

**A statistical forecasting framework and models for
the determination of starting price adjustments
for default price-quality paths**

A report undertaken for the
Electricity Networks Association
by

Peter Thomson
Statistics Research Associates Ltd
Wellington, New Zealand

Tony van Zijl
Sapere Research Group
Wellington, New Zealand

8 December 2010

Contents

Summary of findings and recommendations	1
1 Background	2
2 Stochastic properties of ROI data	3
2.1 ROI data	4
2.2 Exploratory analysis	8
2.3 Summary of findings	15
3 A forecasting framework for ROI assessment	16
3.1 Forecasting framework	18
4 Examples	19
4.1 Constant level model	19
4.2 Random walk model	20
4.3 Local level model	21
4.4 Forecast error analysis	22
5 Future developments	24
Acknowledgements	25
Appendix	26
References	26

Summary of findings and recommendations

This report considers the statistical forecasting framework and models needed to underpin the starting price adjustment framework proposed in the Commerce Commission discussion paper ComCom (2010). It is informed by an exploratory analysis of the statistical properties of ROI (return on investment) data for New Zealand electricity distribution businesses (EDBs).

The stochastic properties of the historic ROI data for 16 non-exempt EDBs were explored using suitably adjusted or normalised data sets. Key findings are that:

- non-exempt ROI time series are non-stationary, even after taking account of any trend movements common to all EDBs;
- non-exempt ROI time series can be assumed to follow a simple stochastic trend plus error model (a non-stationary *local level* model closely related to the widely used *exponential smoothing* model) that is governed by two key parameters which control the smoothness of the trend and the size of the error;
- the key parameters can be assumed to be the same for all EDBs and can be reasonably estimated from historic ROI data within a DPP.

In addition to other more general comments, the commentary makes the following specific recommendations.

Recommendation 1 *To underpin the reliability of statistical forecasting models of ROI time series, it is recommended that suitable industry standards and reporting conventions for the systematic formation and normalisation of ROI time series remain unchanged over 5 year DPP periods, or preferably longer, so that ROI data are consistently and accurately measured over time.*

Recommendation 2 *Given a suitable forecasting model fitted to ROI time series over all but the last year of a DPP, it is recommended that*

- optimal forecasts of the underlying ROI for the last year of the DPP be determined using this model;*
- the decision that the underlying ROI for an EDB exceeds the target ROI is made when its forecast exceeds the control limit*

$$\text{target ROI} + z_{\alpha} \text{ standard deviation of the forecast error}$$

where z_{α} controls the false positive rate α and z_{α} should be at least 2.

The non-stationary local level model identified during the exploratory analysis is a suitable forecasting model for ROI panel time series that is 'fit for purpose'.

Further details are given in the following sections together with comments on possible future developments.

1 Background

On 5 August 2010 the New Zealand Commerce Commission (Commission) released a discussion paper (ComCom, 2010) giving its preliminary views on a framework for adjusting the starting prices of gas transmission businesses, gas distribution businesses and non-exempt electricity distribution businesses (EDBs) where the starting prices apply from the beginning of the regulatory period covered by a default price-quality path (DPP). A DPP is a regulatory instrument provided for under Part 4 of the Commerce Act 1986.

The Commission's proposed framework is based on assessing a supplier's "current and projected profitability" using a return on investment (ROI) measure calculated from suitably normalised historical data. Since starting prices will typically need to be in place before the start of a new DPP, a supplier's ROI for the last year of the old DPP is not known in advance and must be forecast from past available data. As a consequence, the ROI measure of current and projected profitability necessarily has a forward view that involves forecasting ROI.

If the ROI profitability measure falls between specified ROI limits based on an industry-wide weighted average cost of capital (WACC) estimate, then the Commission proposes that starting price adjustments would not generally be made, but would be made otherwise. ComCom (2010) advises that the setting of ROI limits be informed by statistical analysis (ComCom, 2010, paragraph X.5) and, to the extent that consistent time series information is available, it would be used for calculating suppliers' returns (ComCom, 2010, paragraph 4.18).

Once starting prices are set, the annual growth rate of a supplier's prices over the DPP is capped at the corresponding annual growth rate of the consumer price index (CPI) less an efficiency factor X (so-called CPI- X indexation) where X is common to all suppliers. If a supplier considers the DPP to be unfair or inappropriate, then that supplier has the option of applying to the Commission for a customised price-quality path (CPP) that is specifically tailored to their situation. However, it is understood that the design of the DPP and CPP is intended to see relatively few regulated suppliers required to apply for CPPs due to their high compliance costs. As a consequence, primary objectives of the DPP include low compliance costs and widespread acceptance by suppliers.

The Commission's view on the basic purpose of starting prices is given in paragraph 2.13 of ComCom (2010). Downward adjustments to prices may be appropriate "in cases where suppliers are considered to be earning, or likely to earn, excessive profits and/or to promote the sharing of efficiency gains with consumers", whereas upward adjustments to prices may be justified "where current prices are likely to be too low to allow a supplier to earn a normal rate of return and facilitate efficient investment".

This report, commissioned by the Electricity Networks Association (ENA), considers the statistical forecasting framework needed to underpin the starting price adjustment framework proposed in ComCom (2010). It is informed by an exploratory analysis of the stochastic properties of ROI data for New Zealand EDBs both over time and across EDBs. Only when these are known, can a suitable statistical forecasting model be determined and statistically valid ROI limits and ROI assessments made.

The stochastic properties of ROI data are explored in Section 2. A statistical forecasting framework is developed in Section 3 for one ROI limit (an upper limit) and a given target ROI (the Commission’s WACC point estimate). Section 3 also gives a forecasting model appropriate for ROI time series from non-exempt EDBs that can reasonably be used for fitting and forecasting short panel time series. Section 4 gives examples of special cases of this model applied to ROI data, and topics for further research and development are discussed in Section 5.

2 Stochastic properties of ROI data

The starting price adjustment framework for DPPs proposed in ComCom (2010) crucially depends on the stochastic (random) properties of the historic ROI data on which any ROI-based assessment of current and future profitability might be based. Once these properties are known, both over time (dynamic or time series properties) and over the EDBs (static or cross-sectional properties), a suitable statistical forecasting model can provide forecasts of future ROI data and, in particular, the standard deviations of the resulting forecast errors. The latter turn out to play a key role in the setting of suitable ROI limits (see paragraph 6.19 of ComCom, 2010) and are dependent on the model adopted for the ROI.

The model fitted in paragraph 6.19, ComCom (2010), implicitly assumes that each EDB’s ROI data can be regarded as uncorrelated with constant mean and standard deviation that do not vary over time. Of these two time-homogeneous parameters, the mean is assumed to be EDB specific, but the standard deviation has a value common to all EDBs. Furthermore, for any given year, ROIs are assumed to be uncorrelated across EDBs and not influenced by common effects such as long-term weather and economic cycles. This simple model is an example of a *stationary* mean-reverting time series.

Firth (1982) considered the time series properties of corporate earnings in New Zealand for 110 firms over a period from the early 1960s to 1979. This study showed that these annual time series were, for the most part, well represented by *random walk* models. This finding is consistent with a smaller study of 42 New Zealand listed public companies by Caird and Emanuel (1981) and results from other countries. Thus, for these time series, it was the year-on-year differences (increases/decreases) that were stationary and mean-reverting, rather than the original times series. Random walk models are non-stationary with quite different time series properties to those of stationary mean-reverting models. Since each model gives quite different forecasts and forecast standard deviations, choice of model is important.

Suvas (1996) examines the time series forecasting properties of accrual accounting annual income figures over the period 1974 to 1989 for 134 Finnish firms comprising 51 listed firms, 42 limited liability firms subsequently declared bankrupt, and 41 smaller privately held firms. The 26 models considered included stationary mean-reverting models, random walk models and exponential smoothing models among others. As discussed later, exponential smoothing models can be seen as more general forms of random walk models. In terms of short-term forecasting, Suvas (1996) finds that the best of the exponential smoothing models almost always outperformed the others, for all time series considered and across

all three firm categories, and always outperformed the stationary mean-reverting models. Fama and French (2000) look at profitability and earnings for 2343 NYSE, AMEX and NASDAQ firms (small firms, financial firms and utilities excluded) over the period 1964 to 1996. Their model, in essence, adjusts the data by subtracting an estimate of the cross-sectional mean which will take out any common trend across all firms due to business or longer-term economic cycles. Their main finding was that the adjusted data were generally mean-reverting, but highly non-linear, with the mean-reversion factors varying depending on whether profitability was below its mean or far away from its mean in either direction. The behaviour of any common trend was not examined in this study, nor were out-of-sample forecasts explicitly evaluated.

These studies, and the others they reference, give a guide as to the nature of the stochastic properties of ROI data for New Zealand EDBs that might be expected. In general the data they consider, and the ROI data itself, are examples of *panel time series* which is a rapidly developing and evolving area of time series econometrics (see Baltagi, 2008, and Choi, 2006, for example). A distinguishing characteristic of the ROI data for New Zealand EDBs is that both the number of time points available and number of EDBs is very limited, especially the former. This will necessitate very simple models which, nevertheless, will still need to capture the major sources of variation in the historical ROI data in order to forecast satisfactorily at least one year ahead.

2.1 ROI data

Good forecasts from statistical models of ROI time series need data that is consistently measured over time. To the extent that such data sets are available, the Commission has indicated that it intends to use them to assess current and future profitability of EDBs (see paragraph 4.18, ComCom, 2010). However, to produce suitable ROI data entails careful *prior adjustment* (referred to as normalisation in ComCom, 2010) to ensure that the ROI data is consistent with accepted accounting conventions and consistently measures the desired return on investment over time.

PricewaterhouseCoopers (PwC) New Zealand have prepared a number of ROI data sets to assist the ENA and the exploratory analyses undertaken in this report. Over a number of years, PwC has prepared detailed databases comprising information disclosed by EDBs in accordance with the various regulatory requirements. This knowledge and experience have informed the procedures and methods they have used to prepare these data sets which are fully documented in PwC (2010). Note that all the data sets have been prepared on a post-tax basis.

Normalised ROI data sets have been created corresponding to six scenarios, three of which cover the 12 years ending 31 March 1999 to 31 March 2010, and three of which cover the 6 years 31 March 2005 to 31 March 2010. A brief description of the six scenarios is as follows.

Scenario 0 comprises the disclosed ROIs for the years ending 31 March 1999 to 31 March 2010. Apart from naming conventions, these data sets have not been adjusted (nor-

malised) and reflect the relevant regulatory Information Disclosure Requirements (IDRs) for each year concerned.

Scenario 1 is an adjustment of Scenario 0 that smooths the impact of the irregular revaluations that took place over 1999–2004, and uses CPI-based revaluations over 2005–2007 in line with those for 2008–2009.

Scenario 2 adjusts Scenario 0 by replacing all the disclosed revaluations over 1999–2007 by CPI-based revaluations. All valuations over 1999–2010 are now made on the same basis.

Scenario 3 is an adjustment of Scenario 0, but only over the years ending 31 March 2005 to 31 March 2010. Significant changes to the construction of ROI data were introduced in 2008 with the publication of the 2008 IDRs. Scenario 3 ROI restate the regulatory asset base of each EDB over 2005–2007 in a consistent manner to that used over 2008–2010.

Scenario 4 is a modification to Scenario 3 with revenue adjusted for any price path breaches over 2005–2009.

Scenario 5 is a modification of Scenario 4 that excludes capital contributions from regulated revenue and deducts them from the regulatory asset base. This is in accord with the draft input methodologies issued by the Commission on 22 October 2010.

Further details can be found in PwC (2010).

Figure 1 shows boxplots of ROI data by year for non-exempt and exempt EDBs and all six scenarios. Since the ROI data has been prepared on a post-tax basis, a target ROI of 7.57% has also been superimposed for reference where this is the post-tax value given in Appendix B, ComCom (2010). In each boxplot, the central box shows the lower quartile (lower line), median (middle line), upper quartile (upper line) and the whiskers (dashed lines extending above and below the boxes) give the effective range of the data with isolated points (asterisks) indicating potential outliers. For large samples from a Gaussian distribution, roughly 1% of values would be classified as potential outliers. The boxplots clearly show that normalisation is very important as is the regulatory definition of ROI itself.

Consider first Scenarios 0–2. The progressive adjustments leading to Scenario 2 have moderated the impact of anomalous years (2004 for Scenario 0) and smoothed the general level of ROI (as measured by the median) over 1999–2007. The variability of ROI by year, as measured by the interquartile range (the length of the central box), is more homogeneous for Scenario 2 than Scenarios 0–1. Nevertheless there is a marked jump in the median ROI and a reduction in variability from 2008, presumably reflecting the impact of the 2008 IDRs.

For Scenarios 3–5 there is much greater homogeneity in terms of level and variability over 2005–2010. The impact of the restatement of the regulatory asset base for each EDB over 2005–2008 is clearly shown with marked differences evident between the boxplots for Scenarios 2 and 3 over those years. Scenarios 3–5 generally present very similar

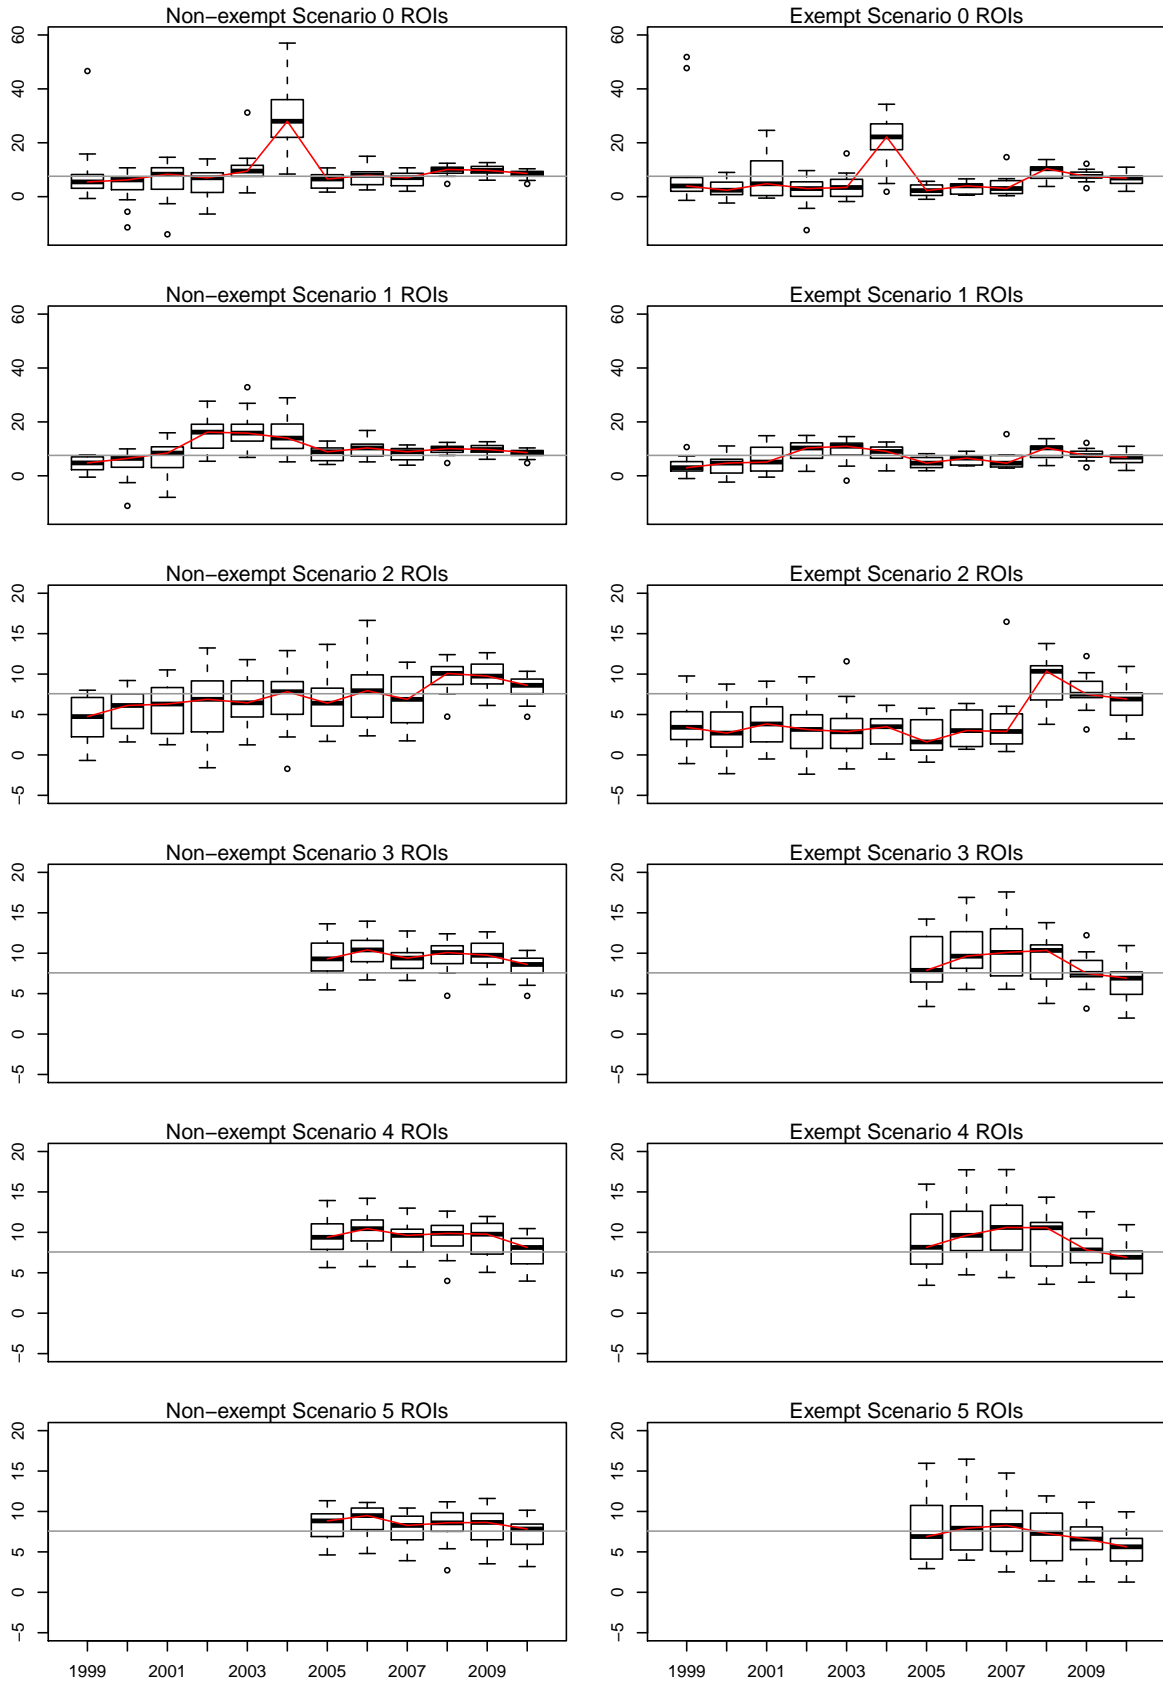


Figure 1: Boxplots of ROI data (%) by year for non-exempt and exempt EDBs and all six scenarios, with annual median ROIs (red) and ComCom (2010) target ROI (grey) superimposed.




Figure 2: Plots of Scenario 2 ROI data (%) by year for non-exempt EDBs with the ComCom (2010) target ROI (grey) superimposed.




Figure 3: Plots of Scenario 5 ROI data (%) by year for non-exempt EDBs with the ComCom (2010) target ROI (grey) superimposed.

descriptions of ROI, although there is a general drop in levels for Scenario 5 by comparison to Scenarios 3–4. This is consistent with the exclusion of capital contributions in the definition of ROI under the revised specification of the 2008 IDRs.

Examples of the plots of individual ROI time series for non-exempt EDBs are given in Figure 2 (Scenario 2) and Figure 3 (Scenario 5). The individual ROI time series for both scenarios appear to be evolving relatively smoothly and predictably over time, with trends that are generally increasing, in the case of Scenario 2 data, and decreasing in the case of Scenario 5 data. The general movement of these apparent trends over time is modest and consistent with the evolution of each Scenario’s annual median non-exempt ROI shown in Figure 1.

Figure 1 shows that the cross-sectional variability of the Scenario 5 time series by year is much the same over time and comparable with the last three years of the Scenario 2 time series. This is also evident in Figures 2 and 3 and contrasts with the cross-sectional variability of the Scenario 2 time series over 1999–2007 which, although common across this period, is markedly greater.

The smoothness of these ROI time series and the fact they are not dissimilar in appearance suggests that a simple non-stationary trend plus error model may well be appropriate. The common general movement of these trends implies that they may well be correlated, possibly due to a common trend that reflects longer-term common factors (movements in the CPI, economic and climate cycles, for example) that impact on all EDBs. These and other issues are directly addressed in the following section.

2.2 Exploratory analysis

In this and subsequent sections attention is restricted to ROI data for non-exempt EDBs since the latter are subject to DPP regulation.

All the available non-exempt ROI time series are short (6 or 12 observations for each EDB) which makes the determination of their time series properties challenging, in the case of 12 observations, and problematic in the case of 6 observations. Of the longer time series, Scenario 2 has the most homogeneous data and, as a consequence, will be used to explore the dynamic and static stochastic properties of ROI data for non-exempt New Zealand EDBs in order to identify suitable statistical forecasting models.

On the other hand, the shorter time series are the most homogeneous and cover a time span that is commensurate with the short time span (as little as 4 years within a 5 year DPP) needed to fit the chosen time series models. These fitted models can then be used to generate forecasts and the standard deviations of forecast errors. Since non-exempt Scenario 5 ROIs are the most consistent with current definitions of ROI, they will be used to illustrate this general approach in Section 3.

For the reasons discussed in the previous paragraphs, the basic strategy adopted throughout this report is to use Scenario 2 data for model identification and Scenario 5 data for estimation and forecasting. The exploratory analysis that follows now considers only Scenario 2 ROI data for non-exempt EDBs.

Since the identified models will need to be fitted to data collected up to, but not including, the last year of a 5 year DPP, the models will need to be suitable for individual time series of length 4 years and a cross-section of 16 non-exempt EDBs. The very small time span dictates time series models with at most 4 parameters, but ideally less. If these parameters were all EDB specific then they would not be able to be reliably estimated. However estimation is significantly more reliable when each of the 16 non-exempt EDBs follows the same model, but otherwise evolves independently, and some or all of the model's parameters are common to all EDBs. For example, if all non-exempt EDBs followed random walks with mutually independent Gaussian increments, each with zero mean and common standard deviation, then the 48 increments available would lead to a relatively accurate estimate of the common standard deviation. This leads to the following underpinning assumptions.

Key assumptions. *The ROI time series for non-exempt EDBs are each assumed to*

- (a) *follow a simple stochastic model that is common to all EDBs, with parameter values that are largely common to all EDBs;*
- (b) *evolve independently of the ROI time series for other EDBs, after possible adjustment for a common trend;*

at least to a reasonable approximation.

These key assumptions are quite stringent and relatively hard to check with any degree of certainty. Nevertheless, the plots in Figures 1, 2 and 3 do not seem inconsistent with these assumptions. It might also be expected that a common regulatory environment, similar economic conditions, geographic separation, and largely comparable businesses, are all factors that will generally support these key assumptions. However, if present, issues such as variability dependent on the size of EDBs (scale), and more complicated cross-sectional dependence, may not be well-accounted for by such simple models.

A first step in the determination of a suitable family of time series models is to determine whether the non-exempt Scenario 2 ROI time series are stationary or non-stationary. Figures 2 and 3 suggest that the time series are non-stationary (trending is evident, for example), but this hypothesis needs to be better supported by appropriate statistical tests. Two such tests are considered.

The Augmented Dickey-Fuller (ADF) test of the hypothesis that a single time series $Y(t)$ follows a (non-stationary) unit root process (a random walk with stationary increments) is based on a model of the form

$$Y(t) = \alpha + \beta t + \rho Y(t - 1) + \epsilon(t) \quad (t = 1, \dots, T) \quad (1)$$

where $\epsilon(t)$ is a stationary autoregressive process. The ADF test procedures formally test the null hypothesis that $Y(t)$ follows a unit root process ($\rho = 1, \beta = 0$), against the alternative that it follows a linear trend with stationary error. The simplest of these procedures was applied to each of the non-exempt Scenario 2 time series and the notched boxplot of p-values (a p-value is the probability under the null hypothesis that the test statistic is at least as extreme as its observed value) are shown in Figure 4. Here the




Figure 4: Notched boxplots of the p-values of the ADF and NH (Nyblom and Harvey, 2000) tests for Scenario 2 ROI data for each non-exempt EDB (ADF and NH), and for the adjusted Scenario 2 ROI data (ADFa and NHa). A 5% horizontal line (grey) has been superimposed for reference.

ends of the vertical notches give the upper and lower limits of a confidence interval for the median. All p-values are large (greater than 5% for example) and support the assumption of non-stationarity.

Given the short time series involved ($T = 12$) and the general lack of power of the ADF test in small samples, a more powerful single test of the common null hypothesis can be constructed by combining the p-values of each of the $N = 16$ separate (approximately independent) tests into one test statistic. One approach (see Choi, 2006) is to calculate

$$Z = \frac{1}{\sqrt{N}} \sum_{i=1}^N \Phi^{-1}(p_i)$$

where $\Phi(\cdot)$ is the standard normal function, $\Phi^{-1}(\cdot)$ is its corresponding quantile function, and the p_i are the p-values for the N individual tests. Under the null hypothesis, Z has a standard normal distribution and this hypothesis will be rejected when the value of Z is large negative (1% of values from a standard normal are less than -2.326 for example). For the non-exempt Scenario 2 ROI time series the computed Z value is positive (1.37) which again supports the hypothesis that all the individual ROI are non-stationary and follow unit root processes.

Nyblom and Harvey (2000) consider tests based on the *local level* model

$$Y(t) = \mu(t) + \epsilon(t), \quad \mu(t) = \mu(t-1) + \eta(t) \quad (t = 1, \dots, T) \quad (2)$$

where $\mu(0)$ is a parameter to be estimated and the $\epsilon(t)$, $\eta(t)$ are mutually independent Gaussian errors with zero means and standard deviations σ_ϵ , σ_η respectively. The random walk $\mu(t)$ can be thought of as a stochastic trend, so that the observations $Y(t)$ follow a simple trend plus error model. This model is non-stationary unless $\sigma_\eta = 0$ when $Y(t)$




Figure 5: Plots of the cross-sectional means (black) and standard deviations (red) by year of Scenario 2 ROI data for non-exempt EDBs with the ComCom (2010) target ROI (grey) superimposed.

reduces to a stationary time series, a sequence of independent Gaussian observations with constant mean and standard deviation σ_ϵ . This leads to a test of the null hypothesis that $Y(t)$ is stationary ($\sigma_\eta = 0$) against the alternative that it is non-stationary. When this test is applied to the 16 non-exempt Scenario 2 time series it yields p-values whose notched boxplot is shown in Figure 4. Most p-values are small (typically less than 5%) implying strong evidence in favour of the alternative hypothesis of non-stationarity. This is confirmed by the combined Z test which is very large and negative (-7.14).

The combined Z test depends on p-values that are exact, or well approximated, and on independent component tests. The former is the case for the Nyblom and Harvey (2000) test since the null distribution was able to be generated exactly. However, the ADF test relies on asymptotic approximations for its null distribution and so its p-values may be less secure. The two Z tests also have different interpretations. A rejected null hypothesis for the ADF based Z test implies that at least one non-exempt Scenario 2 ROI time series has a linear trend plus stationary error whereas, for the Nyblom and Harvey (2000) based Z test, it implies that at least one of the time series follows a stochastic trend plus error model. Perhaps the greatest weakness of the Z tests is the assumption that the components of each Z test are independent.

To minimise the impact of any cross-sectional dependence, one simple procedure is to mean-correct each non-exempt Scenario 2 ROI time series by subtracting the annual cross-sectional means from the individual ROI time series. In terms of Figures 1 and 2 this will remove the impact of any common industry-wide trend and make the individual time series less dependent. However, mean correction alone will not account for the marked reduction in cross-sectional variability over 2008–2010 due, presumably, to the impact of the 2008 IDRs. This suggests that, in addition to mean correction, the non-exempt Scenario 2 ROI time series should be standardised so that their cross-sectional standard deviations are more alike.

Plots of the time series of annual cross-sectional means and standard deviations of the non-exempt Scenario 2 ROI data are given in Figure 5. The smooth and predictable




Figure 6: Plots of adjusted Scenario 2 ROI data (%) by year for non-exempt EDBs with a zero line (grey) superimposed.

movement of the cross-sectional means over time reflects a common industry-wide trend due to economic, regulatory, climate and other factors impacting all EDBs. Over 1999–2008 the trend has increased slowly from around 5% to 10% before declining slightly over 2008–2010. The marked change in the cross-sectional standard deviations from 2008 is also evident in Figure 5. To compensate for the evolving cross-sectional mean and the structural break in cross-sectional variability, the Scenario 2 ROI data were mean corrected and then standardised by applying a scale factor of just over 50% to the mean-corrected ROI data prior to 2008. The resulting average cross-sectional variance over 1999–2007 is now the same as that over 2008–2010. This mean-corrected, re-scaled version of the non-exempt Scenario 2 ROI data will be referred to as the *adjusted Scenario 2 ROI data* for non-exempt EDBs.

Figure 6 plots the adjusted Scenario 2 ROI data for non-exempt EDBs. By comparison to Figure 2, the time series show little evidence of any common trend, and cross-sectional variability is now similar across all years. The individual adjusted ROI time series appear, as before, to be evolving relatively smoothly and predictably over time. They are similar in appearance to the corresponding Scenario 5 ROI time series plotted in Figure 3 and, apart from a level shift, can be regarded as a reasonable proxy for longer Scenario 5 ROI time series. More importantly, cross-sectional dependence has been reduced.

Applying the ADF and Nyblom and Harvey (2000) tests to the adjusted Scenario 2 ROI time series for non-exempt EDBs yields individual p-values whose boxplots are plotted in Figure 4. For the ADH tests, the combined Z test has value 1.40, and for the Nyblom and Harvey (2000) tests it is -5.26. These results are little changed from those reported for

the original Scenario 2 ROI data. This provides further and stronger evidence in favour of non-stationary models for the ROI data from non-exempt EDBs.

The general time series model (2) underpinning, and favoured by, the Nyblom and Harvey (2000) test provides a simple class of non-stationary models with a useful structural interpretation as a (stochastic) trend plus error model. This non-stationary model involves 3 parameters, some of which (σ_ϵ and σ_η) may be able to be assumed to be common to all ROI time series for non-exempt EDBs, and is commonly used for short-term prediction. The model also includes the simplest of the null models underpinning the ADF test. For these reasons, the non-stationary model (2) provides a useful and flexible starting point for modelling ROI time series from non-exempt EDBs.

Within the framework of model (2), the issue of which parameters can be considered common across ROI time series is now considered. From the plots of the ROI time series in Figures 2, 3 and 6, the initial trend values $\mu(0)$ would seem to almost inevitably need to be EDB specific. Nevertheless, for future DPP regulatory periods, it may be possible to reduce the number of these parameters, or eliminate them, by requiring them to be drawn from a prior distribution with mean equal to the target ROI. The parameters σ_ϵ and σ_η are more likely to be common across ROI time series and it is these which are now focussed on.

If a time series $Y(t)$ follows (2) then it is non-stationary, but the time series differences

$$Y(t) - Y(t - 1) = \eta(t) + \epsilon(t) - \epsilon(t - 1) \quad (3)$$

are stationary with theoretical standard deviation $\sigma_\epsilon\sqrt{2 + \lambda^2}$ and lag-one autocorrelation $-1/(2 + \lambda^2)$ where $\lambda = \sigma_\eta/\sigma_\epsilon$ can be interpreted as a signal to noise ratio. Thus the theoretical lag-one autocorrelation is negative and lies between 0 and -0.5, depending on the values of σ_ϵ and λ . If the theoretical standard deviations of the differenced ROI time series are the same, and the theoretical lag-one autocorrelations are also the same, then it can be concluded that the values of σ_ϵ and σ_η are common to all. This observation forms the basis for the following analysis.

Boxplots of the sample standard deviations and lag-one sample autocorrelations for the differences of the Scenario 2 ROI time series, and the adjusted Scenario 2 ROI time series, are shown in Figure 7. They are all relatively tightly clustered about their medians with the lag-one sample autocorrelations largely falling between 0 and -0.5 as expected. To visually assess whether this sampling variability could arise from panel time series following model (3) with homogeneous parameters, 16 independent simulations of length 11 years were generated from (3) with theoretical standard deviation 1 and lag-one autocorrelation -0.25. The latter figures are approximately the medians of the sample standard deviations and lag-one autocorrelations for the differences of the adjusted Scenario 2 ROI data. As can be seen from Figure 7, there is very good agreement between the boxplots for the simulated data and the differences of the adjusted Scenario 2 data. This provides strong informal evidence in support of the hypothesis that the values of σ_ϵ and σ_η can be regarded as common across ROI time series.

To formally check for common variances, standard homogeneity of variance tests can be applied to the differenced ROI time series. Conover et al. (1981) examined 56 such tests,




Figure 7: Notched boxplots of the standard deviations (`sd`) and lag-one autocorrelations (`rho`) for the differences of the Scenario 2 ROI time series for each non-exempt EDB, and for the differences of the adjusted Scenario 2 ROI time series (`sd (adj)` and `rho (adj)`). Boxplots of these quantities (`sd (sim)` and `rho (sim)`) determined from 16 independent simulations of the same model (3) with standard deviation 1 and lag-one autocorrelation -0.25) are also shown with a horizontal zero line (grey) superimposed.

both parametric and nonparametric, for small to medium sample sizes. Only a few were found to be robust and have good power. Of these, the Fligner-Killeen test was among the three best with well-controlled error rates under the null hypothesis of homogeneity. This test was applied to the first and second differences of the unadjusted and adjusted Scenario 2 ROI time series. Testing the homogeneity of variance for each of these data sets is equivalent to testing the homogeneity of common values for σ_ϵ and σ_η across ROI time series. In the case of the Scenario 2 ROI time series, the homogeneity hypothesis was unequivocally retained (the p-value for the test based on the first differences was 24.6%; for the second differences it was 39.1%), but was marginally rejected at the 5% level in the case of the adjusted Scenario 2 ROI time series (the p-value for the test based on the first differences was 4.2%; for the second differences it was 2.8%). When the ROI time series with the smallest standard deviation (Orion) was omitted, the two tests based on the first and second differences of the adjusted Scenario 2 ROI time series were again retained (the p-values for both tests exceeded 13%).

However the Fligner-Killeen and other tests for homogeneity of variance are typically based on time series of independent observations which does not apply here. The individual differenced series are also short (11 and 10 observations respectively) so the results are more indicative than definitive, but do show that there are no gross departures from the homogeneity assumption. A better approach would be to test for homogeneous parameters directly using the multivariate version of model (2). However this has not been done. In practice, any small departure from cross-sectional homogeneity will typically produce negligible changes in one-year-ahead forecasts, especially when these are based on short time series such as 4 annual ROI observations within a DPP. In short, the cross-sectional

homogeneity of variance assumption seems both reasonable, as illustrated by Figure 7 and the homogeneity of variance tests, and relatively costless.

2.3 Summary of findings

The stochastic properties of the historic ROI data for 16 non-exempt EDBs have been explored using suitably adjusted or normalised data sets. Key findings are as follows.

(a) **Non-exempt ROI time series are non-stationary.**

Non-exempt ROI time series were found to be non-stationary, even after adjusting for the cross-sectional mean which removes the impact of any industry-wide trend due, perhaps, to economic, regulatory, climate and other factors impacting all EDBs. This finding is consistent with Suvas (1996) who considered 26 forecasting models for annual earnings of 134 Finnish firms over a 16 year period, but not entirely consistent with Fama and French (2000) who found that the profitability of 2343 US firms (excluding small firms, financial firms and utilities) over a 33 year period, adjusted for cross-sectional averages, typically followed a non-linear mean-reverting process.

(b) **Model for non-exempt ROI time series.**

A suitable model for each non-exempt ROI time series $Y(t)$ is the stochastic trend plus error model

$$Y(t) = \mu(t) + \epsilon(t), \quad \mu(t) = \mu(t-1) + \eta(t) \quad (t = 1, \dots, T)$$

where the stochastic trend $\mu(t)$ has initial value $\mu(0)$ that is a parameter to be estimated, and the $\epsilon(t)$, $\eta(t)$ are mutually independent Gaussian errors with zero means and standard deviations σ_ϵ , σ_η respectively. This structural time series model is parsimonious (only 3 parameters) and is a subset of a wider class of trend plus error models that are commonly used in practice for short-term prediction. This family includes the exponential smoothing models favoured by Suvas (1996). If it is found necessary to allow for a common trend across ROI time series, as well as individual trends, then a further development would be to allow the $\eta(t)$ to be cross-sectionally correlated.

(c) **Parameter homogeneity and estimation.**

The parameters σ_ϵ , σ_η can be reasonably assumed to be common across all non-exempt ROI time series, but the initial trend values will be EDB specific. For fitting the model to non-exempt ROI time series over 4 years of a DPP, the initial trend values for this DPP could be fixed at the last trend estimates from the previous DPP, where these are estimated by conventional time series smoothing techniques. This leaves only two parameters σ_ϵ , σ_η to be estimated from 64 observations in the case of 16 ROI time series. This should be sufficient data to provide reasonable and relatively reliable estimates of these key parameters.

For future DPPs it is likely that the initial trend values will be more concentrated around the target ROI and so could also be modelled by a prior distribution whose

mean is the target ROI. However initialisation of the trend is unlikely to be a major issue for prediction based on 4 years of data since the optimal forecast will typically place greatest weight on the most recent ROI value and negligible weight on the initial trend value.

A number of adjusted or normalised data sets were provided for this exploratory data analysis. The transition from Scenario 0 (unadjusted) to Scenario 3 showed marked changes in the variability of the ROI time series, but those from Scenario 3 to Scenario 5 were much more homogeneous. This underscores the importance of normalisation. Accurate and reliable statistical forecasts of ROI times series are dependent on data that is consistently and accurately measured over time. This observation leads to the following recommendation.

Recommendation 1 *To underpin the reliability of statistical forecasting models of ROI time series, it is recommended that suitable industry standards and reporting conventions for the systematic formation and normalisation of ROI time series remain unchanged over 5 year DPP periods, or preferably longer, so that ROI data are consistently and accurately measured over time.*

3 A forecasting framework for ROI assessment

At first sight, statistical process control (SPC) or stochastic control would seem a suitable statistical framework that is consistent with, and would support, the Commission’s general approach to the setting of starting prices using ROI data. SPC uses statistical methods and models to monitor and control a time series of measurements to ensure that they continue to stay on track and close to a given target, in this case a given target ROI (see Montgomery, 2009, for example). Such an approach is consistent with the use of past ROI values to determine present actions in such a way that the future course of the process is as near as possible to the desired one. This general approach informs the statistical model-based framework adopted here.

The findings in Section 2.3 suggest that a suitable panel time series model for annual ROI data from non-exempt EDBs is

$$Y_j(t) = \mu_j(t) + \epsilon_j(t) \quad (j = 1, \dots, p; t = 1, \dots, T) \quad (4)$$

where t indexes years, T denotes the number of years of data available (typically $T = 4$) and $Y_j(t)$ denotes the annual ROI data for EDB j with smoothly evolving trend $\mu_j(t)$. The measurement errors $\epsilon_j(t)$ are assumed to be mutually independent Gaussian random variables (independent over time and EDB) with zero means and common standard deviation σ_ϵ .

The stochastic trend $\mu_j(t)$ can be thought of as the supplier’s underlying ROI free of measurement errors and it is this trend that should ideally be compared to any target ROI. The $\mu_j(t)$ can be modelled in many ways, including dependence on suitable covariates

such as weather and business cycle indices, or time t itself. In line with Section 2.3, one of the simplest time series trend models is adopted where

$$\mu_j(t) = \mu_j(t - 1) + \eta_j(t)$$

and the $\eta_j(t)$ are assumed to be mutually independent Gaussian random variables with zero means and common standard deviation σ_η . Thus each ROI time series $Y_j(t)$ is assumed to follow the local level model (2). The initial value $\mu_j(0)$ can be set in a number of ways, depending on the circumstances, but could be an unknown constant or a random variable with known mean and standard deviation. This trend model (a random walk) will yield a suitably smooth and evolving level running through the data provided σ_η is small enough.

Model (4) assumes that the ROI time series for each EDB evolves independently of the ROI time series for the other EDBs. Although all EDBs are influenced by the same longer-term economic, regulatory and climate cycles, these common factors are likely to be less important than local factors such as the physical separation of the EDBs, their different topographies and weather, etc. From Figure 1, the non-exempt Scenario 2 ROIs do show evidence of a common trend over time, although this is much less evident in the later scenarios, particularly the non-exempt Scenario 5 ROIs. Nevertheless, if necessary, a common trend could be incorporated in model (4) by allowing the $\eta_j(t)$ to be contemporaneously correlated across EDBs.

The family of models (4) is more general than it might seem at first sight and is, in turn, a subset of a larger class of structural time series models that allow for drift, covariates, and more flexible trends (see, for example, Durbin and Koopman, 2001). If $\sigma_\eta = 0$ then $\eta_j(t) = 0$ and $\mu_j(t)$ is a constant for all t , so that the $Y_j(t)$ are independent Gaussian random variables with common mean and common standard deviation. This model is considered in more detail in Section 4.1. If $\sigma_\epsilon = 0$ then the $Y_j(t)$ form a dependent random walk whose differences $Y_j(t) - Y_j(t - 1)$ are independent Gaussian random variables with zero means and common standard deviation σ_η . This model is considered more fully in Section 4.2. The cases where σ_ϵ and σ_η are both greater than zero span a large range of possibilities between these two extremes. All these models and their variants are widely used in practice for short-term forecasting.

Consider the situation where historical ROI data are available from a larger number of EDBs ($p = 16$ say) over a smaller number of years ($T = 4$ years of the current DPP) giving a total of pT observations (64 in this case). In line with the findings of Section 2.3, model (4) can be fitted to the historical data and estimates of σ_ϵ , σ_η determined. The fitted model can be used to determine optimal forecasts of the ROI, or their underlying trend, for year $T + 1$ and the standard deviations of the respective forecast errors. Finally, these forecasts can now be compared to the target ROI taking proper account of the standard deviations of the forecast errors.

These observations lead to the following general forecasting framework which is consistent with the approach advocated in ComCom (2010) and, in particular, formalises the need to base starting price adjustments on estimates of current and projected profitability.

3.1 Forecasting framework

Consider forecasting the underlying annual ROI for each EDB in year $T + 1$ given that the panel time series model (4) has been fitted to the available annual ROI data for all EDBs over years $t = 1, \dots, T$ (typically $T = 4$). For EDB j

- (a) determine the optimal forecast $\hat{\mu}_j(T + 1)$ of its underlying ROI $\mu_j(T + 1)$ from the fitted model and the ROI data up to and including time T , and the standard deviation $V(T + 1)$ of its forecast error;
- (b) determine the *control limit*

$$\tau(\alpha) = \tau + z_\alpha V(T + 1)$$

where τ is the target ROI, $\Phi(z_\alpha) = 1 - \alpha$, and $\Phi(z)$ is the standard normal distribution function;

- (c) decide that $\mu_j(T + 1)$ exceeds the target ROI τ if its forecast $\hat{\mu}_j(T + 1)$ exceeds the control limit $\tau(\alpha)$ and take appropriate action.

Note that the standard deviation of the forecast error is the same for all EDBs since the parameters $\sigma_\epsilon, \sigma_\eta$ are assumed to be common across EDBs.

Given knowledge of the history of ROI data up to and including time T , the probability that the future underlying ROI value $\mu_j(T + 1)$ is less than the target ROI is

$$\Phi\left(\frac{\tau - \hat{\mu}_j(T + 1)}{V(T + 1)}\right)$$

for each EDB j . This probability is less than α when the forecast $\hat{\mu}_j(T + 1)$ exceeds the control limit $\tau(\alpha)$ so that

$$P(\text{underlying ROI at time } T + 1 < \tau | \text{its forecast} > \tau(\alpha)) \leq \alpha$$

and so α can be interpreted as a *false positive rate*.

In general α should be small since it measures the accuracy of the classification procedure. Just how small α should be depends on the relative costs of these errors to suppliers and consumers with the former preferring small values of α and consumers the reverse. In the SPC literature z_α is often taken to be 3 (or even 6) leading to the so-called three sigma limits which correspond to $\alpha = 0.13\%$. If $z_\alpha = 1$ then $\alpha = 15.87\%$ and if $z_\alpha = 2$ then $\alpha = 2.28\%$. For example, if 4 EDBs had forecasts that breached the control limit $\tau(\alpha)$, then approximately 50% of the time when $z_\alpha = 1$, approximately 9% of the time when $z_\alpha = 2$, and approximately 0.5% of the time when $z_\alpha = 3$, at least one of these EDBs would incorrectly have its starting prices adjusted. Selecting $z_\alpha \geq 2$ seems a reasonable choice here.

Finally, it is noted that the forecasting framework proposed above is generic in the sense that the same general procedure would be followed regardless of which panel time series model was fitted. This leads to the following recommendation.

Recommendation 2 *Given a suitable forecasting model fitted to ROI time series over all but the last year of a DPP, it is recommended that*

- (a) *optimal forecasts of the underlying ROI for the last year of the DPP be determined using this model;*
- (b) *the decision that the underlying ROI for an EDB exceeds the target ROI is made when its forecast exceeds the control limit*

$$\text{target ROI} + z_\alpha \text{ standard deviation of the forecast error}$$

where z_α controls the false positive rate α and z_α should be at least 2.

4 Examples

For model (4) and a given EDB j , it can be shown that the optimal forecast of the underlying ROI at time $T + 1$ based on ROI observations $Y_j(t)$ over years $t = 1, \dots, T$ is a simple weighted average of the form

$$\hat{\mu}_j(T + 1) = w_0\mu_j(0) + w_1Y_j(1) + w_2Y_j(2) + \dots + w_TY_j(T) \quad (5)$$

where the weights w_j sum to one and are determined using the recursions given in the Appendix. It is also the case that $\hat{\mu}_j(T + 1)$ is the optimal forecast of $Y_j(T + 1)$, although the standard deviations of their respective forecast errors are different.

In this section three special cases of (4) are considered which will lead to different patterns of weights and therefore different forecasts. The adjusted Scenario 2 ROI data for non-exempt EDBs will be used to check the accuracy of the forecasts $\hat{\mu}_j(T + 1)$ as forecasts of $Y_j(T + 1)$. The non-exempt Scenario 5 ROI data will be used to illustrate the procedures.

4.1 Constant level model

Setting $\sigma_\eta = 0$ in (4) and initialising appropriately yields the *constant level* model

$$Y_j(t) = \mu_j + \epsilon_j(t) \quad (j = 1, \dots, p; t = 1, \dots, T) \quad (6)$$

where the $\epsilon_j(t)$ are independent Gaussian measurement errors, each with zero mean and standard deviation σ_ϵ . This stationary mean-reverting model was rejected in Section 2 as a model for ROI data, but is included here since it is implicitly assumed in Section 6 of ComCom (2010) when calculating control limits.

Given the historic ROI data $Y_j(1), \dots, Y_j(T)$ for each EDB j , the standard estimates of μ_j and σ_ϵ^2 are

$$\hat{\mu}_j = \frac{1}{T} \sum_{t=1}^T Y_j(t), \quad \hat{\sigma}_\epsilon^2 = \frac{1}{p} \sum_{j=1}^p s_j^2, \quad s_j^2 = \frac{1}{T-1} \sum_{j=1}^T (Y_j(t) - \hat{\mu}_j)^2.$$

Note that the estimate of σ_ϵ differs from that used in paragraph 6.19 of ComCom (2010). The latter is non-standard and less efficient than $\hat{\sigma}_\epsilon$ given above.

For EDB j and the central case where $T = 4$, the optimal forecast of $\mu_j(T + 1) = \mu_j$ and the standard deviation of its forecast error are given by

$$\hat{\mu}_j(T + 1) = \hat{\mu}_j = \frac{1}{4}(Y_j(1) + Y_j(2) + Y_j(3) + Y_j(4)), \quad V(T + 1) = \frac{\sigma_\epsilon}{\sqrt{T}} = \frac{\sigma_\epsilon}{2}$$

and the estimated control limit is given by

$$\tau(\alpha) = \tau + z_\alpha \frac{\hat{\sigma}_\epsilon}{2}.$$

Note that this limit is not the same as that suggested in ComCom (2010), Section 6, which appears to be based on forecasting $\mu_j(T + 1) = \mu_j$ by $Y_j(T)$, the last available ROI value. Clearly $\hat{\mu}_j$ is the better forecast since it has smaller standard deviation.

Here the weights in (5) are

$$w_0 = 0, w_1 = w_2 = w_3 = w_4 = 0.25$$

so that all past ROI values are equally weighted.

For non-exempt Scenario 5 data over the DPP from 2006–2010, the estimate of σ_ϵ calculated from the years 2006–2009 is $\hat{\sigma}_\epsilon = 0.85\%$ and the control limit is 8.42% for $z_\alpha = 2$ and the ComCom (2010) target ROI of 7.57% . Ten EDBs out of 16 were assessed by this model as having exceeded the control limit. For these EDBs, the differences between the forecasts of the underlying ROI and the control limit are 0.25% , 0.38% , 0.38% , 0.45% , 0.60% , 0.94% , 1.53% , 2.11% , 2.40% and 2.41% .

4.2 Random walk model

Setting $\sigma_\epsilon = 0$ in (4) yields the *random walk* model

$$Y_j(t) = Y_j(t - 1) + \eta_j(t) \quad (j = 1, \dots, p; t = 1, \dots, T) \quad (7)$$

where the $\eta_j(t)$ are independent Gaussian measurement errors, each with zero mean and standard deviation σ_η . Unlike the static model (4.1), the random walk model (7) is a dynamic time series model (the simplest) which models the ROI in year t as the previous year's ROI plus random error.

Given the historic data $Y_j(1), \dots, Y_j(T)$ for each EDB j , the standard estimate of σ_η is given by

$$\hat{\sigma}_\eta^2 = \frac{1}{p} \sum_{j=1}^p \hat{s}_j^2, \quad \hat{s}_j^2 = \frac{1}{T-1} \sum_{j=2}^T (Y_j(t) - Y_j(t-1))^2.$$

Moreover, the optimal forecast of $\mu_j(T + 1) = Y_j(T + 1)$ and the standard deviation of its forecast error are given by

$$\hat{\mu}_j(T + 1) = Y_j(T), \quad V_j(T + 1) = \sigma_\eta$$

and the estimated control limit becomes

$$\tau(\alpha) = \tau + z_\alpha \hat{\sigma}_\eta.$$

For the random walk model, the decision to adjust starting prices is based on the last available ROI so that the weights in (5) are now

$$w_0 = w_1 = w_2 = w_3 = 0, w_4 = 1$$

for $T = 4$. Although this decision rule is similar to that given in Section 6 of ComCom (2010), it is based on $\hat{\sigma}_\eta$, the standard deviation of the differenced ROI data, rather than $\hat{\sigma}_\epsilon$, the standard deviation of the mean corrected ROI data.

For non-exempt Scenario 5 data over the DPP from 2006–2010, the estimate of σ_η calculated from the years 2006–2009 is $\hat{\sigma}_\eta = 1.11\%$ and the control limit is 9.79% for $z_\alpha = 2$ and the ComCom (2010) target ROI of 7.57% . Four EDBs out of 16 were assessed by this model as having exceeded the control limit. For these EDBs, the differences between the forecasts of the underlying ROI and the control limit are 0.17% , 0.57% , 1.21% and 1.82% .

4.3 Local level model

Now consider (4) for the general case where both σ_ϵ and σ_η are strictly positive. These two parameters are typically estimated using maximum likelihood estimation techniques applied to the historic data (see Durbin and Koopman, 2000, for details). When the model is fitted to non-exempt ROI time series over the first 4 years of a DPP, the initial trend values $\mu_j(0)$ can be fixed at the last trend estimates from the previous DPP, where these are estimated by conventional time series smoothing techniques.

However, there has not been sufficient time to write and test the computer routines needed to fit this model, although this should be a relatively straightforward exercise. Instead, the parameters σ_ϵ , σ_η adopted for the simulated data shown in Figure 7 were used. These particular values were chosen to match the properties of the adjusted Scenario 2 ROI data which can be regarded as a proxy to the non-exempt Scenario 5 data over the period 1999–2010. Although not optimal, these values of σ_ϵ , σ_η should provide indicative results.

Given $T = 4$ and these values of σ_ϵ , σ_η , the optimal forecast of $\mu_j(T + 1)$ for EDB j is given by

$$\hat{\mu}_j(T + 1) = 0.01\mu_j(0) + 0.01Y_j(1) + 0.05Y_j(2) + 0.20Y_j(3) + 0.73Y_j(4)$$

where the formula for the weights is given in the Appendix. In terms of (5) the w_j are now given by

$$w_0 = 0.01, w_1 = 0.01, w_2 = 0.05, w_3 = 0.20, w_4 = 0.73$$

which places most weight on the more recent observations, but not as extreme as the random walk.

Again consider the non-exempt Scenario 5 data over the DPP from 2006–2010. Using the formula given in the Appendix and the four years of ROI data from 2006–2009, the

standard deviation of the forecast error is $V(T + 1) = 0.83$ and the control limit is 9.22% for $z_\alpha = 2$ and the ComCom (2010) target ROI of 7.57%. Five EDBs out of 16 were assessed by this model as having exceeded the control limit. For these EDBs, the differences between the forecasts of the underlying ROI and the control limit are 0.48%, 0.50%, 1.18%, 1.72% and 2.18%.

4.4 Forecast error analysis

For model (4), the optimal forecast of the ROI for EDB j at the end of year year $T + 1$ given ROI data for the previous T years is exactly the same as $\hat{\mu}(T + 1)$. So, for each EDB j , the optimal forecast of $Y_j(T + 1)$ given the past ROI observations $Y_j(t)$ over years $t = 1, \dots, T$ is

$$\hat{Y}_j(T + 1) = \hat{\mu}(T + 1)$$

but now the standard deviation of the resulting forecast error is given by

$$\sqrt{V(T + 1)^2 + \sigma_\epsilon^2}.$$

This is larger than $V(T + 1)$ as expected, since $Y_j(T + 1)$ is the sum of $\mu_j(T + 1)$ and the random error $\epsilon_j(T + 1)$. In particular, if the ROI observations $Y_j(t)$ follow model (4), then the bias and root mean squared error (RMSE) of the forecast error $Y_j(T + 1) - \hat{Y}_j(T + 1)$ are

$$\text{forecast bias of } \hat{Y}_j(T + 1) = 0, \quad \text{forecast RMSE of } \hat{Y}_j(T + 1) = \sqrt{V(T + 1)^2 + \sigma_\epsilon^2}$$

which are functions of the model parameters.

The previous theoretical results can be benchmarked against the actual forecast errors that arise when model (4) is fitted to ROI data. Furthermore, these forecast errors can also be used to empirically check the forecasting performance of the three cases of model (4) considered earlier (the constant level, random walk and local level models). This has been done for the adjusted non-exempt Scenario 2 data where the fitted models gave forecast errors over the period 2003–2010 (8 years). For each year in this period and each EDB, the ROI for that year was forecast using the ROI data available over the previous 4 years, and then the sample mean and RMSE of the 8 forecast errors computed. The results are shown in Figure 8.

All the models considered appear to have no significant forecast bias (it is effectively zero), but the local level model shows better agreement between the typical (median) empirical RMSE for each EDB and its theoretical counterpart. The boxplots of the forecast errors from the non-optimal local level model are also better (more tightly clustered about a lower median) than the random walk model which, in turn, is slightly better than the constant level model. Since the forecasts for each EDB are based on just 4 past observations and yield only 8 forecast errors, the differences are modest. Nevertheless, the results shown in Figure 8 confirm the appropriateness of the non-stationary models over the stationary constant level model, with the non-stationary local level model preferred. Since the latter was calibrated informally to the adjusted Scenario 2 data, further improvements in the




Figure 8: Notched boxlots of the forecast bias (left plot) and root mean squared error (right plot) for the adjusted Scenario 2 ROI time series (%) for each non-exempt EDB. The one year ahead ROI forecasts were based on the four preceding years and calculated using the constant level model (**c1**), the random walk model (**rw**) and a non-optimal local level model (**ll**). For each model the theoretical forecast bias and root mean squared error are superimposed (red).

forecasting performance of the local level model might be expected when it is fitted using optimal statistical estimation methods.

Finally, the previous analysis is applied to the Scenario 5 ROI data for non-exempt EDBs. Here we only have 2 forecast errors for each EDB, corresponding to the years 2009 and 2010. Summary figures for the average bias and RMSE across EDBs for these years are given in Table 1. Although the RMSE figures are in general agreement with those for the adjusted Scenario 2 ROI data (the non-stationary models perform best), the forecast biases are more considerable. From the plots of the non-exempt Scenario 5 data in Figure 3, it can be seen that most ROI time series have a slight downwards trend, particularly over 2009–2010, and it is this common movement that has biased the two forecasts concerned. When the cross-sectional means are subtracted from the non-exempt Scenario 5 data to adjust for this common trend, the biases vanish and the 2009–2010 RMSE figures reduce to 0.96% (constant level model), 0.86% (random walk model), and 0.87% (non-optimal local level model) with the same conclusions as before.

Year	Bias			RMSE		
	c1	rw	ll	c1	rw	ll
2009	-0.22	-0.22	-0.15	0.94	0.75	0.77
2010	-1.01	-0.85	-0.89	1.42	1.30	1.32
2009–2010	-0.61	-0.53	-0.52	1.20	1.06	1.08

Table 1: Summary figures over the period 2009–2010 for the forecast bias and root mean squared error (RMSE) across EDBs for the non-exempt Scenario 5 ROI time series (%). The one year ahead ROI forecasts were based on the four preceding years and calculated using the constant level model (**c1**), random walk model (**rw**) and a non-optimal local level model (**ll**).

5 Future developments

A basic statistical forecasting framework has been established that provides a basis for the starting price adjustment strategy proposed in ComCom (2010) for ROI time series from non-exempt EDBs. A suitable panel time series model for the ROI time series has also been identified and illustrative examples of its use have been considered. Due to time constraints some developments have yet to be undertaken and, in addition, further issues were identified during the course of the analysis that may warrant further investigation. These developments and issues include, but are not limited to, the following.

Fitting the model

No attempt was made to optimally fit model (4) to the ROI time series for non-exempt EDBs. Rather, the values of the two key parameters σ_ϵ and σ_η were informally calibrated to the adjusted non-exempt Scenario 2 data. In practice these two parameters would typically need to be estimated using maximum likelihood (a well-known statistical estimation method with optimal statistical properties) or equivalent, following the developments given in Durbin and Koopman (2000) for example. The necessary computer programs needed to fit model (4) need to be written and tested on both simulated as well as actual ROI data.

Model extensions including allowance for a common trend

Given that the model will typically be fitted to the first 4 years of a DPP, there is little room for further extension of the model, particularly the dynamics. For example, a drift parameter could be included to allow the ROI time series to show more secular linear time trends, either common to all EDBs or EDB specific. However this will be costly in terms of parameters leading to imprecise and unreliable estimates, and such trend projections often prove too volatile in practice. This is likely to be the case here where a target ROI is in view.

If the non-exempt Scenario 5 data shown in Figure 3 are typical of what might be expected over future DPPs, then it might be advantageous to include a common (industry-wide) stochastic trend in the model in addition to the current EDB specific trend. This is achieved by assuming that the $\eta_j(t)$ are equi-correlated across EDBs which increases the number of model parameters that must be estimated to three (σ_ϵ , σ_η and the common correlation) which, although marginal, does not seem unreasonable (the 3 parameters will typically need to be estimated from 64 observations). This model takes advantage of additional cross-sectional structure and will lead to individual EDB forecasts that are based on a compromise between EDB specific movements in the past data, and the movements of the industry-wide average. If a common stochastic trend is present, then this model should lead to improved forecasts with smaller forecast error standard deviations.

However the imposition of a target ROI on EDBs may make future ROI data less prone to industry-wide movements, at least over relatively short time periods such as 4 or 5 years. If this is the case, then the inclusion of a common trend may be debatable. Nevertheless,

the impact of common economic, regulatory and climate cycles on EDBs and their ROI time series makes the inclusion of a common trend in model (4) a worthwhile development to explore.

Model issues

The assumption of common parameters across EDBs has been checked and found to be reasonable. Even if they are different and the parameters were known, this may not necessarily lead to results that are very different from the homogeneous case because of the short time spans involved. However some EDBs are very large, compared to a greater number of EDBs that are of comparable size (as measured by their regulatory asset base for example). Dividing the EDBs into groups based on size and undertaking more comprehensive statistical tests based on optimal estimation of the models in each category would be one way to further check the homogeneity assumption. However data limitations are likely to make such an analysis more indicative than definitive.

The 5 year DPP cycle inevitably leads to parameters being primarily estimated over short time spans (typically 4 years) which is a major limitation. However past data from previous DPPs are able to inform and ameliorate this estimation, especially over future DPPs, either by formally incorporating such prior information using Bayesian methods, or more informal procedures. The use of covariates within a common industry-wide trend model fitted over a longer history may also be used to prior adjust ROI data prior to checking for any breach of the +target ROI.

The CPI-X factor potentially limits the movements of the errors $\epsilon_j(t)$ and $\eta_j(t)$ making the Gaussian assumption more of an approximation. However this approximation should be reasonable and still lead to optimal linear forecasts provided the CPI-X factor is not overly restrictive.

Forecast horizon

The statistical forecasting framework that has been established estimates “current and projected profitability” by the forecast of the underlying ROI at the end of the current DPP. However the discussion could equally well have been framed in terms of forecasting other functions of future values of the underlying ROI. For example, one might consider the underlying ROI for the first year of the new DPP, or a discounted average of the underlying ROIs over the new DPP, as the forward variable of interest to target. Given the simplicity of the panel time series model used (the forecast more than one year ahead is the same as the one year ahead forecast), the main impact of such a development would be to increase the root mean squared error of the forecast leading to a larger control limit.

Acknowledgements

The authors are grateful to Stuart Shepherd (Sapere Research Group), Nathan Strong (Unison) and Lynne Taylor (PwC) for advice and guidance provided over the course of this research project.

Appendix

For the local level model defined by (4) with parameters σ_ϵ , σ_η that are common to all EDBs, the optimal forecast of the underlying ROI $\mu_j(t+1)$ for EDB j from historic data over years $1, \dots, t$ is given by the recursion

$$\hat{\mu}_j(t+1) = \omega(t)Y_j(t) + (1 - \omega(t))\hat{\mu}_j(t) \quad (t = 1, \dots, T)$$

where the $\omega(t)$ satisfy

$$\omega(t) = \frac{\omega(t-1) + q}{\omega(t-1) + q + 1}, \quad q = \frac{\sigma_\eta^2}{\sigma_\epsilon^2} \quad (t = 1, \dots, T)$$

and both recursions are suitably initialised. These recursions lead directly to forecasts of the form (5) where the w_t are functions of the $\omega(t)$. In particular, the standard deviations of the forecast errors $\mu_j(t) - \hat{\mu}_j(t)$ are given by

$$V(t+1) = \sigma_\epsilon \sqrt{\omega(t) + q} \quad (t = 1, \dots, T)$$

so that $V(T+1)$ and the control limit $\tau(\alpha)$ are readily computed.

When t is large (not the case here) and $\sigma_\epsilon > 0$ then these recursions reduce to *simple exponential smoothing*, a forecasting procedure widely used in practice. Specific examples of these forecasts are given in Sections 4.1, 4.2 and 4.3.

References

- Baltagi, B.H. (2008) *Econometric analysis of panel data (3rd edition)*. John Wiley, UK.
- Caird, K.G. and Emanuel, D.M. (1981) Some time series properties of accounting income numbers. *Australian Journal of management* **6**, 7-16.
- Choi, I. (2006) Nonstationary panels. In *Palgrave Handbook of Econometrics Volume 1: Econometric Theory* (eds. Mills, T.C. and Patterson, K.). Palgrave Macmillan, 511-539.
- ComComm (2010) *Starting price adjustments for default price-quality paths*. Commerce Commission Discussion Paper.
- Conover, W.J., Johnson, M.E. and Johnson, M.M. (1981) A comparative study of tests for homogeneity of variances, with applications to the outer continental shelf bidding data. *Technometrics* **23**, 351-361.
- Durbin, J. and Koopman, S.J. (2001) *Time Series Analysis by State Space Methods*. Oxford University Press, Oxford.
- Fama, E.F. and French, K.F. (2000) Forecasting profitability and earnings. *Journal of Business* **73**, 161-175.
- Firth, M. (1982) Some time series properties of corporate earnings in New Zealand. *Journal of Business Finance and Accounting* **9**, 353-360.

- Montgomery, D.C. (2009) *Introduction to Statistical Quality Control, Sixth Edition*. John Wiley, New York.
- Nyblom, J. and Harvey, A.C. (2000) Tests of common stochastic trends. *Econometric Theory* **16**, 176-199.
- PwC (2010) Electricity Networks Association return on investment data. PwC technical report, October 2010.
- Suvas, A. (1996) An examination of the time series properties of listed, privately held and failing firms. *Journal of Business Finance and Accounting* **23**, 617-639.