**Joint Glasgow Conference of the European Real Estate Society and the Cutting Edge Conference of the RICS 4th to 7th June 2002**

# Title: "Forecasting Values of Commercial and Residential Property Using Non-Linear Mathematical and Statistical Techniques"

**Chris Satchwell, Howard James and Mandy Bradley**
**(Technical Forecasts Limited)**

## Abstract:

Past data series are commercially available both for residential and commercial property. They are used by property professionals as a benchmark for investors. Attempts have also been made in the past to forecast these series to try to identify future rises, falls or turning points in the property market. The methods available to practitioners have been of limited use because of their inability to detect and model the past patterns in property data. A major problem has been the relative lack of reliable past data. Typically, the series from which forecasts might be possible do not go back very far in time. Additionally, data may often be temporally sparse; i.e annual or quarterly rather than monthly. The effect of these restrictions is to limit the amount of information in the series that could be used to forecast them into the future. Furthermore, past methods have tended to confine forecasting to linear models that restrict output behaviour and usually fail to represent the complexities of real property markets.

The work described in this paper shows the feasibility of producing better property forecasts by compensating for the limitations of property series by using 'associated series' in forecasting models. In addition, there is a principled method for matching a model to the complexity of the target data series, and thence to produce the most probable forecasts, in a Bayesian sense, of the past property data. The methods contain principled procedures for the reduction of errors in the final forecasts.

These new methods show high levels of accuracy for property forecasts, which allow them to push much farther into the future than previously. Accuracy assessments are shown both for forecasts of commercial and residential property series.

# Contents

References

Appendices

# 1 Introduction to data series information and forecasting

## 1.1 Relevance of property data series

Property data is available from a variety of private and public sources. The data is useful for assessing the performance of property as a key to future investment. Whilst there is a general acceptance that it is useful to review past performance as a clue to the future, the extent to which past performance can be a guide to the future in the context of property has never been properly assessed. Disclaimers that past performance is ***not*** necessarily an indicator of future behaviour understandably breeds some caution that past data ***can be*** a guide to the future. Our premise is that where representative data for past outcomes from starting condition(s) similar to the present exist, with other things being equal, those same outcomes will tend to repeat and provide a basis for a forecast. Where other things are not equal, such a forecast can provide a datum from which the effects of differences in conditions can be (subjectively) assessed. This paper presents evidence that past property in conjunction with associated data series can be projected well into the future with useful accuracy, though risk can never be entirely eliminated.

Those individuals who are known to be good at making investment decisions take past performance into account as part of their understanding of market processes, consciously or otherwise. They draw on a wealth of past experience as well as knowledge of current and local factors.

## 1.2 Data complexity

Market knowledge only helps partially in the making of investment decisions. Property markets are so complex that it is impossible to have all the necessary knowledge to make a perfect decision. Even if complete knowledge were possible, occasional shocks in the market disrupt expectations as recent tragic events demonstrate only too well. The effects of imperfect knowledge combined with occasional shock events contribute to risk in all business decisions. Investors live with risk. The goal of this work is to reduce the risk that arises from imperfect market knowledge; the extent to which market shocks may be predicted is a matter that will take a long time to resolve.

The research described in this paper is part of an ongoing effort within Technical Forecasts Ltd (TFL) to develop and commercialise good practice for the projection of property data into the future using two major sources of property time series data available in the UK, namely the Investment Property Databank for commercial property and HM Land Registry for residential property. Accurate forecasting of this type of data should constitute a major resource for property decisions, contributing to the ability to manage, possibly to understand, the complexity of the property markets.

Forecasts can be a tool to assist property professionals. They stand to gain materially from combining their knowledge with clear mathematical prognostications of market directions and turns.

## 1.3 External influences on data series

The values of a property time series such as retail rents in the south east of England vary in time because of a multitude of influences that make the rents rise and fall (demand, the economy, supply and so on). In a perfect market, a single time series should react to every change (ref. 8). Consequently, a single series can act as an imprint of all the major influences in the market. Given this information encoded in the fluctuations of a time series, it is possible to make a mathematical model that reconstructs the market along with all its influences. Such a model is capable (in theory) of making a good forecast of the time series.

In reality, however, data series are beset by problems: they may not react correctly to market influences due to delays in the transfer of market information; the data, though collected on a regular basis, may not be sufficiently frequent to catch all of the changes caused by market influences; the data series may have been collected over too short a time, the Land Registry data being a case in point. All these factors and others conspire to make it difficult if not impossible to make good mathematical forecasting models of single time series of property data.

# 2 The use of mutual information

## 2.1 Earlier work

An early pointer to a solution to the single series problem in this project was to associate the target series with another time series that was known to be a leading indicator to the target series (Ref. 1). Those models provided forecasts for horizons up to the limit of the 'lead' provided by the leading indicator, but did not provide a basis for continuing beyond that point since they could not forecast their leading indicator. An essential difference between that and the present work is that the 'lead' of the other series has ceased to be important. What is important is the information that another series can contribute to the forecasting process and for this reason the 'other' series are now referred to as 'associated' series. The authors describe a composite non-linear mathematical model that uses information from both target and associated series. The results of forecasts using the composite models give better results than non-linear models of the target series alone.

## 2.2 Information theory

It is possible to measure the information that an associated series can give to a target series. This measure of information is a number on an information scale. A basis for this can be found in information theory where a quantity known as 'entropy' -- which measures the degree of disorder in data -- is defined. It is possible to assess the disorder in data alone, and then in conjunction with an associated series. The reduction in disorder brought about by the associated series defines the mutual information. The greater the reduction, the greater the information given to the target series by the associated series. A formal description of this technique can be found in (3).

## 2.3 Novel application of mutual information

In the methods described here, the mutual information between every target series is measured against a hundred or more candidate associated series. Following this measurement stage, every target series is paired in turn with the n best associated series ($2 < n < 12$) to produce multiple double series models that are then used to project the data into the future. The net result is to provide an objective means of selecting the best combinations of time series for forecasting models.

It is believed that some users of property data apply traditional linear regression models to project the data into the future. Our work has shown that data characteristics unique to the property markets dictate that non-linear modelling methods are essential if accurate forecasting models are to be achieved. Property markets are highly non-linear, so linear regression models can only be approximate at most, even if multiple series are used.

Non-linear methods have the potential to pick up the complex relationships in property data but they need to be correctly applied. In the following sections of this paper, the methods are discussed and results of thousands of forecasts on real data assessed. Section 3 contains a discussion of how models are developed and adapted to market complexity; Section 4 explains how forecasting errors are minimised; Section 5 assesses the accuracy of forecasts; finally, Section 6 closes the paper with a review of the main results of the research.

# 3 Forecasting Models, and matching complexity to available data

Complexities of non-linear data modelling are well beyond the scope of this paper but for those who are interested, (2) and (3) provide a good introduction.

## 3.1 Forecasting approach

The essential problem is to find a relationship between 'windows' of past values of target series (usually prices of interest) and associated series, to predict the value of a target series one step ahead.

Provided forecasts of associated series are available, the ‘windows’ can then slide one step ahead to encompass that forecast and use it as part of the raw data for the next one.  In principle, this process can continue forever.  In practice, forecasts can become rapidly unstable as errors compound to drive answers rapidly upwards or downwards.  Delaying the onset of such instabilities demands accurate forecasting methods.

With regard to the associated series, these are forecast as a batch, one step ahead at a time, referencing others as their own ‘associated series’.  These series reflect a microcosm of the economy and are usually smoother and longer than target property series.  Their forecasts tend to be good, and provide an anchor to deter instabilities in subsequent forecasts of the target property series with which they will later be associated.

## *3.2 Defining the problem of complexity*

There is a fundamental problem in non-linear data modelling of separating a repeatable ‘signal’ from unrepeatable ‘noise’.  If a model is too simple, data is not optimally modelled.  If it is too complex then some of the noise is inadvertently modelled and when the model is used on previously unseen data, forecasts are usually poor because noise characteristics of the new data differ from those on which the model was developed.

Most non-linear data modelling methods use a non-linear transformation (which we will not concern ourselves with) of the input data to create a ‘feature’ space, which is then linearly related to the target.  Thus, between the feature space and the target, we are back into the more familiar realm of linear regression.  The modeller has control over the number of features to use in this space and there are established techniques for determining how many to use.

The most common of these involves dividing the data into ‘training’ and ‘testing’ sets, and plotting graphs of modelling errors against complexity – as measured by the number of features used.  The number of features consistent with a minimum error in the ‘test’ data set is usually used for the non-linear model.  Essentially, this provides an estimate of the optimal number of features required to separate signal and noise for that particular division of the data into ‘training’ and ‘testing’ files.

Regrettably, this technique requires more data than is usually available for forecasting property time series – but its description helps one to appreciate the problem.  Finally, it needs to be appreciated that this is essentially a manual process and unsuitable for mass forecasting.

## *3.3 Regularisation and incorporation of Bayesian techniques*

The Bayesian techniques used are described in (2).  They offer some important advantages over all other non-linear data modelling techniques.

Before describing them, we need to introduce the idea of regularisation.  Where too many features are related to the target, noise is represented and a model’s output can be ‘spiky’; i.e. not as smooth as it should be when plotted against the inputs.  Regularisation is a technique that can adjust the regression coefficients (or weights) between features and target, to penalise curvature in the output and suppress spikes.

Another minor diversion is needed to describe what happens to the weights when too many features are used.  In magnitude, weights become quite large – to produce the effect of a large positive weight cancelling out that of a large negative weight.  When a small number of features are used, weights tend to be smaller and reflect an intuitive expectation of what their values should be to relate features to a target.  To stop this ‘large weight’ syndrome, the Bayesian technique imposes a prior condition that all possible weights are clustered around some mean value according to a Normal (Gaussian) distribution.  When the maths are worked through, the effect is to apply the right amount of regularisation to correct an inherently over complicated model to one of appropriate complexity for the data.  More importantly, its results are independent of arbitrary divisions of the data into training and testing files and can be shown to be the most probable model between a given feature space and its target.

For the forecasting of property series, this provides a tool to wring the most inference out of the limited data available. This method has been developed to make it suitable for mass forecasting, and TFL currently use it to produce around 1.5 million non-linear models monthly to support their on-going forecasting activities.

## 4 Reducing Forecasting Errors

In section 3, issues of non-linear data modelling were discussed which may have led to the misconception that a perfect model can result.

In practice, nearly all models have errors – typically through the choice of inputs or features used. The approach we have adopted is to accept that (for the moment) we cannot produce 'the perfect model' but to try to approach the correct answer via the imperfect models that we can produce.

Suppose each model consists of the right answer plus an error. If results from a number of statistically independent models are averaged, then their average should approximate to the right answer; because the property of statistical independence means errors should self-cancel. Using this road to the correct answer the issue therefore is one of how to build statistical independence into the models rather than one of how to produce 'the perfect model'.

Our legacy technology was the two-series model (1). This is not defended in terms of building 'the perfect model', as better individual models can undoubtedly be produced if more series provide more information to the forecasting process. The two-series approach remains competitive when used in this 'committee' environment, as their simple inputs makes statistical independence more visible, the selection of appropriate associated series easier ,and more models can be produced for a given computing effort.

Statistical independence can also be improved by varying the choice of features used in the non-linear models. Typically some 50 models, consisting of 10 different sets of associated series and 5 different sets of features for each, are used in a forecast.

## 5 Discussion of the results

### *5.1 Tables of historic forecastability- Appendix 1*

In Section 6.2 on accuracy assessments, methods used in this research for obtaining measures of historic forecast accuracy are presented. The system adopted in that section allows an A/B/C/U classification, described in section 6.2.4, of forecasts of series which have been cut back, by two years for residential, or three years for commercial, and are projected up to the currently known situation. Forecast results are compared with these known values, and their accuracy assessed. Since checks by TFL have shown that forecasts which have historically proven accurate tend to continue accurate for some time, it is a useful guide for users of the forecasts as an indication of current reliability, which, along with market knowledge, forms an element of the decision-making process.

Tables 1 & 2, presented at Appendix 1, give the characteristics of classifications A, A and B combined, and all classifications (A+B+C+U) for residential and commercial forecasts respectively. The differences between residential (Table 1) and commercial (Table 2) performances stand out clearly: commercial forecasts over three years are better than the residential forecasts over two, because the residential data series tend to be short and volatile whilst the commercial data series are long and smooth. This is most clearly visible in the first column of each table, in which 56% to 86% of residential properties (postcode Districts) and 38% to 68% (postcode Sectors) are in the A and B classes, whereas for commercial properties, 80 % to 100% are in same classes. The A and B class of both types of properties are a useful guide to forecast performance. In both sets of series, the average correlation figures indicate that the vast majority of A and B class forecasts are in the right direction.

The performance of residential properties is divided into two sections (Table 1). It can be observed that postcode Districts offer a significantly better forecast than Sectors. The Districts are much larger than the Sectors (average around 20,000 dwellings, as opposed to average 2000 in a Sector) so that there are more recorded sales. The result is that the District series are smoother (less volatile) than the Sectors, so that better forecast performance can be expected; and is in fact the case in practice.

The last three columns in Tables 1 and 2 give the percentage of forecasts that are within, respectively 5, 10 and 15% of the true values at the end of the forecast period. For the residential forecasts, the 15% figure is a helpful guide when it is considered that local valuations of properties at the current date should be within about 10% of the value that is realised in a sale. Since the 15% measure is the value after ***two years*** in the case of these forecasts, then the 15% measure seems reasonable.

The corresponding figures for commercial properties are significantly better than the residential series, since they are taken over three years rather than the two years for the residential properties.

The remaining columns of the two tables concern average and absolute errors. They confirm that in general, the rises in the markets over the last two to three years have been underestimated in the case of residential and overestimated in the case of commercial properties. The reason for this effect is currently under investigation.

## *5.2 Graphical samples of forecasts - Appendix 2*

The graphs provided at Appendix 2 are intended to give an idea of how the forecasts perform, but in fact many thousands of forecasts are produce each month so it is difficult within the confines of this paper to give an overall picture. The samples are chosen from ‘A’ class forecasts, and illustrate some features of interest.

### 5.2.1 Residential districts

Figure A1 (Appendix 2) shows a forecast of a rise in District CH7 in Essex. The rise has been anticipated but in line with many residential forecasts, the extent of the rise in this District was not quite modelled. It is possible that local demand in many of these cases was not modelled sufficiently well probably due to lack of inclusion of local data in the model.

The Sheffield District S71 in Figure A2 is modelled reasonably well given the high volatility of this series. The forecast captures the more or less flat performance of this District.

The Southampton District in Figure A3 is well captured by the forecast, in spite of the volatility of the series. The forecast captures the initial drop and then the gradual rise of the prices.

The TN13 District in Figure A4 has quite a low visible volatility. The gradient of the rise is well captured by the forecast model.

District S25 in Figure A5 is well modelled apart from the initial drop. It is possible that local factors were responsible for this drop which the model missed but which seemed to have a lasting effect. Alternatively, the start of the forecast could have coincided with a peak in local Land Registry values, which returned to normal in the next quarter.

### 5.2.2 Commercial series

Figure A6 is the first of the commercial examples. Here, the three year rise is well captured though the break in the gradient just before March 2000 is not modelled in this case.

In Figure A7 the forecasts anticipates a maximum in the value rather earlier than it actually occurred. It would seem that the market was heading for a turn but the onset was delayed.

In Figure A8, a rather volatile series is well modelled. The forecast has anticipated the minimum at the right time and at more or less the correct level.

Finally, in Figure A9, the accurate match between a smooth series and its forecast can be seen. The downturn is anticipated by a few months, but warning of its onset was given over 2 years in advance.

# 6 Accuracy of the forecasts

## 6.1 *Significance of historic accuracy in data-led forecasting*

The value of a forecast must lie in the confidence which can be placed in it, in order to assess the potential risk/reward of decisions based upon it.

### 6.1.1 Explanations and subjective forecasting

Where forecasts are produced by experts, along with substantial explanations of why the forecasts are as given, then users of these forecasts feel justified in making apparently informed decisions on that basis, whatever happens subsequently. In contrast, where they are produced by computerised calculation solely on the basis of historic data, without the associated verbal explanation of the future, then a track record of accurate forecasts is essential in order to offer sufficient confidence in future predictions to persuade users to act on these recommendations, without the comforting base of words on which they have acted previously.

A psychological perspective on this issue provides an interesting insight into the acceptability of different types of forecasts. In (7), experimental evidence from social psychologists shows that providing a reason for a proposition tends to generate a fixed-action response in its favour; resulting in a weapon of influence to have it believed. What is usually overlooked in forecasting is that such explanations require subjectively-formulated models, and that there are usually a great many of such models with very different results that could be devised, with no guarantee offered that the one chosen is necessarily better than any of the others.

From the psychologists' observations in (7), it seems likely that many will choose to ignore the fundamental subjectivity (and consequential inaccuracies in results) of the 'explained' model selection process. The consequence is that they give greater credence to results from the particular model chosen, in preference to those of an objective model based on the best interpretation of the relevant data, but lacking a plausible tale to justify them

### 6.1.2 Objectivity and data-led forecasts

In the longer term it is likely that the kind of objective results produced by TFL will gain acceptance and be used as a basis for selecting a subjective model producing co-incidental outputs, in order to maximise the chances of getting the right answers – if not the right explanations.

As a significant factor in its research and development of data-led forecasting methods, TFL ensures that there is a logical, mathematical basis for historic reliability of its forecasts. It must be on this that confidence in any forward forecast must be based – unless one waits to see how accurate it was, in which case any value in the forecast will have ceased to exist.

This consideration of accuracy in data-led forecasting will be examined in the remainder of this section.

## 6.2 *Measures for assessing accuracy*

### 6.2.1 Measures adopted

Assessment of the accuracy of a forecast should be made on the underlying behaviour of the data being forecast, without necessarily matching the short-term variability in that data, termed 'noise', which at any specific point may deviate significantly from the underlying trend. Thus accuracy should be assessed primarily on three factors:

- how well the shape of the forecast matches the actual behaviour of a series (the ***correlation*** between the curves);
- how accurately the actual and forecast behaviours match (the ***RMS*** or root mean square error between the series);

- how erratically the series being forecast behaved in the past – the ***volatility***.

Having assessed accuracy on this basis, it is then useful to consider how closely the forecast matches the final actual value of the series. This is not a prime measure of accuracy, since, due to the volatility of any data series, the value at any particular point may lie away from the underlying series – see section 6.2.2.

### 6.2.2 Accuracy and volatility

The volatility of a data series can render the actual accuracy of a forecast at one specific time period as a totally misleading indicator of the overall accuracy of a forecast data set.

In Fig. 1 below, the jagged curve illustrates the average selling price of detached houses in one postcode sector in the YO postcode area, using actual Land Registry data. The smoothed curve indicates the probable underlying behaviour of the average detached house price in that postcode sector.

As can be seen, the quarter's average price at any point can be widely different from the overall average: at Dec00, the actual average price from the Land Registry is 50% up on what a forecast value following the underlying curve would be, yet the following quarter it would be fractionally down on forecast.

Overall, performance would be quite acceptable for such a volatile series, yet at one instant, such as Dec00, accuracy could be considered highly suspect – until the next data point is observed.

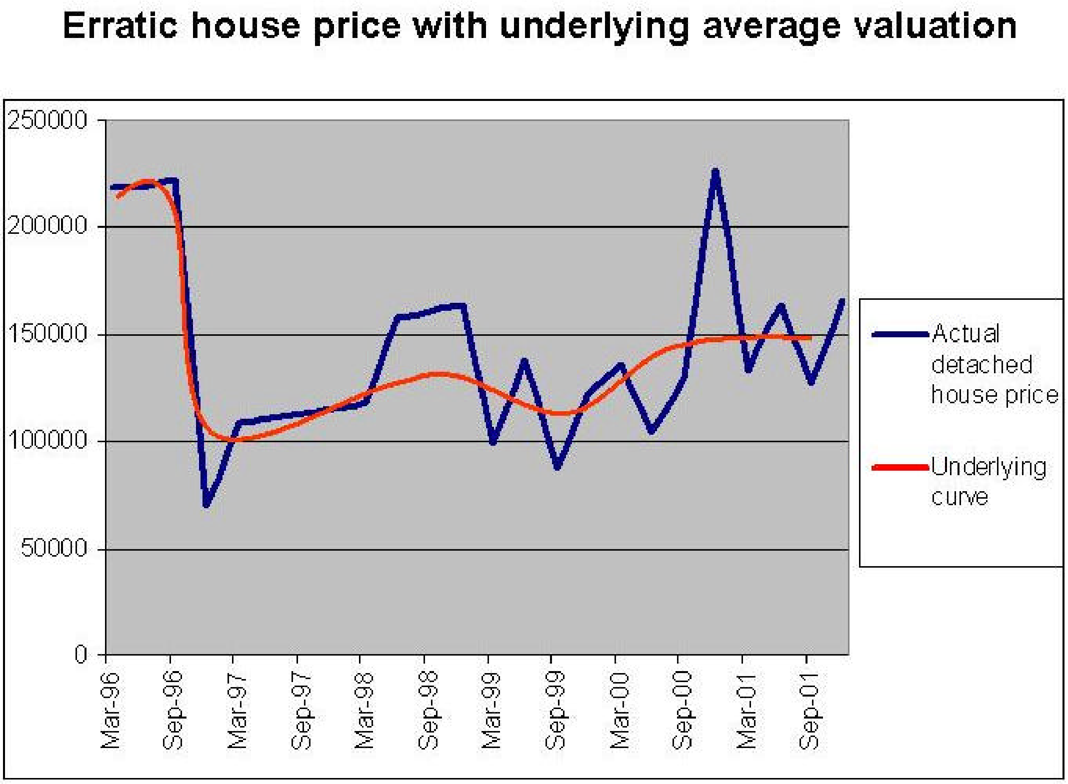

Fig 1 Accuracy and volatility

### 6.2.3 Volatility: TFL definition

As its name indicates, the volatility of any series is a measure of how erratic it is over a given number of data points. It is often used to assist in the measure of risk/reward ratio in evaluating the “right price” of a stock or bond option (Black-Scholes model, Ref 4). In TFL’s case, it is used to assist in evaluating the historic forecastability of residential and commercial property data series.

In the language of the Black-Scholes model, the volatility $s$ for a house price series is defined as:

$$s = \frac{s}{\sqrt{t}} \ ,$$

where

$$s = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(u_i - \bar{u})^2}$$

and

$$u_i = \ln\left(\frac{S_i}{S_{i-1}}\right)$$

$s$ = standard deviation of average house price of given type over specified no. of periods in postcode sector under observation;

$S_i$ = average price of house of given type at end of ith period in postcode sector under observation (i = 0, 1, … n);

$(n+1)$ = number of periods – eg one year’s quarterly data would result in (n+1)=5, including start and finish points, hence n=4;

$\frac{S_i}{S_{i-1}}$ = return on house value in ith interval;

$\bar{u}$ = mean of all $u_i$ ;

$t$ = time interval in years. (Ref. 5)

Similar considerations also apply for commercial property series, such as IPD’s monthly Regional and City Pages, and Segments, which are also forecast by TFL in its IndexEye publication.

### 6.2.4 Historic forecastability assessment

It is axiomatic that the past history of data will enable a limited capability to forecast its future behaviour. It has further been shown (Ref. 1, 6) that forecasting of a property series can be dramatically improved by jointly forecasting associated series alongside the required series. Logically, this can be accepted; after all, it is well-known that if oil supply from the Middle East is reduced, petrol prices will rise, and vice versa.

In order to confirm the forecasting capability of TFL’s models, it has been found that, where a model can forecast well historically, then it is likely to continue to forecast well, over a measurable time-frame. However, in order to achieve this to best effect, a long history of data is required. The shorter the time series available, the less feasible realistic forecastability assessment becomes. It may be that forecasts from short time series are ‘good’, but no proof or history of forecastability can be offered in such cases.

TFL’s standard method of assessing historic forecastability is to cut back series being forecast by 2 – 3 years, including all parallel data, such that **a model will not have any information beyond the cut-off date**. The trimmed series are then forecast to the latest date for which data is actually known, and performance compared between forecast and actual, using the criteria defined in section 6.2.1

A further refinement is offered by evaluating $\Psi$, the ratio between RMS and volatility, $s$ :

$$\Psi = RMS/s \approx 1$$

- if >> 1, the model might be considered insufficiently complex for the data set it is modelling;
- if << 1, the model might be attempting to model the noise in the data.

No research has currently taken place on the correctness of the above assumptions, although the ratio value of 1 must always by definition of RMS and $\sigma$ be a target. The exception is where $\sigma$ is very small, which would of course influence the ratio to excess.

Assessment of historic forecastability is made on the basis of an A / B / C / U classification, as below:

if ($RMS < 0.1$) AND (($0.9 < \Psi < 1.10$) OR ( $\sigma < 0.05$ )) then historic forecastability = A

if ($RMS < 0.2$) AND (($0.8 < \Psi < 1.25$) OR ( $\sigma < 0.08$ )) then historic forecastability = B

if ($RMS < 0.35$) AND (($0.6 < \Psi < 1.5$) OR ( $\sigma < 0.12$ )) then historic forecastability = C

otherwise: historic forecastability = U

The definition of A-class is very tight, which reflects in the relatively few A allocations, particularly in the Residential forecast series, but also in the significantly higher accuracy that these achieve.

A U (‘Unclassifiable’) category can be awarded for a number of reasons. Typically, this will be either because a series is highly volatile, or ‘noisy’, so that extracting underlying information from the noise is not feasible; or that a series is too short to enable a judgment to be made on its historic forecastability. It does not necessarily imply that the forecast offered is inaccurate (although this could well be the case!); merely that more care should be taken in its use.

## 6.3 Accuracy figures

The characteristics of historic residential house-price series are very different from typical commercial series, in that commercial series tend to be smoother, because they take account of more contributions per series than can be the case in Land Registry data at the postcode sector level. We will therefore consider the two separately.

### 6.3.1 Residential forecast accuracy

In the Land Registry data, four house-types are listed: detached, semi-detached, terraced and flats / maisonettes. House price sales data is provided to postcode sector level, averaging around 2000 dwellings per postcode sector, but varying widely in fact, from under 100 to over 7000. A price is not provided for any house type which does not achieve more than 3 sales in a quarter. Neither does the classification of house types distinguish between size of dwelling: number of bedrooms, plot size, condition etc will dramatically affect the sale price of individual properties, and can bias the average price significantly in any postcode sector.

Land Registry data series are quite short, since data exists only back to March 1995. Thus trials can only be performed on 2-year forecasts, as 3-year trials would leave insufficient data for adequate training of the model.

Historic accuracy figures for residential property are shown in Table 1, Appendix 1. Please note that use of the term ‘error’ is strictly mathematical: it is the percentage difference between forecast and actual values at the end of the two-year period.

As can be seen from the table, the accuracy of forecasts for districts and sectors in the ‘A’ category appears to justify the classification: these are significantly more accurate than the remainder, although, considering the (A+B) category – usually well over 50% of all districts and sectors – still one third of forecasts are within 10% of the actual value after two years. Except for semi-detached houses, more than 25% of ‘A’ districts and sectors are correct to within 5% of actual value after 2 years. More than half are within 10%, and over 80% within 15% of actual, after 2 years.

The model has consistently under-estimated the increase in property prices over the two-year period (along with all experts in the property field), although fairly consistently by around 10-12%, with accuracy in the ‘A’ categories again being greater, at 5-8% under-estimate. The average accuracy, the average of the absolute accuracy value, is within 10% for ‘A’ categories, and within 15% for A and B categories in combination.

### 6.3.2 Commercial forecast accuracy – IPD Monthly series

Historic accuracy figures here are calculated along similar lines, using Investment Property Databank’s (IPD) Regional Pages of commercial property indices. These series are much longer than those for residential property, so that more indicative checks can be made on forecasting reliability.

In this case, the data series for March 2002 was taken back to March 1999, along with all parallel series, and forecast for the (known) three years. Results were then compared with actual; a summary is given in Table 2, Appendix 2.

Using similar A / B / C / U classifications to those for residential property, it is useful to note that a much higher proportion of series are categorised as historically ‘A’ forecastable, with the vast majority of the remainder as ‘B’ forecastable. None was classed as ‘U’ – unclassifiable.

Overall, the models have over-estimated index values, by around 4 – 6%, slightly more for Industrials. It is notable that half as many Offices series are forecast to within 5% after 3 years than any other, yet relaxing the standard to within 10% brings them back into line with other series forecasts. Of the individual categories of commercial property, Retail seems to have been the most accurately forecastable over the last 3 years, although for all categories and classes, almost half are forecast to within 15% over 3 years.

## 7 A forecast to March 2005 – Appendix 3

Fig. A10 in Appendix 3 illustrates a genuine TFL forecast, for 3 years from March 2002 to March 2005, of an IPD Regional Pages series.

The series chosen is for All Properties in Central London; Total returns.

Please note that any fair comparison of performance of that forecast must be carried out via that specific IPD data series, as all series from different IPD datasets perform differently, due to the differing constituents of each series.

Our forecast starts from a base of 100 at March 2002; thus to obtain the forecast value of the IPD index at any particular point, our forecast value should be multiplied by the March ’02 IPD index value, and divided by 100.

TFL’s forecast indicates a flattening in total returns during the first half of 2003, but returning to steady growth from around Sept 2003. We welcome your feedback over the next 3 years!

# 8 Summary of the research

This paper describes an ongoing research effort aimed at developing forecasts of property rents, capital and yields using time series data. The property data that is the target of this research does not in itself contain sufficient information to give acceptable forecasts, but with the addition of mutual information from associated series, demonstrably improved results are possible. The target and associated series are used in conjunction to produce multiple models, each model being adjusted to the right complexity using Bayesian statistical principles. The resulting multiple models are finally averaged in order to maximise probability of error cancellation in the forecasts.

The forecasts were assessed using tests on commercial property data from Investment Property Databank (IPD) and house price data published by HM Land Registry. The forecasts were made for three years ahead for the commercial data, and two years for the residential data – the difference being due to relative shortness of the historical residential data series. Results were compared with actual data and tabulated in the paper. In addition, some graphs of forecasts and actual data are presented to give a visual perspective of the results. One example of a genuine forecast is also presented.

The question of the confidence that can be placed in a forecast has been addressed from a perspective of historical accuracy and inherent volatility, and remains crucial to the ability to use these forecasts in a commercial environment. In terms of scale, we believe these methods have led to high-quality mass forecasts becoming available to the property profession for the first time, leading to a more rational basis for many of their decisions. We are conscious that economic decisions taken without good forecasts have been likened to driving a truck whilst only being able to look in the rear view mirror. We try to provide an insight into the road ahead. Our ongoing research effort will therefore continue with the aims of providing better forecasts, confidence measures for those forecasts and finally technologies to enable those forecasts to be optimally used. We hope to have the opportunity to report our progress in these endeavours at future conferences.

# References

1. Connellan, O. & James, H. (1998) ‘Estimated Realisation Price (ERP) by neural networks: forecasting commercial property values’, Journal of Property Valuation & Investment, Vol. 16 No 1 pp71-86

2. Bishop, C. M. ‘Neural Networks for Pattern Recognition’, pp 164-191, 385-439, Oxford University Press, 1985, ISBN 0-19-853864-2

3. Haykin, S ‘Neural Networks – A Comprehensive Foundation’, pp 256 – 312, 484-499, Prentice Hall, 1999, ISBN 0-13-908385-5

4. Black, F. & Scholes, M. ‘The Pricing of Options and Corporate Liabilities’, Journal of Political Economy, June 1973, pp 637-654

5. Gasquet, V. ‘Option pricing, the Black-Scholes Model, and Neural Networks’, M.Sc. thesis, Aston Univ. Dept of Computer Science & Applied Mathematics, Sept 1994

6. Connellan, O. & James, H. ‘Forecasts of a Small Feature in a Property Index’, RICS Cutting Edge Conference 2000

7. Cialdini, R. B. ‘Influence - Science and Practice’, Fourth Edition, Allyn and Bacon, ISBN 0-321-01147-3, 2001, pp4.

8. Taken F. (1981) Lecture Notes in Mathematics 898. Ed Rand DA and Young L-S. Springer-Verlag: Berlin

# Appendices

## Appendix 1 Tables of accuracy results

### Table 1: Historic forecastability of residential forecasts

| | All figs at end 2yrs (Dec 99-01) | % of total | Average error: | Av. Abs(error) | +ve correl: | Av. correl: | Inside 5%: | Inside 10%: | Inside 15%: |
|---|---|---|---|---|---|---|---|---|---|
| | | **Postcode Districts** | | | | | | | |
| **A only** | **Detached** | 9.75% | -6.37% | 8.40% | 91.28% | 50.95% | 33.72% | 61.63% | 87.21% |
| | **Semis** | 26.17% | -7.46% | 9.50% | 89.63% | 51.56% | 23.33% | 55.29% | 82.94% |
| | **Terraced** | 16.69% | -7.31% | 9.04% | 85.71% | 51.26% | 25.97% | 59.74% | 87.34% |
| | **Flats/Mais** | 9.36% | -6.49% | 7.76% | 96.43% | 70.17% | 41.67% | 64.29% | 86.90% |
| **A + B** | **Detached** | 56.80% | -10.21% | 14.06% | 86.63% | 44.46% | 16.07% | 36.93% | 56.59% |
| | **Semis** | 81.57% | -11.88% | 14.32% | 87.60% | 50.95% | 14.07% | 31.67% | 55.51% |
| | **Terraced** | 69.70% | -11.29% | 13.56% | 87.79% | 51.77% | 15.79% | 36.47% | 58.09% |
| | **Flats/Mais** | 55.96% | -13.09% | 15.15% | 92.63% | 62.84% | 16.93% | 29.48% | 48.41% |
| **ALL** | **Detached** | 100.00% | -11.79% | 17.05% | 81.35% | 35.91% | 13.27% | 29.93% | 45.86% |
| | **Semis** | 100.00% | -12.94% | 16.02% | 85.70% | 47.67% | 12.66% | 28.55% | 49.92% |
| | **Terraced** | 100.00% | -11.78% | 15.76% | 84.99% | 45.62% | 14.69% | 32.85% | 51.76% |
| | **Flats/Mais** | 100.00% | -12.53% | 20.77% | 84.73% | 47.51% | 12.93% | 23.97% | 37.79% |
| | | **Postcode Sectors** | | | | | | | |
| **A only** | **Detached** | 4.42% | -5.63% | 9.55% | 93.30% | 53.59% | 25.36% | 56.46% | 77.99% |
| | **Semis** | 15.50% | -7.59% | 9.69% | 91.32% | 53.09% | 25.52% | 55.96% | 80.70% |
| | **Terraced** | 8.39% | -6.87% | 9.62% | 90.74% | 55.38% | 28.24% | 52.78% | 79.63% |
| | **Flats/Mais** | 10.14% | -5.49% | 8.15% | 97.70% | 66.42% | 34.56% | 62.67% | 88.02% |
| **A + B** | **Detached** | 38.04% | -10.04% | 15.06% | 85.71% | 41.54% | 15.79% | 35.19% | 54.25% |
| | **Semis** | 67.54% | -11.69% | 14.46% | 88.65% | 47.99% | 16.14% | 34.35% | 54.44% |
| | **Terraced** | 53.70% | -11.48% | 14.64% | 86.80% | 48.96% | 13.67% | 32.95% | 52.84% |
| | **Flats/Mais** | 55.70% | -11.67% | 15.16% | 92.87% | 57.08% | 16.69% | 32.80% | 51.01% |
| **ALL** | **Detached** | 100.00% | -10.14% | 22.63% | 78.73% | 31.07% | 12.39% | 25.80% | 39.61% |
| | **Semis** | 100.00% | -13.27% | 18.14% | 86.29% | 41.93% | 13.17% | 27.88% | 44.42% |
| | **Terraced** | 100.00% | -11.62% | 18.64% | 81.03% | 39.28% | 13.40% | 27.52% | 43.79% |
| | **Flats/Mais** | 100.00% | -10.99% | 23.19% | 85.98% | 45.92% | 12.34% | 24.16% | 37.06% |

# Appendices

## Table 2 Historic forecastability of IPD Regional Pages indices

| | All figs at end 3yrs (Mar 99- Mar02) | *% of total* | *Average error (end 3rd year)* | *Av. Abs(error) (end 3rd year)* | *+ve correl over 3 yrs* | *Inside 5% after 3 yrs* | *Inside 10% after 3 yrs* | *Inside 15% after 3 yrs* |
|---|---|---|---|---|---|---|---|---|
| **A only** | **All Regional Pages** | 60.00% | 5.46% | 7.97% | 94.95% | 41.41% | 67.68% | 83.84% |
| | **All property** | 57.50% | 4.27% | 6.28% | 95.65% | 60.87% | 82.61% | 86.96% |
| | **Industrial** | 51.11% | 8.01% | 8.72% | 100.00% | 39.13% | 52.17% | 82.61% |
| | **Offices** | 60.00% | 4.85% | 10.05% | 95.24% | 19.05% | 61.90% | 71.43% |
| | **Retail** | 71.11% | 4.87% | 7.29% | 90.63% | 43.75% | 71.88% | 90.63% |
| **A + B** | **All Regional Pages** | 95.76% | 6.95% | 14.06% | 81.65% | 25.95% | 43.67% | 57.59% |
| | **All property** | 100.00% | 4.07% | 12.98% | 85.00% | 35.00% | 47.50% | 55.00% |
| | **Industrial** | 88.89% | 12.74% | 17.61% | 77.50% | 22.50% | 30.00% | 47.50% |
| | **Offices** | 94.29% | 6.85% | 16.32% | 78.79% | 12.12% | 39.39% | 45.45% |
| | **Retail** | 100.00% | 4.43% | 10.21% | 84.44% | 31.11% | 55.56% | 77.78% |
| **ALL** | **All Regional Pages** | 100.00% | 5.16% | 14.95% | 77.71% | 24.70% | 41.57% | 54.82% |
| | **All property** | 100.00% | 4.07% | 12.98% | 85.00% | 35.00% | 47.50% | 55.00% |
| | **Industrial** | 100.00% | 7.46% | 19.52% | 68.89% | 20.00% | 26.67% | 42.22% |
| | **Offices** | 100.00% | 4.41% | 17.44% | 74.29% | 11.43% | 37.14% | 42.86% |
| | **Retail** | 100.00% | 4.43% | 10.21% | 84.44% | 31.11% | 55.56% | 77.78% |

# Appendix 2 – Performance of forecasts against actuals

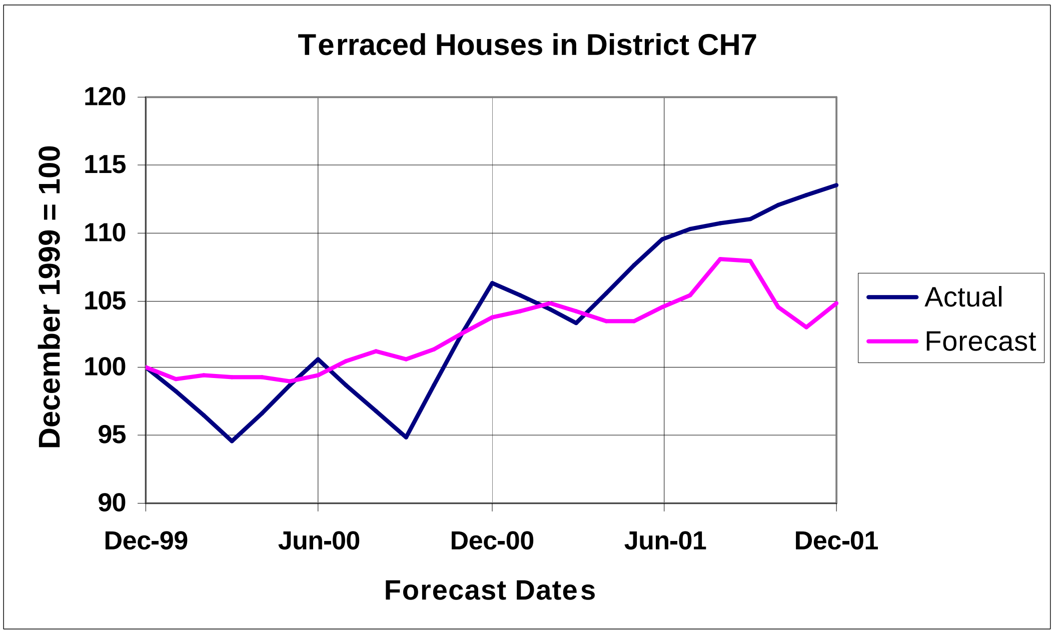

Figure A1 showing a forecast of a rise in house values.

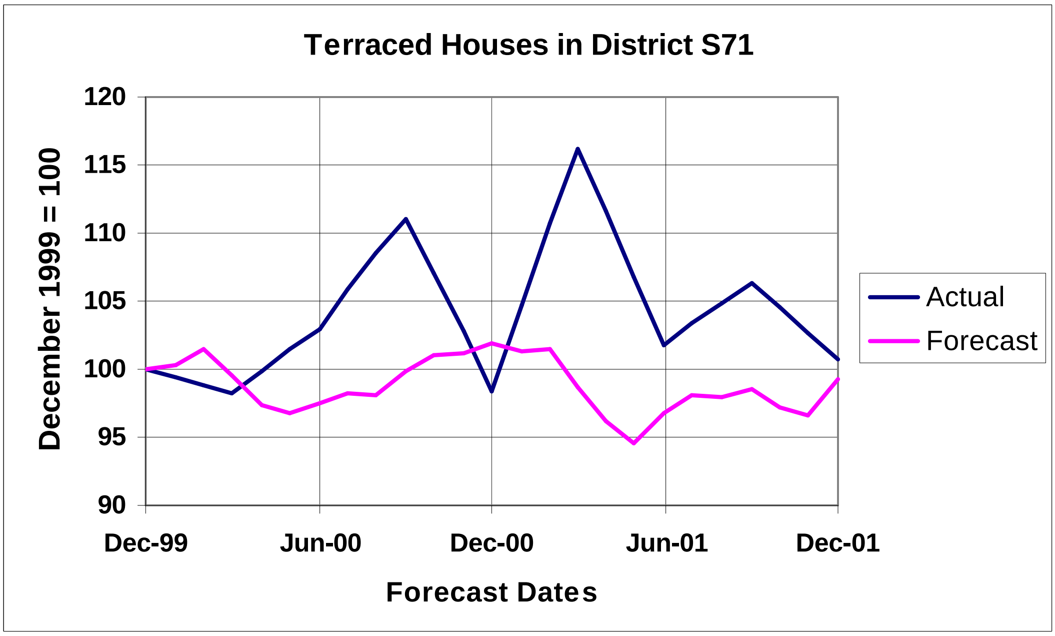

Figure A2. This forecast of a volatile series captures the somewhat lack-lustre performance of this District.

# Appendices

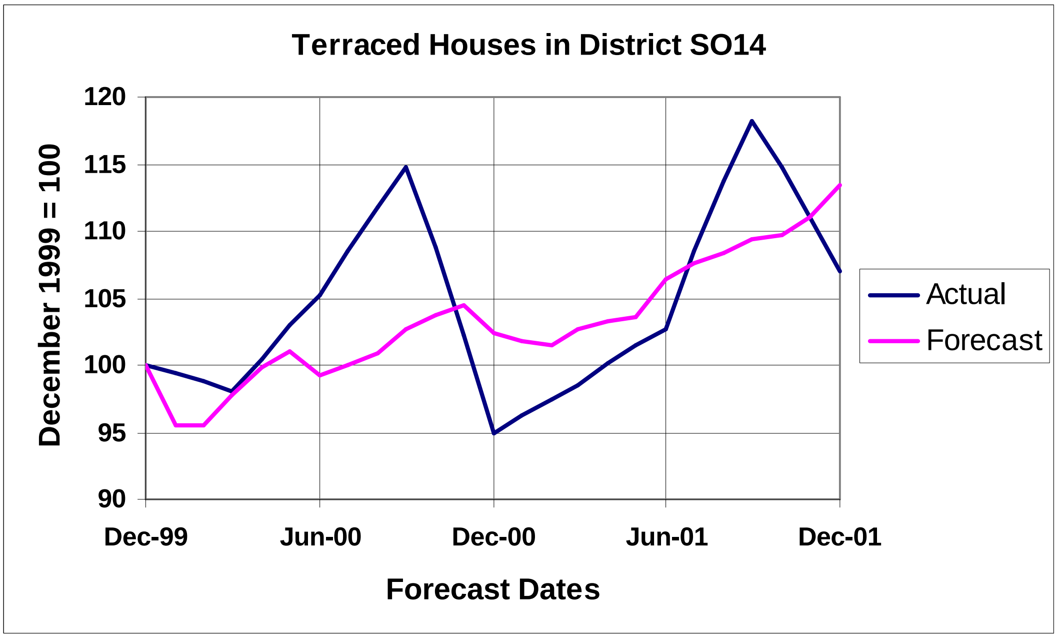

Figure A3. In spite of the volatility, the forecast manages to capture the apparent trend in this District.

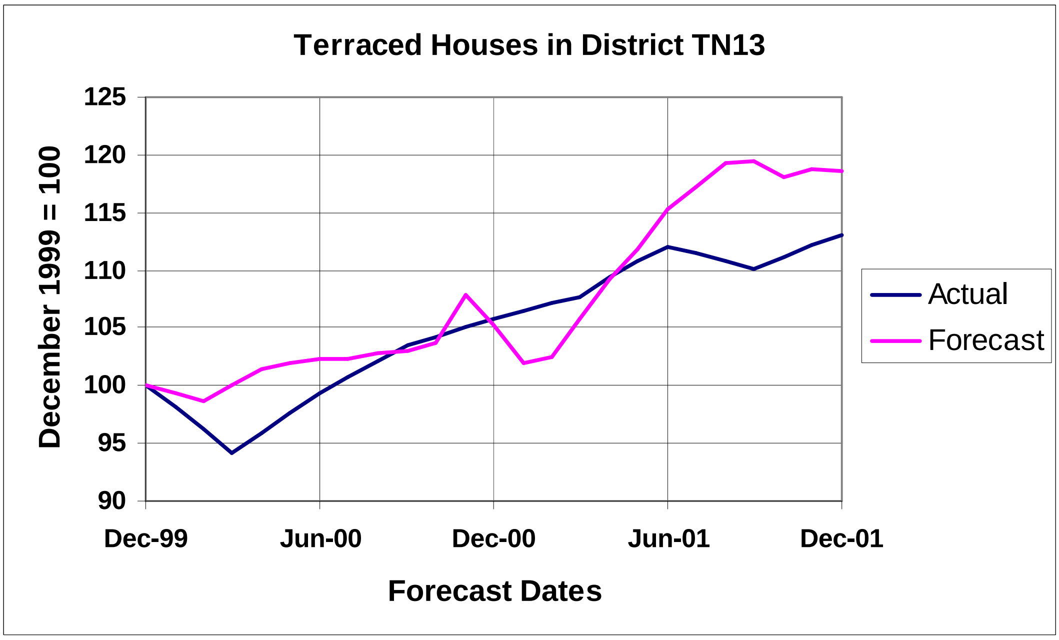

Figure A4. This District has relatively smooth series that is well captured by the forecast line.

# Appendices

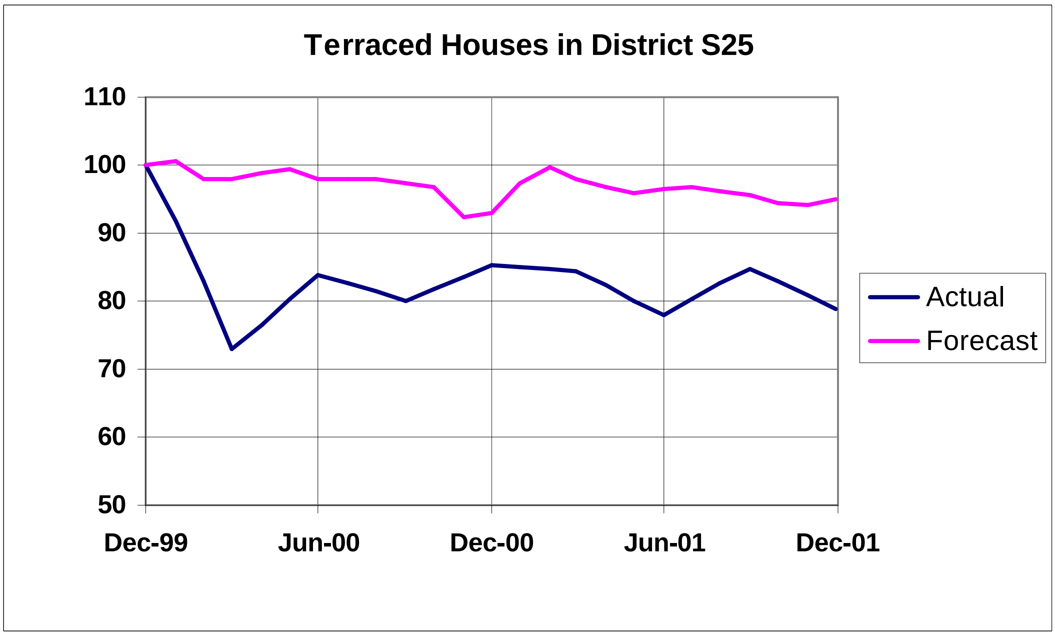

Figure A5. The rapid drop in the actual values is not captured by the model, but the forecast line captures the underlying flatness of this District.

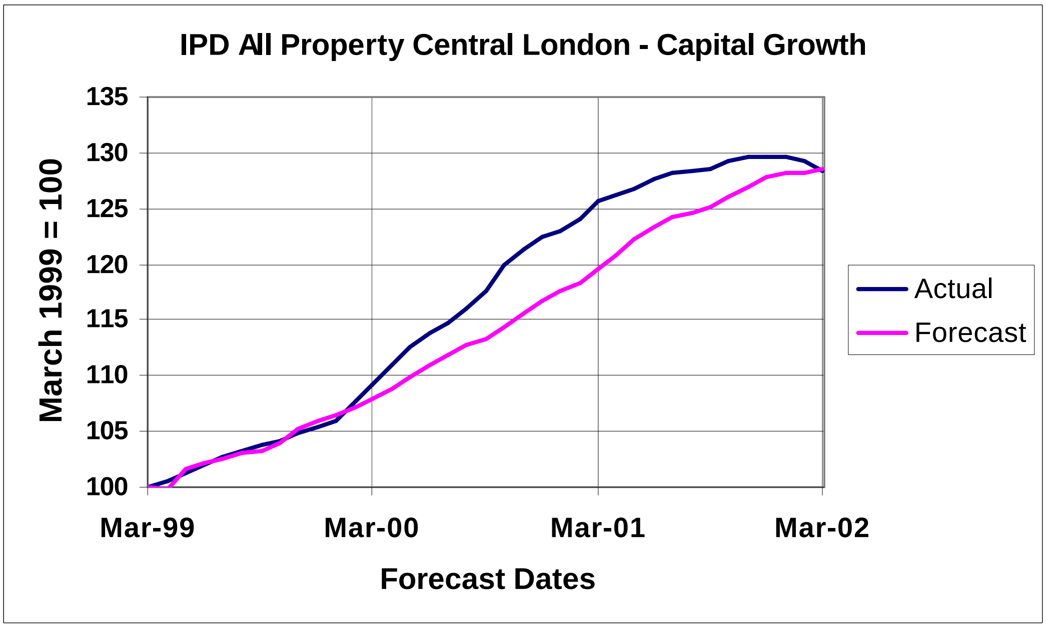

Figure A6. The forecast captures the rise in this series though it misses the two year level by about 4%.

# Appendices

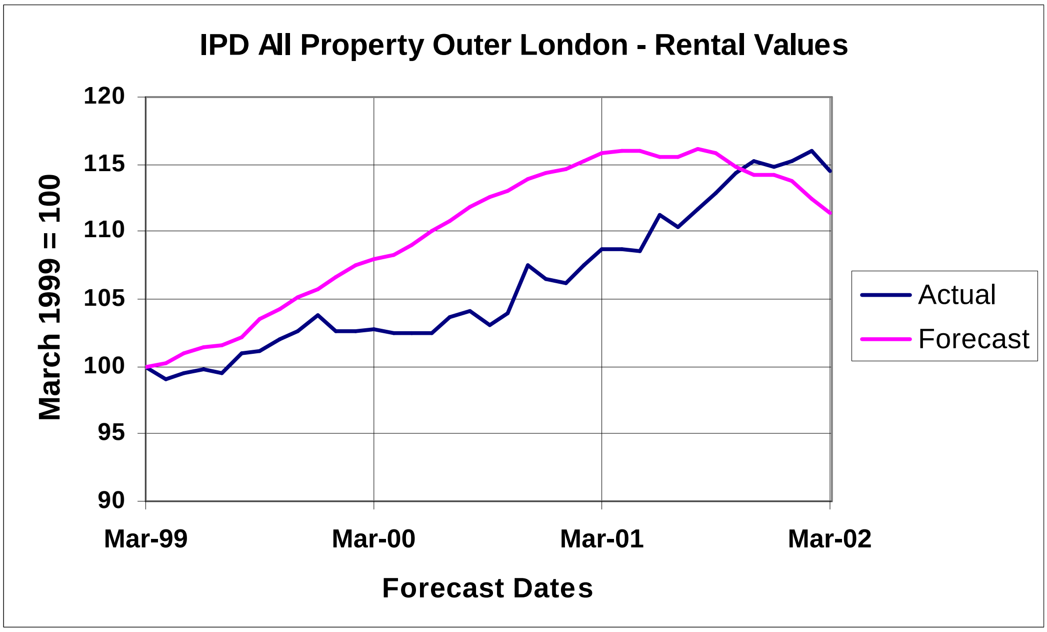

Figure A7. The turning point of this rather spiky rental series is anticipated by the forecast.

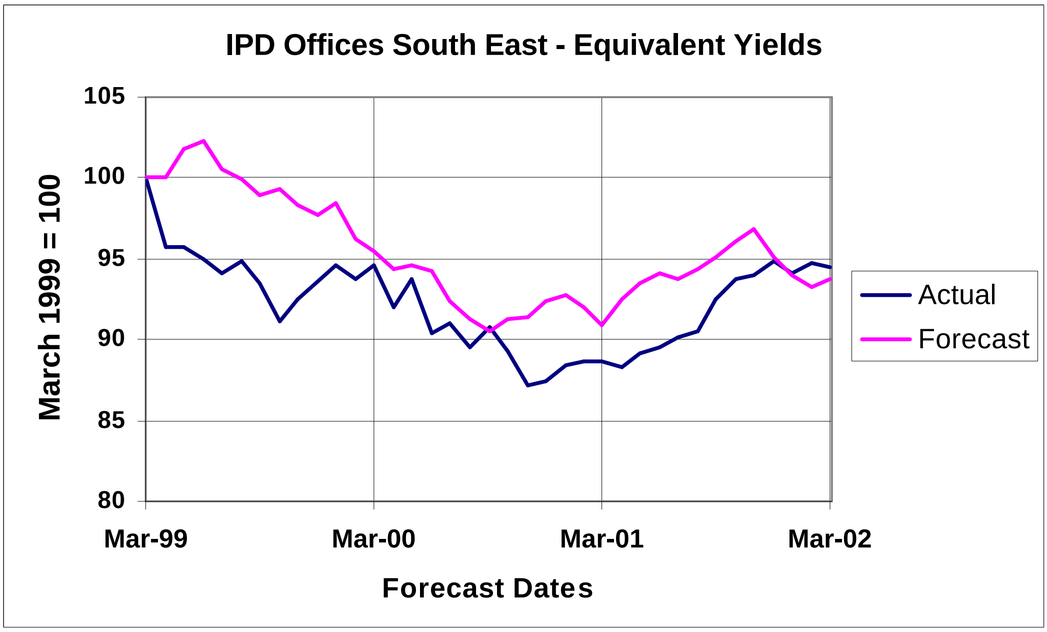

Figure A8. The minimum of this equivalent yield series is well captured by the forecast.

# Appendices

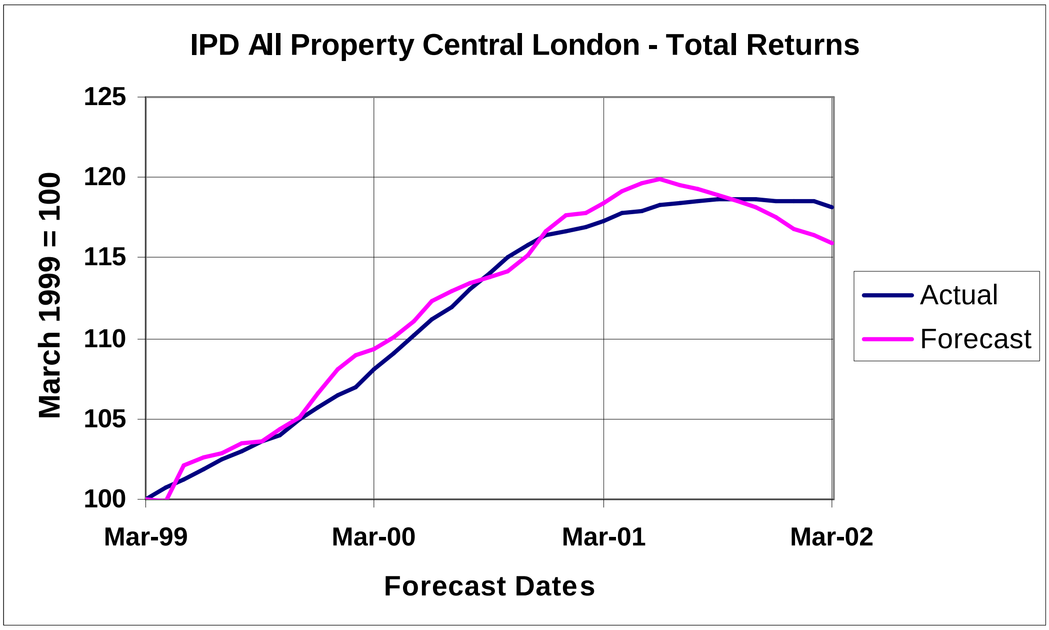

Figure A9 The forecast follows the actual values very closely throughout most of the period, anticipating the downturn by a few months at the far end of the forecast.

# Appendix 3 – a 3-year forecast

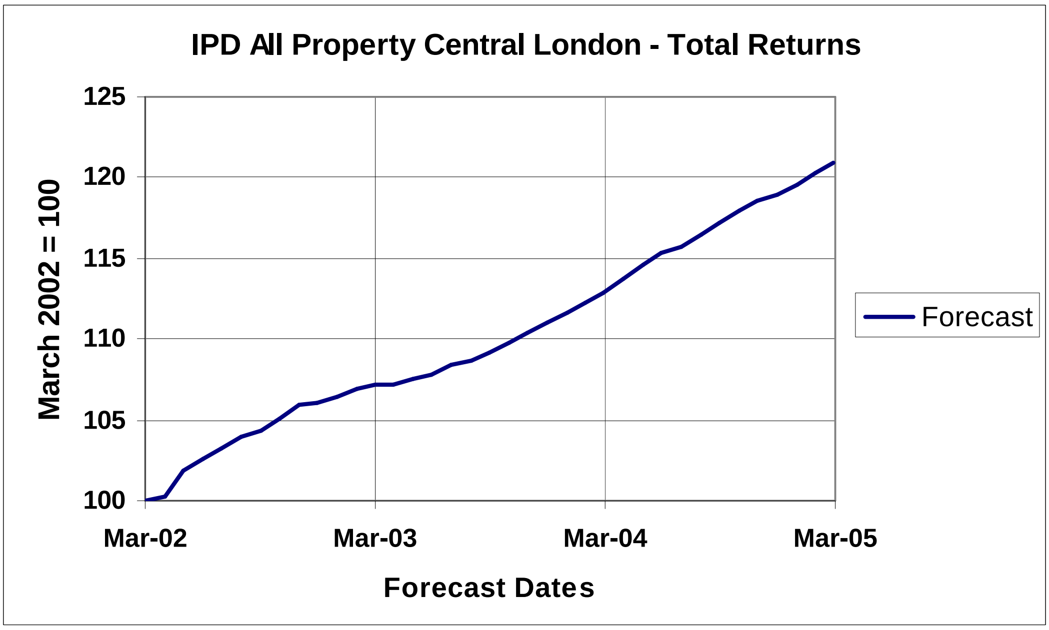

Figure A10 TFL’s forecast indicates a flattening in total returns in the first half of 2003, but returning to steady growth from around Sept 2003.