will not make a large difference. Therefore, only once a year the sub-sample is expanded and the parameters are re-estimated.

Further, it is chosen to use for all sub-samples the same variables that were significant in the full model, so the same variables will be used as in Tables 4 and 7. Not all variables will be significant in all sub-sample-based models, but to make it possible to compare all models, it is chosen not to remove insignificant variables.

Define i_{t+1} as the one-month-ahead forecast of the target value, based on information on time t, captured in Υ_t, which contains all explanatory variables at time t (z_t and w_t). Then the expectation of i_{t+1} can be written as

$$\mathrm{E}(i_{t+1}|\Upsilon_t) = (1 - h_{t+1})i_t + h_{t+1}\sum_{j=1}^{4}(i_t + s_j) \cdot [\Phi(c_j - \pi' w_t) - \Phi(c_{j-1} - \pi' w_t)] \quad (19)$$

where h_{t+1} is calculated from (6), s_j are as given in Table 6, $c_0 = -\infty$ and $c_4 = \infty$. All parameters in h_{t+1}, as well as the parameters in π and all c_j's are re-estimated for each sub-sample. The parameter estimates for all models can be found in Table 10 in the Appendix.

After doing this, Figure 2 displays the actual target value along with the predicted target values for each month. As we can see, in many cases the forecast predicts a target value for month $t + 1$ that is very close to the observed target value in month t. Only in month $t + 1$, when it became clear what the real target change in that month was, the forecast predicts for the next month a value that is close to the target value after its change in month $t + 1$.

However, in some cases the forecast predicts a large increase or decrease in the target very well. As this is a bit difficult to see from the figure, some calculations are performed on the forecast errors, which are defined as the differences between the actual target values and the predicted target values for all months. First of all, the mean forecast error is equal to -0.036 percent. This means that on average, the predicted value of the target is 0.036

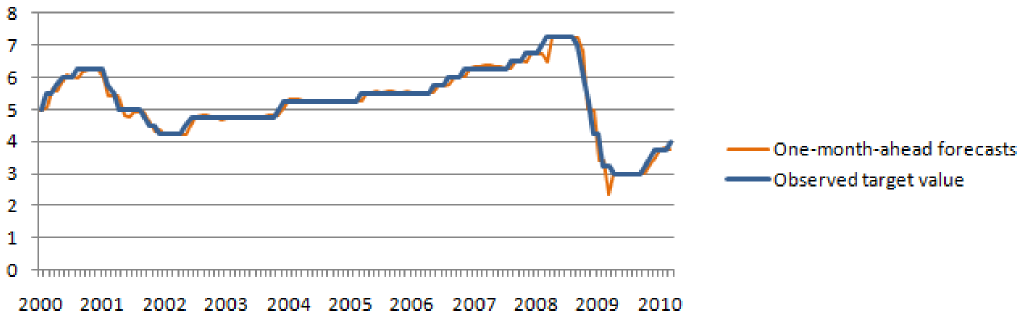

Figure 2: One-month-ahead forecasts, 2000-2010

percent lower than the actual value. After testing whether this value differs significantly from zero, it turned out that the p-value of this test is equal to 0.06, so the null hypothesis cannot be rejected. This means that the forecast errors are unbiased.

Next, the mean absolute error and the the mean squared prediction error are calculated. They are equal to 0.118 percent and 0.044 percent, respectively, which are both not very large. However, these values cannot be compared to other values such as the variance of the residuals, because the forecasts are made with different models. So, by using these criteria, we cannot conclude whether or not the forecasts are "good"; we can only see that the forecast errors are not that large.

By performing a Mincer-Zarnowitz regression, that is, regressing the actual value of the target on a constant and on the forecast of the target, it turns out that the coefficient of the intercept is significantly different from zero, and that the coefficient for the forecast is slightly different from one (the null hypothesis should be rejected, because the t-value $= -2.696$). This means that the actual value of the target is significantly different from the forecasted target value.

5.2 One-year-ahead forecasts

One-year-ahead forecasts are much more difficult to perform. This is because with the models from Section 4, we always need information at time t to make a forecast for the target value on time $t + 1$. So, if we want to forecast the target value on time $t + 12$, we need the values of the explanatory variables at time $t + 11$, which are obviously not yet available on time t. To make one-year-ahead forecasts, we have therefore a couple of possibilities that will be discussed in this section.

The first possibility is to re-estimate the models from Section 4 by using the effect of the explanatory variables of one year ago. So, instead of using Υ_{t-1}, we can now use Υ_{t-12}. Re-estimating the models for the whole sample and eliminating the insignificant variables one at a time, the models in Tables 11 and 12 in the Appendix remain.

It can be expected that those models perform much worse, because it seems logical that information on, for example, the GDP or the US federal funds rate from last month can influence the target this month, but does information of a year ago really influence the target value of next month? This can be seen by making one-year-ahead forecasts. Again, once a year the parameters are re-estimated by using an expanding window and using the same variables in all sub-sample models. In this case, we use the data from 1990 until 1999 to make forecasts for December 2000 until November 2001; for December 2001 the models based on 1990-2000 are used. We have made forecasts for the years 2000 until 2004, because it was expected that these years were the least difficult to forecast; in these years, the economy was much more stable than in the last few years, due to the international financial crisis. The forecasts can be found in Figure 11 in the Appendix.

As can be seen from this figure, this method is not able to predict the value of the target based on information of a year ago. The forecasts look more like the actual value of the target 12 months before, so this method does not need any further review.

Another way to perform these one-year-ahead forecasts is the way in which Hamilton and Jordà performed their 12-month-ahead forecasts. This can be done by using the same models as before (that is, use information in month $t-1$ to estimate the target value on time t), and predicting the values of the explanatory variables for the next eleven months by simulation. The biggest problem from this approach is that it is very hard to predict the value of, for instance, the GDP growth in a year. Hamilton and Jordà did this by using an autoregressive distributed lag (ARDL) model for each explanatory variable. For the GDP growth it can be written as follows:

$$GDP_t = \beta_1 + \beta_2 i_t + \beta_3 i_{t-1} + \beta_4 GDP_{t-1} + \varepsilon_t, \tag{20}$$

where i_t denotes the target value in month t, i_{t-1} denotes the target value in the month before and ε_t is the error of the regression. By iterating we can each time forecast the target value for the next month and using that predicted value, we can make a forecast for all explanatory variables for that next month in the way of (20). Using these predictions, the target value can be forecasted one month further, and so on, until the forecast for the target value for month $t+12$ is reached. The forecasts can be found in Figure 12 in the Appendix.

From this figure it can be seen that apparently, this method also does not work well. The forecasts are too far away from the observed target values, and therefore, no further calculations on this method are needed. The main problem with this method is that it is very difficult to predict future values of GDP, CPI, unemployment and so on.

The third way to make one-year-ahead forecasts is less convenient, because this way makes use of data that is not yet observed. However, it is a method to evaluate the estimated models. In this way, the explanatory variables do not need to be simulated because we use the explanatory variables that are observed from month t until month $t+11$ to make a forecast for the target value in month $t+12$. This is actually what was done in Section 5.1; because we updated the model only yearly, the forecast for December 2000 was based on the model for 1990-1999, and because we use the data of November 2000 this can be seen as a kind of one-year-ahead forecast. Of course, at the end of 1999 we do not know what value the GDP will have in November 2000, so these are not very convenient one-year-ahead forecasts. Therefore they are used as one-month-ahead forecasts.

6 Conclusion and discussion

The main question of this thesis was:

How can changes in the Australian cash rate target be modeled, and is it possible to predict future changes with these models?

First of all, it turned out that past durations between target changes do not have any significant influence on future changes in the cash rate target. We found that, with the use of the autoregressive conditional hazard model and the ordered probit model, the target changes can be well modeled according to the hit rate tables; many target changes are well predicted. Also relatively good one-month-ahead forecasts can be made and therefore these models are useful to model changes in the Australian cash rate target. However, one-year-ahead forecasts were much less convenient, but there are two problems for that. The first problem is that the target value is very difficult to predict, given past target changes; we have seen that there is no autocorrelation between them. Therefore, decisions of the Reserve Bank Board are very hard to predict. The second main problem is the issue on forecasting the explanatory variables. It turned out that it is very hard to forecast for example the Australian GDP or the US federal funds rate, and therefore one-year-ahead forecasts are difficult to make. We therefore suggest to take a look at these models; how can the explanatory variables be modeled in a better way? This was beyond the scope of this thesis but would be a good topic for further research.

Another problem that arose, was on the fact that there is no board meeting in January and therefore the board has two months of new information at their disposal at the meeting in February. Although we captured a dummy for this in the model, the model is not fully able to compensate for this. This is a large disadvantage of this approach and could be another topic for future research.

One more thing that has to be mentioned, is that in the models in this thesis, we used explanatory variables in month $t-1$ to model the target change in month t. But, it can be questioned whether the data of month $t-1$ is already available at the board meeting at the beginning of month t; maybe some information, for example about the GDP or unemployment rate in a particular month, may not be available until month $t+1$ or $t+2$. This was not taken into account in this thesis, but it is expected that at the board meeting in month t the members of the board already have some information about month $t-1$, although the precise data might not be available yet.

To conclude, the models perform relatively well and can be used to predict the target value at least one month ahead, although some further research would be desirable. It can also be concluded that there are some differences between the target changes in the US federal funds rate target and the Australian cash rate target; in the latter one there is no significant autocorrelation between the durations between changes and therefore the autoregressive part of the model is not useful for the Australian cash rate target.

References

Burda, M.C. and Wyplosz, C. *Macroeconomics.* New York: Oxford University Press, 2005.

Easton, S. and Pinder, S. "Predicting Reserve Bank of Australia Interest Rate Announcements: Beware of the Target Rate Tracker." *The Australian Economic Review* 40, no. 1 (2007): 119-22.

Engle, R.F. and Russell, J.R. "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data." *Econometrica* 66 (1998): 1127-62.

Hamilton, J.D. and Jordà, Ò. "A Model of the Federal Funds Rate Target." *Journal of Political Economy* 110, no. 5 (2002): 1135-67.

Reserve Bank of Australia. "The Reserve Bank's Open Market Operations." *Reserve Bank of Australia Bulletin* (June 2003): 1-7.

A Appendix

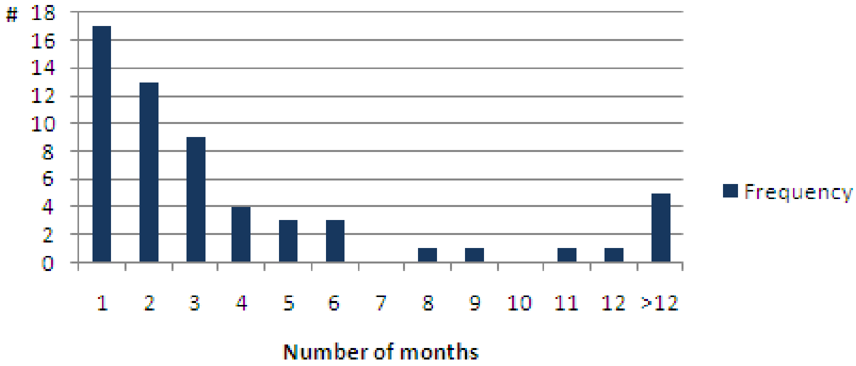

Figure 3: Frequency of durations between target changes in months (1990-2010)

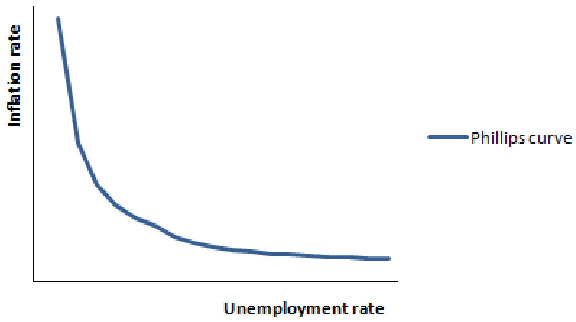

Figure 4: Example of a short-run Phillips curve

Variable	Target change	GDP growth	CPI	Unemp.	Budget	Govern. yield	Cons. sent.	US$ /AUS$	US fed. funds	Japan call	Last change
Target change	1*	0.35*	0.04	-0.10	0.07	-0.22*	0.04	0.23*	0.29*	0.04	0.38*
GDP growth		1*	-0.31*	0.08	0.12*	-0.35*	0.03	0.08	0.22*	-0.01	0.31*
CPI			1*	-0.25*	0.09	0.28*	-0.13*	-0.11*	-0.11*	0.15*	0.20*
Unemp.				1*	0.01	0.41*	0.07	-0.07	0.01	-0.36*	-0.39*
Budget					1*	0.04	-0.10	-0.14*	0.01	-0.01	0.05
Govern. yield						1*	-0.05	-0.04	-0.03	0.04	-0.25*
Cons. sent.							1*	0.18*	0.04	-0.05	-0.06
US$/AUS$								1*	0.11*	0.02	0.13*
US fed. funds									1*	0.13*	0.29*
Japan call										1*	0.09
Last change											1*

Table 9: Correlations between variables (* = significant)

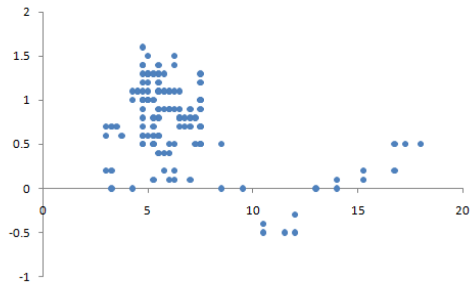

Figure 5: Scatter for target (x) against GDP growth (y)

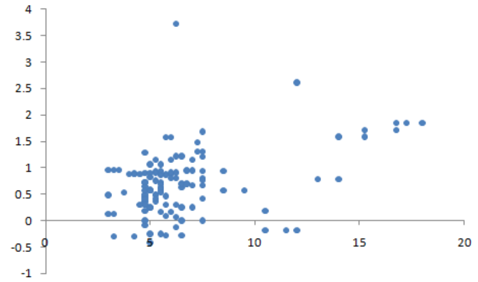

Figure 6: Scatter for target (x) against CPI change (y)

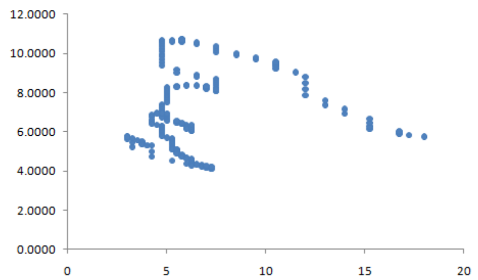

Figure 7: Scatter for target (x) against unemployment rate (y)

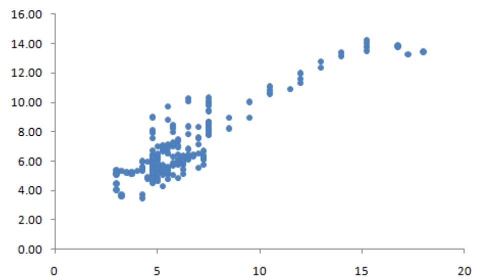

Figure 8: Scatter for target (x) against government yield (y)

	'90-'99	'90-'00	'90-'01	'90-'02	'90-'03	'90-'04	'90-'05	'90-'06	'90-'07	'90-'08	'90-'09
δ_1	1.23	-1.83	-0.97	-0.63	-0.07	-0.05	0.04	0.06	0.15	-2.49*	-2.23*
δ_2	2.63*	2.48*	1.73*	2.15*	2.18*	2.38*	2.42*	2.40*	2.59*	2.51*	2.17*
δ_3	2.17*	1.95*	1.97*	2.08*	2.14*	2.40*	2.52*	2.36*	2.10*	1.85*	0.98*
δ_4	0.26*	0.26*	0.30*	0.25*	0.25*	0.27*	0.28*	0.28*	0.29*	0.38*	0.47*
δ_5	1.47	1.23	4.87*	3.47*	3.64*	4.19*	4.57*	4.46*	4.77*	6.36*	4.75*
δ_6	-0.24	-0.21	0.21	0.01	0.01	0.00	-0.01	-0.02	-0.05	0.24*	0.21*
δ_7	-11.26*	-10.38*	-10.09*	-10.40*	-10.69*	-12.20*	-12.82*	-12.10*	-10.53*	-8.71*	-6.33*
π_1	1.30*	1.16*	1.32*	1.36*	1.38*	1.38*	1.37*	1.28*	1.25*	1.23*	1.20*
π_2	4.02*	4.12*	1.76	1.89	2.02*	2.02*	2.14*	2.43*	2.27*	1.15	1.29*
π_3	-0.10	-0.00	-0.06	-0.03	-0.03	-0.03	-0.02	-0.01	0.02	0.09	0.12*
π_4	0.70	0.92*	1.49*	1.51*	1.54*	1.54*	1.57*	1.66*	1.73*	1.89*	1.87*
c_1	-0.22	-0.49	-0.83	-0.87*	-0.91*	-0.91*	-0.96*	-1.12*	-1.16*	-1.09*	-1.19*
c_2	2.00*	1.45*	1.59*	1.44*	1.38*	1.38*	1.32*	1.08*	0.98*	0.95*	0.82*
c_3	2.36*	2.42*	2.47*	2.60*	2.77*	2.77*	2.81*	2.83*	2.87*	2.89*	2.95*

Table 10: Parameter estimates for the ACH and ordered probit model for each sub-sample, used to forecast (* = significant)

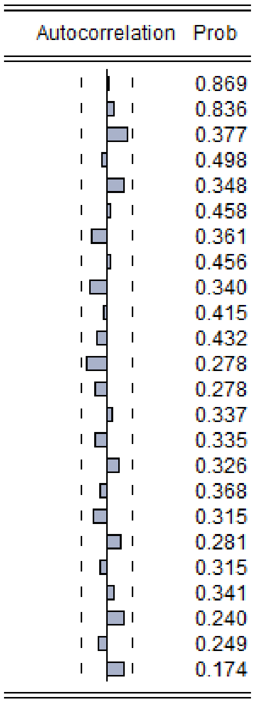

Figure 9: Correlogram for autocorrelations between durations

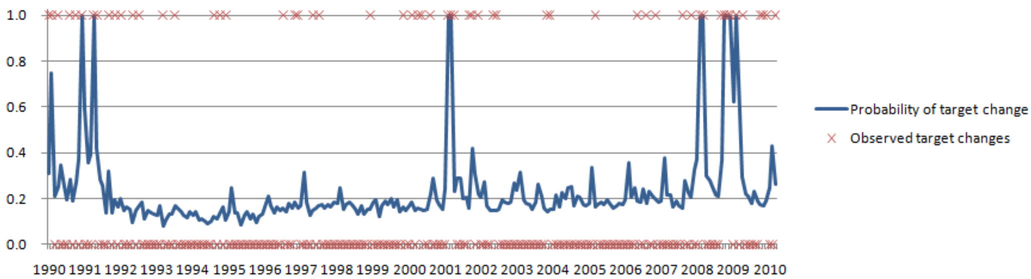

Figure 10: Probabilities of target changes over time

Parameter	Variable $(t-12)$	Estimate	Std. error	P-value
δ_1	GDP growth	2.3558	0.8988	0.0129
δ_2	Unemployment rate	0.4204	0.0868	0.0000

Table 11: Parameter estimates for ACH model, monthly data 1990-2010 (240 observations)

Parameter	Variable $(t-12)$	Estimate	Std. error	P-value
π_1	GDP growth	0.9756	0.3388	0.0063
c_1		-0.1903	0.3012	0.3268
c_2		0.8323	0.3144	0.0120
c_3		2.1977	0.3768	0.0000

Table 12: Parameter estimates for the ordered probit model, monthly data 1990-2010 (240 observations)

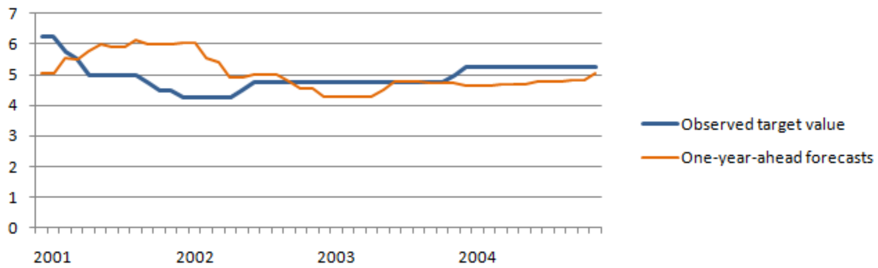

Figure 11: One-year-ahead forecasts, 2000-2004

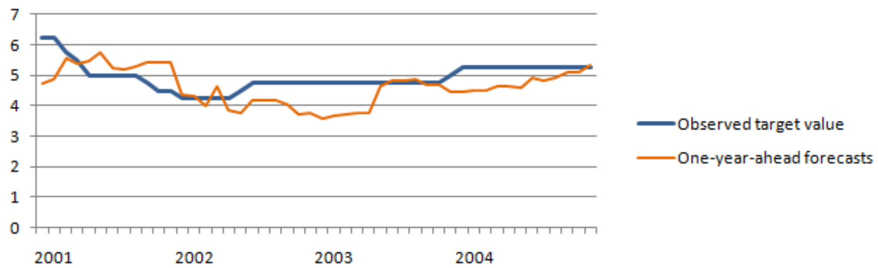

Figure 12: One-year-ahead forecasts, 2000-2004