Estimating the volatility of property assets

Problem presented by

Stewart Richie

Actuarial Profession

Executive Summary

When an investor is allocating assets between equities, bonds and property, this allocation needs to provide a portfolio with an appropriate risk/return trade-off: for instance, a pension scheme may prefer a robust portfolio that holds its aggregate value in a number of different situations. In order to do this, some estimate needs to be made of the volatility or uncertainty in the property assets, in order to use that in the same way as the volatilities of equities and bonds are used in the allocation. However, property assets are only valued monthly or quarterly (and are sold only rarely) whereas equities and bonds are priced continuously and recorded daily. Currently many actuaries may assume that the volatility of property assets is between those of equities and bonds, but without quantifying it from real data. The Study Group was challenged to produce a model for estimating the volatility or uncertainty in property asset values, for use in portfolio planning.

The Study Group examined contexts for the use of volatility estimates, particularly in relation to solvency calculations as required by the Financial Services Authority, fund trustees and corporate boards, and it proposed a number of possible approaches.

This report summarises that work, and it suggests directions for further investigation.

Version 1.2 (Final)
14 August 2008

Report coordinator

Melvin Brown (Smith Institute)

Contributors

Jeff Dewynne (Universities of Oxford and Wollongong)
Tom Fischer (Heriot-Watt University)
Sam Howison (University of Oxford)
Kamil Kulesza (Isaac Newton Institute for Mathematical Sciences)
Andrew Lacey (Heriot-Watt University)
Karina Piwarska (University of Warsaw)
David Wilkie (Heriot-Watt University)
Maxim Zyskin (University of Texas)
Angus MacDonald (Heriot-Watt University)

Contents

1   Introduction

1.1   Background and motivation

(1.1.1)   Estimates of the volatility or uncertainty are currently used in property assets, as in equities and bonds, to evaluate the solvency of insurance companies and large pension schemes. In this context, solvency is defined as the excess of assets (such as equities, bonds, property and cash) over liabilities (such as insurance policies or pension payments) expressed as a fraction of liabilities.

(1.1.2)   Solvency estimates are required by the Financial Services Authority (FSA), fund trustees, and corporate boards. The timescale for which solvency is evaluated may exceed 35 years. Since 2004, the FSA has required a market consistent approach to the evaluation of solvency. In this regime, solvency estimates are required to certain levels of confidence, and so they require the use of variances and covariances between different types of portfolio holdings; these can be calculated using the respective volatilities in their rates of return.

(1.1.3)   In a similar way, volatilities are also required for activities such as portfolio planning. Pension fund trustees and insurance company investment managers may review their asset allocation strategy typically every three to five years. Within such periods they may also adjust their portfolios as conditions change. In both these cases, volatility estimates inform decisions in portfolio planning.

(1.1.4)   For pension schemes, although there are many variable elements, there is a guaranteed element in the final pension. For life assurance companies, there is the guaranteed sum assured, but for with profits policies there are also bonuses. Property and bonds mainly are held against promises which have a guaranteed element. The tension between the guarantees and the uncertainties in the investment assets generates insolvency risk. Monte Carlo simulations are run to assess this risk, by modelling a variety of futures for the benefits i.e. the liabilities and for the growth in assets, and so evaluate the probability of insolvency, and identify asset mixes that minimise this probability. It is in this context, that estimates of the volatilities of and correlations between various asset types are required.

(1.1.5)   Estimates of solvency are highly sensitive to correlations between asset types, and they may change. For example, if equities and property act in concert during extreme economic shocks then the diversification plan for less extreme situations may not work. However, it is volatility estimates that determine the scale of uncertainties, as represented in variance-covariance matrices.

(1.1.6)   A key feature of commercial property assets is that they are only valued
           monthly or quarterly (and are sold only rarely) whereas equities and bonds
           are transacted and thus priced with high frequency and recorded daily.
           Currently, many actuaries will assume that the volatility of property assets
           is between those of equities and bonds, but without quantifying it from
           real data.

(1.1.7)   Commercial property indices are published, for instance by IPD, which
           use surveyors estimates. However, the volatility in such an index may not
           correctly represent the long-term risk, because the sale price of a property
           is subject to various unpredictable factors that mean it will not be directly
           linked to the index. This is similar to the 'thin trading' problem for
           equities with small capitalisation ('small cap' equities). They appear to
           have a good risk-adjusted return, because infrequent trading means that
           the volatility of the shares is understated. This issue has been addressed by
           Dimson (1979), and by Roll (1981) but there is no corresponding analysis
           for property assets.

1.2   Challenges for the Study Group

(1.2.1)   The challenge for the Study Group was to produce a model for estimating
           the volatility or uncertainty in property asset values, for use in portfolio
           planning and solvency assessments.

(1.2.2)   The following questions were put to the Study Group:
           (a) What information do the surveyors estimates use? Are they based
               on commercial rents, or do they use information from property sales
               when available?
           (b) Might other information such as returns on real estate investment
               trusts (REITs) be useful surrogates for estimating property portfolio
               volatilities.
           (c) Can a model for a sale price be obtained from using the IPD index,
               but then also multiplying by a random factor F at the time of sale,
               where F has a distribution over perhaps the interval from 0.8 to 1.2?
               If one made such a model, is there data available that could be used
               to validate it?
           (d) Is there data available to find what other economic variables F is
               correlated with?
           (e) How does F change with the time horizon?
           (f) Should a model for property asset values look at the extent to which
               similar properties sold at nearby times and locations can be used to
               give a more specific measure of volatility?
           (g) What can be inferred from the change in variance over different time
               horizons, allowing for the extent to which there is serial correlation
               that might have an impact? (Booth & Mercato 2004)

2   Current understanding and definitions

2.1   General observations

(2.1.1)   Discussion with actuarial experts at Heriot-Watt University, who were invited to participate in the Study Group provided further background to the challenges posed. These are recorded in this section.

(2.1.2)   The usefulness of statistical indices for portfolio analysis is highly dependent on the diversity and size of the portfolio of properties held; they are most apt for large portfolios with a spectrum of properties that matches the index.

(2.1.3)   The distribution of pension fund sizes is very wide, and includes a large number of small funds. However, the distribution of insurance companies sizes is distributed towards the large sizes. For smaller companies it is important to analyse the details of heterogeneities in individual property holdings with respect to lease and rental arrangements, some of which may extend over many years. In these cases the statistical methods based on volatilities are not applicable.

(2.1.4)   An indication of the levels of property holding as a proportion of total assets can be found from a sample pension fund and insurance company balance sheets over 2006-7: 6%, 7%, 3.5%, 0.3%. This small sample is at the lower end of the range $5-20\%$ which was suggested in the presentation of the problem to the Study Group.

(2.1.5)   A perceived distinction between pension and life assurance companies was that the former are more likely to be interested in total property value, whereas the latter might respond differently to capital and rental values, being more able to rapidly change the equities portfolio in response to changes in capital value.

(2.1.6)   UK funds do not hold residential property. But building societies have large amounts of high frequency transactional data and hedonic models (Booth & Mercato 2004) for indices which express the transaction value in terms of property attributes such as number of rooms, location, land, etc, through statistically fitted parameters. Such indices have the advantage of being based on actual transactions and do not suffer biases and smoothing effects of valuation by agents.

(2.1.7)   In addition to their commercial property valuation index, IPD (the publisher of real estate indices) have a commercial residential valuation index. This opens a possible way to investigating the degree of volatility smoothing arising from the valuation process, by comparing it with residential transactional data, if allowance could be made for any differences between commercial and private residential properties. The Study Group obtained

Nationwide and Halifax transactional data but was unable to obtain from IPD the residential valuation index, for which a price must be a paid, and so was unable to evaluate this option.

2.2   Definitions

(2.2.1)   The indices considered in this report, give the return R_i (index value) on a holding S_i at time t_i with respect to its value S_0 at a reference time t_0. Thus,

$$R_i = \frac{(S_i - S_0)}{S_0}, \tag{1}$$

which means that

$$S_i = (1 + R_i)S_0. \tag{2}$$

(2.2.2)   There are a very large number of volatility measures, some with free parameters whose values are set 'from experience of use' in given applications. For example consider the 'exponential' form for a zero mean series $z_i, i = 1, ..., \infty$, with parameter λ:

$$E(z) = \frac{1}{(1 - \lambda)} \sum_{n=1}^{\infty} \lambda^n z_n \tag{3}$$

$$E(z^2) = \frac{1}{(1 - \lambda)} \sum_{n=1}^{\infty} \lambda^n z_n^2 \tag{4}$$

In the Study Group, a number of simple measures of variance in the return series were used to compare different real estate property indices and property equities.

(2.2.3)   For clarity, the results presented in this report are based on the following definition of volatility. Let R_i be the compound return (index value) at time t_i in a time series of constant interval Δt starting from some reference time t_0. The volatility v_i is expressed in terms of a variance of the series of incremental returns $r_i, r_{i-1}, r_{i-2}, ..., r_{i-m}$, where $m + 1$ is the number of samples. Thus,

$$r_i = \frac{(R_i - R_{i-1})}{R_{i-1}} \tag{5}$$

$$\bar{r}_i = \frac{\sum_{j=0}^{m} r_{i-j}}{m + 1} \tag{6}$$

$$V_i = 100\sqrt{\alpha_{year} \frac{\sum_{j=0}^{m}(r_{i-j}^2 - \bar{r}_i^2)}{m + 1}}, \tag{7}$$

where V_i is the volatility (%) at time t_i, and $\alpha_{year}\Delta t$ is an interval of 1 year. The factor $\sqrt{\alpha_{year}}$ scales the volatility to 1 year.

(2.2.4)   By a straight forward extension of the definition of volatility, covariances, and thus correlation coefficients, between pairs of holdings, return indices may be estimated. There are a number of ways estimating correlations between pairs of holdings in a portfolio. However, given such correlation coefficients (whether prescribed or derived), volatilities are required to estimate the uncertainties in holdings and to compute confidence levels for their value.

3   Real estate returns modelling

3.1   US Residential Property Model

(3.1.1)   The Study Group briefly reviewed the predictions of a US Federal Reserve index for UK residential properties. See Figure 1. The model in Martin (2005) predicts that the primary force underlying the evolution of real house prices is the systematic and predictable changes in the working age population driven by the baby boom. The model is calibrated to U.S. data and tested on international data. One surprising success of the model is its ability to predict the boom and bust in Japanese real estate markets around 1974 and 1990.
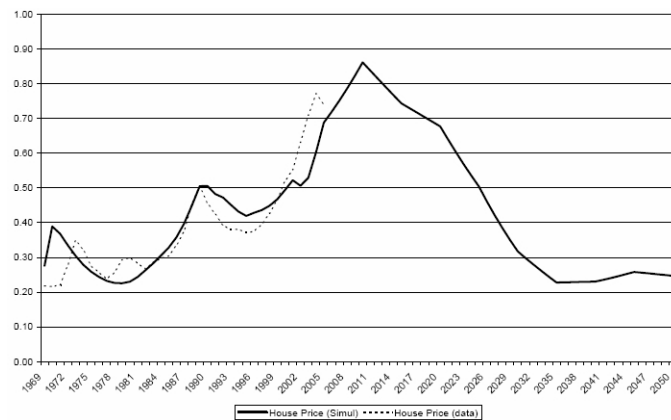
Figure 1: FED demographic model: United Kingdom simulated and real house prices

3.2   Index modelling

(3.2.1)   The Study Group noted the work of Booth & Mercato (2004) in developing auto-regressive models of the IPD property index against bonds, equities and treasury bills, and also reference therein to ways of removing the smoothing present in the IPD index on the basis of assumptions of the underlying processes of smoothing. The best results (no evidence

of serial correlation in the residuals) are obtained for the second order auto-regressive model. Bonds are found to have low or small negative correlation with the index and equities to have a positive correlation. This work is useful in forecasting property index behaviour and would bear further scrutiny, although it is primarily based on valuation indices not on property transaction data.

(3.2.2)   Commercial property is infrequently traded so limiting the information in transactional data. Earlier identified downward biasing of the volatility estimates of infrequently traded shares was addressed by Dimson (1979) and then corrected by Fowler & Rorke (1983). This work is based on actual share transactions of long periods compared to normal transaction periods. The Study Group was not certain that this methodology could be applied given the low frequency and irregularity of transactions in the UK commercial property market. Further work would be required to establish this. The method could be applied to the IPD property total returns index, to remove the biasing of volatility caused by the use of monthly evaluation data. However, that data would still be subject to various other sources of evaluation bias.

(3.2.3)   The Study Group began preliminary work to develop a stochastic model to simulate the IPD total returns index (IPD-TRI). The model is based upon drawing random samples from a distribution of incremental logarithmic returns $y_i | i - s = 1, 2, ...$ of the IPD-TRI. The variable y_i sampled from the series of IPD-TRI returns $R_i | i = 1, 2, 3, ...$ is

$$y_i = \log\left(\frac{R_i}{R_{i-1}}\right) - \langle y \rangle_{(i,s)} \tag{8}$$

where the local average $\langle y \rangle_{(i,s)}$, is defined by

$$\langle y \rangle_{(i,s)} = \frac{1}{s} \sum_{j=i-s}^{i-1} y_j. \tag{9}$$

Here, s is the size of the averaging window.

(3.2.4)   The probability density function of y for $s = 3$ is shown in Figure 2, and multiple series of successive random samples from $\rho(y)$ are shown in Figure 3. The corresponding simulations of the IPD total return are obtained by successive iterative solution of equation 8, examples of which are shown in Figure 4.

(3.2.5)   The work of the Study Group on simulating the IPD index is at an exploratory stage. The simulation described above assumes no correlation between the incremental changes in the IPD total return index. Further work might investigate the effects of window size s, and the implications and developments of this model.
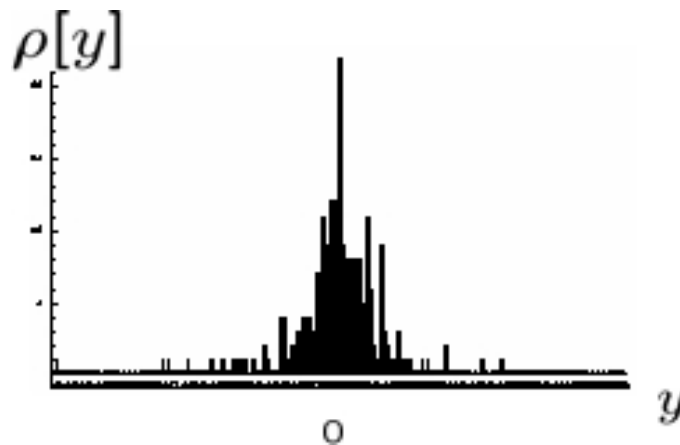
Figure 2: The probability probability density function of IPD incremental returns for averaging window size $s = 3$
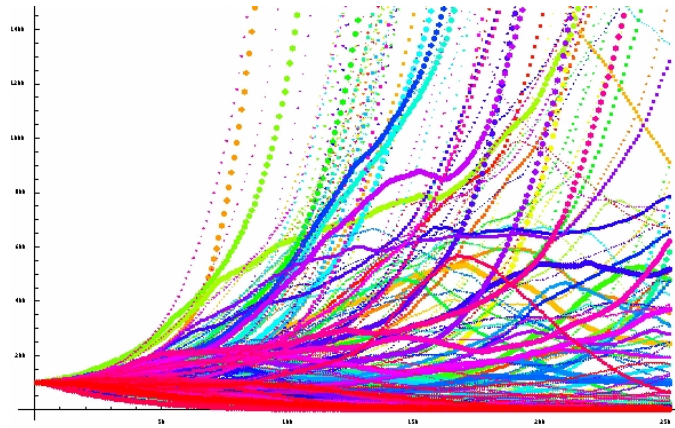

Figure 3: Simulation of incremental IPD total returns for averaging window size $s = 3$ over 250 months

4   Real estate index trends and volatilities

4.1   IPD commercial property valuation indices and property equities

(4.1.1)   The IPD commercial property valuation indices that were supplied to the Study Group, are for every month from December 1986 to December 2007, and are broken down into retail, office and industrial sectors as well as being provided for the whole market. Details of how the indices are computed from various incomes, costs, and capital growth may be found in Booth & Mercato (2004), IPD - Investment Property Databank (1985). Figure 5 shows the IPD total return index, while Figure 6 shows

Figure 4: Simulation of IPD total returns for averaging window size $s = 3$ over 250 months

the return indices for capital value and rental, which are components of the total index. The latter are based on the estimated values of individual properties were they to come on the market at the time of the evaluation. (The nomenclature in the legend of these figures is that of the IPD property index.)
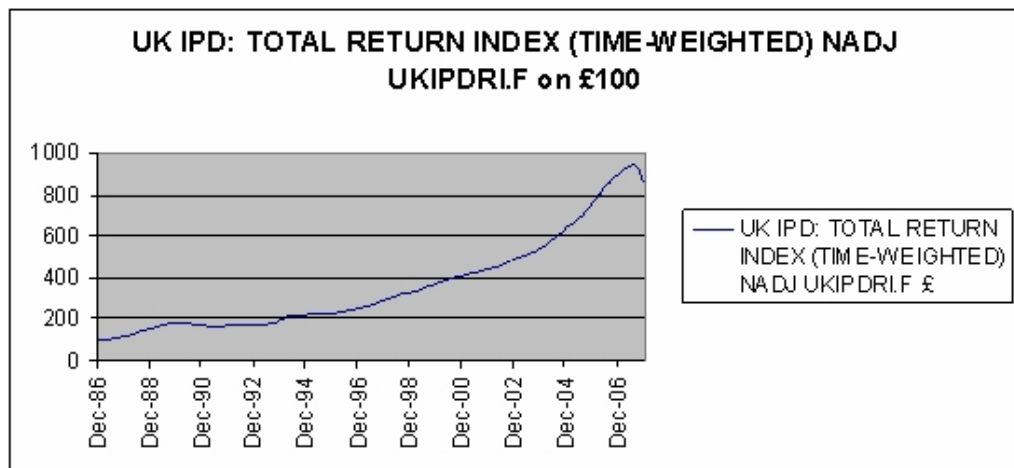
Figure 5: IPD total return value index

The recent changes in the property market are reflected in the down-turn in the total return index and in a period during which the capital return index has exceeded the rental return index, having long previously been below it.

(4.1.2)   The Study Group chose to focus upon the IPD total return index and to compare it with the share indices of the property companies, Land Secu-
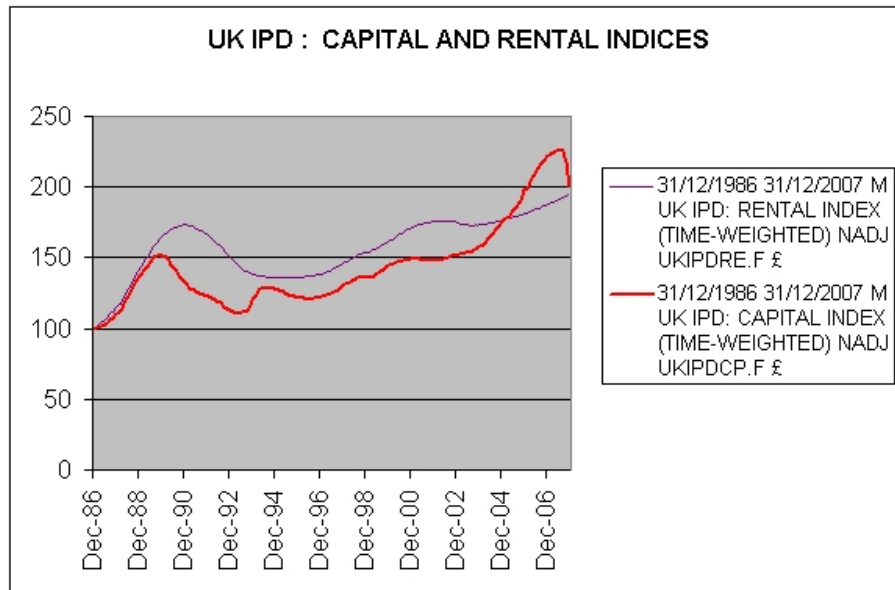
Figure 6: IPD capital and rental value indices

rities Group and British Land, and also with the FTSE Real Estate and Development (FTSE-RED) index. The data for these indices was obtained from the companies' websites. The motivation for this comparison was the conjecture that while the share indices would reflect overall sentiment in the equities market they might also, in part, reflect underlying volatility in the property market, which is smoothed out by the evaluation process on which the IPD indices are based. Such companies carry portfolios of properties and have frequently traded shares. Figure 7 shows the indices of the property companies and FTSE-RED over the period January 2000 to December 2007. Both property companies correlate well with the FTSE-RED index.

(4.1.3)   Figure 8 exemplifies how over the same period, the IPD total return index not only has lower volatility and amplitude than property equities; its peak lags the peak in property equities. A comparison of property shares with the FTSE-100 index, shown in Figure 9, shows that the FTSE-100 index also is less volatile than and lags the peak in property shares.

(4.1.4)   The above analysis suggests that, at least in the period 2000-2007, the overall trend of the IPD total property return index is better represented by the FTSE-100 index, than the FTSE-RED and individual property holding shares. The volatility of the FTSE-100 index appears to be greater than the IPD index, although the Study Group did not have time to make a numerical comparison of the respective volatility profiles. This is an item for further work, although it is difficult to argue why the FTSE-100
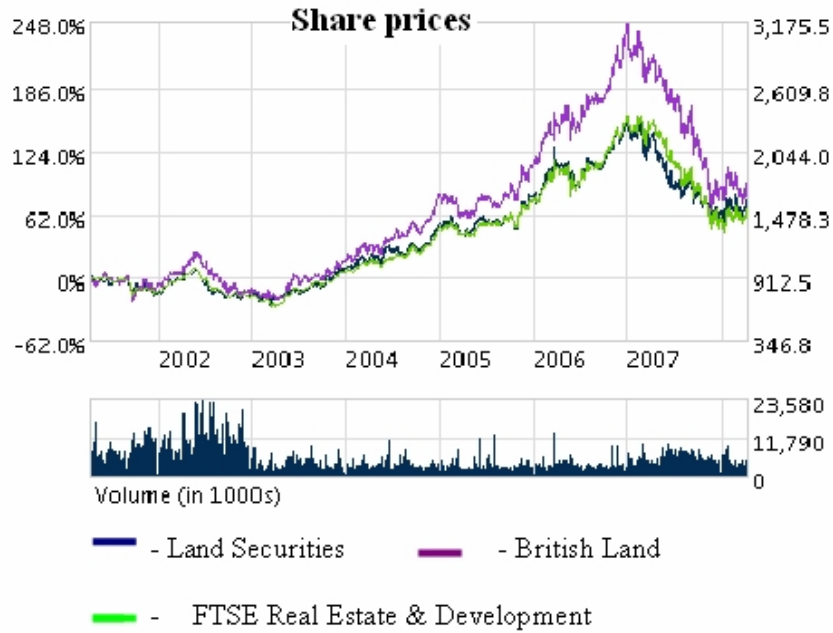
Figure 7: The property company share indices. Ref: Land Securities Group website.
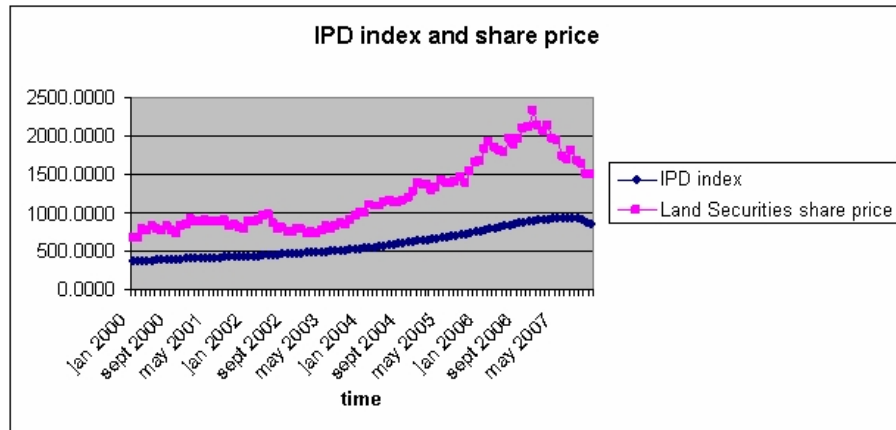
Figure 8: IPD total return index and Land Securities Group shares.

volatility should represent the underlying volatility in property and rent prices.

(4.1.5)   The above analysis of indices also shows that property equities are substantially more volatile than the FTSE-100 and IPD total return index.
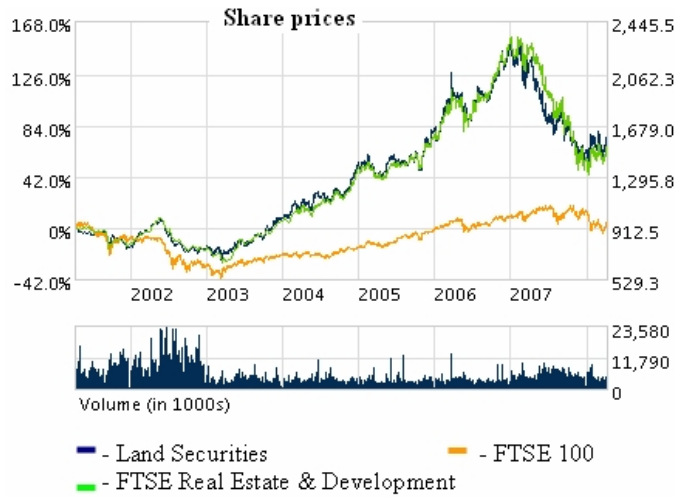
10

Figure 9: Property company share indices and the FTSE-100 index

4.2 Volatilities of the IPD property valuation indices and property equities

(4.2.1)   The Study Group computed the volatilities, defined in subsection 2.2, using 11 samples for variance estimation i.e. $m = 10$, for the IPD total return index (Figure 10) and for the shares of Land Security Group (LSG)
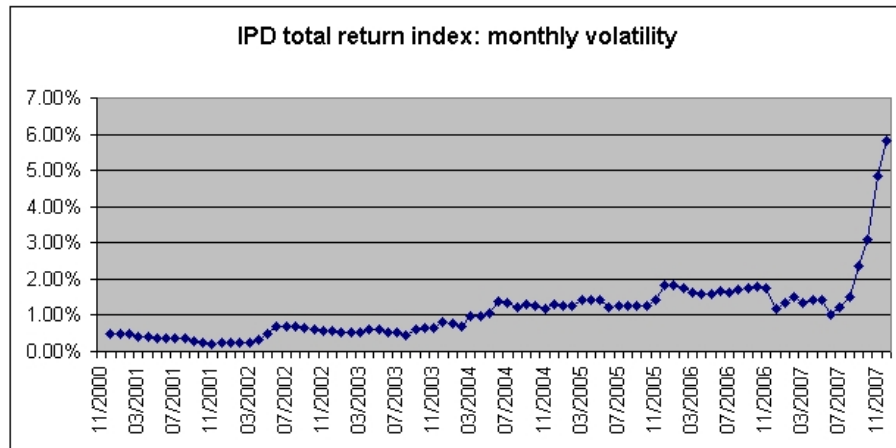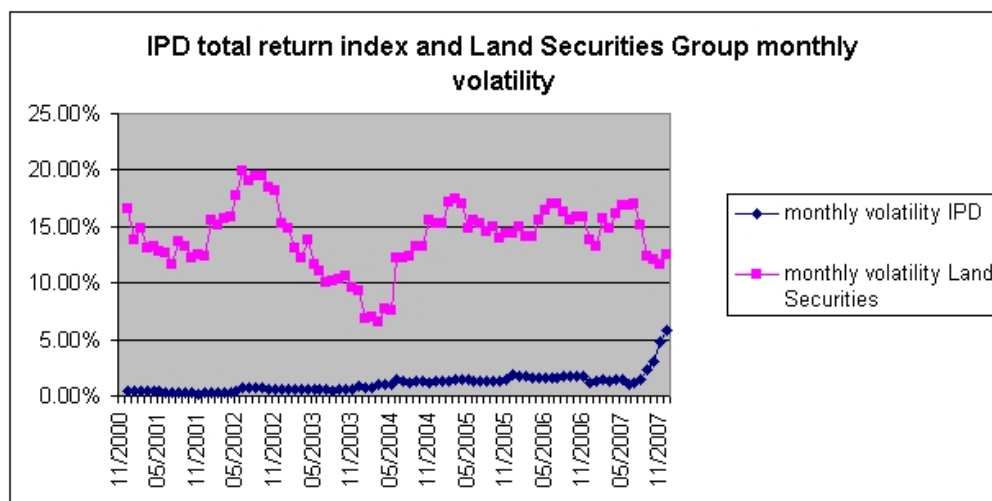

Figure 10: Property company share indices and the FTSE-100 index

(Figure 11).   The difference in volatilities is significant, with the LSG shares reflecting market sentiment, company activity and any underlying value of commercial properties in its portfolio. Further work would also include comparison of the volatilities for FTSE-RED and FTSE-100.

Figure 11: Property company share indices and the FTSE-100 index

4.3 Extracting the underlying volatility in real estate equities

(4.3.1)   Can overall effects in equities markets be removed from real estate equities to yield the bare real estate volatility? The Study Group considered the possibility of taking a basket of equities which are variously distant from real estate e.g. Vodafone, IT companies, energy, etc, in the sense that they are not strongly affected by real estate equities. A collection of real estate equities might then be added to the basket, and principal component analysis (PCA) carried out on the covariance matrix of the time series of the share indices in the basket to identify and extract bare real estate volatility.

(4.3.2)   A preliminary PCA with a small basket of equities was carried out by the Study Group, but the results were inconclusive. Further work, carried out after the Study Group and summarised in the Annex, has strongly undermined the idea of constructing a proxy property index from property company shares.

5   Treating property differently

5.1   Worst case factoring

(5.1.1)   Discussion repeatedly turned to the fact that property is different from equities and bonds, because it is so rarely transacted. Under such conditions it is not possible to obtain reliably volatilities of the transaction price of commercial property. This strongly suggests that such property

should be treated differently in solvency analysis and portfolio management decision-making.

(5.1.2) The Study Group considered the approach whereby the FSA would define non-stochastic worst case factors on property valuation indices for use in solvency analysis and portfolio management decision-making. This would avoid the need for volatilities to support a probabilistic analysis for the return on property holdings. Thus, the factored forecast value return on property holdings would be matched off against appropriate parts of the liabilities e.g. guaranteed parts, and the 2004 market consistent approach would then be applied to bonds and equities for which the use of covariances and probability distributions is more apt.

(5.1.3) If a factoring approach is taken to an index of commercial property such as the IPD property valuation index, there remains the challenge of choosing the factor F to apply to an index of property value returns. The Study Group considered whether modelling of the post-insolvency or 'rundown' phase would yield realistic values for F, since it is in this phase that the property holdings in question would actually be put on the market and transacted, so yielding transaction prices. The process might be treated in the manner of 'default', through prioritisation and pay-off of creditors. The Study Group considered that given the circumstances of insolvency and the relative ill liquidity of the property market, the factor F may well be considerably lower than 1 and in a range $x < F < 1$. Data on transactions in such circumstances would be required to establish a value and uncertainty for x. However, the FSA might define a maximum value of F that should be used in the absence of such data.

6 Conclusions and further work

6.1 Conclusions

(6.1.1) Much of the Study Group activity was spent in trying to understand the context of the challenge of estimating the volatility of property assets, through review - by no means comprehensive - of the real estate literature, and by data analysis of property indices and equities and some preliminary modelling of the IPD property index.

(6.1.2) Returning to the question posed in Section 1.2, we can provide comments in the light of the activities of the Study Group:

   (a) *What information do the surveyors estimates use? Are they based on commercial rents, or do they use information from property sales when available?*
   The Study Group did not investigate this question. However, Booth & Mercato (2004) notes that valuers have a tendency to use

historical transaction values which introduces bias, and also to use comparable sales which introduces smoothing of differences between properties. This work, particularly its Section 3, is worth further study.

(b) *Might other information such as returns on real estate investment trusts (REITs) be useful surrogates for estimating property portfolio volatilities.*

The Study Group investigated the indices and volatilities of transacted shares of two property companies (Land Securities Group and British Land) and the FTSE Real Estate and Development index. These indices had substantially higher trend amplitudes and volatility in the last 7 years than the IPD property total return index, which is valuation based. Equities in property companies represent market sentiments, activities and influences over a wider range of sectors, in addition to the underlying value of property. The direct adoption of such volatilities for the returns on REITs is hard to justify. The Study Group made a preliminary investigation of the use of principal component analysis to separate out the underlying trends in property transaction prices. However, subsequent investigation using this approach, which is summarised in the Annex, has strongly undermined the idea of constructing a proxy property index from property company shares.

(c) *Can a model for a sale price be obtained from using the IPD index, but then also multiplying by a random factor F at the time of sale, where F has a distribution over perhaps the interval from 0.8 to 1.2? If one made such a model, is there data available that could be used to validate it?*

The Study Group was sceptical about the use of a 'factored' valuation index for the estimation of sale price for the purpose of solvency analysis and real estate portfolio management, particularly if the factor is based on an analysis of property company equities. It considered that use of a factoring approach might be justifiable if F were estimated (or specified by the FSA) for the contexts of insolvency or default, in which case it may well be that $F < 1$.

(d) *Is there data available to find what other economic variables F is correlated with?*

Given the lack of transaction data for commercial property it would seem difficult to establish such a correlation. Correlates with economic variables might be established for property company equities, but their application to solvency analysis may be difficult to justify.

(e) *How does F change with the time horizon?*

Given a lack of transaction data, the Study Group did not address this question.

(f) *Should a model for property asset values look at the extent to which*

14

*similar properties sold at nearby times and locations can be used to
give a more specific measure of volatility?*
The Study Group noted the observations of Booth & Mercato (2004),
regarding the biasing and smoothing of valuations that can arise
when spatially and temporally separated transacted properties are
used as a basis for valuation. These effects may be minimised if such
properties are 'close' in space-time, but it seems that the constraint
may be too severe to get adequate statistics to estimate a volatility.

(g) *What can be inferred from the change in variance over different time
horizons, allowing for the extent to which there is serial correlation
that might have an impact? (Booth & Mercato 2004)*
The Study Group made cursory observations on the change of volatil-
ity over time for property company equities, but it did not make any
inferences regarding the time-dependence of volatility in estimates
of the transaction value of property holding portfolios.

6.2   Further work

(6.2.1)   The Study Group began a number of threads of investigation, which would
be interesting to pursue, but perhaps most useful would be the following:

(6.2.2)   Further work building on that reported in the Annex would be expected
confirm that that principal component analysis cannot usefully extract
estimates of the underlying volatility in the transaction value of property
from property company equities.

(6.2.3)   Studies of commercial property transactions in post-solvency or default
circumstances should be researched and/or carried out, with the aim of
revealing an F-factor on property valuation indices for use in solvency
analysis and property portfolio management.

References

BOOTH, P. & MERCATO, G. (2004) The measurement and modelling of commercial real estate performance. *B. A. J.* **10, I**, 5–61.

DIMSON, E. (1979) Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics* **7**, 197–226.

FOWLER, D. & RORKE, C. (1983) Risk measurement when shares are subject to infrequent trading (comment). *Journal of Financial Economics* **12**, 279–283.

IPD - INVESTMENT PROPERTY DATABANK (1985) http://ipdglobal.com.

MARTIN, R. F. (2005) The baby boom: predictability in house prices and interest rates. *Board of Governors of the Federal Reserve System (U.S.) in its series International Finance Discussion Papers* **847**.

ROLL, R. (1981) A possible explanation of the small firm effect. *The Journal of Finance* **36**, 879–888.

Annex: Follow-on analysis of property company equities

(1.0.1)   The analysis summarised in this Annex was carried out and reported by
Jeff Dewynne, after the Study Group. It investigates whether an under-
lying "property value" risk factor can be revealed in the behaviour of
property companies equities. This work is aimed at the question, posed in
section 4.3, of whether background effects in equities markets be removed
from real estate equities to yield the bare real estate volatility?

(1.0.2)   For share price data retrieved only back to about 2004/5, it was found not
possible to get sensible numbers for the correlation between the property
index (IPD) and property company share prices because the overlap in
monthly data between the index and the share prices was not large enough.
However, a principal component (PCA) analysis was carried out, and the
results were surprising.

(1.0.3)   In the following analysis, risk is the same as standard deviation of returns.
This is standard practice in VaR and most other areas of mathematical
finance. In particular, the risk is daily risk (that is, the standard deviation
of daily returns).

(1.0.4)   We may write the returns on a property company in the form

$$R_i(t) = \sum_{k=1}^{n} \alpha_{i,k} F_k(t) \tag{1}$$

where $R_i(t)$ is the return on company i's shares (over some fixed time
interval), $F_k(t)$ is a random variable associated with the k-th risk factor,
and $\alpha_{i,k}$ is a measure of the correlation between the return on company i
and the risk-factor k. So, assuming the risk-factors are uncorrelated,

$$covar(R_i, R_j) = \sum_{k=1}^{n} \alpha_{i,k} \alpha_{j,k} \sigma_k^2, \tag{2}$$

where σ_k is the standard deviation of risk factor k, *i.e.*, the magnitude of
the riskiness of risk factor k. Now, it might be conjectured that "property
value" is a risk factor that is common to most property companies and
that most property companies' risk would depend rather strongly on this.
So, in taking a sample of property companies, working out their covariance
matrix, then diagonalising it the following might be expected:

   (a) relatively large correlations between the returns (because of the com-
       mon "property value" risk factor)
   (b) an eigenvector with a large eigenvalue relative to the other eigen-
       vectors. This eigenvector, normalised so its components sum to one,

gives the fractions of wealth that should be invested in each individual share in order to construct a portfolio that tracks the largest risk-factor common to all the shares involved, and this would be expected to be identified with the "property value" risk factor in this case.

Since the covariance matrix is positive definite, symmetric, the eigenvectors are orthogonal and we can interpret them as independent risk-factors with the positive eigenvectors being the variance of the corresponding risk-factor. A measure of the total risk of the portfolio is then the sum of the eigenvalues, and the proportion of the total risk due to the i-th risk-factor may be represented by the ratio of the i-th eigenvalue to the sum of all eigenvalues.

(1.0.5)   As a control, the principal component analysis was done with a few sectors: banks, utilities, phone companies as well as property companies. The surprising result is that it worked as expected in all the cases examined EXCEPT the property companies. For property companies, the correlation matrix was quite close to the identity matrix, suggesting the absence of a common underlying risk factor. That strongly undermines the idea of constructing a proxy property index in this manner. It may be that a bad sample of companies was taken, but given the almost total lack of correlation found, it is difficult to see how taking a larger set of companies would help. It is also the case that many of the publicly listed property companies' shares are rather thinly traded (*i.e.*, their price does not change very often) and these were deliberately not used in the analysis.

(1.0.6)   Data and MATLAB files (with supporting notes) are held by Jeff Dewynne (Universities of Oxford and Wollongong). The files comprise data downloaded from Yahoo, stored in OpenOffice and Excel formats, and MATLAB scripts. The latter read in the data, perform the PCA analysis and plot graphs of returns and their auto-correlations.