

Forecasting real estate returns using financial spreads

Article

Accepted Version

Brooks, C. and Tsolacos, S. (2001) Forecasting real estate returns using financial spreads. *Journal of Property Research*, 18 (3). pp. 235-248. ISSN 1466-4453 doi: <https://doi.org/10.1080/09599910110060037> Available at <https://centaur.reading.ac.uk/35970/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

Published version at: <http://dx.doi.org/10.1080/09599910110060037>

To link to this article DOI: <http://dx.doi.org/10.1080/09599910110060037>

Publisher: Routledge

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

This is an Author's Accepted Manuscript of an article published in the *Journal of Property Research* (2001) [copyright Taylor & Francis], available online at: www.tandfonline.com/10.1080/09599910110060037

Forecasting Real Estate Returns using Financial Spreads

Chris Brooks (corresponding author), ISMA Centre, PO Box 242, The University of Reading, Whiteknights, Reading RG6 6BA; e-mail: C.Brooks@reading.ac.uk

and

Sotiris Tsolacos, Jones Lang LaSalle, 22 Hanover Square, London, W1A 2BN

October 2000

Forecasting Real Estate Returns using Financial Spreads*

Abstract

This paper examines the predictability of real estate asset returns using a number of time series techniques. A vector autoregressive model, which incorporates financial spreads, is able to improve upon the out of sample forecasting performance of univariate time series models at a short forecasting horizon. However, as the forecasting horizon increases, the explanatory power of such models is reduced, so that returns on real estate assets are best forecast using the long term mean of the series. In the case of indirect property returns, such short-term forecasts can be turned into a trading rule that can generate excess returns over a buy-and-hold strategy gross of transactions costs, although none of the trading rules we develop could cover the associated transactions costs. We therefore conclude that such forecastability is entirely consistent with stock market efficiency.

J.E.L. Classifications: C32, C52

Keywords: real estate returns, spreads, forecasting, time series models, vector autoregressions

* The authors acknowledge comments from two anonymous referees on a previous version of the paper. The usual disclaimer applies.

1. Introduction

There is considerable evidence in the real estate literature that the behaviour of real estate returns, that is returns on real estate-backed assets, are related to economic trends and changing business conditions, and that they can be modelled, at least in part, using publicly available economic information (Chan *et al*, 1990; McCue and Kling, 1994 and Ling and Naranjo, 1997). This information is contained in a number of well publicised macroeconomic financial time-series used in financial markets. Two variables commonly used in the existing literature as indicators of future economic and monetary conditions are the term structure of interest rates and the gilt-equity ratio. Research in the area of real estate return predictability has been extended to investigate the possibility of superior investment performance of real estate portfolios utilising the predictive power of real economy variables and of financial aggregates.

Since the large majority of these studies examine the predictability of real estate returns in the context of the US, further research is warranted to provide empirical evidence on the linkages between the macroeconomy and real estate returns in other markets. This need is particularly pronounced in European markets where equity investors are increasingly considering real estate backed stocks as an investment vehicle in their portfolios (see Ball *et al* (1998) for the increasing importance of the indirect property vehicles). Although all equities are valued according to the expected future profit stream of the firms, it can be argued that the fundamentals underpinning these profit streams and returns on real estate backed stocks may differ from those on other types of stocks, resulting in distinct time series behaviour of returns between these asset classes. The main reason for this feature is the fact that the former are expected to reflect conditions in the underlying direct property market. This, coupled with the evidence that real estate

returns seem to be more predictable than returns on other assets (Liu and Mei, 1992), makes it possible that investors might be able to exploit this predictability in order to pursue more profitable investment strategies, or more efficient portfolio diversification in the mean-variance sense than portfolios which exclude real estate as an asset class. In addition, although previous studies have explained the historical variation in real estate returns using a set of predetermined variables within pre-specified frameworks, such as multifactor regression models and vector autoregressions, they have not tested for the forecasting ability of the different methodologies and variables. Therefore, the question that arises is how accurately the real estate return generating models in the existing literature predict the direction in future values of real estate returns and how suitable they are for the purpose of short-term and longer term forecasting.

The main thrust of the present study is to examine the performance of alternative econometric methodologies in forecasting real estate returns in the short-run in the UK. Real estate returns are defined as the growth rate in the Primark Datastream price index of real estate stocks. An analyst can study the time series properties of past data on real estate returns and utilise this information to make short term predictions. Alternatively, these forecasts can be made based upon a more general modelling methodology of real estate return determination which includes information contained in lagged macroeconomic or financial variables. The particular aim of this research is to forecast the short- and long-run variation in real estate returns utilising both the time series properties of the returns series alone and then a multivariate model that incorporates the term structure of interest rates and the gilt-equity yield ratio to exploit the economic information content of these variables. It is very important that investors have evidence on whether the time series properties of real estate returns can be used for forecasting

purposes or whether they should employ more general models and frameworks to carry out the forecasting analysis. This study provides the first evidence on comparative forecasting research work in the field of real estate returns in the UK based on four different prediction methodologies. The study also evaluates whether the relative forecasting accuracy of these methodologies is invariant as the forecasting horizon becomes longer.

The remainder of the paper is organised in five sections. A discussion of the methodology followed and data employed in this paper is presented in section 2. Section 3 reports the empirical findings while section 4 presents a comparative analysis of the forecastability of direct and indirect property returns. Section 5 offers a trading rule to exploit potential forecastabilities in the indirect returns series, and finally, section 6 concludes.

2. Forecasting frameworks, data, and production of forecasts

The simplest approach to forecasting real estate returns is to utilise the time series properties of the returns series itself. For this purpose, forecasts are derived from two models: (i) a forecasting model that simply comprises the long term mean of the returns series, estimated as the unconditional mean of a trailing sample (in this case of size 200 observations) and (ii) an ARMA model which is fitted to the property returns series, with the optimal model order being chosen by Akaike's information criterion, Schwarz's Bayesian criterion, and the Hannan-Quinn criterion.

An third approach to forecasting real estate returns is to consider additional explanatory variables that explicitly incorporate analysts' expectations about business conditions in a vector autoregressive (VAR) model framework. The set of variables for inclusion in the

VAR is determined by the conclusions of recent research investigating the modelling of stock returns or real estate returns. It is assumed that changing economic conditions are ultimately responsible for the variation in property share prices. There is ample evidence in applied economics research that the term structure of interest rates and the gilt-equity yield ratio (GEYR) adjust when market participants revise their expectations about the state of the economy. Therefore, they can be used as leading indicators of future business conditions (Estrella and Hardouvelis, 1991; Hardouvelis, 1994; Davis and Fagan, 1997).

Several authors have argued that the term structure of interest rates contains information about the future course of the economy (Campbell, 1987; Chen, 1991; Estrella and Hardouvelis, 1991; Hardouvelis, 1994). In particular, Estrella and Hardouvelis (1991) demonstrate that the yield curve has predictive power beyond that contained in the short term interest rate, and can help predict changes in national output up to four periods ahead. The term structure of interest rates is defined as the difference between the long-term and the short-term rates of interest. This spread is usually positive so that the yield curve is upward sloping. A tightening of monetary policy in order to curb inflation or inflationary growth will be reflected in higher short-term interest rates. If the financial markets believe that future economic activity will slow down and therefore that the future short-term rate of interest will fall, the yield curve will flatten out or even decline. High short rates in the current period can lead to lower aggregate investment expenditure that in turn results in a decline in future economic activity. Therefore, market participants expect lower prices and returns on traded property assets during the forthcoming economic contraction. Chan *et al* (1992) provide supporting evidence to the view that the term structure affects real estate returns but Liu and Mei (1992) did not find such

evidence. On the other hand, Ling and Naranjo (1997) found that the spread variable could become important in particular periods.

The predictive power of the gilt-equity yield ratio has not been examined in the literature on real estate returns, although research on general equity market indices suggests that it can have additional explanatory power beyond that contained in other macroeconomic or financial variables (see Levin & Wright, 1999; Brooks and Persaud, 2000). The GEYR ratio is measured as the ratio of the long gilt yield to the equity dividend yield. Levin and Wright argue that the GEYR has a normal “long run” level, reflecting a long run no-arbitrage relation between government bond and equity markets. Movements in this ratio are strongly affected by changes in the stock price and dividend yield (see Davis and Fagan, 1997). If, due to higher expected profitability, the dividend yield falls, the GEYR may become too high so that equities are thought expensive relative to bonds. Levin and Wright also state that bond yields do not fall, the low value of equity yields cannot be sustained. Therefore, equity prices, and in the context of the present study, the prices of property stocks, must fall to restore the long run equilibrium. When the GEYR is too low, equity prices are expected to rise in order to restore the equilibrium relationship.

Other macroeconomic variables, such as changes in unemployment, inflation, short term interest rates are not included in the results given here for two reasons. First, the results obtained by using these variables (not reported but available from the authors upon request) were inferior to those reported here using the two financial variables. Second, recent research by Brooks and Tsolacos (1999) has shown that the other variables do not even have any in-sample predictive power for real estate returns, evidenced by the joint

lack of significance of the coefficients in the VAR representations when general stock market effects are removed.

The data employed in this study comprise monthly observations on the Primark Datastream UK Property Index, the term spread (measured as the difference between the yields on a 20 year government bond and the three month Treasury bill rate), and the gilt-equity yield ratio (calculated by taking the ratio of the yield on 20 year government bonds and the dividend yield on the FTSE 100). All data were obtained from *Datastream International*, and cover the period January 1968 until January 1998, yielding a total of 361 sample points. The property index employed comprises a market value-weighted index constructed by Primark Datastream, based upon the top 26 property stocks traded on the London Stock Exchange. The relative weightings given to the component stocks are updated on a monthly basis.

Some summary statistics for the data are presented in table 1. The property returns series shows significant autocorrelation at the first lag, but none thereafter, and is both skewed and leptokurtic. Meanwhile, the GEYR series is leptokurtic but not skewed, and the spread series seems to show little departure from normality. Both the GEYR and spread series are very strongly autocorrelated, but there is no evidence that any of the series are non-stationary.

The effect of the above variables is examined within the context of an unrestricted reduced-form vector autoregressive (VAR) model, with three equations (one for each of the three variables: real estate returns, the term structure, GEYR), which is described by:

$$Y_t = \beta_0 + \beta_1 Y_{t-1} + \dots + \beta_m Y_{t-m} + u_t \quad (1)$$

where Y is the set (or 3×1 vector) of variables included in the system, the β terms give the sets of coefficient vectors (β_0 is a 3×1 vector of constants, β_1, \dots, β_m are 3×3 matrices of coefficients on the lagged variables, m represents the number of lags of each variable in each equation), and u_t is a vector of error terms (or innovations) which are assumed to be mutually uncorrelated and independent of the Y s. The number of lags of each variable to be included in the VAR is chosen using multivariate generalisations of Akaike's Schwarz's Bayesian, and the Hannan-Quinn information criterion (see, for example, Enders, 1995).

The forecasts are constructed as follows. The sample is split roughly in half, with the first 200 observations being used for in-sample model estimation. Then a series of out of sample forecasts up to six steps ahead are generated. The sample is then rolled forward by one observation, the models re-estimated, and a new series of forecasts constructed. This procedure is repeated until 160 such forecasts are generated. The UK economy has been the subject of numerous changes in monetary and fiscal policy over our 30 year sample period (see below), and thus the use of shorter windows used in a rolling fashion helps to minimise the possibilities of structural breaks whilst ensuring sufficient in-sample data to estimate the models and to produce the forecasts.

The forecasts for the different models are evaluated and compared on the grounds of mean squared error (MSE), mean absolute error (MAE) and the proportion of times that the model correctly predicts the return's sign. The "best" model is defined as the one with the lowest mean squared or mean absolute error, since this would indicate the model whose forecasts are closer to the realised values of the series, and whose forecasts are therefore the most accurate. However, it has also been shown (Gerlow *et al.*, 1993) that

the accuracy of forecasts according to traditional statistical criteria may give little guide to the potential profitability of employing those forecasts in a market trading strategy, so that models which perform poorly on statistical grounds may still yield a profit if used for trading, and vice-versa. Models that can accurately forecast the sign of future returns, or can predict turning points in a series have been found to be more profitable (Leitch and Tanner, 1991). Hence the proportion of times that the model correctly predicts the sign of the real estate return is also calculated.

In table 2, we present the results of Granger-causality tests for the three variables employed in the analysis, used to determine whether the VAR formulation seems sensible or not. A detailed description of the operation of this test and its interpretation can be found in Brooks and Tsolacos (1999). Changes in the property index are found to be caused by previous changes in its own values, and by previous changes in the values of the GEYR and the interest rate spread. Meanwhile, changes in the property index do not seem to Granger-cause changes in either of the other two series. Thus we conclude that the other two variables are weakly exogenous, and that the additional information contained in the GEYR and the term spread could be useful for modelling and forecasting changes in the property index.

3. Results

Initially the order of the ARMA and the VAR models was established as accomplished by the three information criteria. The in-sample minimisation of Akaike's criterion suggested that four lags of the autoregressive part and two lags of the moving average term should be used, while the other two criteria both favoured an ARMA(1,1) model. For the VAR model the same criteria indicated the inclusion of four lags, one lag and one

lag respectively. We employ the same number of lags for each variable so that the VAR is completely unrestricted. The results for the out of sample forecasts generated using the different methods, with ARMA and VAR models being estimated using the numbers of lags suggested by all the criteria, are given in Table 3 for 1 and 6 month prediction horizons.

In this application, the forecast accuracy evaluation measures give differing orderings of models. The VAR models, which attempt to incorporate the macroeconomic influences of the term spread and the gilt-equity yield ratio, give surprisingly poor overall performances for a forecasting horizon of one or six steps ahead, on MSE or MAE grounds. When evaluated using these criteria, the VAR(4) does particularly badly, and presumably represents an over-parameterisation, given that there are 12 right-hand side variables plus a constant, in each equation. The VAR(1), on the other hand, whilst not the best model, produces 1-step ahead forecasts which are among the most accurate. Even the small VAR model rapidly loses forecasting power, however, as the forecast horizon is extended. This is not surprising, since to produce multi-step ahead forecasts from a VAR, the values of the other variables in the system must also be forecasted (within the system), which introduces another source of uncertainty and potential error. In order to further evaluate why the VAR model (which contains all of the information in the univariate autoregressive framework, plus additional variables) did not produce superior forecasts compared with simpler models, we compute and plot in Figures 1 to 3 the impulse-response functions associated with unit shocks in each of the three explanatory variables, which are subsequently tracked for 24 months. Considering the signs of the responses, changes in the value of the property index are positively correlated over a short horizon, since the impulse response is large positive for lags 1 and 2, but the

impulse-responses thereafter are negative, although the two standard error bands (the dotted lines) span the horizontal axis, indicating that the impulse-response coefficients are no longer significant. The response of the property index to changes in the GEYR, in Figure 2, shows no clear pattern over time. The unit shock of GEYR leads first to a rise, then a fall, and then a rise in the value of the property index, although overall the effect is positive. Similarly for changes in the term spread, the effect upon the real estate index is difficult to determine, with first a fall, and then a rise, and then a fall in the value of the index following a unit shock. For all three plots, however, it is clear that after perhaps 5 or 6 periods, the effects of the shocks for the current value of the property index returns, are negligible.

The best models, for short-term forecasts are the long term mean on MSE grounds, and the profligate ARMA(4,2) on MAE grounds, while the VAR(4) produces the highest proportion of correct sign predictions. For the 6-month ahead forecasts, the long term mean produces the most accurate forecasts under all evaluation measures.

4. The Forecastability of Direct Versus Indirect Property Investment Returns

In this section, we compare the relative forecast accuracies of the various models presented above using the returns on both the indirect property investment vehicle (the equity quoted property index) and the returns on a direct property investment index - the IPD monthly total return series. The latter is compiled from valuation and management records for individual buildings in complete portfolios. The IPD total return is defined as the overall return on capital employed and is the sum of income return and capital growth. The Income return is income receivable net of management and irrecoverable costs divided by capital employed through the month. Capital growth is the change in

capital value from one valuation to the next, net of any capital flows, divided by capital employed (more information can be found in the IPD UK monthly index publications). Since the IPD series is only available from December 1986, we also re-compute the forecasts for the indirect property series analysed above for the same sample periods in order to facilitate comparisons. Thus forecasts are now generated using a moving window of 5 years' (60 months') worth of data, with forecasts up to 6 steps ahead being generated and then the sample rolled forward 1 observation until 74 such forecasts are produced. The results are presented in tables 4 and 5.

Again, model orders are chosen using Akaike's, Schwarz's Bayesian, and the Hannan-Quinn information criteria. These criteria chose for indirect property series, and for the ARMA models, a (2,5) - all criteria; for the direct (IPD) series the selected models were (6,3) - AIC, (3,4) - SBIC and HQIC. All three criteria chose a VAR(1) for the indirect series, while AIC chose a VAR(2), but SBIC and HQIC chose a VAR(1) for the direct series. The fact that the selected lag-lengths and sizes of models are larger for the direct property returns is indicative that there is potentially more forecastable structure in this series than that based on equities.

The results in table 4 are broadly similar to those in table 3, although there are some differences due to the change in sample length and the timing of the estimation and forecasts. In table 4, the VAR model for the indirect returns is almost universally superior. It provides the lowest MSE and MAE, and the highest proportion of correct sign predictions, except at the longest forecasting horizon, where the VAR is the worst model for sign prediction, getting only half of the 6-step ahead signs right.

Table 5 presents the results for the appraisal-based IPD series. In this case, it is clearly the ARMA approach, based upon only lags of the property returns themselves, which is the winner, while the VAR model, containing additional information regarding the financial markets, provides the least accurate forecasts, except in terms of sign predictions. The parsimonious VAR(1) model provides a marginally higher proportion of correct sign predictions at the short forecasting horizon, although the ARMA(1,1) and ARMA(1,3) models are better as the prediction horizon increases.

In general, we can state that the use of “generous” information criteria, such as Akaike’s, will lead to the selection of models which are over-parameterised, leading to poor forecasts. Those concerned with forecasting real estate returns should be concerned to use small models which do not use up valuable degrees of freedom, and which can generalise.

It seems also that there is a great deal more forecastability in the appraisal-based, as opposed to the equity-based series. This is evidenced by the fact that some models are able to correctly predict the direction of change of the former series up to 6 months ahead with over 90% accuracy. In financial markets, this accuracy would be considered phenomenal. The mean absolute and mean squared errors of the time series and VAR models are still smaller than those based on the long term mean, even at the 6 months ahead horizon. This is, however, hardly surprising, and the apparent predictability of the direct returns could be merely a statistical artefact, resulting from appraiser-induced smoothness of the series. This smoothness in the series is exacerbated by the frequency of the valuations. Monthly valuations of commercial properties do not provide sufficient time for economic and other information to induce a significant variation in total returns.

Therefore, current returns are bound to be linked to previous returns since they are not traded at the frequency of the property shares. On the other hand, this long memory in direct property returns may reflect characteristics of the property market that result in cyclical shocks having long-term effects on rental and capital values.

4. Production and testing of trading rules

Although the evidence for the apparent predictability of property returns, both direct and indirect, may be considered *prima facie* evidence of stock market informational inefficiency. However, according to modern definitions of stock market efficiency, a market may only be considered inefficient with respect to a particular information set if it is possible to make abnormal profits from trading on the basis of that information set. In order to investigate whether we can find any evidence of property market informational inefficiency, we develop a set of trading rules based on the forecasts produced as discussed above. For the reasons outlined above, the trading rules are developed only in the context of indirect, equity-based property returns. The results of a trading rule based on appraisal-based returns are likely to be misleading and unrealistic since such returns are not based on transactions data, and therefore it may not in fact have been possible to buy or sell property at the prevailing valuation prices. Additionally, investing in direct property is likely to occur large transactions costs and market orders to buy or sell direct property will only be executed with a long lag.

Although there are an infinite number of ways to operationalise a trading rule, we consider a simple approach where an investor is either long the equity-based property index or has no position, with the funds being held in treasury bills and therefore earning the short term “risk-free” rate of interest. The 1-step and 6-step ahead forecasts are

produced, and a decision about whether to be in property is made on that basis. We employ two rules: first, invest in the property stock index if the property return during the next period is forecast to be positive; second, invest in property only if the return forecast for the next month is greater than the average return over the sample period (i.e. greater than 0.87% per month). In each case, if a return is predicted to be sufficiently large to signal a buy, the property index notionally purchased. Then, the sample is rolled forward one observation, and if the next return predicted positive again, the index continues to be held, and so on. If the return is then predicted to be negative, however, the index is sold, with the proceeds invested in the risk free rate. It is assumed that the investment began at the start of our out of sample period and continued for the full 160 observations (i.e. a 13 years and 4 month period). We can effectively ignore risk in these calculations, since the benchmark will be a “buy-and-hold the property index” approach, and this is likely to represent more risk than our trading strategies, which involve being out of equities and in treasury bills for part of the period.

Table 6 presents the trading profitability results, denominated in simple annualised percentage returns, for each of the 5 models. The long-term mean model produces property return forecasts that are always (small and) positive. Therefore, the long term mean rule would be identical to passively holding the property index throughout the whole period. The simple average annualised return to simply holding the property index throughout is 8.07%. Overall, some of the models are able to modestly improve on this profitability. The best performing models in aggregate are the VAR models, which produce returns in excess of 9% per annum for the 1 step ahead forecasts with holdings in equities being based on any positive forecasted returns. In general, as one might have anticipated, trading on the basis of the 1-step ahead forecasts is more profitable than

trading on the basis of the 6-step ahead forecasts. Such a result ties in with the statistical evaluation, which showed that the one step ahead forecasts are typically more accurate and give a higher percentage of correct sign predictions.

A further cautionary note is required, however, in interpreting these results. The returns presented in table 6 are calculated gross of transactions costs. Sutcliffe (1997) suggests that an appropriate “round-trip” figure for transacting in companies contained in the FTSE-100 is 1.7% of the value of the purchase/sale per transaction for an investor. This figure is made up of bid /ask spread (0.8%), stamp duty (0.5%) and commission (0.4%). Even ignoring the fact that some of the companies in the property index are not in the FTSE 100, and are therefore likely to be less liquid with higher spreads, these transactions costs are likely to wipe out any profits made by the trading rules. The best performing rule generates a return of 9.26% per annum, but would generate approximately 3-4 trades per year, resulting in transactions costs of around 6% per annum. A filter which states that investors should only buy into the property index when the index is predicted to rise by more than its historical average, typically results in fewer trades, but also generates lower gross returns. For example, the best performing model in the context of the strict filter is the VAR(1) using one-step ahead predictions, giving an average annual return of 8.70%. This rule still requires 37 trades over the out of sample period (approximately 2-3 trades per year), again wiping out any excess gross profitability over a pure buy-and-hold equities strategy. Under modern definitions of market efficiency, forecastability in returns can only be considered evidence for market informational inefficiency if those forecasts can be turned into a trading rule that generates abnormal returns net of costs. Since our trading rules could not generate excess returns net of transactions costs, we would infer that the apparent forecastability of

indirect property returns is entirely consistent with stock market efficiency and the absence of possibilities for arbitrage.

6. Conclusions

The present study examined the performance of three different forecasting techniques that are available to analysts and investors to forecast real estate returns in the UK. Two of these techniques utilise the time-series properties of the real estate return series. The third methodology is a VAR model that allows the testing of the role of the term structure of interest rates and the gilt-equity yield ratio in forecasting real estate returns. Forecasts were produced for one to six months ahead and evaluated on the basis of standard forecast evaluation criteria.

This investigation has found that a vector autoregressive model, incorporating lagged values of term structure of interest rates and the gilt-equity yield ratio, provides superior one step (month) ahead predictions of real estate returns than other simpler univariate and naive forecasting models. However, at longer predictive horizons (that is up to six months), the relative advantage of the more complex multivariate model for forecasting equity-based returns disappears. The implication of this finding is that the multivariate asset pricing models developed in existing work to explain the intertemporal variation in real estate returns may only be suitable for immediate or short-term forecasts (one to two steps ahead). The recommendation in the present paper is that analysts can do no better than to forecast the long term behaviour of indirect, equity-based real estate returns series based on its long term mean. The short run variation around the mean can be forecast using multi-factor frameworks. Meanwhile, returns to direct property investments appear *prima facie* to be forecastable using time series models at least 6 periods into the future.

In the case of indirect property returns, such short-term forecasts can be turned into a trading rule that can generate excess returns over a buy-and-hold strategy gross of transactions costs, although none of the trading rules we develop could cover the associated transactions costs. We therefore conclude that such forecastability is entirely consistent with stock market efficiency.

References

- Akaike, H.(1974) A New Look at the Statistical Model Identification *IEEE Transactions on Automatic Control* **AC-19**(6), 716-723
- Ball, M., Lizieri, C. and MacGregor, B. (1998) *The Economics of Commercial Property Markets*, Routledge, London.
- Brooks, C. and Persaud, G. (2000) The Trading Profitability of Forecasts of the Gilt-Equity Yield Ratio *International Journal of Forecasting* forthcoming
- Brooks, C. and Tsolacos, T. (1999) The Impact of Economic and Financial Factors on UK Property Performance *Journal of Property Research* **16**(2), 139-152
- Chan, K.C., Hendershott, P.H. and Sanders, A.B. (1990) Risk and return on real estate: evidence from equity REITs, *American Real Estate and Urban Economics Association Journal* **18**, 431-52.
- Davis, E.P. and Fagan, G. (1997) Are Financial Spreads Useful Indicators of Future Inflation and Output Growth in EU Countries? *Journal of Applied Econometrics* **12**, 701-714.
- Enders, W.A. (1995) *Applied Econometric Time Series* Wiley, New York.
- Estrella, A. and Hardouvelis, G.A. (1991) The Term Structure as a Predictor of Real Economic Activity. *Journal of Finance* **46**, 555-576
- Gerlow, M.E., Irwin, S.H. and Liu, T-R. (1993) Economic Evaluation of Commodity Price Forecasting Models *International Journal of Forecasting* **9**, 387-397
- Hardouvelis, G. (1994) The term structure spread and future changes in long and short rates in the G7 countries, *Journal of Monetary Economics* **33**, 255-83.
- Leitch, G. and Tanner, J.E., (1991) Economic Forecast Evaluation: Profit Versus the Conventional Error Measures *American Economic Review* **81**(3), 580-590
- Levin E.J. and Wright, R.E., (1998) The Information Content of the Gilt-Equity Yield Ratio *Manchester School Supplement*, 89-101
- Ling, D. and Naranjo, A. (1997) Economic risk factors and commercial real estate returns, *Journal of Real Estate Finance and Economics* **14**, 283-307.
- Liu, C.H. and Mei, J. (1992) The predictability of returns on equity REITs and their comovement with other assets, *Journal of Real Estate Finance and Economics* **5**, 401-18.
- McCue, T.E. and Kling, J.L. (1994) Real estate returns and the macroeconomy: some empirical evidence from real estate investment trust data, 1972-1991, *Journal of Real Estate Research* **9**(3), 5-32.
- Sutcliffe, S. (1997) *Stock Index Futures: Theories and International Evidence*, Second Edition, Thomson Business Press, London.

Table 1. Summary Statistics for the Variables 1968-1998

Statistic	Indirect Property Returns	Log-Changes in the GEYR	Log-Term Spread
Mean	0.869	-0.814	1.182
Variance	64.720	0.015	5.000
Skewness	-0.659**	-0.061	-0.233
Excess Kurtosis	5.048**	1.174**	0.0346
Minimum	-43.905	-1.199	-4.820
Maximum	38.869	-0.3512	6.730
Bera-Jarque	408.226**	20.909**	3.280
Acf(1)	0.125**	0.873**	0.960**
Acf(2)	-0.099	0.752**	0.923**
Acf(3)	-0.027	0.670**	0.883**
Acf(4)	-0.022	0.562**	0.845**
Acf(5)	-0.071	0.479**	0.810**
LB-Q*(12)	11.957	1063.998	1650.394**

Notes: Bera-Jarque is a normality test statistic, asymptotically distributed as a $\chi^2(2)$ under the null; Acf(x) denotes the autocorrelation coefficient at lag x; * and ** denote significance at the 5% and 1% levels respectively; LB-Q*(12) denotes the Ljung-Box statistic for autocorrelation of order up to 12, which is distributed asymptotically as a $\chi^2(12)$ under the null.

Table 2. Granger Causality Tests for VAR(4)

<u>Causality to</u> <u>(Dependent</u> <u>Variable)</u>		<u>Causality from (independent variable)</u>		
		Indirect Property Index	GEYR	Term Spread
Indirect Property Index		0.0232	0.0193	0.0957
GEYR		0.4954	0.0000	0.2978
Term Spread		0.5144	0.1154	0.0000

Note: Cell entries are *p*-values.

Table 3. Out of sample forecast accuracies for indirect (DS) property return forecasting models 1968 - 1998

Steps	Forecasts using					
	Long term mean	ARMA(4,2)	ARMA(1,1)	ARMA(1,3)	VAR(1)	VAR(4)
Panel A: Mean Square Error						
1	98.9625	143.8746	100.6335	100.7223	99.6245	114.0778
6	98.8341	153.9693	100.8244	100.6871	102.4357	109.0905
Panel B: Mean Absolute Error						
1	7.2197	7.1819	7.3820	7.3307	7.2049	7.5951
6	7.2197	7.8485	7.3590	7.3500	7.3528	7.3528
Panel C: % Correct sign predictions						
1	58.09	51.86	55.62	57.50	56.88	58.25
6	58.13	46.88	50.00	48.13	55.62	52.50

Table 4. Out of sample forecast accuracies for indirect (DS) property return forecasting models 1986-1998

Steps	Forecasts using		
	Long term mean	ARMA(2,5)	VAR(1)
Panel A: Mean Square Error			
1	56.8638	76.0518	55.5625
6	56.8842	72.9506	46.4988
Panel B: Mean Absolute Error			
1	5.8497	6.9969	5.7702
6	5.8521	6.6779	5.5091
Panel C: % Correct sign predictions			
1	58.11	51.35	59.46
6	58.11	51.35	50.00

Table 5. Out of sample forecast accuracies for direct (IPD) property return forecasting models 1986-1998

Steps	Forecasts using					
	Long term mean	ARMA(4,2)	ARMA(1,1)	ARMA(1,3)	VAR(1)	VAR(2)
Panel A: Mean Square Error						
1	0.6717	0.2381	0.1729	0.1756	0.7052	0.7061
6	0.6724	0.8661	0.2668	0.2739	0.8917	0.9332
Panel B: Mean Absolute Error						
1	0.5459	0.3598	0.2808	0.2999	0.6338	0.6527
6	0.5456	0.6212	0.3556	0.1766	0.7447	0.7677
Panel C: % Correct sign predictions						
1	93.24	87.84	91.89	90.54	94.60	91.89
6	91.89	75.68	91.89	90.54	89.49	85.14

Table 6: Profitability of Indirect Property Index Trading Rules based on Out of Sample Forecasts

Steps ahead forecasts produced for:	Trade when forecast next return > 0.87% per month		Trade when forecast next return positive	
	Number of trades	Average annual return	Number of trades	Average annual return
		<u>ARMA(1,3)</u>		
1 step ahead	23	2.69%	74	7.91%
6 steps ahead	1	2.95%	68	6.19%
		<u>ARMA(1,1)</u>		
1 step ahead	22	4.13%	66	8.93%
6 steps ahead	1	2.50%	76	1.87%
		<u>ARMA(4,2)</u>		
1 step ahead	48	4.55%	77	7.48%
6 steps ahead	34	8.50%	79	5.18%
		<u>VAR(4)</u>		
1 step ahead	37	7.82%	50	9.26%
6 steps ahead	34	2.28%	51	4.99%
		<u>VAR(1)</u>		
1 step ahead	37	8.70%	41	9.11%
6 steps ahead	56	3.50%	41	4.93%
Buy and hold equities			1	8.07%
Buy and hold treasury bills			1	2.50%

Figure 1: Impulse Responses and Standard Error Bands for Innovations in the Property Index




Figure 2: Impulse Responses and Standard Error Bands for Innovations in the GEYR




Figure 3: Impulse Responses and Standard Error Bands for Innovations in the Term Spread

