



Reserve Bank of New Zealand Analytical Notes

Inflation expectations curve: a tool for monitoring inflation expectations

AN2016/01

Michelle Lewis

March 2016

Reserve Bank of New Zealand Analytical Note Series
ISSN 2230-5505

Reserve Bank of New Zealand
Po Box 2498
Wellington
NEW ZEALAND

www.rbnz.govt.nz

The Analytical Note series encompasses a range of types of background papers prepared by Reserve Bank staff. Unless otherwise stated, views expressed are those of the authors and do not necessarily represent the views of the Reserve Bank.

Non-technical summary¹

Inflation expectations play a key role for inflation targeting central banks. In particular, well-anchored inflation expectations demonstrate that a central bank has credibility and make it easier than otherwise for policymakers to achieve inflation targets. Hence, central banks monitor inflation expectations closely to understand current pricing dynamics and assess the credibility of achieving the target.

This *Analytical Note* describes an approach for monitoring and assessing inflation expectations using yield curve modelling techniques, where a term structure is estimated from various surveys of inflation expectations.

These expectations curves use information from a wide range of inflation expectation surveys, rather than relying on a few individual surveys or a few horizons. This information is condensed into two intuitive concepts. The first is long-run inflation expectations, which can be interpreted as the perceived inflation target focus and used to assess whether inflation expectations are well anchored. The second concept is ‘expected time to target’, where the shape of the curve shows how survey participants expect inflation to converge to the perceived inflation target focus.

The results show that surveyed inflation expectations in New Zealand are currently well anchored and have been throughout the inflation targeting regime. Long-run expectations have always been well within the relevant official target band and the expected time to target has, on average, been around two-and-a-half years.

Different *Policy Targets Agreements* have had a detectable effect on the expectations curves, with the entire curve shifting higher when the inflation target mid-point was increased in 2002 (the target range was also narrowed). Since the increased focus on the mid-point of the target range was introduced in 2012, the long-run inflation expectation has shifted lower and closer to the two percent mid-point.

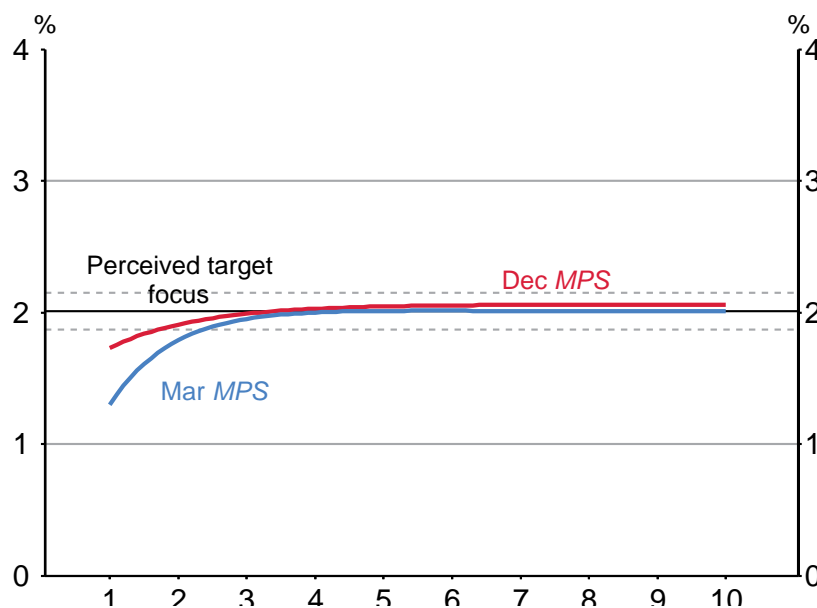
¹I would like to give special thanks to Leo Krippner for assistance with the modelling approach and advice. I am also grateful to Güneş Kamber, John McDermott, Geordie Reid and other colleagues for helpful discussions.

1 INTRODUCTION

Inflation expectations play a crucial role in an inflation targeting framework, where well-anchored inflation expectations make it easier than otherwise for a central bank to achieve the target. Hence, central banks put in a lot of effort to monitor and understand inflation expectations from a variety of sources.²

In this *Note* I describe an approach for monitoring and assessing inflation expectations, applying yield curve modelling techniques to inflation expectations from many different surveys and range of horizons. The Bank has used these techniques to help monitor inflation expectations and has published the expectations curve in the *Monetary Policy Statement (MPS)* since June 2015. Figure 1 shows the most recent expectations curve from the March 2016 *MPS*.

Figure 1: Inflation expectations curve in March 2016 *MPS*



Note: The perceived target focus is where the inflation expectations curve trends towards in the long run. The dotted grey lines reflect a band that is close to the perceived target focus (calculated as ± 1 standard deviation of the perceived target focus since the 2012 PTA).

²For example, in the Bank's core macro model, NZSIM, inflation expectations play a crucial role. Kamber et al. (2015) show that observing particular inflation expectations surveys has a material impact on the inflation dynamics within the model.

Using survey data and a term structure model provides a way of combining information from various survey data with different time horizons and of overcoming issues that are involved with using financial market data. Typically, when estimating inflation expectations from financial market data, term structure models are applied to nominal and inflation-indexed bonds. However, this approach is problematic for two key reasons. First, there are few inflation-indexed bonds on issue in New Zealand (this is also the case for many countries). The second issue is conceptual, where nominal and inflation-indexed bonds are required to calculate inflation expectations and these two bond types carry various risk premia that are different for each bond (for example, inflation uncertainty and liquidity). This makes interpretation of implied inflation expectations problematic in practice.³ Using many surveys of expectations negates these issues since the surveys relate directly to expectations, so there are no contaminating risk premia.

Applying a term structure to inflation expectation surveys is akin to model averaging for forecasting, where expectations of future outcomes are taken from a range of sources. Each of the surveys may have some bias in terms of forecasting future inflation but combining the surveys, in principle, improves the representation of inflation expectations over various horizons. The term structure approach also provides continuous curves, so inflation expectations for any standard horizon can be calculated.

I apply the Nelson-Siegel (1987) model to surveyed measures of inflation expectations, with horizons out to 10 years. This provides an intuitive approach to monitoring numerous inflation expectation surveys, condensing the wide range of information into two key concepts: the long-run level of inflation expectations (perceived target focus) and the expected time for inflation to return to target.

Recent studies that have applied a term structure to inflation expectation surveys have used it in the context of analysing the effect of central bank policy changes on inflation expectations. Lewis and McDermott (2015) use expectations curves to investigate how expectations changed following changes to the *Policy Targets Agreement* (PTA). Aruoba (2015) uses a term structure of inflation expectations to analyse the effects of the Federal Reserve's monetary policy following the global financial crisis of 2008. In related work, Mehrotra and Yetman (2014)

³See Bauer and Rudebusch (2015) for further discussion on the difficulties with interpreting inflation expectations from inflation indexed bonds.

assess the degree of anchoring in inflation expectations for 44 countries using survey data out to two years and a decay function.

2 INFLATION EXPECTATIONS SURVEYS

For New Zealand, there are over a dozen surveyed measures of inflation expectations, measuring headline CPI and wage inflation. I use 12 surveys of headline CPI inflation, with survey horizons ranging from one year to 10 years ahead (see table 1 and further discussion on the data is in the appendix). The full sample period is from 1996Q1 until 2016Q1, reflecting data availability.

Table 1 shows summary statistics for the survey measures, with sub-sample periods reflecting different *PTA* objectives. The average for all expectations has been within the inflation target range in each of the periods, although there has been some variation depending on which *PTA* prevailed. The average inflation expectation across all horizons rose from around 1.88 percent in the 1996Q4-2002Q2 period to 2.51 percent in the 2002Q3-2012Q2 period, and has since fallen to 2.12 percent.⁴

The standard deviation across horizons in each of the sample periods evolves as would be expected for well-anchored inflation expectations. The shorter horizon expectations show a larger variance, reflecting that shorter horizons are more influenced by the business cycle and the transmission lags for monetary policy. For instance, any transitory/temporary shock that hits the economy and is expected to effect inflation in the near-term (for example, a spike in oil prices, a one-off tax change, or a demand shock), should be reflected in the volatility of inflation expectation measures, where the survey measures fluctuate to show how inflation is expected to respond. As the horizon increases, the effect of the shock dissipates and hence the longer-term measures should show lower volatility relative to shorter-term measures, if expectations are well-anchored. This pattern is evident in table 1, where the standard deviation at the 10-year horizon is around half of the one-year horizon.

⁴The inflation target range in the 1996Q4-2002Q2 period was 0 to 3 percent. In the 2002Q3-2012Q2 period the target range was 1 to 3 percent. In 2012Q3 the *PTA* was changed to have more focus on the mid-point of the 1 to 3 percent target range.

Table 1: Survey measure summary statistics

	1996Q1-2016Q1		1996Q4-2002Q2		2002Q3-2012Q2		2012Q3-2016Q1	
	$\bar{\pi}$	$\hat{\sigma}$	$\bar{\pi}$	$\hat{\sigma}$	$\bar{\pi}$	$\hat{\sigma}$	$\bar{\pi}$	$\hat{\sigma}$
RBNZ 1yr	2.27	0.63	1.98	0.47	2.69	0.51	1.66	0.33
ANZ 1yr	2.68	0.46	2.53	0.40	2.96	0.31	2.14	0.39
Aon 1yr	2.19	0.70	1.76	0.42	2.63	0.68	1.87	0.29
RBNZ 2yr	2.29	0.36	2.01	0.23	2.57	0.22	2.10	0.24
Consensus 2yr	2.23	0.60	1.83	0.49	2.59	0.54	2.07	0.15
Consensus 3yr	2.20	0.33	1.85	0.14	2.40	0.24	2.33	0.16
Aon 4yr	2.15	0.35	1.76	0.10	2.40	0.19	2.24	0.14
Consensus 4yr	2.16	0.35	1.73	0.12	2.40	0.17	2.30	0.18
Consensus 5yr	2.14	0.33	1.73	0.13	2.39	0.12	2.21	0.21
Consensus 6yr	2.14	0.29	1.77	0.08	2.36	0.13	2.21	0.16
Aon 7yr	2.14	0.29	1.82	0.14	2.34	0.14	2.25	0.09
Consensus 10yr	2.15	0.27	1.81	0.07	2.37	0.11	2.19	0.13

$\bar{\pi}$ is the average expectation for each survey and $\hat{\sigma}$ is the standard deviation for each survey.

Note latest data available for Consensus is October 2015.

3 MODEL

The Nelson-Siegel model is a natural approach for modelling the term structure of inflation expectations based on survey data. Firstly, it is a parametric approach, which means the underlying components can be given economic interpretations. This feature, along with the stable fit over time, makes the Nelson-Siegel class of models a popular modelling choice in the term structure literature and in practice at central banks.⁵ Secondly, while the Nelson-Siegel model needs an arbitrage-free adjustment when applied to bond data, the original form can be used in this analysis since the survey data relate directly to expectations and do not need to be adjusted for the market price of risk.⁶ However, I do assume consistency among the inflation expectation surveys to the extent that respondents, regardless of the specific survey, are giving their best assessment of future price inflation at the specified horizon. Survey-specific bias is accounted for in the Nelson-Siegel model, with that particular survey being treated more as an

⁵See Krippner (2015a) for a theoretical foundation for the Nelson-Siegel model and see Diebold and Rudebusch (2012) for a full discussion on Nelson-Siegel model applications.

⁶For comprehensive discussion on yield curve modelling concepts and the market price of risk, see Krippner (2015b).

anomaly.⁷

The Nelson-Siegel framework represents the term structure of forward rates or interest rates at a given point in time using a latent factor model with three coefficients and one parameter. The analogous Nelson-Siegel functional form applicable to annual inflation expectations, which is derived from the original Nelson-Siegel forward rate function presented later in equation 2, is:⁸

$$\pi_{NS}^e(t, \tau - 1, \tau) = L(t) + S(t) \cdot [\text{Slope loading}] + B(t) \cdot [\text{Bow loading}] \quad (1)$$

where $\pi_{NS}^e(t, \tau - 1, \tau)$ is the Nelson-Siegel annual inflation expectation, t is the observation date, τ is the expectation horizon, and $L(t)$, $S(t)$, and $B(t)$ are the Level, Slope, and Bow (sometimes called Curvature) coefficients. Those coefficient names refer to the shapes of the respective factor loadings that the coefficients multiply into. The Level factor loading has a value of one that does not vary by horizon, and is hence referred to as a long-run factor. The Slope and Bow factor loadings are respectively:⁹

$$[\text{Slope loading}] = \frac{1}{\phi} \{ \exp(-\phi[\tau - 1]) - \exp(-\phi\tau) \}$$

$$[\text{Bow loading}] = \frac{1}{\phi} \{ \exp(-\phi[\tau - 1]) - \exp(-\phi\tau) \} + \{ [\tau - 1] \exp(-\phi[\tau - 1]) - \tau \exp(-\phi\tau) \}$$

where ϕ is the decay parameter that controls how quickly the Slope and Bow factor loadings decay to zero as the horizon τ increases. The loading on the Slope factor starts near one and decays monotonically to zero as $\tau \rightarrow \infty$, and this loading is typically referred to as a

⁷The main alternative approach in the term structure literature is to use empirical techniques such as splines. These techniques are popular among, for example, traders that want to closely replicate all observed points on the curve. A key trade-off is in interpretability, where economic meaning can be given to the fundamental components of the Nelson-Siegel model but can't be given to non-parametric methods such as splines.

⁸I give special thanks to Leo Krippner for the derivation of the expressions in this section and for comments to improve the exposition.

⁹The Slope and Bow loading expressions differ from the typical Nelson-Siegel factor loadings used in Aruoba (2015) or for fitting yield curve data, because I am modelling surveys of expected annual rates at a future horizon. The Nelson-Siegel interest rate factor loadings would only be applicable to surveys of the average of expected inflation up to each horizon.

short-term factor. The Bow factor starts near zero, rises for short- and medium-term horizons, and then decays monotonically to zero as $\tau \rightarrow \infty$. Hence, this loading is typically referred to as a medium-term factor. The economic interpretations of these factors from the perspective of inflation expectations are discussed below.

Equation 1 is estimated as in Lewis (2015), i.e. by minimising the squared residuals of the fitted and actual inflation expectations over the entire sample of expectations data, where the optimal value for ϕ is globally estimated and values for the Level, Slope, and Bow are estimated for each expectations curve observation. The global optimisation of ϕ ensures the expectations curves are consistent across time as well as cross-sectionally.

In this *Note* I use two perspectives from the estimated model to monitor inflation expectations. First, the Level factor may be interpreted as a measure of long-run inflation expectations, which in turn provides a proxy for the inflation target focus as perceived by survey participants. Second, the combination of the Slope and Bow factors show how expectations converge to the perceived target. Hence, I can use those components to indicate the ‘expected time to target’ implied by the survey participants.¹⁰

4 RESULTS

In this section I assess the model’s fit to New Zealand data and show how inflation expectations curves have behaved since 1996.

4.1 MODEL FIT

The model fits the data well, with an average absolute error (difference between estimated and observed survey data) of only 13 basis points.¹¹ This compares well with the literature of fitting bond yields (see, for example, Lewis 2015 for New Zealand and Gurkaynak, Sack, and Wright 2007 for the US). Table 2 shows the average absolute difference between the fitted values and

¹⁰In practice, while the mid-point of the target band is important under the current PTA, the time taken for reaching the target band, and being comfortably inside, is also relevant.

¹¹A basis point is 100th of a percentage point.

actual data for the individual surveys and figure 2 provides an illustration.

These results show the Nelson-Siegel model fits some surveys better than others. For example, at the one-year horizon the absolute error for the ANZ and AON surveys are 36 and 24 basis points respectively, while the RBNZ survey is 13 basis points. As the time horizon increases, the fitted error generally decreases. This reflects the greater stability in expectations as the time horizon increases (illustrated in figure 2).

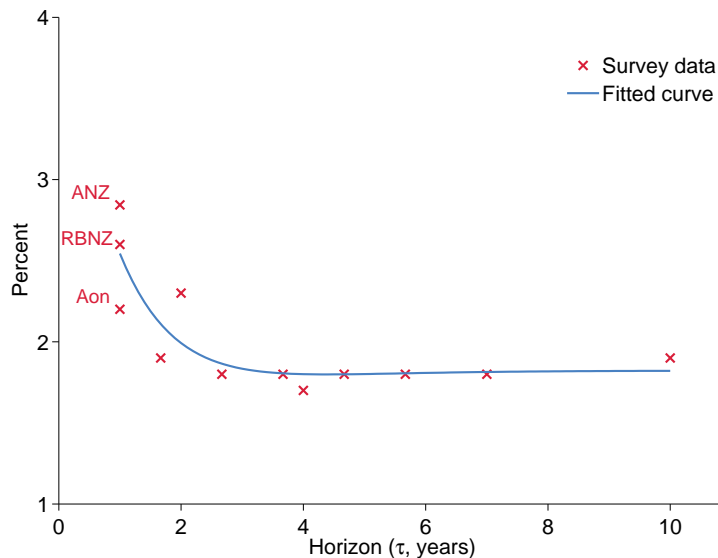
Table 2: Absolute error statistics (basis points)

	Average	Minimum	Maximum		Average	Minimum	Maximum
RBNZ 1yr	13	0	42	Aon 4yr	7	0	29
ANZ 1yr	36	2	98	Consensus 4yr	8	0	27
Aon 1yr	24	1	96	Consensus 5yr	7	0	26
RBNZ 2yr	12	1	39	Consensus 6yr	6	0	18
Consensus 2yr	16	0	80	Aon 7yr	7	0	26
Consensus 3yr	10	0	37	Consensus 10yr	6	0	16

The notation used for Consensus Economics surveys are for convenience. The moving horizons are controlled for in the Nelson-Siegel model.

Figure 2 illustrates how the fitted yield curve smoothes through the individual survey measures. In this example the one-year ahead RBNZ and Aon surveys are below the ANZ one-year survey and above the rest of the yield curve. The yield curve smoothes through the three observations, with the RBNZ survey still having significant influence due to the weighting in the model and taking all the other survey information into account.

Figure 2: Inflation expectations curve: actual and fitted



Note: The solid line is the fitted yield curve from the Nelson-Siegel model.

The crosses are the actual survey data. The example used here is for 2002Q2.

4.2 MONITORING INFLATION EXPECTATIONS

There are several ways to analyse the information in the inflation expectations curves that are relevant for monetary policy. In this *Note* I focus on two key aspects: the perceived inflation target focus and the expected time to target.

PERCEIVED INFLATION TARGET FOCUS

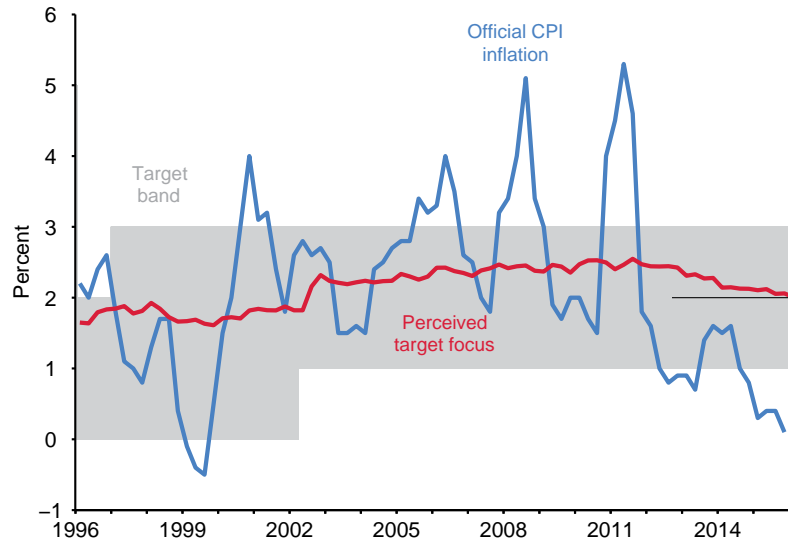
As previously discussed, the Level component can be interpreted as the perceived inflation target focus, which can be used to assess whether expectations are well-anchored to the official inflation target. The benefit from using this approach is that information from all the surveys is used rather than a single survey measure.¹²

Figure 3 shows the perceived inflation target focus along with the official inflation target range and headline CPI inflation since 1996. The perceived target focus has remained well within the target range since 1996 and has followed changes to the PTA (see Lewis and McDermott 2015

¹²The longer-run surveys certainly play an important role in determining the Level component but the surveys from other horizons also have influence.

for formal tests on how changes to the *PTA* affected inflation expectations). The figure also shows that the perceived inflation target focus has not always been strictly at the mid-point of the official inflation target. More recently, the perceived target has drifted lower from 2.5 percent in 2011 to around 2 percent.

Figure 3: New Zealand's inflation target and long-run inflation expectations



Note: The grey band represents the inflation target range and how it has changed over time.

The solid black line represents the focus on the mid-point of the band since September 2012.

EXPECTED TIME TO TARGET

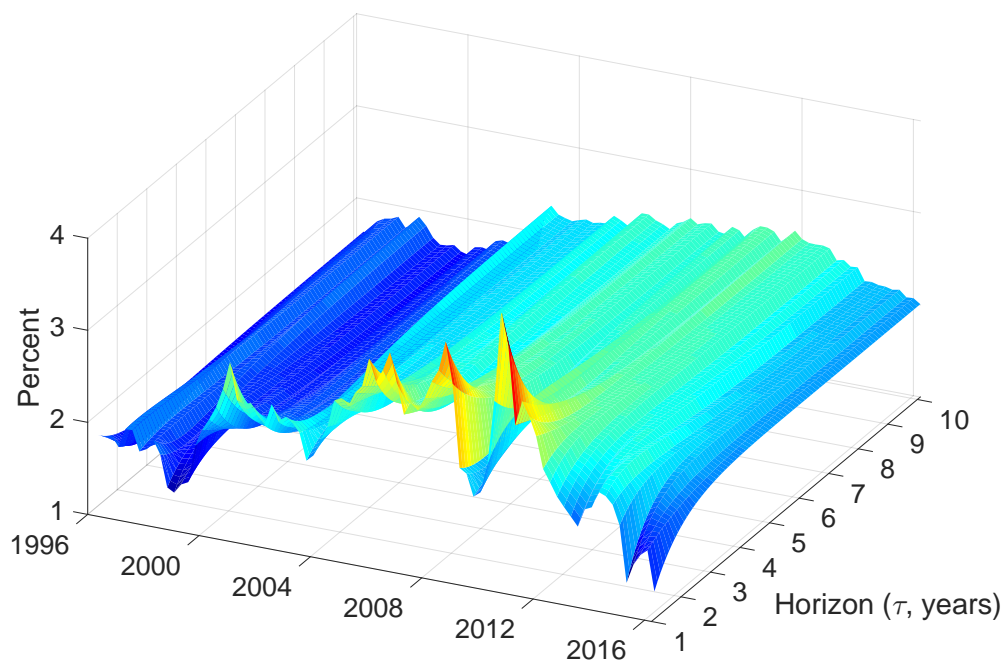
The second aspect for monitoring inflation expectations is how expectations converge to the perceived target. An inflation expectations curve consistent with flexible inflation targeting would have a smooth transition back to target following any deviations from target. In addition, the expected time to target would be consistent with the medium-term notion in the *PTA* (see Ford et al. 2015 for further discussion on the *PTA*). Further to this, the expected time to target can help with understanding the impact that inflation expectations will have on wage- and price-setting behavior at horizons relevant for forecasting inflation and setting monetary policy.

Figure 4 gives a heuristic view of how inflation expectations have evolved over time. The curves have had a variety of shapes but all curves show expectations smoothly transitioning back to the perceived target focus. During the mid-2000s when inflation outturns were outside the

target band, the expectations curves showed inflation returning inside the official target band within two years. The spike in inflation following the increase in Goods and Services Tax in 2010 is also evident, where inflation expectations increased markedly at short-horizons (to four percent) and returned within the official inflation target band in two years.

The different *PTAs* are also evident in figure 4. The entire yield curve lifted following the 2002 *PTA* change (light blue region) and has since shifted lower following the 2012 *PTA* change (darker blue region).

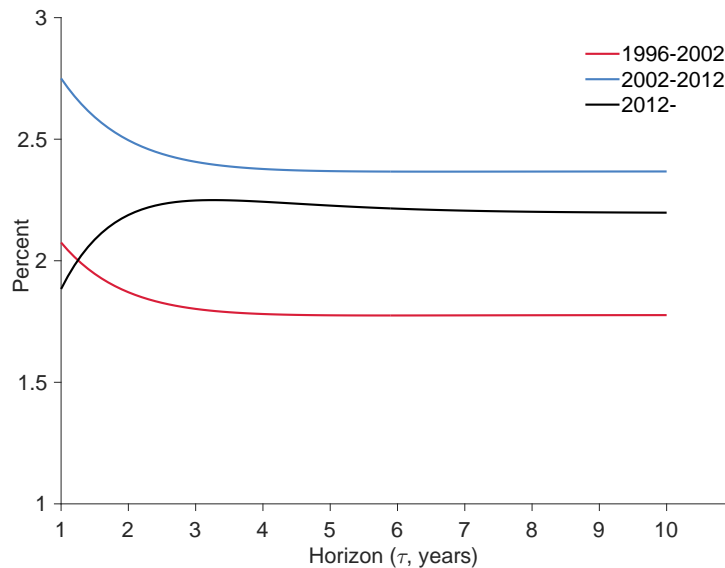
Figure 4: Evolution of inflation expectations curves



Note: The colours in the figure reflect a heatmap, where warmer colors represent higher inflation expectations.

Figure 5 shows the average inflation expectations curves during each *PTA*. Between 1996 and 2002, when the inflation target range was 0 – 3 percent, the expectations curves were slightly downward sloping and converged to around 1.8 percent. Between 2002 and 2012, when the target range was narrowed to 1 – 3 percent, the curve had a similar downward slope but the level convergence was around 2.5 percent. Since 2012, when the emphasis on the mid-point was introduced to the *PTA*, the long-end of the curve has shifted lower to around 2.2 percent on average.

Figure 5: Average inflation expectations for different PTAs



Note: The curves are the averages for each horizon during the different inflation targeting regimes.

An expected time to target can be directly calculated using the forward rate version of the Nelson-Siegel model.¹³ This approach provides a future point estimate for inflation at each horizon. Equation 2 gives the forward rate representation of equation 1:

$$\pi_f^e(t, \tau) = L(t) + S(t)e^{-\phi\tau} + B(t)\phi\tau e^{-\phi\tau} \quad (2)$$

In this exercise, I calculate how long it takes for the expectations curve to return and remain within a specified threshold around the perceived target focus. The threshold used is one standard deviation of the Level component in each of the different PTAs (1996 – 2002, 2002 – 2012, and 2012 – present).¹⁴ Using the standard deviation of the Level component reflects that

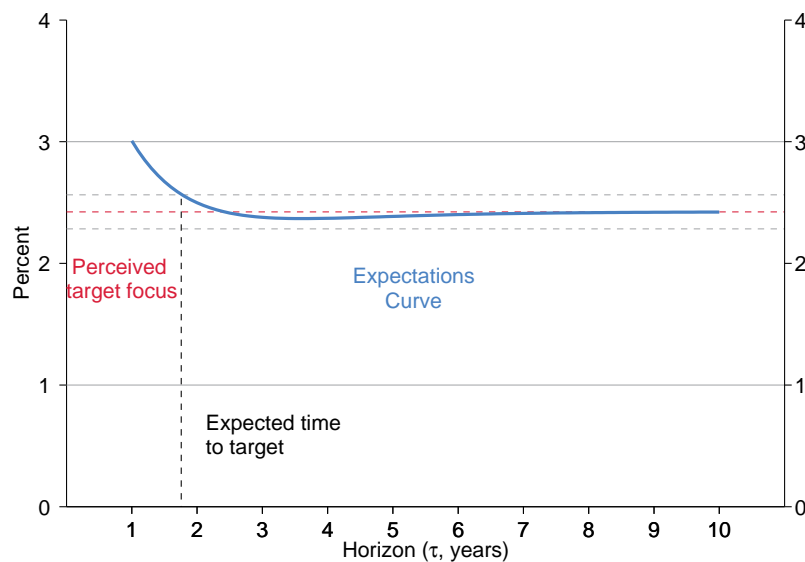
¹³Another approach could have been to calculate the half-life based on the decay parameter, ϕ . The Nelson-Siegel model has an exponential function on the Slope component and a polynomial function on the Bow component. If it were the case that the Bow was insignificant then the half-life calculation would be $\tau_H = -\frac{1}{\phi} \log(0.5)$. For example, with $\phi = 0.8$, $\tau_H = 10$ months. However, the Bow component is often statistically significant and thus this calculation omits an important dynamic.

¹⁴Different thresholds could be applied, which would give different results. For example, using the official target bands would show that the expectations curves are usually within the threshold, while using the point estimate of the perceived target (and no threshold) would result in an increase in the expected time to target given the often-slow decay to precisely reaching the Level.

the perceived target focus is not necessarily a fixed point and conceptually takes into account that the *PTA* has specified the inflation target as a range, rather than a point, through the Bank's history. The current *PTA* specifies the target as keeping future CPI inflation between 1 and 3 percent on average over the medium term, with a focus on keeping future average inflation near the 2 percent target mid-point. The choice of splitting the sample periods reflects the distinct changes in inflation expectations behavior in each period. The standard deviation of the Level component was 9 basis points in the 1996 – 2002 period, 11 basis points in the 2002 – 2012 period and 14 basis points since 2012.

Figure 6 illustrates how the expected time to target is calculated: the time horizon is traced out from where the expectations curve crosses the target threshold and remains within it.

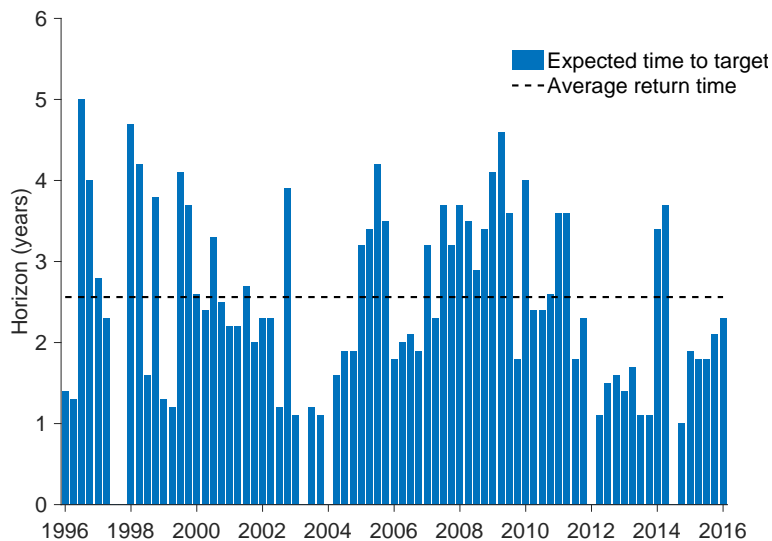
Figure 6: Stylised expected time to target



Note: solid grey lines represent the current 1-3 percent *PTA* inflation target range, the dashed red line is the perceived target at a point in time, and the dashed grey lines are the one standard deviation around the perceived target.

Figure 7 shows the expected time to target since 1996. On average, when inflation expectations are outside the threshold, survey participants expect inflation to return close to the perceived target focus in around two-and-half years.

Figure 7: Expected time to target



Note: zero values show the forward curve was inside the threshold at all horizons.
The average return time only includes instances where inflation expectations are outside the threshold.

5 CONCLUSION

Applying a term structure to inflation expectation surveys can be a useful tool for monitoring inflation expectations. The yield curve modelling approach allows a wide range of information to be condensed into intuitive concepts. These concepts can be useful for assessing a central bank's credibility for achieving its inflation target, understanding current pricing dynamics, and how expectations are affected following major events or policy changes.

This *Note* shows that inflation expectations in New Zealand have been well anchored since 1996. In particular, the perceived inflation target focus has been stable and remained well within the official inflation target band. Despite some periods of actual inflation being outside the official target band, expectations have shown a smooth transition back to target, with expectations returning close to the perceived target focus in around two-and-half years on average.

References

Aruoba S (2014), Term structure of inflation expectations and real interest rates: the effect of unconventional monetary policy. University of Maryland and Federal Reserve Bank of Minneapolis. Staff Report 502.

Bauer M and G Rudebusch (2015), Optimal policy and market-based expectations. Economic Letters, Federal Reserve Bank of San Francisco.

Ford D, E Kendall, A Richardson (2015), Evaluating monetary policy. Reserve Bank of New Zealand Bulletin Vol. 78, No. 7, November 2015.

Kamber G, C McDonald, N Sander, and K Theodoridis (2015), A structural forecasting model for New Zealand: NZSIM. Reserve Bank of New Zealand, Discussion Paper. DP2015/05.

Krippner L (2015a), A theoretical foundation for the Nelson and Siegel class of yield curve models, *Journal of Applied Econometrics*.

Krippner L (2015b), Term Structure Modeling at the Zero Lower Bound: A Practitioner's Guide. *Palgrave-Macmillan*.

Lewis M (2015), Forecasting with macro-finance models: Applications to United States and New Zealand. Master of Commerce in Economics Thesis, Victoria University of Wellington, School of Economics and Finance.

Lewis M and C J McDermott (2015), New Zealand's experience with changing its inflation target and the impact on inflation expectations. Reserve Bank of New Zealand, mimeo.

Mehrotra A and J Yetman (2014), 'Decaying expectations: what inflation forecasts tell us about the anchoring of inflation expectations'. BIS Working Paper No 464.

Nelson C and A Siegel (1987), Parsimonious modelling of yield curves. *Journal of Business*, 473-489.

Appendix

A INFLATION EXPECTATIONS DATA

Some of the survey data used are from Consensus Economics and some of these data are not constant horizon. The questions refer to the expectations for future inflation at the end of the current year plus two, three, four, five, and six years.¹⁵ For simplicity, in this *Note*, labels refer to the ‘plus year-ahead’ (for example, the two-year ahead measure is the end of current year plus two years). Correct treatment for the changing horizon is used in the term structure curve model. Also, these longer-term Consensus Economics surveys are conducted half-yearly, while the other expectations surveys used are undertaken quarterly and are constant horizon. The treatment used here is to interpolate to quarterly frequency holding the previous value constant. A more robust treatment would be to use the Kalman filter to estimate the missing observations. This will be explored in future research. Note, the retail and household surveys of inflation expectations are not included in this study given their substantial level biases. However, future work, using the Kalman filter, would help overcome this issue.

¹⁵The 10-year measure is the five-year, five-year forward inflation expectations rate and is thus treated as a 10-year rate.