YALE
INTERNATIONAL
CENTER
FOR FINANCE

Yale ICF Working Paper No. 06-04
May 2006

Simulating Real Estate in the Investment Portfolio: Model Uncertainty and Inflation Hedging

William N. Goetzmann, ***Harvard Business School and the Yale School of Management***

Eduardas Valaitis, ***American University***

Simulating Real Estate in the Investment Portfolio:
Model Uncertainty and Inflation Hedging

William N. Goetzmann, *Harvard Business School and the Yale School of Management*
Eduardas Valaitis, *American University*

March 8, 2006

Abstract: Real estate is regarded as an inflation hedge, however the autocorrelation of property return indices and the autocorrelation of changes in the CPI pose serious problems of inference. In this paper we address these problems in two ways. First, we use robust methods to test of changes in the relationship between property returns and inflation. Second, we perform simulations of sample investment portfolios using vector autoregressions to study the ability of commercial and residential housing to hedge inflation. Despite the relatively short sample period available, we find that property is likely to hedge inflation well.

JEL Classification R3

Acknowledgements: We thank the participants in the housing session of the Econometrics Society Meetings in 2006, Jeff Fisher for suggestions and help with data access, the MIT Real Estate Center workshop for helpful suggestions and Rick Peiser for his support. Funding was provided by the Real Estate Academic Initiative at Harvard University and the International Advisory Board.

Contact: William N. Goetzmann, Visiting Professor, Harvard Business School, Baker Library 279, Boston, MA 02163, wgoetzmann@hbs.edu.

I. Overview

There is considerable research suggesting that real estate is a good hedge against inflation, since property returns and inflation measures are historically correlated.[1] However, these correlations may be misleading when they are based on a limited historical experience. More generally, inter-temporal econometric estimation is ultimately subject to temporal constraints. This small-sample problem is even greater when observations are not serially independent, and when there is some predictable time variation in the expected value. Given the strong historical correlation between real estate and inflation, who would have predicted that the past five years would yield such robust real estate returns in a low inflation environment? Despite the reassuring historical correlations, is it possible to forecast future relationships between the two series' with reasonable accuracy?

Most widely-used measures of real estate performance extend back merely to the mid-1970's, capturing just a fraction of the economic dynamics of the 20th and early 21st centuries. Assumptions about the stationarity of asset class correlations based on such a short time series are questionable.[2] Thus, when the structure of the time-series model itself is conditional upon possibly unknown state variables, and a forecast of future relationships is vital, it is prudent to explicitly model not only what is known, but what is unknown.

In this paper we revisit the relationship between real estate returns and inflation in the context of a robust simulation that allows for serial and cross-sectional inter-temporal dependencies in asset return series' and related economic variables. We explicitly address the issue of meta-uncertainty through Monte-Carlo simulations. The results of these simulations suggest that, despite the evidence for historic structural shifts in the relationship between real estate and macroeconomic variables, commercial and residential property returns are predicted to remain significant inflation hedges.

More generally, our results suggest that the expected variation in the long horizon real estate returns is broad, due to the autocorrelation of returns and the meta-uncertainty about econometric models. This leads us to the conclusion that we are more confident about the long-run inflation-hedging capacity of real estate than about its long-run expected return.

II. Paper Structure

Section III discusses the background and previous research on the issue of real estate as an inflation hedge. Section IV describes the data and the test for structural changes in the

---

[1] See, for example, Fama and Schwert (1977), Hartzell, Heckman and Miles (1987), Brueggerman, Chen and Thibodeau (1984), Ibbotson and Siegel (1984), Rubens, Bond and Webb (1989)Goetzmann and Ibbotson (1990), Graff and Cashdan (1990), Fisher and Goetzmann (2005).

[2] C.f. Goetzmann, Rouwehorst and Li (2005) and Li (2004).

correlation between real estate and macroeconomic variables. This test requires the use of econometric adjustments to the standard errors to account for lagged dependencies. In the spirit of classic studies of spurious correlation, we use simulations under the null to estimate the probability of incorrect inference about the co-movement of real estate series' and inflation. We use these simulations to compare alternative procedures for adjusting the critical values for tests of shifts in correlation, and we estimate the size and power of these tests. Using this information, we construct a test for structural breaks in the correlation between real estate and other economic variables of interest, and we reject the null of stationary bivariate historical correlations between real estate and GDP changes. Evidence for the break motivates the use of sub-sample estimation.

In Section V, we examine the role of real estate in the asset allocation decision in which the investor confronts an inflation-indexed liability. We use a vector autoregression [VAR] model to simulate the inter-temporal dependency across variables, and then simulate long-term holding period returns for a range of different asset portfolios as well as inflation. We augment the VAR with other macro-economic time series to take into account the potential structural relationship between commercial property, housing, home construction, inflation and production. A major issue in the use of a VAR model for simulation is that it is conditioned upon one particular historical realization, which itself is subject to sample variation. If the simulation is to be used for long-term financial planning, it is important to consider the possibility of structural changes in the model, and allow for a range of outcomes not covered by the historical experience. We use sub-period sampling in a hierarchical simulation procedure to address this issue. The simulation model introduces meta-variation in the VAR model via random sampling of the error variance-covariance matrix from a Wishart distribution. This robust sampling is highly conservative as it produces significantly fatter than normal tails in the distribution of future portfolio outcomes. These tails allow us to examine the ability of real estate to hedge inflation under extreme conditions that might exceed historical experience. Section VI shows that under these conservative assumptions, commercial and housing real estate remain good hedges against inflation. Section VII concludes.

III. Background

The relationship between real estate returns and inflation has long interested the academics and the practitioners alike. From a theoretical perspective, it is not immediately evident that real estate is a good inflation hedge. For example, the value of a building that is leased at a fixed rent to a tenant who does not pay maintenance, taxes and utility costs is adversely affected by inflation shocks. On the other hand, a building leased net of maintenance, taxes and utilities, with contractual rent increases tied the CPI might represent an excellent inflation hedge.[3] This duality motivated Graff and Cashden (1990) to propose a decomposition of real estate returns into income and capital appreciation components with differing inflation exposures. For residential property, Sinai and Souleles (2003) argue that the demand for housing is driven by the demand for hedging rent inflation. Ultimately, the question of whether real estate is a good inflation

[3] See Graff and Cashdan (1990).

hedge is an empirical one, and depends upon the availability and quality of data for testing.

Fama and Schwert (1977) were the first to uncover strong evidence that residential property returns hedged inflation shocks and expected changes in inflation over the period 1953 through 1971. Their tests used the housing component of the CPI, constructed as a three month moving average of hedonically-adjusted U.S. housing price per square foot. They explicitly addressed the problematic time-series issues with the data – particularly the fact that the moving average in effect concealed a potentially higher correlation to monthly measures of inflation. They found that the housing return series was explained by lags of an interest-rate-based measure for unanticipated inflation and thus they concluded that real housing appreciation returns were empirically unrelated to the inflation rate. Several authors since have used the Fama and Schwert estimation framework, and its more general extensions.[4] For example, Linneman and Gyourko (1988) performed a comprehensive study of the relationship between every available time-series measure of real estate and inflation measures based on Fama and Schwert instruments for expected inflation, as well as measures based upon time-series forecasts of expected inflation. They found that commercial real estate is a good inflation hedge, residential properties slightly less so, and real estate investment trusts are perverse inflation hedges – behaving more like other publicly traded equities in this regard.[5]

A specific drawback to the use of correlation and regression estimates of the relationship between real estate and inflation is that model misspecification and non-stationarity in the variables can lead to type I error more frequently than expected.[6] More generally, long-term relationships between real assets and inflation may not be reliably detected through short-term econometric models, particularly when variables are characterized by sustained but ultimately transitory deviations in performance. These problems have motivated a co-integration approach to measuring the relationship between real estate and inflation. The logic being that, even if inflation and real estate returns are themselves non-stationary (or close to it), their linear combination might itself be stationary. Barkham and Ward (1996) and Ganesan and Chiang (1998) adopt this technology, reaching differing conclusions about real estate as an inflation hedge – the latter finding

[4] E.g. Ibbotson and Siegel (1984), Hartzell et al (1987), Sirmans and Sirmans (1987), Rubens et al (1989), Miles (1996), Hoesli (1994) and Newell (1996).
[5] We do not explore the REIT puzzle in this paper. Many authors have observed that REITS tend to be "overly" correlated to the stock market and not well correlated to other measures of real estate returns. Glascock, Lu and So (2002) argue that the perverse relationship between REITs and inflation shocks is explained by a joint correlation to other macroeconomic factors, in particular to monetary policy shocks.
[6] The history of research on spurious correlations and regressions is itself of interest. See Yule (1926), James and Orcutt (1948), Walker (1950), Granger and Newbold (1974), Phillips (1986, 1998), Aldrich (1995). More recently this topic has been revisited by Granger, Hyung and Jeon (2001), Ferson, Sarkissian and Simin (2003) both of which provide an excellent summary and simulations of the magnitude of the problems and the utility of proposed solutions.

that Hong Kong real estate is not effective as such a hedge. Chaudry, Meyer and Webb (1999) apply co-integration methods and find that inflation appears to play a role in the structural relationship between real estate and financial assets. The implication of these and related studies is that it appears useful to account for a wider set of macroeconomic variables and financial variables in modeling real estate returns, but they do not generally reverse previous conclusions about real estate and inflation.

IV. Data and Tests for Sample Correlation Breakdown

One of the goals of this paper is to construct a robust simulation technology based on VAR that includes commercial and housing real estate along with a few macroeconomic variables. The simulated series' then are used to forecast the distribution of future outcomes of portfolio allocation decisions, as well as for examining issues related to risk – in particular how variables co-move. As a foundation for the model, however, it is important to understand the characteristics of the data. The quarterly variables selected for VAR simulation and further analysis span the period from 1978 Q2 to 2004 Q2 and are provided in Table 1.

IV.1 Real Estate Series'

The most comprehensive database of commercial real estate properties is the NCREIF National Property Index. The National Council of Real Estate Investment Fiduciaries [NCREIF] has pooled property-level investment data among its membership for 27 years.[7] Its database consists of properties that are held by investment managers on behalf of tax-exempt investment funds. Members of NCREIF contribute data quarterly on individual properties, including the acquisition price, net operating income, capital expenditures and market value of each property. The values and cash flows of the properties in the database are reported on an unlevered basis. The index is constructed as a value-weighted average of the property appraisals (or prices) on an end-of-quarter basis, and then the percentage change from last quarter is calculated to construct the capital appreciation return to NCREIF Property Index. Similarly, the income flows from the properties are cap-weighted and averaged, and percentage returns reported.

When a property is sold, the sale price for the property is reported and that value is used in the index construction; however, most valuations are based upon appraisals. It is common practice among institutions to seek external appraisals for properties on a regular basis. Some properties are appraised annually, others at higher or lower frequencies. These external appraisals are not synchronized, in the sense that any given quarter may include valuations based upon recently appraised properties and properties appraised as much as five quarters ago. In between external appraisals, investment managers update the appraisal values internally. Ross and Zisler (1990) and Geltner (1993) show how this process effectively creates a moving-average process in the reported returns; given a reliable specification of the lag, it is possible to "unsmooth" the

[7] This description is adapted from Goetzmann and Fisher (2005).

index.[8] It is also possible to calculate returns using only external appraisals. Geltner and Goetzmann (2000) for example, omit internal appraisals from the calculation and construct a “repeated-measures” index from NCREIF data that also includes income returns. They find that this index does not differ dramatically at longer horizons than the widely-used appraisal-based NCREIF.

A key issue for the current paper is the question of how this updating occurs. If, for example, the firm increases or decreases the appraised value by the most recent quarterly change in the CPI, this could induce a spurious short-term co-movement between real estate returns and inflation. Even if the internal appraisals were not tied to inflation, the moving average component induced by infrequent external appraisal could significantly influence statistical tests about correlation with other variables

The Office of Federal Housing Enterprise Oversight [OFHEO] was created to monitor the agencies that insure and securitize home mortgages. A key concern of OFHEO is the calculation of meaningful capital requirements for its subject agencies. This necessitates a measure of the value of the firms’ collateral – i.e. the U.S. housing stock the mortgages of which are held by the agencies. OFHEO uses a valuation process based upon the repeat-sale regression developed by Case and Shiller (1987) with a bias adjustment proposed in Goetzmann (1992). As inputs to the repeat-sales regression, OFHEO constructs price relatives for houses that represent agencies’ collateral. These price relatives are based on actual home sales or professional appraisals when homes are re-financed. Since these events occur infrequently, the repeat-sales regression estimates the log-returns to an equal-weighted index of the agencies’ housing portfolio, by minimizing sums of squared errors between the change in index value over the interval between sales, and the actual price-relatives of the properties. An adjustment is required to estimate the return index, as opposed to the log-index, due to Jensen’s inequality.

Although it is a best, unbiased estimator, the repeat-sales regression has some undesirable characteristics. It is not chosen to most efficiently estimate cross-correlations to other series,’ for example. The estimation procedure is also known to introduce negative serial correlation in returns when the data are sparse and to smooth returns when data are dense. This latter effect can simply be thought of as error in the ability of the regression to completely identify whether a shock occurred in one period or a neighboring period.[9] These two features of the housing indices can both manifest themselves in the same index, since data may be sparse in the early intervals and dense in the middle intervals. Thus, although it represents the most comprehensive measure of housing values available, and it is used in a regulatory process to control agency risk exposure, it has time-series features that adversely affect inference about its co-movement with other variables, and potentially influence long-term forecasts of housing returns.

[8] Cf. Ross and Zisler (1990) and Geltner (1993)

[9] Goetzmann (1991) for simulations of the reversal. Geltner (1993) for discussion of persistence.

Other variables used in our VAR analysis may have peculiarities that affect inference, but these are more generally understood outside of the real estate literature and in the interest of space will not be addressed. The variables included in our VAR model are: percentage change in completed housing units nationwide, percentage change in the consumer price index, percent return on U.S. long-term government bonds (from Ibbotson associates [IA]), percentage change in U.S. gross national product, percentage change in the S&P 500 Index, and the percentage change in the national unemployment level.

IV.2 Sub-Period Correlation Differences

Meaningful forecasts from a VAR model rely on the assumption that the underlying sample correlation structure is constant. This is a testable proposition. To address it, we split the short available sample of 106 quarterly returns into two equal 53-quarter periods and compute correlations between the pairs of variables; Table 2 reports the arithmetic difference between the sample correlation in the first period and second period $\rho_{XY,1..53} - \rho_{XY,54..106}$. These differences suggest plausible structural shifts in the relationship between GNP and NCREIF returns, and inflation and NCREIF returns: the differences are .63 and .40 respectively.

This, in turn, suggests that a full-sample VAR model might be potentially miss-specified, as the correlation structure might not be constant. Indeed, other correlation differences are also extreme: for example, S&P 500 and government bonds have a difference of .55. Thus, there are two questions that need to be answered before a VAR model can be applied. First, how can we judge whether the computed correlation differences are significant, and second, how should we split our sample to ensure the most stable correlation structure within each sub-period.

The autocorrelation or non-stationarity of variables is a major problem when modeling the distribution of the returns in a VAR when the correlation structure itself is potentially non-constant. It is well known that two autocorrelated variables may appear to be spuriously positively or negatively correlated and this makes a general test for a change in correlation between any two time series difficult to construct. Table 3 shows that when AR(1) model is fit to the variables considered in the study, all of them, except for stocks and bonds, are autocorrelated. Housing returns have a first lag autocorrelation coefficient ϕ_1 of .92, while NCREIF's is ϕ_1 = .68. When the Akaike Information Criterion (AIC) is used to determine the number of optimal lags k in the AR, some variables like inflation have k = 12. Even then, ϕ_1 remains high for the housing construction and housing returns. Thus, the pairwise correlations between real estate series' and other variables are likely to have a spread-out distribution. In Appendix 1 we develop a test of the null hypothesis of no change in the pair-wise correlations of time-series variables in two disjoint sub-periods; we then use this test to assess the presence of structural breaks in sample correlation.

IV.3 Test for Structural Change in Sample Correlation

Using the test presented in Appendix I, we calculated autocorrelation-corrected *t*-statistics for the change in pairwise correlations across variables that we will employ in our VAR. A line search is then used to identify break-points that best separate the sample data according to different correlation regimes. Specifically, in step one, we start with two sub-periods with observations indexed 1..30 and 31..106; for each pair of variables in our data we calculate the difference in correlations between the sub-periods and find the standard deviation for the correlation difference using our corrected formula for the standard error. We compute and record the *t*-value for each difference. In steps i=2..36: we repeat Step 1 one but take sub-periods 1..(30+i) and (31+i)..106 instead. Note that sub-sample size is an issue – tests for correlation breaks have less power near the beginning or the end of the series'. Thus we arbitrarily chose to have at least 30 observations in the smallest sub-period, hence the selection of values for *i*. Table 4 reports the results for the pairwise correlation difference test between commercial real estate returns and other variables.

The evolution of the *t*-statistic in Table 4 is mostly smooth. This implies that the estimation of the standard deviation of the difference between the sample correlations is fairly precise even though we might have as few as 30 observations. Only two variables exhibit *t*-values that are higher in absolute value than 2: split points 52 through 59 produce significant differences in correlation between NCREIF returns and GNP, and split points larger than 65 do the same for NCREIF and HPI. That said, it might seem disconcerting that the t-value for the correlation difference between national commercial real estate and GNP drops dramatically from 3 to .4 when one goes from split point 50 to 60. This corresponds to the dramatic crash in property values beginning in 1990, and a known regulatory change with respect to bank lending on real estate. Thus, the tests may in fact capture an economic shift in macroeconomic relationships. Figure 1 shows GNP and NCREIF times series' with a vertical dashed line representing the split point. As one can see, right after that split point, there is a large positive GNP change, while NCREIF returns drop a lot. In fact, after the 59th observation (which corresponds to the Q4 in 1992) the time series no longer move together; hence, the sizeable drop in the t-value from 3 to .4 is reasonable. Interestingly, this period also marks the beginning of the emergence of Real Estate Investment Trusts [REITs] as major owners of commercial investment property in the U.S.

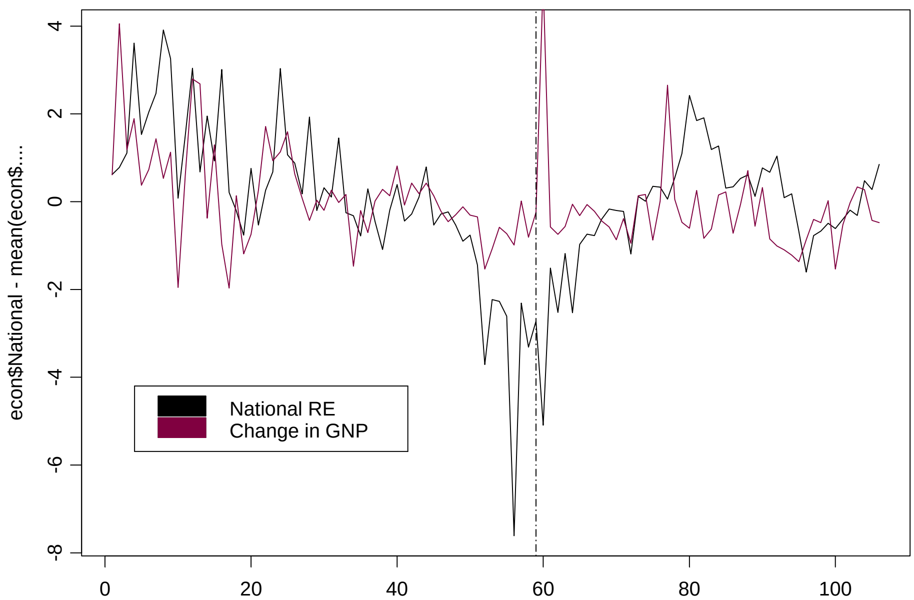

Figure 1: Time series of NCREIF returns and change in GNP with the optimal split point maximizing the significance of the difference between the sub-period correlations.

Note that the difference between sub-period correlations ranges between 1.3 and 1.7 for inflation and NCREIF returns, which suggest that there might be a benefit to dividing the estimation period into *any* two sub-periods that separate the beginning and end of the series'. For example, a division 1..40 and 67..106, yields a *t*-value of 2.3 for sub-period correlation difference between NCREIF and inflation. Of course, one must be mindful of data-mining in the interpretation of this *t*-value. A standard test for a single structural change via a line search over a *t*-statistic (or its equivalent) does not lead to rejection of the null that the correlation between NCREIF returns and inflation has remained constant over the period of study. Thus, despite large measured correlation differences, the standard errors around the sample difference statistic are equally large as a result of the autocorrelation of the two series'.

In sum, tests for changes in correlation suggest that even fairly large historical shifts in correlations sometimes will not lead to a rejection of the null, given the short historical record. This in turn suggests that it is important to introduce such uncertainty into the simulation model in addition to utilizing sub-samples. In basic terms, if the correlation structure of the variables used to represent assets and liabilities in an investment forecast is estimated with imprecision, it is important to introduce such parameter uncertainty into the VAR model if the goal of the simulation is to replicate the range of future joint potential realization of the variable set.

V. VAR Modeling and Simulations

The VAR is widely used in economics to model the evolution of the economy, and has been particularly useful in analyzing the effects of policy choice (Bernanke, Boivin and Eliasz, 2004). In financial economics, it is frequently employed to capture the inter-temporal behavior of asset returns. Campbell and Viciera (2002), for example, provide and excellent overview of the application of the VAR technology to solving inter-temporal portfolio decision problems. The VAR structure can also be used as the basis for simulating the effects of portfolio choice in the presence of a variety of factors, including intermediate cash flows, tactical rebalancing decisions and non-normal error distributions. Although the literature about the VAR is considerable, there are two broad challenges to its application to modeling financial asset returns. The first, as we discussed above, is the problem of conditioning upon a particular historical realization which may not be representative of the future. The second is that financial variables are typically assumed to have an equilibrium relationship with each other. Stocks are presumed to have expected returns higher than bonds, due to the equity risk premium, expected future short-term interest rates should be consistent with the current yield curve, stock prices and dividends should not stray too far from one another. These and other structural relationships implied by economic theory – and to some extent tested and verified empirically in specifications more complex that a VAR – lead to increasingly complex models of evolution of financial asset returns. Although it is important to recognize that a simple VAR model may not capture the longer-term co-integration of two processes, in this paper, since we are focused on the question of model error more generally, we maintain a fairly simple specification. We include long-term bond, stocks and two measures of real estate returns. We augment these with related macroeconomic variables: home construction, inflation[10], unemployment and change in GNP. With these, we obtain marginal distributions of asset returns that are mostly consistent with historical realizations. The following describes the procedure in more detail.

We do not separate the nominal inflation into the expected and shock components because the median of the expected inflation provided by the Michigan University housing survey is the least biased and least MSE when predicting nominal inflation and it confirms the adaptive expectations in predicting inflation.

[10] We do not separate nominal inflation into expected and shock components in the subsequent analysis. When we performed the separation, we used the median of the expected inflation provided by the Michigan University housing survey as the measure of the expected inflation, as it is the least biased and the least MSE predictor of nominal inflation (see Baghestani, 1992, Mehra, 2002, and Valaitis, 2006). The shock component produced by using this measure confirms the adaptive expectations theory, as the expected inflation measure depends heavily on the previous realizations of the nominal inflation series. This produces a high in-sample correlation between the expected and shock components, which in turn, influences the VAR model loadings. Therefore, when we simulated the full sample VAR model with the nominal inflation replaced with the expected and shock components, there were only minor differences in the simulated return correlations between the two components and other variables.

V.1 Full Sample Simulation

In this section we describe the details of our hierarchical sampling procedure. To capture the meta-uncertainty in the VAR estimates, we sample the covariance matrix from a Wishart distribution. That is:

Step 1: We calculate the VAR matrix of coefficients $\hat{\beta}$ for the aforementioned variables. The VAR is a single lag specification, and the beginning values are the last in the historical time-series'. Application of the AIC criterion suggests that lag one captures most of the predictability in the time-series, however it is important to note that previous authors have found that, even in the presence of modest predictability at longer lags, overfitting may lead to spurious predictability and correlation. To address this, we simulated the system under a null of no lagged link between real estate and inflation. The results of this test are reported in Appendix 2. A simulation under the null of no relationship between commercial real estate and inflation leads to a (statistically) significant drop in correlation from .7 to .21. The remaining positive correlation is likely due to the influence of the other autocorrelated macroeconomic variables. This suggests that we can reject the null of a spurious economic relationship between real estate and inflation.

Next, since the estimation of $\hat{\beta}$ is strongly conditioned on the sample covariance matrix $\hat{\Sigma}$, we must address the question of whether the uncertainty in $\hat{\Sigma}$ affects $\hat{\beta}$.

Step 2: We use a Wishart distribution to draw *k* updates for Σ from

$$\Sigma^{-1} \sim Wishart\left(\left\{n\hat{\Sigma}\right\}^{-1}, n-m\right),$$

where *n* is the sample size and *m* is the number of estimated parameters.

Step 3: For each Σ obtained in *Step 2* we update the $\hat{\beta}$ by drawing *j* random samples from:

$$\beta | \Sigma = N\left(\hat{\beta}, \Sigma \otimes \left\{X' X\right\}^{-1}\right).$$

Step 4: We then use each updated β in Step 3 to generate *i* paths of *t* quarterly returns for each of the eight variables in VAR.

For the full sample simulation, we chose k = 1,000, as most variation in simulated returns will be due to uncertainty in estimating the sample covariance matrix; j = 10 and i = 5. We then first examine the distributions of the simulated annualized returns for the eight variables and certain portfolios achieved over the holding period t = 40 (10 years). Thus, each variable (and portfolio) has 50,000 annualized returns produced. Panel A in table 5 provides summary statistics for the distributions of the eight variables used in our VAR model. Notice that the null of Kolmogorov-Smirnoff test for normality is rejected for all variables, and the Q-Q plots reported in Figure 2 imply that the tails of the distributions of the variables are quite heavy. It is well known that non-constant variance will generate fat-tailed distributions.

The historical values for most variables, particularly the macroeconomic ones, are relatively close to their simulated means. However, the relative means for certain variables are potentially contradictory to some economic equilibrium expectations and to history. For example, commercial real estate has historically exceeded inflation by 400 basis points per year, not 600 bps as suggested by the simulation – this may be due to the high autocorrelation of the series and conditioning upon recent real estate performance, coupled with the low recent level of inflation. On the other hand, long term bonds have historically had a return equal to commercial real estate, which does not make sense given that most real estate is levered at rates higher than the long term bond yield. The simulation forecasts a future premium of real estate returns over bonds. Finally, the mean of the simulated S&P 500 returns is about 240 bps lower than historical values, suggesting lower forecasted future equity returns.

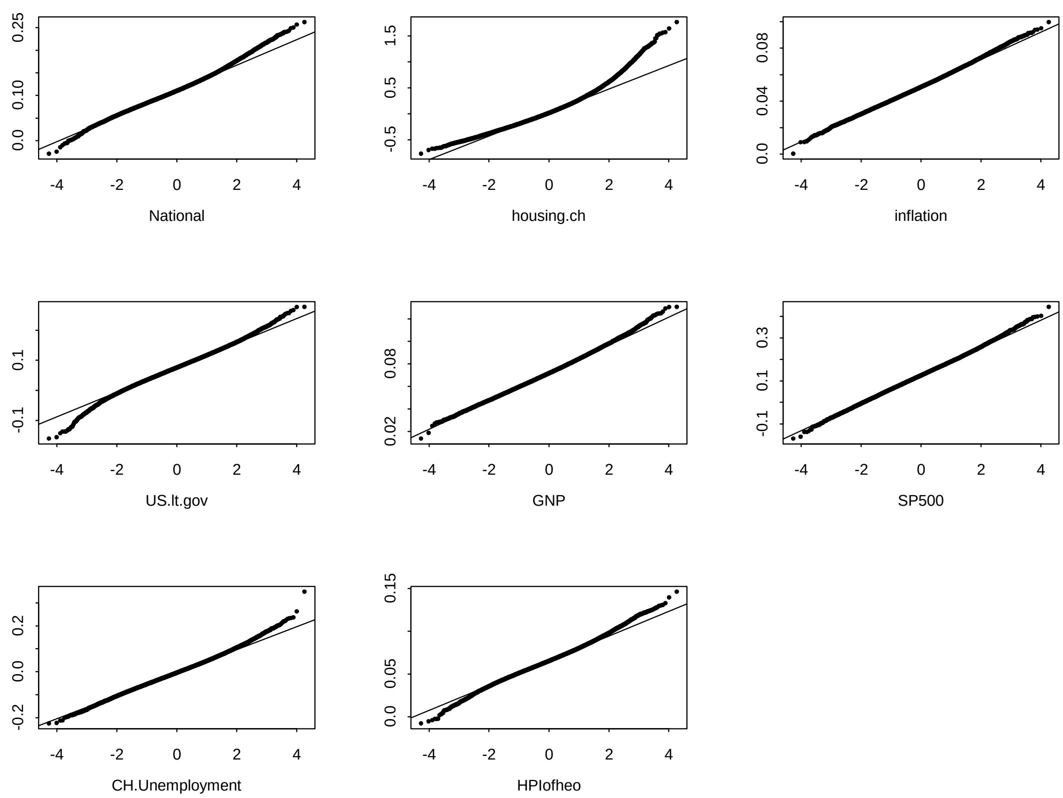

Figure 2: Q-Q plots for annualized return distributions for the eight variables used in the VAR.

The first question we ask is what happens if one invests in only one of the following four asset classes: bonds, stocks, commercial real estate, or housing. Figure 3 exhibits a random sample of 5,000 returns from our simulated 50,000 annualized returns for each variable class as plotted against inflation in the same paths.

Housing and commercial real estate returns appeared to be excellent hedges against inflation in our full-sample simulations; while stocks had a very low positive association with inflation, and U.S. government bonds were strongly negatively associated with inflation. Hence, combining bonds, stocks and one of the real estate asset classes might not offer protection from inflation.

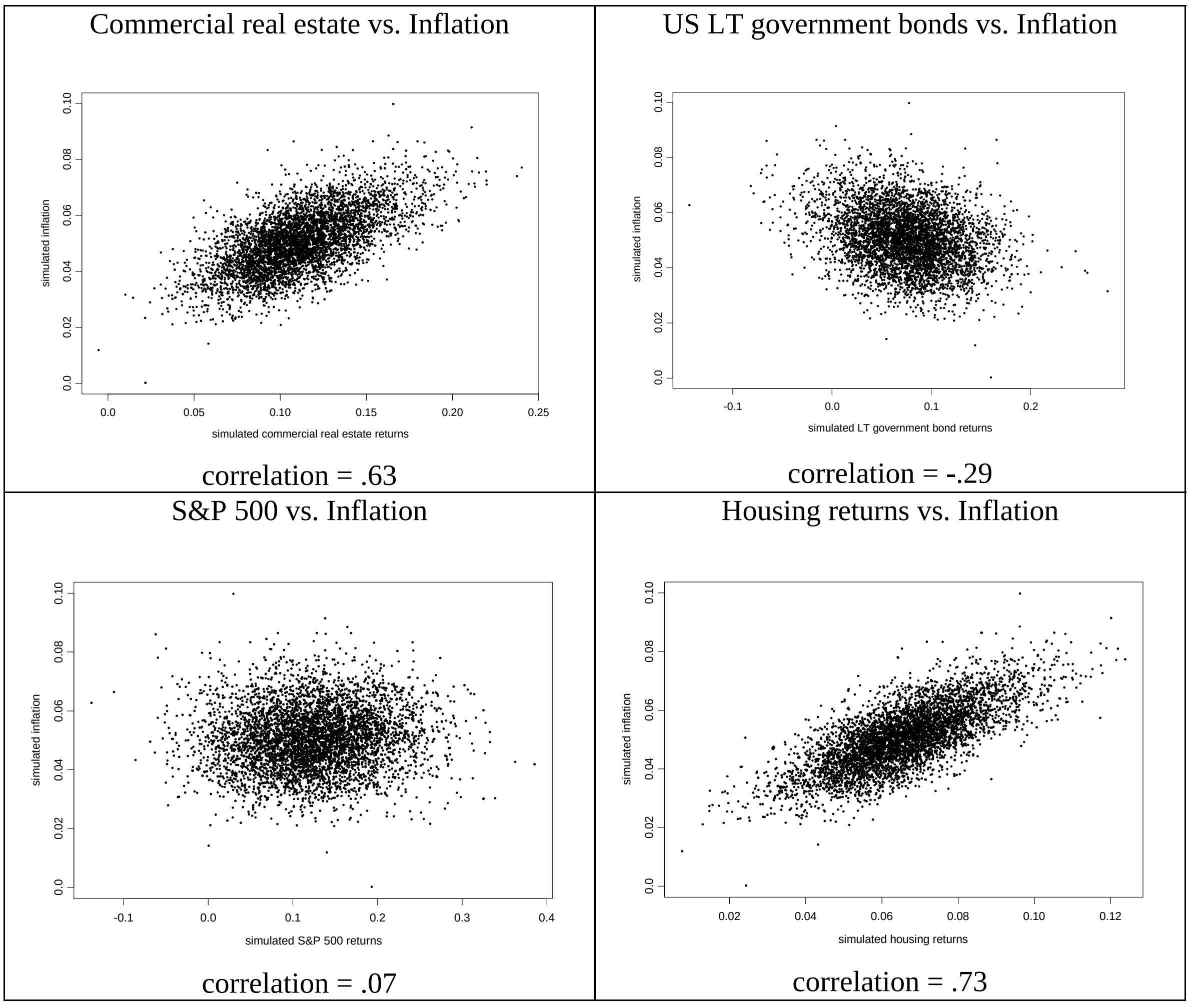

Figure 3: Relationship between inflation and the four asset classes in the VAR produced time series.

V.2 Portfolio Performance for Full Sample

We next construct eight different portfolios and examine how they perform along each simulated path. Every portfolio comprises of assets in 3 classes: stocks, U.S. government bonds and real estate (commercial or housing):

C All	1/3 in stocks, 1/3 in bonds and 1/3 in commercial real estate
C Stocks	1/2 in stocks, 1/4 in bonds and 1/4 in commercial real estate
C Bonds	1/4 in stocks, 1/2 in bonds and 1/4 in commercial real estate
C CommRE	1/4 in stocks, 1/4 in bonds and 1/2 in commercial real estate
H All	1/3 in stocks, 1/3 in bonds and 1/3 in housing real estate
H Stocks	1/2 in stocks, 1/4 in bonds and 1/4 in housing real estate
H Bonds	1/4 in stocks, 1/2 in bonds and 1/4 in housing real estate
H HousingRE	1/4 in stocks, 1/4 in bonds and 1/2 in housing real estate

Table 6 presents the annualized summary statistics for the portfolios. Not surprisingly the portfolio with the highest average annual return for the 40 quarter-period contains commercial real estate, and is dominated by stocks. All the portfolios with housing underperform the respective portfolios with commercial real estate. Table 6 also provides the correlations of the portfolios with the simulated inflation along each path. It appears that the best protection against inflation is offered by a portfolio that contains 50% commercial real estate, 25% bonds and 25% stocks. This portfolio returned an average of 11.3% a year and had a .34 correlation with inflation. The correlation of .34 can be achieved when investing in housing, stocks, and bonds by raising the proportion of housing to 69% and diving the remaining 31% between stocks and bonds evenly. But such a portfolio is forecast to return 8.26% annually – 300 bps lower than the portfolio dominated by the commercial real estate. While a portfolio dominated by housing does indeed appear to provide an inflation hedge, it does so with a lower expected return. Because the NCREIF index is comprised of properties held by tax exempt institutions – particularly pension funds – individuals are only indirectly benefited by this investment vehicle as an inflation hedge.

V.3 Sub-Sample Analysis

Recall from our earlier tests for structural change in the pairwise correlations of the VAR variables that there were strong indications of a structural change roughly in the middle of the sample period. One concern is that we are currently in a regime that may not continue. To address this issue, we can split the data into two sub-periods and re-run the simulations. This has two effects. The first is to widen the Wishart distribution for the sample covariance matrix, since the variation it introduces is based upon the degrees of freedom reduced by the smaller sub-sample size. The second is that it conditions upon a different empirical estimate of the VAR coefficients.

V.3.a Sample Period 1977 through 1991

The first sample period ranges from quarter 1 to 58, and the second from 59 to 106. Summary statistics for the simulated VAR variables in the first sample period are presented in panel B of Table 5. Both real estate indices are below their historical mean levels for the first period, and other variables – notably changes in housing construction and change in unemployment – differ significantly from historical averages. The first sub-sample can also be used to see if the variables distributions obtained by simulations contain the actual realizations of the eight variables in sub-period indexed by 59:106. In other words, how does our VAR technology perform out of sample, under the assumption of stationarity?

Figure 4 shows that the out-of-sample means of most variables fall well within the standard variation of the simulations. Inflation changed a lot from the first sub-period to the second, consistent with the notion that the lower inflation in that period was a shock. However, it was still captured in the left tail of our simulated distribution, meaning that the VAR allows for such shocks and shifts appropriately. What is the relationship between simulated inflation and the 4 asset class returns for sub-sample 1:58? NCREIF and OFHEO had correlations of .66 and .70 respectively. Stocks had a slightly positive correlation (.10) and bonds a significantly negative correlation (-.62) with NCREIF returns. It is important to point out that we do not model expected and unexpected inflation separately in the VAR, so our simulations are silent on the issue of whether the correlations are driven by inflation shocks or expectations. Indeed, by omitting short term rates from the simulation, we are removing a key instrument that has been used to differentiate the two. On the other hand, the VAR effectively models inflation expectations via an AR(1) process, along with lagged cross-correlations with other variables. As a result of the offsetting positive and negative inflation correlation with real estate and bonds, most of the sample portfolios had low or negative correlations to inflation in the first sample period. Panel B in Table 6 wxhibits these results.

V.3.b Sample period 2: 1992 through 2004

In the interest of space we do not report all the details of the simulations based on the second sample period. It is instructive, however, to look at the distributions. The figures suggest that the simulation allows for some highly skewed and kurtotic distributions. This is particularly interesting, as the return generating process is a simple VAR with sample variation introduced by a Wishart. Take, for example, the S&P 500 return outcomes in Figure 5. They display surprising deviations from log-normality, suggesting that the range of future outcomes, even for stock market investing is strongly affected by the introduction of the uncertainty parameter.

The pairwise correlations between inflation and the four asset classes of interest are lower for this simulated period: NCREIF and OFHEO have correlations of .29 and .13 respectively; stocks and bonds have correlations of .14 and -.19 respectively. Thus, while the commercial real estate is still the best hedge against inflation, its positive association with inflation has decreased from .66 in the sub-sample to half that value in the second sub-period. The decrease in positive association with inflation is even more pronounced for the housing returns: down to .13 from .70. The correlations of the sample portfolios to inflation, reported in Table 6 are likewise significantly reduced.

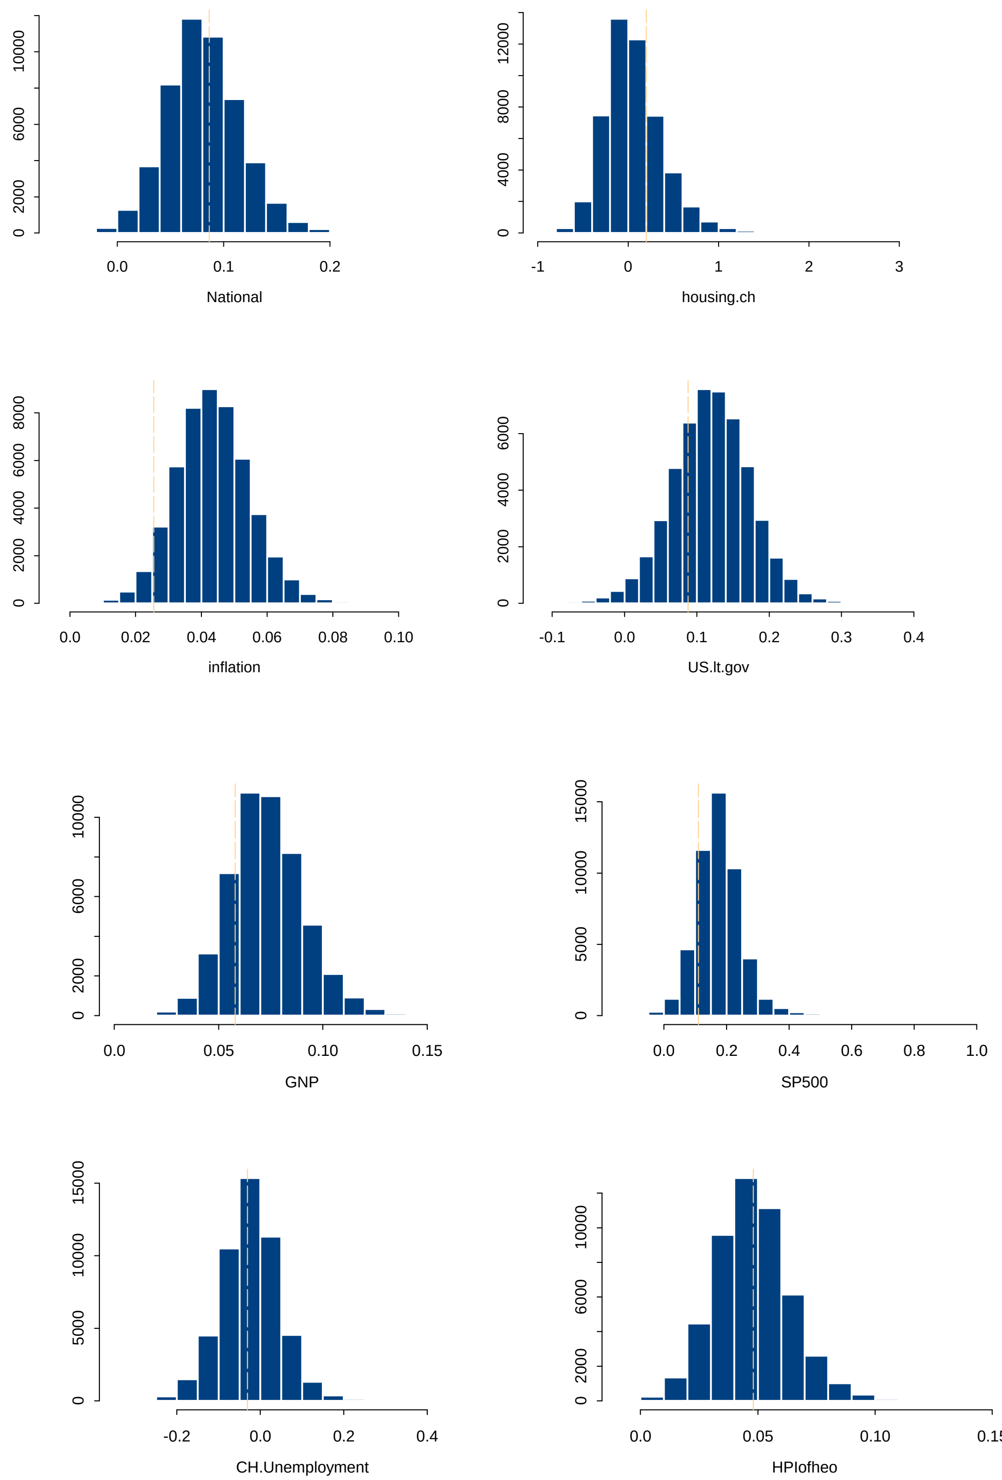

Figure 4: Simulated distributions using sub-sample 1977:1991. The vertical line corresponds to the annualized return produced by each variable in the actual period 1992:2004.

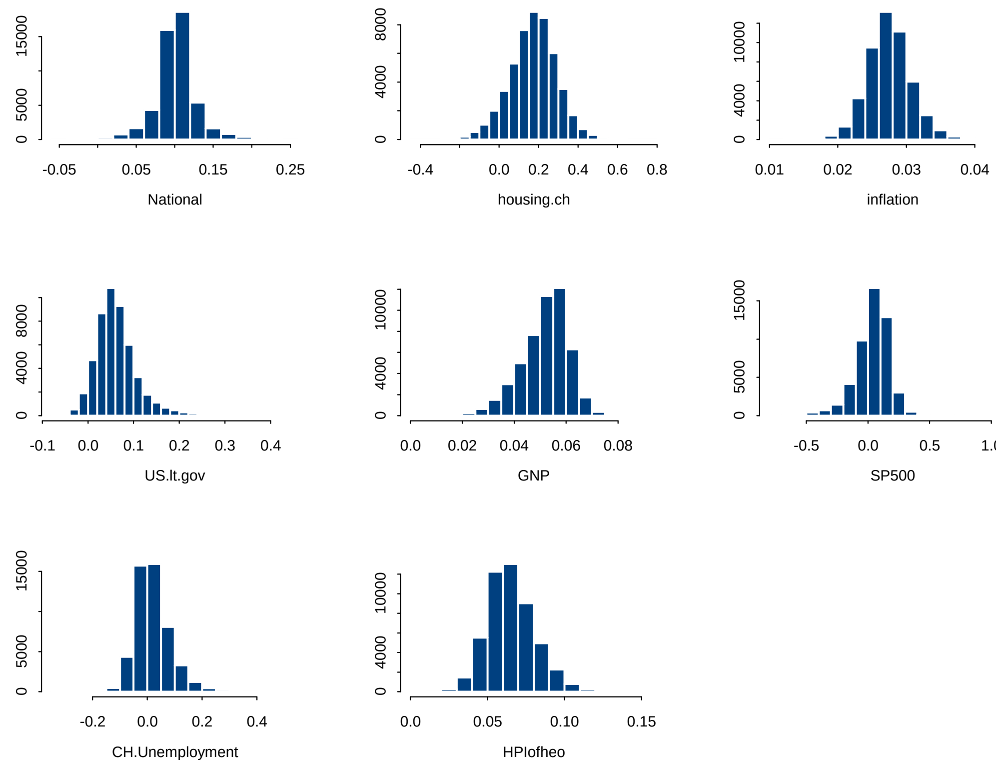

Figure 5: Simulated distributions using sub-sample 59:106.

VI. Shortfall Analysis

In this section we examine the frequency of portfolio shortfall with respect to inflation. Figures 6 through 8 present histograms of the inflation-adjusted annual returns for the sample portfolios consisting of three assets (stocks, bonds, and commercial real estate or housing). The number reported in Figure 6 is the percent of these inflation-adjusted portfolios that fall below 0. The investment horizon is ten years. In general, the probability of a shortfall in real terms for any portfolio is relatively low. For example, when a portfolio is dominated by commercial real estate (50% of money invested in NCREIF) only 2 in 1000 paths have an annualized negative real return. When housing is included instead of commercial real estate, the shortfall probability is larger. When a portfolio is dominated by bonds and includes housing ("H Bonds"), almost 9.5% of all inflation-adjusted portfolios returned less than 0 annually. It is also interesting to note how much tighter the real estate-dominated portfolio distributions are than other portfolios – this is driven by the fact that real estate is correlated to inflation. There are two ways to generate positive real returns in a portfolio. One is to select assets with high absolute returns and no correlation to inflation and the other is to select assets with high correlation to inflation and modest absolute returns. We see from the distributions of real portfolio returns dominated by real estate the effects of the latter.

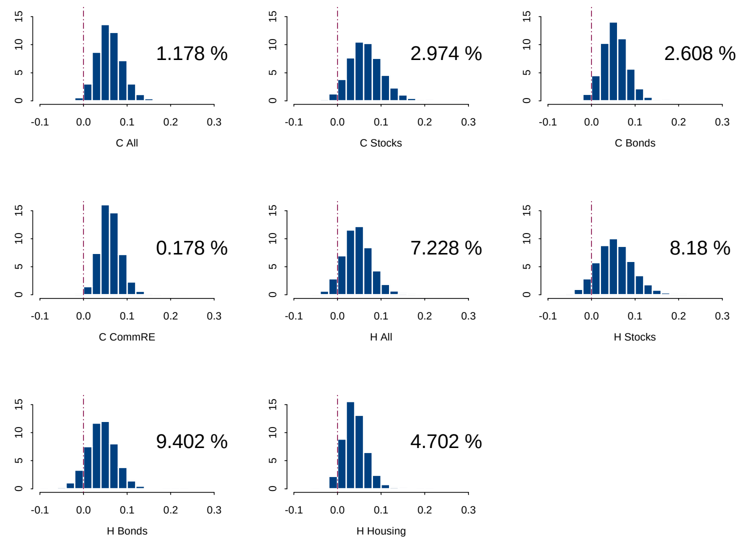

Figure 6: Full sample inflation –adjusted ten year simulated portfolio returns

The sub-period shortfall results indicate a broader spread, consistent with the introduction of greater parameter uncertainty. The more recent sub-period suggests a much greater likelihood of shortfall in real terms than does the first sub-period (see Figures 7 and 8). Thus, if the structure change in the early 1990's was genuine, it suggests that the capacity going forward to hedge inflation over the longer term is more difficult, at least given the few basic instruments we include in this study. It is worth noting that studies of U.S. housing sub-markets also suggest non-trivial probabilities of real loss in home values over long time-periods.[11]

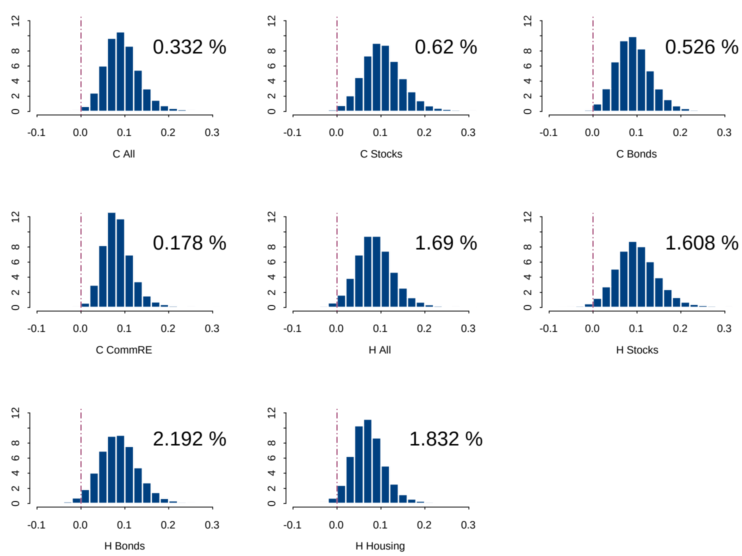

Figure 7: Portfolio shortfalls in sub-period 1977-1991.

[11] See, for example, Goetzmann and Spiegel (2002) and Caplin et al. (2003).

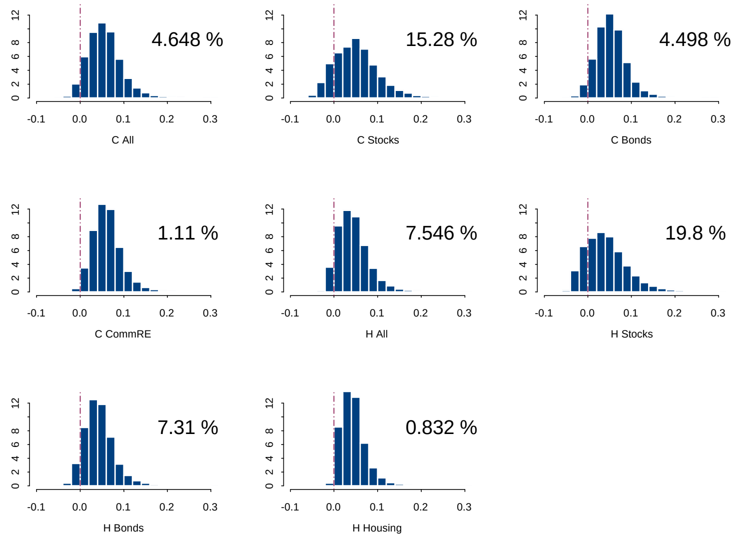

Figure 8: Portfolio shortfalls in sub-period 1992-2004.

Table 7 summarizes the shortfall probabilities for the sample portfolios over three different time horizons (5, 10 and 20 years), using both the whole period and sub-periods to estimate the VAR and simulate uncertainty. The shortfall probabilities decrease with the longer investment horizon, reflecting the positive real expected returns for all assets. As previous simulations suggest, the commercial real estate-dominated portfolios are the best at reducing shortfall probabilities. The interesting variation is across the sub-periods, with the later sub-period suggesting a significant probability of shortfall at all but the longest horizon.

VII. Conclusion

One of the primary goals of many institutional investors is the preservation of capital in real terms, and for individual investors it is building a portfolio that keeps up with the cost of living. The challenge is that some of the most promising assets from this perspective are also those whose correlation to inflation is inherently difficult to forecast with precision. Indeed, is should not be surprising that any asset that correlates well to an autocorrelated variable is itself highly autocorrelated. While past researchers have addressed this problem by decomposing inflation into an expected and an unexpected component, with expected inflation presumably hedgeable through the fixed income markets, in the final analysis, investors care about total realized inflation. In this paper we examine the relationship between commercial and residential property returns through a VAR model that includes macroeconomic variables as well as stocks, bonds and real estate.

Much of the paper is focused on understanding the potential weakness of reliance on historical correlations between real estate and inflation as the basis for long term financial planning. In the first part of the paper we developed a method for testing for structural change in correlations between two autocorrelated variables and examined the performance of this test under three different specifications driving the correlation. We found that our proposed test worked well, however in general the power of a test in the presence of serial correlation is low. We used this test as the motivation for considering two sub-periods as potentially representative of different regimes for real estate and the macro-economy. Our VAR simulation was performed using estimates from the whole sample period as well as estimates from the two relevant sub-periods. We also introduced parameter uncertainty into the simulations through the use of a meta-sampling of the covariance matrix for the system. The effect of this hierarchical sampling was to dramatically fatten the tails of the distribution of simulated outcomes. This effect was more pronounced when the sample period used to estimate the covariance matrix was short. We found fairly convincing evidence that real estate is a relatively good asset to use as an inflation hedge, particularly over the long term. However the sub-period analysis should motivate caution – the data since 1992 suggest a somewhat weaker connection between property returns and inflation, and this may in fact be a structural change. Even under this pessimistic scenario, real estate is a better inflation hedge than stocks or a long term bond portfolio.

Our study is far from definitive. There are many other asset classes we omit. There are other real estate vehicles we could include. A much more sophisticated model of the economy might be warranted – one in which equilibrium relationships among inflation, bonds, stocks and real estate are assumed and modeled. Expected and unexpected inflation could be separately considered. Nevertheless, the analysis we provide in this paper should offer some guidance to investors interested in preservation of real spending power over the long term.

Appendix I. A Test for Change in Correlation in the Presence of Autocorrelation

I.a Background

The correlation coefficient for variables that are independent and not autocorrelated is distributed normally with mean 0 and standard deviation $T^{-1/2}$:

$$\rho_{XY} \sim N\left(0, T^{-1/2}\right) \qquad \textit{(Equation 1)}$$

where T is the number of periods in a time series. Unfortunately, this is not true when both variables are autocorrelated. To illustrate this, in the spirit of Yule (1947), we generate two highly autocorrelated ($\phi_x = \phi_y = .9$) AR(1) processes, X and Y, independently from each other with T = 100 observations. Note that these are close to a unit root. The left panel in Figure I.1 shows the distribution of the correlation coefficient between X and Y. One can see that it ranges from –.8 to .8. If Equation 1 holds, one would expect the standard deviation of the coefficient to be .1, hence these extreme observations would be around 8 standard deviations away from the mean. Right panel in Figure I.1 shows the *t*-values produced by diving the correlation coefficient by the standard deviation of .1: 50% of all absolute *t*-values are larger than 1.96, hence independence between X and Y would be rejected at a 5% significance level for those *t*-values. So we confirmed that when variables X and Y are highly autocorrelated, the correlation coefficient is not distributed as suggested by Equation 1.

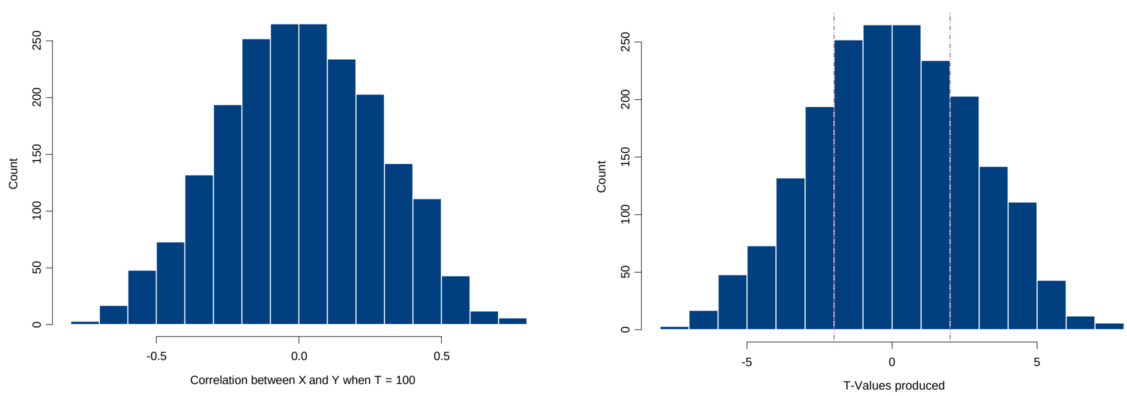

Figure I.1: Distribution of the correlation coefficient between X and Y and t-values produced assuming that Equation 1 can be employed to estimate the coefficient's standard deviation.

To answer the question of when the difference between the correlations for two sub-periods is significantly different from zero, one first needs to have a model for the correlation coefficient's distribution that accounts for autocorrelation. Bartlett's formula (Bartlett, 1949) can be used to find the standard deviation of the zero lag correlation coefficient between two autocorrelated (dependent or not) variables X and Y:

$$\sigma^2_{\rho(0),XY} = T^{-1}\sum_{k=-\infty}^{\infty}\left\{\rho_{xx}(k)\rho_{yy}(k)+\rho_{xy}(k)\rho_{yx}(k)-2\rho_{xy}(0)\times\left[\rho_{xx}(k)\rho_{xy}(k)+\rho_{yy}(k)\rho_{yx}(k)\right]\right\}$$
$$+T^{-1}\sum_{k=-\infty}^{\infty}\left\{\rho_{xy}{}^2(0)\times\left[\frac{\rho_{xx}{}^2(k)}{2}+\rho_{xy}{}^2(k)+\frac{\rho_{yy}{}^2(k)}{2}\right]\right\} \qquad \textit{(Equation 2)}$$

where $\rho_{xy}(k)$ denotes lag k correlation between X and Y, and $\rho_{xx}(k)$ denotes autocorrelation for the variable X at lag *k*. Note that Equation 2 simplifies to $\sum_{k=-\infty}^{\infty}\left\{\rho_{xx}(k)\rho_{yy}(k)\right\}\Big/T$ when the variables are assumed to be independent and to $1/T$ when the variables are independent and not autocorrelated. Box and Jenkins (1976) suggest that at most *k* = T/4 number of lags should be used when calculating lagged autocorrelations and correlations; they also insist that $k > 50$. Our data only have 106 observations total for each time series, hence, if we split them into two equal sub-periods, we are working with 53 observations; when calculating lagged autocorrelations and correlations we can not use lags that are as high as 50 which make the estimate of the variance of the correlation coefficient less precise. Also note that when variables X and Y are not independent and autocorrelated, the mean of the correlation coefficient between them is no longer 0.

I.b Hypothesis Statement for Testing the Significance of the Difference of the Correlation Coefficients

Suppose that we have two time series X and Y with observations recorded over the time period 1…T. Suppose that we would like to test for the correlation breakdown for two disjoint sub-periods 1…K, and K+1…T. That is:

$$H_o : \rho(X,Y,1..K) = \rho(X,Y,K+1..T)$$
$$H_A : \rho(X,Y,1..K) \neq \rho(X,Y,K+1..T).$$

Moreover, we can think of the previous hypothesis as testing for the breakdown in the generating mechanism for these two variables going from one sub-period to another, hence we could use the following asymptotic results

$$\rho(X,Y,1..K) \sim N\left(\mu_1, \frac{1}{\sqrt{K}}B(X,Y,1..K)\right)$$
$$\rho(X,Y,K+1..T) \sim N\left(\mu_2, \frac{1}{\sqrt{T-K}}B(X,Y,K+1..T)\right),$$

where $\frac{1}{\sqrt{K}}B(X,Y,1..K)$ is the standard deviation given obtained from Equation 2. Since under the null, we assume that the correlation does not change from one sub-period to another, we have $\mu_1 = \mu_2$ and if the generating mechanism stays the same, $B(X,Y,1..K) = B(X,Y,K+1..T)$. Then the latter two expressions can be estimated more precisely using the whole sample $B(X,Y,1..T)$. Then, the difference between correlations will have the following distribution:

$$\Delta\rho \sim N\left(0, \sqrt{\frac{1}{K}B^2(X,Y,1..T)+\frac{1}{T-K}B^2(X,Y,1..T)}\right) = N\left(0, \sqrt{\frac{T}{(T-K)K}} \times B(X,Y,1..T)\right). \quad \textit{(Equation 3)}$$

I.c Testing the Performance of the Correlation Difference Test

We will first test the asymptotic performance of the proposed test with T = 5,000 and K = 2,500; we will use lags up to k = 100 to estimate $B(X,Y,1..T)$, and will cycle through 4 different X and Y generating models:

> *Model 1*: no autocorrelation and no correlation (X and Y are both random normal);
>
> *Model 2.x*: the variables are both first-order autocorrelated but independently generated:
> $X_t = \phi_{X,1}X_{t-1} + \varepsilon_x$ and $Y_t = \phi_{Y,1}Y_{t-1} + \varepsilon_y$ with errors as independent normal;
>
> *Model 3.x*: the variables are both first-order autocorrelated and dependent through a "direct source":
> $X_t = \phi_{X,1}X_{t-1} + \varepsilon_x$ and $Y_t = \phi_{Y,1}X_{t-1} + \varepsilon_y$ with errors as independent normal;
>
> *Model 4*: the variables are both first-order autocorrelated and dependent through an "indirect source":
> $X_t = \phi_{X,1}X_{t-1} + \varepsilon_x$ and $Y_t = \phi_{Y,1}Y_{t-1} + \varepsilon_y$ with errors having covariance $\sigma_{\varepsilon_y,\varepsilon_y} \neq 0$.

The autoregressive parameters for the models and sub-models are:

Method	***Autoregressive Coefficients***
1	Random Noise
2.1	$\phi_x = \phi_y = .9$
2.2	$\phi_x = \phi_y = .6$
3.1	$\phi_x = \phi_y = .9$
3.2	$\phi_x = \phi_y = .7$
3.3	$\phi_x = \phi_y = .5$
3.4	$\phi_x = .9; \phi_y = .5$
4	$\sigma_{\varepsilon_y,\varepsilon_y} = .7$

Recall, that under the null, in both sub-periods each variable is generated using the same model. The rejection percentages were obtained based on a *t*-statistic of

$$t = \Delta\rho \Big/ \left\{\sqrt{\frac{T}{(T-K)K}} \times B(X,Y,1..T)\right\}.$$

Table I.1 shows that when Bartlett's formula is employed through *Equation 3* and the *t*-values for correlation differences are computed, the rejection percentages are close to the nominal alpha levels: they range from the low of 4.65% to the high of 5.75% for the 95% significance level, and from 8% to 10% for the 90% one.

alpha = .05			alpha = .10	
Method	% Rej. Using Bartlett's Formula	% Rej. Using 2/sqrt(n) as St. Dev.	% Rej. Using Bartlett's Formula	% Rej. Using 2/sqrt(n) as St. Dev.
1	5.75%	5.80%	10.55%	10.55%
2.1	5.75%	52.45%	9.80%	58.00%
2.2	5.05%	17.10%	8.35%	23.75%
3.1	4.65%	0.75%	9.65%	2.45%
3.2	4.75%	4.35%	8.95%	8.30%
3.3	4.65%	4.85%	8.00%	8.05%
3.4	5.00%	6.25%	8.50%	10.70%
4	5.20%	12.65%	9.25%	19.20%

Table I.1: Percentage of rejected correlation differences at 95 and 90% significance levels using Equation 1 and Equation 2 for the computation of the standard deviation of the correlation coefficient.

I.d Performance in Small Sample

The asymptotic rejection levels when the null is true are satisfactory for our tests of the significance for the difference in the correlation coefficient between sub-periods. But our data is limited; hence it is imperative to examine the behavior of the proposed test in a small sample setting. We assumed that T = 100, and each model had 5,000 X and Y pairs generated. Table I.2 includes 2 new sub-models 2.3 and 3.5 that generate X and Y so that they are AR(3) instead of just AR(1). The rejection percentages given in Table I.2 are much further away from the nominal alpha levels for the Bartlett-corrected measure when compared to the asymptotic case in Table I.1. Model 3.1 (when variables are strongly autocorrelated and directly dependent) has more than 10% of all correlation differences rejected at 5% level.

	alpha = .05		alpha = 10	
Data Generating Method	Bartlett	2/sqrt(n)	Bartlett	2/sqrt(n)
1	5.74%	5.28%	11.42%	11.04%
2.1	7.10%	42.92%	14.28%	50.68%
2.2	5.58%	16.64%	10.98%	24.48%
2.3	5.84%	37.80%	12.22%	45.80%
3.1	10.44%	4.22%	16.25%	7.94%
3.2	5.38%	5.38%	10.80%	10.44%
3.3	5.32%	5.50%	9.80%	10.20%
3.4	8.00%	7.28%	13.82%	13.16%
3.5	8.50%	13.20%	15.22%	20.30%
4	9.18%	27.32%	15.23%	35.66%

Table I.2: Percentage of rejected samples at 95 and 90% significance levels using Equation 1 and Equation 2 when T = 100.

I.e Other Tests for Correlation Difference

We also considered other tests that assess the significance of the differences between correlation coefficients. Some, like differencing the times series, were discarded as when the data were just random normal, differencing made the original series autocorrelated. The three alternative tests examined were:

> *Test 1*: stationary bootstrap using Politis and Romano (1994) algorithm;
>
> *Test 2*: pre-whitening both series by fitting an AR(1) model then using Equation 1 for the estimate of the standard deviation of the correlation coefficient;
>
> *Test 3*: pre-whitening both series by fitting an AR(k) model to each, where k is determined using AIC; then use Equation 1 for the estimate of the standard deviation of the correlation coefficient.

Table I.3 summarizes the results when alternative tests are used to reject the null of zero difference between correlations: stationary bootstrap performs the worst of the three with 7 out of 10 generating methods rejecting further away from the nominal 5% level than the proposed Bartlett test. On the other hand, pre-whitening using k number of lags determined by AIC looks hopeful; the test's largest over-rejection is 6.65% (Bartlett's one is 10.44%) and 7 out of 10 data generating methods for this pre-whitening test have closer rejection levels to 5% than the Bartlett test.

alpha = .05	T = 100	Test 1			Test 2	Test 3
Data Generating Method	Bartlett	Stationary Bootstrap q = .025	Stationary Bootstrap q = .1	Stationary Bootstrap q = .2	AR(1) Pre-whitened	AR(k) Pre-whitened
1	5.74%	8.70%	6.30%	4.60%	5.42%	5.78%
2.1	7.10%	8.90%	9.50%	16.80%	5.74%	5.76%
2.2	5.58%	9.00%	5.90%	6.30%	5.20%	5.36%
2.3	5.84%	8.00%	10.10%	19.10%	8.12%	6.56%
3.1	10.44%	5.50%	5.80%	13.10%	1.88%	2.02%
3.2	5.38%	8.00%	4.10%	6.80%	1.98%	2.36%
3.3	5.32%	8.10%	4.40%	4.20%	3.54%	3.58%
3.4	8.00%	5.60%	6.20%	8.90%	3.14%	3.56%
3.5	8.50%	6.10%	7.60%	13.80%	2.62%	3.54%
4	9.18%	8.20%	9.30%	14.30%	1.52%	1.92%

Table I.3: Percentage of rejected samples at the 95% significance level when T = 100.

Unfortunately, a test based on pre-whitening might produce nonsensical results when the null is not true. Suppose that in the first sub-period X and Y are generated as independent random normal variables and in the second sub-period they are generated using method 2.1, that is, they are both highly autocorrelated yet independently generated. Also suppose that T = 100 and that the correlation between X and Y in both periods is .4. Then the difference between the equally-sized sub-periods is .4 - .4 = 0. When one pre-whitens the series in the first sub-period, nothing happens as the series are random normal with no autocorrelation, hence the correlation stays the same at .4; the correlation between the pre-whitened X and Y in the second sub-period will be close to 0. So the difference between the sub-period correlations will be approximately .4 – 0 = .4. Since now both time periods have data series that are neither dependent nor autocorrelated, Equation 1 can be used to get an estimate of the standard deviation for the correlation coefficient that is $1/\sqrt{50}$; hence the difference will have a standard deviation of $\sqrt{2/50} = .2$. So the difference of .4 will be rejected as significant. But the original difference before pre-whitening between the periods was 0; such instances lead the test based on pre-whitening to rejecting really small or zero differences between sub-periods as significant. Therefore we maintain that Equation 3 should be used when testing for the difference between the sub-period correlations.

Appendix II. Simulation under the Null of no Relationship Between Real Estate and Inflation.

Figure II.1 illustrates the associative relationships among the economic variables in included in our VAR model. Each arrow corresponds to a significant predictive power (*p*-value less than .05), with the variables on the top being predicted by variables below them. Red arrows indicate negative association, while black ones – positive.

For example, national real estate [NCREIF] returns are positively associated with HPI [OFHEO] ones, while in turn, HPI can is higher when inflation is high. Thus if we want to study the hedging properties of the portfolios under the null of no association between inflation and national real estate returns, we need to set to 0 not only the lagged coefficient on inflation in the VAR coefficient matrix when simulating current national real estate returns, but also the lagged coefficient for HPI, as inflation is acting through it to increase national real estate returns. When modeling inflation, such consideration does not matter as inflation depends mainly on itself; however, we also set the lagged coefficient of real estate on inflation to zero.

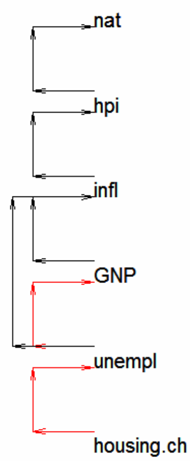

Figure II.1: Graphical representation of the associative relationships between the variables included in the VAR model significant at a 5% alpha level. Negative association is marked with red arrows and positive one with black.

We then simulate 10,000 paths for each variable with three entries in the VAR coefficient matrix equal to zero as described above. We used the full sample data and assumed a 20-year holding period. Table II.1 shows that the correlation between commercial real estate and inflation under the null is .215. While it is significantly different from zero, it is a great reduction from the correlation of roughly .70 obtained using the original VAR coefficient matrix. It does, however, suggest that correlations under the null can be at least as high as .21.

	National	LT Gov B	S&P 500	Housing	Inflation
National	1.00	-0.47	0.08	0.30	0.22
LT Gov B	-0.47	1.00	0.27	-0.47	-0.40
S&P 500	0.08	0.27	1.00	-0.35	-0.16
Housing	0.30	-0.47	-0.35	1.00	0.77
Inflation	0.22	-0.40	-0.16	0.77	1.00

Table II.1: Correlation under the null of independence between five simulated variables.

References

Aldrich, J.H., 1995, “Correlations Genuine and Spurious in Pearson and Yule,” *Statistical Science*, 10(4), 364-376.

Anari, Ali & Kolari, James, 2002, “House Prices and Inflation,” *Real Estate Economics* 30 (1), 67-84.

Ang, A. and Piazzesi, 2003, “A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables,” *Journal of Monetary Economics*. 50, 745–787.

Baghestani, H. 1992. “On the Formation of Expected Inflation under Various Conditions: Some Survey Evidence.” *The Journal of Business*, Vol. 65, No. 2, 281 – 293.

Barkham, R. and C. Ward, 1999, “Investor Sentiment and Noise Traders,” *Journal of Real Estate Research*, 18, 291-312.

Bartlett, M. S., 1949. “Some Evolutionary Stochastic Processes,” *Journal of the Royal Statistical Society*. *Series B (Methodological)*, 11 (2), 211-229.

Belviso, F. and Milani, F. 2003. “Structural Factor-Augmented VAR (SFAVAR),” working paper, Princeton University.

Bernanke, B., Boivin, J., and Eliasz, P. 2004. “Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach,” Finance and Economics Discussion Series No. 2004-3 (Washington: Board of Governors of the Federal Reserve System).

Box, G.E.P., and Jenkins, G.M. (1976). *Time Series Analysis: Forecasting and Control*, Revised Edition. Holden-Say, San Francisco.

Brueggeman, W. B., A. H. Chen and T. G. Thibodeau, 1984, “Real Estate Investment Funds: Performance and Portfolio Considerations,” *Real Estate Economics*, 12(3), 333–54.

Brunnermeier, M. and Christian Julliard, 2005, “Money Illusion and Housing Frenzies” Princeton University Working Paper.

Caplin, Andrew, Eric Hangen, Barry Nalebuff, Elisabeth Prentice, John Rodkin, Matthew I. Spiegel and Tom Skinner, “Home Equity Insurance: A Pilot Project,” Yale ICF Working Paper, 2003.

Case K. and R. J. Shiller, 1987, "Prices of Single Family Homes Since 1970: New Indexes for Four Cities," Cowles Foundation Discussion Papers 851.

Chaudry, Mukesh, P.C. N. Meyer and J.R. Webb, 1999, “Stationarity and Cointegration in Systems with Real Estate and Financial Assets,” *Journal of Real Estate Finance and Economics*, 18(3), 339- 349.

Clive W J & Hyung, Namwon & Jeon, Yongil, 2001, "Spurious Regressions with Stationary Series," *Applied Economics,* 33(7), 899-904.

Fama, E. F., and G. W. Schwert, 1977, “Asset Returns and Inflation,” *Journal of Financial Economics*, 5(2), 115–46.

Ferson, W. , S. Sarkissian and T. Simin,2003, "Spurious Regressions in Financial Economics?" forthcoming*, Journal of Finance.*

Fisher, J.D. and W.N. Goetzmann, 2005, “Performance of Real Estate Portfolios, *Journal of Portfolio Management*, Special Issue.

Ganesan, S. and Y. H. Chiang, 1998, “The Inflation-Hedging Characteristics of Real and Financial Assets in Hong Kong,” *Journal of Real Estate Portfolio Management*, 4(1), 55–67.

Geltner, David & Goetzmann, William, 2000, "Two Decades of Commercial Property Returns: A Repeated-Measures Regression-Based Version of the NCREIF Index," *The Journal of Real Estate Finance and Economics*, vol. 21(1), 5-21.

Geltner, David, 1993, “Temporal Aggregation in Real Estate Return Indices,” *Real Estate Economics*, 21, 141-166.

Glascock, J.L., C. Lu and R.W.So, 2002, “REIT Returns and Inflation: Perverse or Reverse Causality Effects?,” *The Journal of Real Estate Finance and Economics,* 24(3), 301-317.

Goetzmann W.N. and R.G. Ibbotson, 1990, “The Performance of Real Estate as an Asset Class,” *Journal of Applied Corporate Finance,* 3(1), 65-76.

Goetzmann, W. N., Lingfeng, L. and Rouwenhorst, K.G., 2004, “Long Term Global Market Correlations.” *Journal of Business*. 78(1), 1-38.

Goetzmann, W. N.,1992, "The Accuracy of Real Estate Indices: Repeat Sale Estimators," *The Journal of Real Estate Finance and Economics*, 5(1), 5-53.

Goetzmann, W.N. and F. Edwards, 1994, “Short Horizon Inputs and Long Horizon Portfolio Choice,” *Journal of Portfolio Management*; 20(4), 76-81.

Goetzmann, W.N., L. Li & K.G. Rouwenhorst, 2005, "Long-Term Global Market Correlations," *Journal of Business,* 78(1), 1-38.

Goetzmann, William and Matthew Spiegel, 2002, “The Policy Implications of Portfolio Choice in Underserved Mortgage Markets,” in Nicholas P. Retsinas and Eric S. Belsky, eds. *Low Income Homeownership: Examining the Unexamined Goal*, Harvard Joint Center for Housing and Urban Studies Conference and The Brookings Institution, and Brookings Institution Press, 257-274.

Goetzmann, William, 1993, "The Single Family Home in the Investment Portfolio." Journal of Real Estate Finance and Economics, Vol. 6 (3) 201-22.

Gorodnichenko, Y. 2003, “Reduced-rand Identification of Structural Shocks in VARs,” Department of Economics, University of Michigan.

Graff and Cashdan, 1990, “Some New Ideas in Real Estate Finance,” *Journal of Applied Corporate Finance*, 3(1), 77-84.

Granger, C.W.J. and P. Newbold, 1974, Spurious regressions in Economics, *Journal of Econometrics* 4, 111-120.

Hartzell, D., J. Heckman and M. Miles, 1987, “Real Estate Returns and Inflation,” *Real Estate Economics*, 15(1), 617–37.

Hoesli, M., 1994, “Real Estate as Hedge Against Inflation: Learning from the Swiss Case,” *Journal of Property Valuation and Investment*,12(3), 51–9.

Ibbotson R.G., L.B. Siegel, 1984, “Real Estate Returns: A Comparison with Other Investments,” *Real Estate Economics*, 12(3), 219-242.

James and Orcutt, 1948, “Testing the Significance of Correlation Between Time Series,” *Biometrika*, 35(3/4), 397-413.

Li, Lingfeng, 2002, "Macroeconomic Factors and The Correlation of Stock and Bond Returns," Yale University Working Paper.

Linneman, P and J. Gyourko, 1988, “Owner-occupied homes, income-producing properties, and REITs as inflation hedges: Empirical findings,” *The Journal of Real Estate Finance and Economics*, 1(4), 347-372.

Lu, Chiuling & So, Raymond W, 2001, "The Relationship between REITs Returns and Inflation: A Vector Error Correction Approach," *Review of Quantitative Finance and Accounting*, 16(2), 103-15.

Maitland-Smith, J and C. Brooks, 1999, “Threshold Autoregressive and Markov Switching Models: An Application to Commercial Real Estate”, *Journal of Property Research*, 16, 1-20.

Mehra, Y.P., 2002, "Survey Measures of Expected Inflation: Revisiting the Issues of Predictive Content and Rationality," Federal Reserve Bank of Richmond Economic Quarterly, 88, 17-36.

Milani, F. 2003. "Adaptive Learning, Model Uncertainty and Monetary Policy Inertia in a Large Information Environment." Computing in Economics and Finance. 2003. 280. Society for Computational Economics. International Monetary Fund World Economic Outlook (WEO), Chapter 2, September 2004.

Miles, D., 1996, "Property and Inflation," *Journal of Property Finance,* 7(1), 6–20.

Newell, G., 1996, "The Inflation-Hedging Characteristics of Australian Commercial Property: 1984–1995," *Journal of Property Finance,* 7(1), 6–20.

Newey,W., and K.West, 1987, ``A Simple Positive-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica* 55, 703-708.

Phillips, P.C.B., 1986, Understanding spurious regressions in Econometrics, *Journal of Econometrics* 33, 311-340.

Phillips, P.C.B., 1998, New tools for understanding spurious regressions, *Econometrica* 66, 1299-1326.

Piazzesi, M., Schneider, M., and Tuzel, S. 2004. "Housing, Consumption, and Asset Pricing." (Unpublished; Chicago: University of Chicago).

Politis, D.N., and Romano, J.P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 82 (397), 171-185.

Ross, Stephen A & Zisler, Randall C, 1991, "Risk and Return in Real Estate," The *Journal of Real Estate Finance and Economics*, vol. 4(2), 175-90.

Ruben, J., M. Bond and J. R. Webb, 1989, "The Inflation-hedging Effectiveness of Real Estate," *Journal of Real Estate Research,* 4, 45–56.

Sinai, T and N. S. Souleles, 2005, "Owner-Occupied Housing as a Hedge Against Rent Risk," *The Quarterly Journal of Economics*, 120(2), 763-789.

Sirmans, G.F., and Sirmans, C.F., 1987, "Real Estate Returns: The Historical Perspective of Real Estate Returns," Journal of Portfolio Management, 13(3), 22-31.

Valaitis, E. "Expected Inflation and Inflationary Shocks Revisisted," working paper.

Viceira, L. M., and J.Y. Campbell. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press, 2002.

Walker, Sir Gilbert, 1950, “Apparent Correlation Between Independent Series of Autocorrelated Observations,” *Biometrika*, 37(1/2), 184-185.

Wheaton, W. C., Torto, R.G., Sivitanides, P.S., Southard, J.A., Hopkins, R.E., and Costello, J.M. 2001, “Real Estate Risk: A Forward Looking Approach.” Working Paper, Torto Wheaton Research, Boston.

Yule G.U., 1926, “Why do we sometimes get nonsense correlations between time series? A study in sampling and the nature of time series,” *Journal of the Royal Statistical Society* 89, 1-64.

Table 1: Variables Used in VAR

National	Returns on the NCREIF NPI
housing.ch	Change in Completed Housing Units Nationwide
inflation	Change in CPI
US.lt.gov	Returns on U.S. Long-Term Government Bonds
GNP	Change in GNP
SP500	Change in SP500 Index
CH.Unemployment	Change in National Unemployment Level
HPIofheo	Returns of the OFHEO House Price Index

Table 2: Correlation differences between two sub-periods for pairs of variables.

	National	housing.ch	inflation	US.lt.gov	GNP	SP500	CH.Unempl	HPIofheo
National		0.15	0.40	-0.11	0.63	-0.18	0.10	0.09
housing.ch	0.15		0.08	0.12	0.19	0.03	0.01	0.24
inflation	0.40	0.08		-0.28	0.45	0.15	0.37	0.53
US.lt.gov	-0.11	0.12	-0.28		-0.17	0.55	0.28	-0.12
GNP	0.63	0.19	0.45	-0.17		-0.32	-0.13	0.63
SP500	-0.18	0.03	0.15	0.55	-0.32		0.27	0.17
CH.Unempl	0.10	0.01	0.37	0.28	-0.13	0.27		-0.51
HPIofheo	0.09	0.24	0.53	-0.12	0.63	0.17	-0.51	

Table 3: AR(1) and General AR(k) Models

Variable	AR(k=...)	ϕ_1 =	AR(k=...)	ϕ_1 =
National	1	0.68	6	0.40
housing.ch	1	0.77	5	0.76
inflation	1	0.54	12	0.24
US.lt.gov	0	0.00	5	-0.09
GNP	1	0.28	1	0.28
SP500	0	0.00	0	0.00
CH.Unemployment	1	0.34	8	0.29
HPIofheo	1	0.92	1	0.92

Table 4: Results of Pairwise Tests for Correlation Breaks.

End period with	housing.ch	inflation	US.lt.corp	US.lt.gov	GNP	SP500	CH.Unemployment	HPIofheo	TotCompGrowthPrivate
30	-0.3	1.3	-1.4	-1.2	1.2	-1	0.6	0	-0.9
31	-0.3	1.4	-1.3	-1.2	1.2	-0.9	0.7	0	-0.9
32	-0.3	1.4	-1.1	-1	1.2	-0.7	0.7	0	-0.9
33	-0.3	1.5	-1.1	-1.1	1.3	-0.8	0.7	0	-0.9
34	-0.2	1.5	-1.1	-1	1.4	-0.8	0.7	-0.1	-0.8
35	-0.1	1.6	-1	-0.9	1.4	-0.5	0.8	-0.2	-0.7
36	-0.1	1.6	-1	-0.9	1.4	-0.5	0.8	-0.2	-0.7
37	-0.1	1.6	-1	-0.8	1.5	-0.7	0.8	-0.3	-0.6
38	0.1	1.6	-0.7	-0.6	1.4	-0.7	1	-0.4	-0.5
39	0.1	1.6	-0.6	-0.5	1.4	-0.7	1.1	-0.4	-0.5
40	0.1	1.6	-0.6	-0.5	1.5	-0.6	1.1	-0.4	-0.5
41	0.1	1.6	-0.6	-0.5	1.5	-0.6	1.1	-0.4	-0.5
42	0.1	1.6	-0.6	-0.5	1.5	-0.6	1.2	-0.4	-0.5
43	0.1	1.7	-0.5	-0.5	1.5	-0.6	1.2	-0.4	-0.5
44	0	1.6	-0.6	-0.5	1.5	-0.6	1.2	-0.4	-0.5
45	0	1.6	-0.5	-0.4	1.5	-0.6	1.2	-0.4	-0.4
46	0	1.6	-0.5	-0.5	1.6	-0.6	1.2	-0.3	-0.4
47	0	1.7	-0.5	-0.4	1.6	-0.7	1.2	-0.3	-0.4
48	0	1.7	-0.5	-0.4	1.6	-0.6	1.2	-0.3	-0.4
49	0.1	1.6	-0.4	-0.3	1.6	-0.5	1.2	-0.3	-0.5
50	0.1	1.6	-0.4	-0.3	1.6	-0.5	1.2	-0.3	-0.5
51	0.1	1.4	-0.3	-0.2	1.7	-0.2	0.9	-0.1	-0.5
52	0.3	1.4	-0.4	-0.4	**2**	-0.4	0.6	0.2	-0.3
53	0.5	1.4	-0.4	-0.4	**2.1**	-0.6	0.3	0.3	-0.4
54	0.7	1.4	-0.3	-0.3	**2.2**	-0.5	0.3	0.4	-0.5
55	1	1.4	-0.5	-0.5	**2.3**	-0.5	0.3	0.6	-0.6
56	1.6	1.4	-0.8	-0.8	**2.7**	-0.9	-0.1	0.5	-0.5
57	1.7	1.4	-0.6	-0.6	**2.7**	-0.7	-0.2	0.6	-0.6
58	1.7	1.4	-0.7	-0.7	**2.9**	-0.7	-0.5	0.7	-0.3
59	1.6	1.4	-0.8	-0.9	**3**	-0.7	-0.3	0.8	-0.3
60	1.3	1.5	-0.7	-0.7	0.4	-0.9	0.1	1	-0.3
61	1.2	1.4	-0.9	-0.9	0.4	-0.9	0.3	1.2	-0.3
62	1.1	1.5	-1	-1.1	0.6	-0.9	0.3	1.4	-0.4
63	1.1	1.5	-1.1	-1.2	0.6	-0.9	0.5	1.5	-0.3
64	0.7	1.6	-0.9	-1	0.5	-0.9	0.8	1.9	-0.2
65	0.6	1.5	-0.6	-0.8	0.5	-0.8	0.8	**2**	-0.2
66	0.3	1.6	-0.5	-0.7	0.4	-0.8	1	**2.1**	-0.1
67	0.2	1.5	-0.4	-0.6	0.4	-0.8	1.1	**2.3**	-0.2
68	0.1	1.5	-0.4	-0.6	0.4	-0.8	1.3	**2.5**	-0.1
69	0.1	1.5	-0.4	-0.7	0.4	-0.8	1.3	**2.6**	0
70	0.1	1.5	-0.5	-0.8	0.5	-0.9	1.3	**2.7**	0.1
71	0.1	1.5	-0.5	-0.8	0.5	-0.9	1.3	**2.8**	0.1
72	0.3	1.6	-0.8	-1.1	0.6	-1	1.4	**3**	0.4
73	0.3	1.6	-0.8	-1	0.6	-1	1.4	**3**	0.5
74	0.3	1.6	-0.7	-1	0.6	-1	1.4	**3.1**	0.5
75	0.3	1.6	-0.8	-1	0.6	-1	1.4	**3.1**	0.5
76	0.3	1.6	-0.8	-1	0.6	-1	1.5	**3.2**	0.5

Table 5: Simulated Distributions and Historical Means:

A: Entire sample period 1977-2004

Series	Mean	StDev	Min	Max	normal	Historical Mean
NCREIF	0.11134	0.0296	-0.029058	0.2624	0	0.09447
housing	0.03993	0.2452	-0.767641	1.7684	0	0.03440
Infl.	0.05079	0.0106	0.000177	0.0997	0	0.04318
LTGvt	0.07515	0.0425	-0.160922	0.2778	0	0.10341
GNP	0.07188	0.0126	0.013568	0.1307	0	0.06758
SP500	0.12489	0.0653	-0.168912	0.4441	0	0.14816
Unemp.	-0.00388	0.052	-0.225408	0.3486	0	-0.00444
HPIofheo	0.06552	0.0153	-0.007887	0.1463	0	0.05540

B: First sample period: 1977-1991

Series	Mean	Stdev	Min	Max	normal	hist mean
NCREIF	0.0812	0.0346	-0.0573	0.26	0	0.1009
housing	0.062	0.324	-0.9345	2.801	0	-0.0958
Infl.	0.0436	0.0112	-0.00341	0.1	0	0.0579
LTGvt	0.1196	0.0537	-0.11504	0.378	0	0.1127
GNP	0.073	0.0177	0.0015	0.173	0	0.0754
SP500	0.176	0.0731	-0.14946	1.033	0	0.168
Unemp.	-0.0233	0.0704	-0.33858	0.544	0	0.0148
HPIofheo	0.0479	0.0159	-0.00922	0.142	0	0.0614

C: Second Sample Period: 1992-2004

Series	Mean	Stdev	Min	Max	normal	hist median
National	0.1018	0.02652	-0.04189	0.2584	0	0.0959
housing.ch	0.1779	0.11827	-0.40782	0.793	0	0.2185
inflation	0.0274	0.00327	0.01099	0.046	0	0.0248
US.lt.gov	0.0568	0.04534	-0.07866	0.448	0	0.0647
GNP	0.0532	0.00899	0.00213	0.0908	0	0.0495
SP500	0.0534	0.14022	-0.76234	1.1046	0	0.129
Unemp.	0.0118	0.06519	-0.31572	0.549	0	-0.0708
HPIofheo	0.0639	0.01558	0.00388	0.1507	0	0.0507

Table 6: Summary Statistics for Simulated Portfolios.

A: Entire Sample Period

	C All	C Stocks	C Bonds	C CommRE	H All	H Stocks	H Bonds	H Housing
Mean	0.1110	0.1160	0.1060	0.1130	0.0972	0.1060	0.0944	0.0907
St Dev	0.0299	0.0382	0.0282	0.0257	0.0320	0.0404	0.0308	0.0257
Min	0.0059	-0.0137	0.0025	0.0141	-0.0115	-0.0290	-0.0167	0.0031
Max	0.3170	0.3580	0.2900	0.2880	0.3130	0.3560	0.2880	0.2800
Corr with Inflation	0.1810	0.1380	0.0708	0.3400	0.0509	0.0606	-0.0443	0.1440

B: Sample Period 1977-1991

	C All	C Stocks	C Bonds	C CommRE	H All	H Stocks	H Bonds	H Housing
Mean	0.1370	0.1480	0.1340	0.1270	0.1290	0.1430	0.1270	0.1140
St Dev	0.0418	0.0504	0.0393	0.0364	0.0444	0.0524	0.0422	0.0387
Min	-0.0093	-0.0287	0.0002	-0.0057	-0.0164	-0.0368	-0.0248	-0.0114
Max	0.8190	0.8940	0.7680	0.7690	0.8190	0.8940	0.7680	0.7680
Corr with Inflation	-0.0211	0.0307	-0.1940	0.0922	-0.1002	-0.0175	-0.2560	-0.0500

C: Sample Period 1992-2004

	C All	C Stocks	C Bonds	C CommRE	H All	H Stocks	H Bonds	H Housing
Mean	0.0850	0.0804	0.0832	0.0893	0.0735	0.0719	0.0741	0.0723
St Dev	0.0426	0.0551	0.0383	0.0369	0.0406	0.0536	0.0375	0.0326
Min	-0.0391	-0.0625	-0.0243	-0.0353	-0.0093	-0.0290	-0.0199	0.0040
Max	0.8830	0.9610	0.8300	0.8310	0.8830	0.9600	0.8290	0.8290
Corr with Inflation	0.0181	0.0254	-0.0404	0.0723	-0.0344	-0.0084	-0.0849	-0.0126

Table 7: Summary of Shortfall Analysis for Different Holding Periods.

	Portfolio							
Holding Period & Sample	C All	C Stocks	C Bonds	C Comm RE	H All	H Stocks	H Bonds	H Housing RE
5 years / Full	6.38%	11.69%	11.48%	2.14%	21.61%	22.73%	26.59%	16.58%
5 years / 59:106	9.30%	20.63%	10.11%	3.15%	13.98%	26.28%	14.92%	3.47%
10 years / Full	1.18%	2.97%	2.61%	0.18%	7.23%	8.18%	9.40%	4.70%
10 years / 1:58	0.33%	0.62%	0.53%	0.18%	1.69%	1.61%	2.19%	1.83%
10 years / 59:106	4.65%	15.28%	4.50%	1.11%	7.55%	19.80%	7.31%	0.83%
20 years / Full	0.03%	0.24%	0.22%	0.00%	1.18%	1.46%	1.85%	0.54%
20 years / 1:58	0.01%	0.64%	0.28%	0.02%	0.34%	0.35%	0.57%	0.38%
20 years / 59:106	1.78%	11.63%	1.63%	0.19%	3.47%	15.78%	3.09%	0.64%
	Worst Portfolio in RED							
	Best Portfolio in BLUE							