Hedging Housing Risk in London

Matteo Iacoviello*
Boston College

François Ortalo-Magné†
London School of Economics
& University of Wisconsin

October 3, 2002

Abstract

This paper investigates the benefits of allowing households to compensate the portfolio distortion due to their housing consumption through investments in housing price derivatives. Focusing on the London market, we show that a major loss from over-investment in housing is that households are forced to hold a very risky portfolio. However, the strong performance of the London housing market means that little is lost in terms of expected returns. Even households with limited wealth are better off owning their home rather than renting and investing in financial assets, as long as they are willing to face the financial risk involved. In this context, access to housing price derivatives would benefit most poor homeowners looking to limit their risk exposure. It would also benefit wealthier investors looking for the high returns provided by housing investments without the costs of direct ownership of properties. Comparisons with French, Swedish and US data provide a broader perspective on our findings.

*Department of Economics, Boston College, Carney Hall, Chestnut Hill MA 02467-3806; email: iacoviel@bc.edu

†Department of Economics, London School of Economics, Houghton Street, London WC2A 2AE, UK; and Department of Real Estate and Urban Land Economics, University of Wisconsin, 975 University Avenue, Madison, WI 53706; email: ortaloma@lse.ac.uk.

1 Introduction

This paper provides further evidence on the potential benefits of financial instruments linked to the performance of the housing market, using London, England, as a case study. We find that the returns to housing in London have been strong but very volatile compared to other financial assets. Households overinvested in housing due to their housing consumption motive gain from the high returns on their home, but are forced to hold a very risky portfolio. Standard financial assets do not provide much of a hedge against the risk of owning a home. This explains why homeowners pursuing low risk-low return strategies would benefit from the introduction of housing price derivatives. We find that the other major beneficiaries of such derivatives would be investors pursuing high risk-high return strategies.

The first reason for choosing London as the focus of this study is the volatility of the local housing market since the mid-Seventies, the period covered by our data. The second reason is that Londoners now have access to financial instruments which allow them to limit or expand their exposure to housing price risk, independently of their housing consumption. Such instruments have been discussed in the literature for a number of years. The academic literature has attempted to encourage their introduction by demonstrating their potential benefit.[1] Here, at last, we have a city where hedging housing price risk is feasible. Our findings have direct implications for the marketing of the products currently on offer and the development of further financial innovations.

Obviously, the findings of our study are subject to the usual caveats of the standard mean-variance portfolio approach we take. In particular, in determining optimal portfolio weights, we ignore the household's human capital and future housing consumption needs, as well as differences in liquidity between housing and other assets.[2] Despite such drawbacks, it still remains a useful methodology for examining the empirical evidence and assessing the performance of housing relative to other major classes of assets.

Computing optimal portfolio weights and the mean-variance frontier under various combinations of assets is a standard problem. The difficulty with housing as an asset concerns the estimation of the moments of returns on a single home. For this purpose, we adapt the methodology proposed in Englund et al. (2002). We also follow their lead with regards to the portfolio frontiers we compute in order to generate results which are directly comparable to those they obtain with their Swedish data. To gain a broader perspective on the London evidence, we present comparable data for France and the US based on the work of Lagarenne and le Blanc (2002) and Flavin and Yamashita (2002), respectively.

The remainder of the paper is organized as follows. Section 1 explains how we build the time series and derive the moments of investment returns required to compute portfolio allocations. Section 2 considers unrestricted portfolio allocations before focusing on the constrained problem faced by various types of households and how they would benefit from access to housing price derivatives. Section 3 concludes with remarks on the current housing price derivative offerings on the London market and the potential for further developments of housing-related financial products. The Appendix reports the French, Swedish and US data underlying the comparisons in Section 2.

2 Data

The first step of our analysis consists in computing inputs to the standard portfolio optimization problem. We consider two types of housing investments: a UK based and a London based portfolio. We add four standard investment alternatives: general stocks, real estate stocks, long term bonds and t-bills. We use real returns on the FTSE-All-Share index for general stocks, returns from the FTSE-Real-Estate-Stocks index for real estate stocks, the ten-year benchmark government bond yield for long term bonds, and the three-month Treasury Bills yield for short term bonds. We use the Retail Price Index to deflate nominal returns.[3] Figure 1 plots real quarter on quarter returns for the six assets we include in our specification over the period covered by our data, 1977-2000.

2.1 Deriving housing returns

Both UK and London housing investment returns are built from the housing price indexes provided by Nationwide, a major mortgage lender which reports quality-adjusted indexes for the UK. From the housing price index series, we follow the methodology proposed in Englund et al. (2002) to obtain housing investment returns. These authors define the returns, r_t^H, as

$$r_t^H = P_t - P_{t-1} + .01 \tag{1}$$

where P_t is the log of the house price index and .01 is a proxy for the dividends from a housing portfolio net of taxes and depreciation costs.

To solve the portfolio optimization problem of a homeowner, we also need to estimate housing returns on a single housing unit. Again, following Englund et al. (2002), we define the returns on a single housing unit, r_t^h as

$$r_t^h = r_t^H + v_t - v_{t-1}, \tag{2}$$

where v_t is defined as follows

$$v_t = \rho v_{t-1} + \eta_t \tag{3}$$

$$E(\eta_t) = 0, \ E\left(\eta_t^2\right) = \sigma_\eta^2. \tag{4}$$

Within this framework, what makes a single housing unit different from the housing price index are the autocorrelation ρ in prices for individual sales, and the error variance σ_η^2. This specification reflects the implicit assumption that investing in a single home exposes the investor to a higher risk, but the same expected return as the local housing price index.

To compute an individual house price series for a London home, we therefore need numerical values for both ρ, and σ_η^2. We take advantage of the availability of housing price series for London boroughs, hence more "local" than the overall London time series mentioned above. The borough data are only available on a yearly basis for the period 1989 to 2000 from the Halifax, another lender who reports quality-adjusted housing prices in the UK. We transform this annual data into a quarterly series by assuming that local prices follow a quarterly AR(1) process. We calibrate the autocorrelation parameter of this process, ρ to the value of $-.4$. This is the average value estimated by Englund et al. (1998) on a rich data set of individual housing transactions in Sweden.[4] Unfortunately no such data is available for the UK. We experimented with greater and smaller values for ρ. Our main findings are robust to variation around the calibrated value of $-.4$. Trivially, allowing for a smaller (larger) value of ρ in absolute value implies that the individual house becomes less (more) volatile at all horizons, thus reducing (increasing) the benefits of allowing households to invest in the housing indexes, *ceteris paribus*.

The method we use to convert the yearly local prices time series into quarterly time series already goes one step toward providing us with a single-home housing return series, in the sense that it introduces volatility into the borough price series. The next step consists in estimating the volatility of the single home series itself. By definition,

$$r_t^j = P_t^j (1 - L) \tag{5}$$

where $j \in \{h, H\}$ and L is the lag operator. Plugging these definitions in (2) yields:

$$v_t (1 - L) = \left(P_t^h - P_t^H \right) (1 - L) \tag{6}$$

Hence, by (3)

$$\widetilde{\eta}_t \;\; = \;\; P_t^h - P_t^H - \widetilde{\rho} \left(P_{t-1}^h - P_{t-1}^H \right) \tag{7}$$

Given a calibrated value for ρ, equation (7) tells us how to estimate σ_η^2, the sample variance of the series $\widetilde{\eta}$. As shown in Figure 2, various boroughs of London did not experience the same housing return fluctuations over the period of our data. The value of σ_η^2 we choose is the sample average across the six boroughs for which we managed to construct consistent time series for the 1989-2000 period: Croydon, Greenwich, Islington, Hammersmith, Richmond, and South Kensington. That is, for each borough we estimate equation (7), calculate the implied σ_η^2 and then take the average across boroughs.

2.2 Moments of Investment Returns

The first two moments of investment returns for each time series are calculated using vector autoregressions with two lags, running from 1977Q1 to 2000Q4.[5] Table 1 reports expected real returns in percentage terms and variances for all seven asset classes.

UK housing provides returns averaging between those of bonds and general stocks with variances also between those of bonds and general stocks. However, relative to the available evidence on France (Lagarenne and le Blanc, 2002), the US (Goetzman, 1993, Gatzlaff , 2000, and Flavin and Yamashita, 2002) and Sweden (Englund et al., 2002), the London housing market displays very strong returns, higher than the returns from real estate stocks or general stocks, an exception. London housing returns are also very volatile. The studies cited above report standard deviations for housing returns often well below half of that of stocks. Here, the standard deviations are of similar

Table 1: Mean and variances of real quarterly asset returns

	Gen. Stocks	R.E. Stocks	t-bills	Bonds	UK HPI	L HPI	Home
	Expected Returns						
	1.64	1.35	0.85	0.97	1.52	1.91	1.91
	Variances (x100, in quarterly terms)						
Horizon							
1 quarter	0.70	1.14	0.054	0.007	0.07	0.14	8.42
10 quarters	0.50	1.41	0.016	0.020	0.21	0.47	1.06
20 quarters	0.50	1.46	0.018	0.019	0.19	0.46	0.76
40 quarters	0.49	1.47	0.019	0.018	0.17	0.45	0.60

magnitude for all relevant horizons. Note that both the standard deviations reported here and the correlations reported in the next table are remarkably stable once we consider an horizon greater than 1 quarter.

Table 2 presents correlation coefficients between investment vehicles for various time horizons. The returns to housing are positively correlated with the returns to stocks. The correlations between housing and t-bills and between housing and bonds are small, the first is negative, the second is positive. The returns to London housing are positively correlated with the returns to UK housing. The correlations between a single London home and any of the two indexes are low at short horizons, and higher at long horizons.

In contrast, Englund et al. (2002) report smaller correlations of housing and stocks and larger negative correlations between housing and bonds, and between housing and t-bills. Working with French data, Lagarenne and le Blanc (2002) also find negative correlations between housing and bonds and between housing and t-bills, but a negative, albeit small, correlation between housing and stocks. Goetzman (1993), Gatzlaff (2000) and Flavin and Yamashita (2002) report small correlations between housing and financial assets, most of them negative. To facilitate the comparison across countries, we reproduce the comparable data from France, Sweden and the US in the Appendix.

Only the studies of Englund et al. and Gatzlaff isolate real estate stocks as a separate class of assets, as we do here. We find a positive correlation between housing returns and real estate stock returns with a magnitude similar to that found by Englund et

Table 2: Correlation between asset returns at different horizons

Horizon	R.E. Stocks	t-bills	Bonds	UK HPI	L. HPI	Home
Gen. Stocks						
1 quarter	0.5936	0.0968	0.1303	0.0687	0.1449	0.0182
10 quarters	0.5488	0.2649	0.4524	0.3391	0.4515	0.3004
20 quarters	0.5550	0.2285	0.4596	0.4122	0.5229	0.4083
40 quarters	0.5564	0.2152	0.4604	0.4516	0.5584	0.4841
R.E. Stocks						
1 quarter	1	-0.0703	0.0485	0.1969	0.1928	0.0243
10 quarters		-0.2389	0.0640	0.5486	0.5313	0.3535
20 quarters		-0.2792	0.0495	0.5826	0.5543	0.4328
40 quarters		-0.2928	0.0418	0.5941	0.5580	0.4838
t-bills						
1 quarter		1	0.8302	-0.0362	0.0383	0.0048
10 quarters			0.6731	-0.1974	-0.1446	-0.0962
20 quarters			0.6354	-0.1313	-0.1070	-0.0835
40 quarters			0.6238	-0.0820	-0.0756	-0.0655
Bonds						
1 quarter			1	0.0764	0.1814	0.0228
10 quarters				-0.0398	0.1041	0.0693
20 quarters				0.0075	0.1421	0.1110
40 quarters				0.0467	0.1706	0.1479
UK HPI						
1 quarter				1	0.8453	0.1063
10 quarters					0.8959	0.5960
20 quarters					0.8835	0.6899
40 quarters					0.8747	0.7583
London HPI						
1 quarter					1	0.1258
10 quarters						0.6653
20 quarters						0.7809
40 quarters						0.8669

al. for Sweden. This result stands in contrast to the negative correlation reported by Gatzlaff for the US.

Overall, our results provide further support for what is becoming a well-established fact in the literature: housing is a potentially attractive asset within a diversified portfolio due to the low correlations of its returns with those of financial assets. When compared to American, French, and Swedish data, London distinguishes itself by the high first and second moments of housing returns relative to those of other financial assets, as well as the higher correlation of housing and stocks returns.[6]

3 Optimal portfolios

We now turn to the analysis of mean-variance efficient portfolios. First, we report optimal allocations of investment to housing, ignoring housing consumption motives. Second, we study the optimal allocations for households with varying levels of wealth invested in their own home. Third, we investigate the gains accruing to these households through access to investments which mimic housing returns for the UK and London.

3.1 No consumption distortion

As a benchmark, we consider the portfolios chosen when no constraint is imposed by housing consumption. The objective of this exercise is to highlight the implications of the return patterns in the data for portfolio choice.

The results are summarized in Figures 3 to 5. Figure 3 displays portfolio weights and the efficient mean-variance frontier for an investor who has access to all seven assets: general stocks, real estate stocks, t-bills, bonds, UK housing, London housing and a London home. The individual minimum variance portfolio appears heavily geared towards t-bills and long-term bonds. For high risk-high return strategies, the efficient portfolio exploits the non-perfect correlation between bonds and t-bills, going long in bonds and short in t-bills. With regards to housing investments, the single home is trivially dominated by the housing price indexes. The unrestricted investor chooses to invest a greater proportion of his net worth in UK housing than in stocks. A small short position in London housing less than compensates the long position in UK housing.

The computations underlying Figure 3 are based on a 40-quarter investment horizon. Assuming a shorter investment horizon does not change our findings. The only difference concerns the investments in t-bills and bonds. At all horizons, the unrestricted investor would still choose to take an overall long position in housing securities. This suggests that renters would benefit from access to housing investments. In the remainder of the paper, all the figures based on London data are constructed assuming a 40-quarter investment horizon.

When investors have access to general stocks, t-bills, bonds and UK housing only, they choose to go short in t-bills in order to buy stocks, bonds and housing, especially at high risk levels (see Figure 4). The same is true when investors can invest in a single home rather than in UK housing (see Figure 5). The source of the difference between Figures 4 and 5 is the extra volatility of the London home. Comparing the figures indicates that this extra volatility is less than compensated for by a higher return relative to the index. This explains why Figure 5 shows greater weights on general stocks and smaller weights on housing.

3.2 Housing consumption distortion

To account for the consumption distortion to the homeowner's portfolio, again we follow Englund et al. (2002) in order to generate comparable results. We compute optimal portfolio allocations conditional on given fractions of wealth invested in housing. Englund et al. (2002) consider four different cases: (1) renter (0% of net wealth in housing), (2) rich homeowner (100% of net wealth in housing), (3) average homeowner (200% of net wealth in housing), and (4) poor homeowner (400% of net wealth in housing). In addition to the restrictions which define each type of investor, we also restrict portfolio shares to be between +400% and -500% as in Englund et al. (2002). This restriction holds throughout the remainder of the paper. Note that the household is left free to choose an optimal source of debt by taking short positions on any asset on which it is allowed.

The top-left panel of Figure 6 displays the constrained mean-variance efficient frontiers for the four different classes of households at a 40-quarter horizon when they are given access to general and real estate stocks, t-bills and bonds with no short sales on stocks. The greater the consumption distortion (the poorer the household), the higher the risk of the minimum variance portfolio. Relative to the unrestricted portfolio, the minimum variance portfolio of every category of household lies below the unconstrained mean-variance frontier. At high risk levels, investing in a London home is worthwhile: households with small portfolio weights on housing holdings cannot reach as high a return as the poor homeowner. Furthermore, the returns on housing are strong enough to make holding a house worthwhile even for poor homeowners: for a same risk level, the portfolio of a poor homeowner yields substantially higher returns than the portfolio of a renter.

For the sake of comparison, we compute equivalent unrestricted and constrained mean-variance efficient frontiers for France, Sweden and the US using the data reproduced in the Appendix (see Figure 6). Differences in underlying data and methodological approach imply that caution should be exercised when comparing the four different graphs, in particular when comparing levels of returns and risk across markets.[7] In all four markets, poor homeowners are forced to hold very risky portfolios. Poor households are typically better off owning their home than renting it, if they are willing to face the implied portfolio risk, except in Sweden. Sweden is the only market where the constrained mean-variance frontiers for the renter coincide with the unrestricted frontier. In all three other markets, exposure to housing is desirable from an investment point of view except for investors looking for minimal risk exposure.

3.3 Hedging housing risk

Figure 8 displays the mean-variance frontier faced by each category of household for the unrestricted portfolio and portfolios with various degrees of access to each of the assets. The most restricted portfolio allows households to take long positions in general and real estate stocks, and both short and long positions in t-bills and bonds. Next, we expand the household's choice by adding alternatively (1) short sales of both real estate and general stocks, (2) the UK index and (3) the London index. Under these last two scenarios, the household is allowed to take both short and long positions in the relevant index and in stocks.

Figure 7 shows that access to housing derivatives provides substantial financial benefits. Households pursuing low-risk strategies benefit from hedging their housing risk. Investors pursuing higher risk strategies benefit from the ability to gain exposure to housing returns. The London housing price derivative dominates the UK housing derivative as a hedge. The reverse is true when the derivative is used as an alternative to direct investment in properties. At short horizons, the correlation between the two housing returns is sufficiently low that both contribute significantly to generating higher returns at all levels of risk.[8]

Short positions in stocks help homeowners limit their risk exposure when they do not have access to the housing price derivative. However, they are a poor substitute, bringing about a modest reduction in risk. In terms of gaining greater returns, short sales of stock are most useful for poor homeowners who are willing to face the risk imposed by owning their home, yet want to make up the loss of returns due to the consumption distortion. In contrast, short sales of stocks prove much more useful in the Swedish context as reported by Englund et al. (2002). Their findings concerning the benefits provided by housing price derivatives are similar to ours.

Finally, the extent to which London housing derivative securities are a redundant asset depends on the horizon of the investor. The high correlation between London and UK housing returns at high horizons implies that, for a Londoner, a London security would mostly represent a redundant asset. Matters are different for shorter horizons (1 to 10 quarters), when the correlations between London and UK returns are lower. These results provide some justification for offering differentiated housing price index derivatives at short horizons.

4 Concluding Remarks

Housing in London has shown a period of strong returns. If the same pattern of returns were to continue, even homeowners with little wealth would benefit from greater expected returns on their portfolio than renters willing to face the same risk exposure while investing solely in financial assets. This does not mean that households would not benefit from access to financial instruments, allowing them to re-balance their exposure to housing risk. In particular, if London housing has generated high rates of returns, it has been at the expense of a high risk exposure. Here lies the primary cost of the portfolio distortion due to the consumption of housing: homeowners are forced to hold a very risky portfolio, in particular those with heavily leveraged positions, such as young first-time buyers.

How could homeowners decrease their risk exposure? Relevant financial instruments have recently become available in the UK. The firms IG Index and City Index have introduced a range of products which enable investors to bet on the changes in house prices. Both firms offer spread bets; i.e., bets on the change between the future index quoted on the day of the transaction and what the underlying price index will be when published (IG Index relies on the monthly Standardised Average Price as reported by the Halifax; City Index relies on the HM Land Registry Quarterly Residential Price Report).[9] IG index offers bets at most two quarters ahead and only on the UK housing price index. City Index provides bets at various levels of disaggregation over horizons ranging from 1 to 7 quarters. Our results provide support for their approach: over such horizons, households should prefer instruments based on local prices rather than a national index. Longer horizons would be even more desirable.

The interest in these bets by the British public has been minimal so far. British financial intermediaries may benefit from the experience of grain market operators in the US. Few farmers use grains futures directly, although they have been available for a long time. However, grain intermediaries use future contracts to hedge forward contracts with farmers. The take-up of such simpler contract (e.g., I buy your crop at this date at this price) is much larger. British homeowners might benefit from similar types of retail products backed by bets on the index; products such as downpayment-savings accounts or insured mortgages whose interest rates are linked to the performance of the local housing market.

Spread bets on offer at the moment seem much more adequate to serve the second type of potential clientele identified by our study, investors pursuing high risk-high return strategies looking for exposure to housing returns without the costs of direct housing transactions and property management.

Acknowledgements

We wish to thank Min Hwang, John Quigley, Jim Shilling, an anonymous referee, and seminar participants at the University of Wisconsin for helpful discussions and suggestions. We are also grateful for the financial support of the Financial Markets Group and the Suntory and Toyota International Centres for Economics and Related Disciplines, LSE.

Notes

1. E.g., Case, Shiller and Weiss (1993), Shiller and Weiss (1999).

2. For a discussion of these and other related issues, see for example Campbell and Viceira (1999), Cocco (2000), Lin and Vandell (2001) and Ortalo-Magné and Rady (2002).

3. We obtained these time series from Datastream.

4. See Table IV in Englund et al. (1998). The value of ρ that we choose is the sample of their average of their estimates for eight Swedish regions. A negative ρ is likely to reflect idiosyncratic aspects of housing transactions completed after a very short holding period; e.g. unusual profit opportunity arising from the fact that the first seller was uninformed of constrained.

5. See Englund et al. (2002) for a detailed description of the methodology.

6. Results are robust to minor changes in the sample periods.

7. See the Appendix.

8. To economize on space, we do not report the corresponding graphs here.

9. For a discussion of the mechanics of spread bets and the details of the products on offer, we refer the reader to the firms' websites: www.cityindex.co.uk and www.igindex.co.uk.

References

Campbell, John Y., and Luis M. Viceira (1999). "Consumption and Portfolio Decisions when Expected Returns are Time Varying," *Quarterly Journal of Economics*, 114(2), 433–95.

Case, Karl E., Robert J. Shiller, and Allan N. Weiss (1993). "Index-Based Futures and Options markets in Real Estate," *Journal of Portfolio Management*, 19(2):83-92.

Cocco, João F. (2000). "Hedging House Price Risk with Incomplete Markets," Institute of Finance and Accounting Working Paper 317, London Business School.

Englund, Peter, John M. Quigley, and Christian L. Redfearn (1998). "Improved Price Indexes for Real Estate: Measuring the Course of Swedish Housing Prices," *Journal of Urban Economics*, 44(2), 171–196.

Englund, Peter, Min Hwang, and John M. Quigley. (2002). "Hedging Housing Risk," *Journal of Real Estate Finance and Economics* 24(1), 163–197.

Flavin, Marjorie, and Takashi Yamashita. (2002). " Owner-Occupied Housing and the Composition of the Household Portfolio," *American Economic Review* 1, 345–362.

Gatzlaff, Dean H. (2000). "The Effect of Single-Family Housing on Multi-Asset Portfolio Allocations," mimeo, Florida State University.

Goetzman, William N. (1993). "The Single Family Home in the Investment Portfolio," *Journal of Real Estate Finance and Economics* 6(3), 201–222

le Blanc, David, and Christine Lagarenne. (2002). "Owner-Occupied Housing and the Composition of the Household Portfolio: the Case of France," mimeo, CREST-INSEE, Paris, France.

Lin, Zhenguo, and Kerry D. Vandell. (2001). "Illiquidity and Real Estate Assets," mimeo, University of Wisconsin Madison.

Ortalo-Magné, François, and Sven Rady (2002): "Tenure Choice and the Riskiness of Non-Housing Consumption," *Journal of Housing Economics*, 11(3), 266–279 .

Shiller, Robert J., and Allan N. Weiss. (1999). "Home Equity Insurance," *Journal of Real Estate Finance and Economics* 19(1), 21–48.

Appendix

Here we report the returns data we used for the comparisons in Section 3. Expected returns and standard deviations are expressed in quarterly terms. In terms of measurements, the data for London is closest to the Swedish data: before-tax returns are computed assuming a 40-quarter investment horizon. The French and US data are after-tax returns computed assuming a 4-quarter investment horizon. Details about data sources are available in the respective papers, Englund et al. (2002), Lagarenne and le Blanc (2002), and Flavin and Yamashita (2002). For their portfolio computations, Flavin and Yamashita rounded the correlations between the individual home and other assets to zero. We do the same for our US portfolio computations here.

Table 3: France, Real quarterly asset returns

France	Stocks	t-bills	Bonds	Home
Expected	2.8000	0.3000	1.2750	1.0250
Std Dev.	3.1000	0.2500	0.7250	0.5875
		Correlations		
Stocks	1.0000			
t-bills	0.2390	1.0000		
Bonds	0.5610	0.3320	1.0000	
Home	-0.1020	-0.5640	-0.3660	1.0000

Table 4: Sweden, Real quarterly asset returns

Sweden	Stocks	t-bills	Bonds	Home
Expected	3.6500	1.2900	1.8600	1.2700
Std Dev.	1.2923	0.1183	0.3317	0.9327
		Correlations		
Stocks	1.0000			
t-bills	-0.0531	1.0000		
Bonds	0.2334	0.6720	1.0000	
Home	-0.0223	-0.2303	-0.3606	1.0000

Table 5: USA, Real quarterly asset returns

USA	Stocks	t-bills	Bonds	Home
Expected	2.0600	-0.0950	0.1500	1.6475
Std Dev.	3.0189	0.5437	1.0504	1.7803
		Correlations		
Stocks	1.0000			
t-bills	0.0191	1.0000		
Bonds	0.1990	0.6854	1.0000	
Home	-0.0052	-0.0192	-0.0056	1.0000

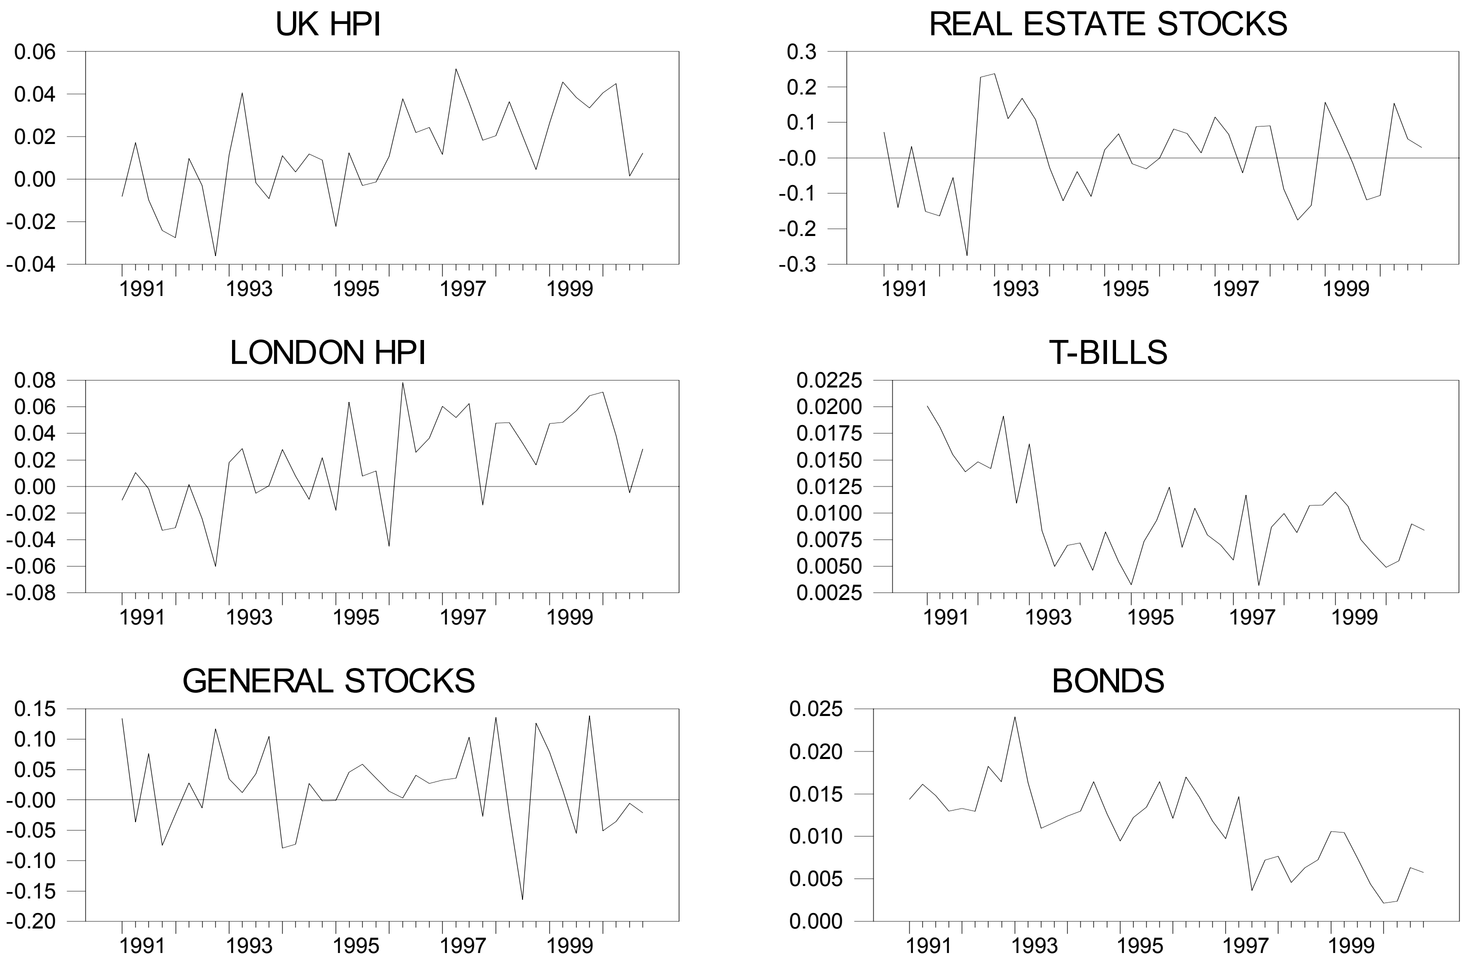

Figure 1. Quarterly Asset Returns

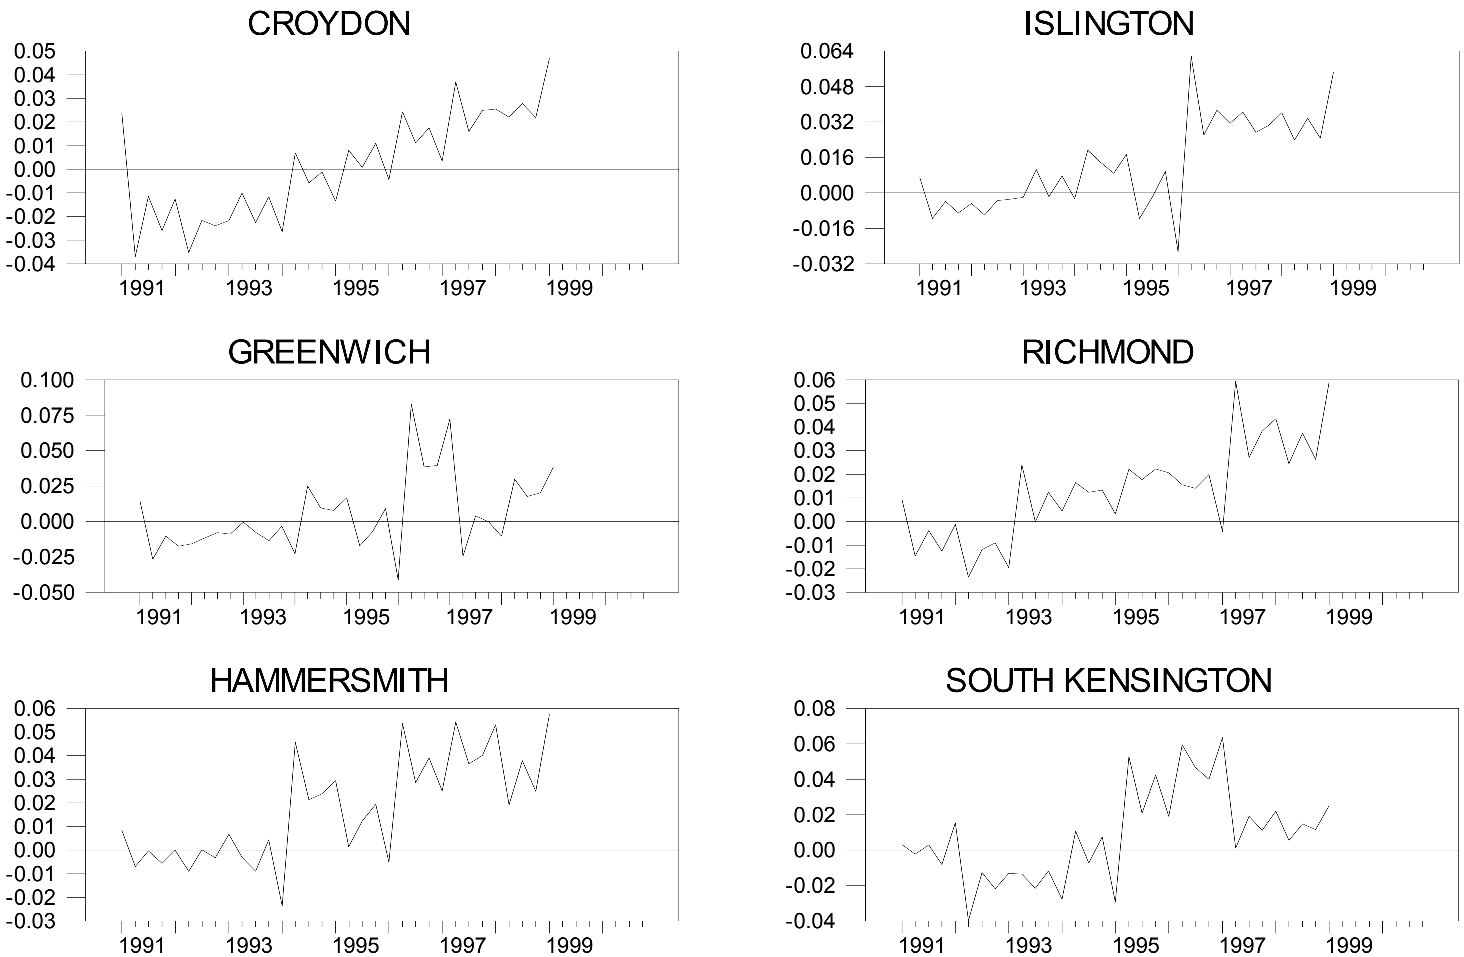

Figure 2. Quarterly Housing Returns

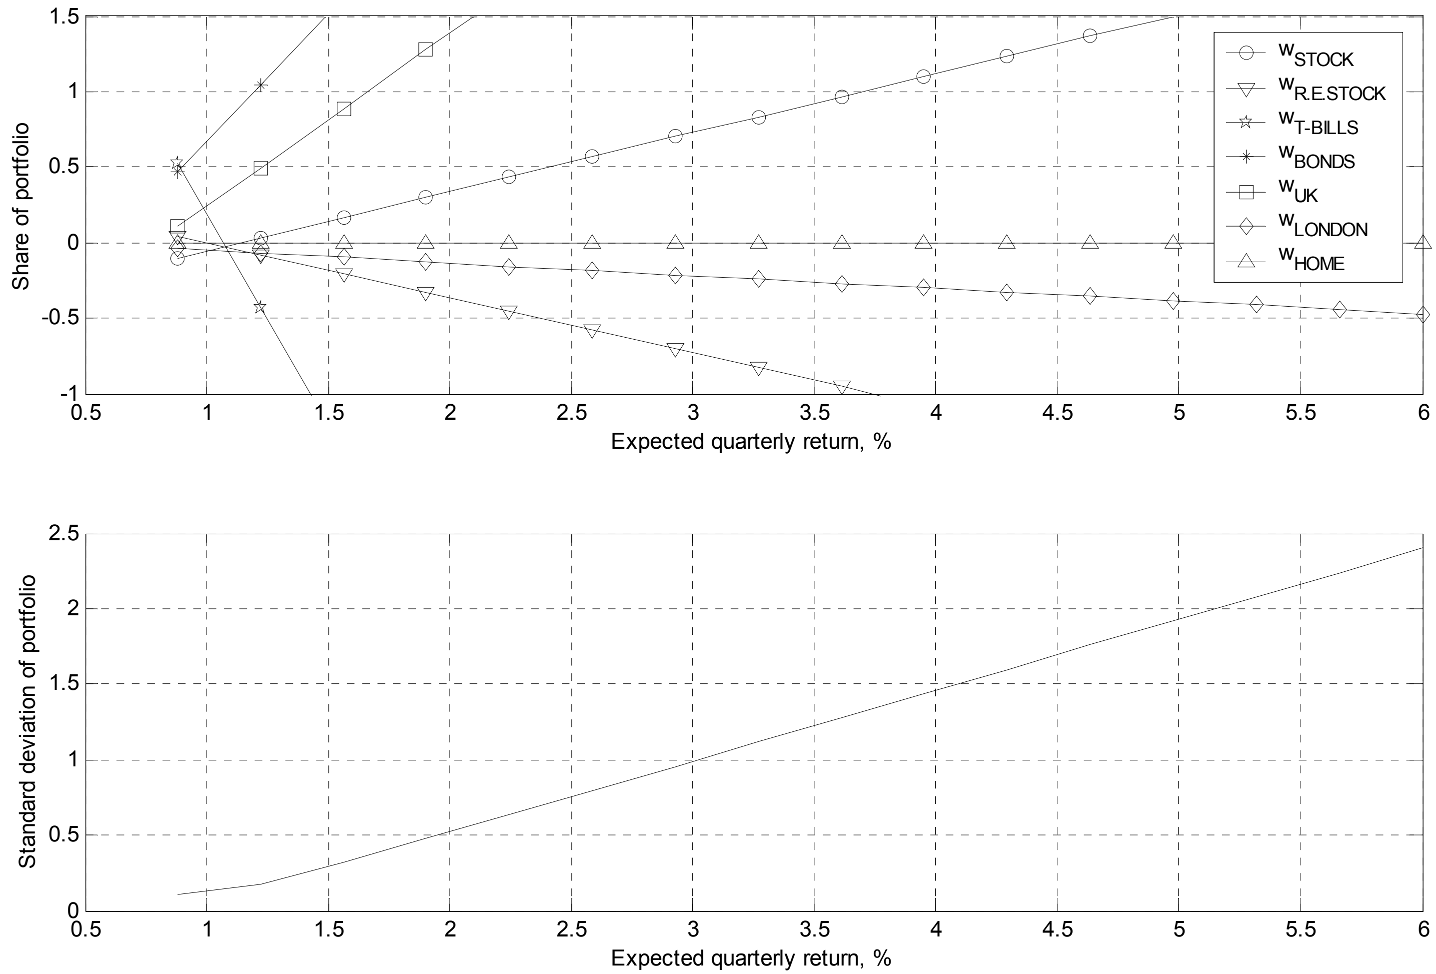

Figure 3. Unrestricted Portfolio, All Assets

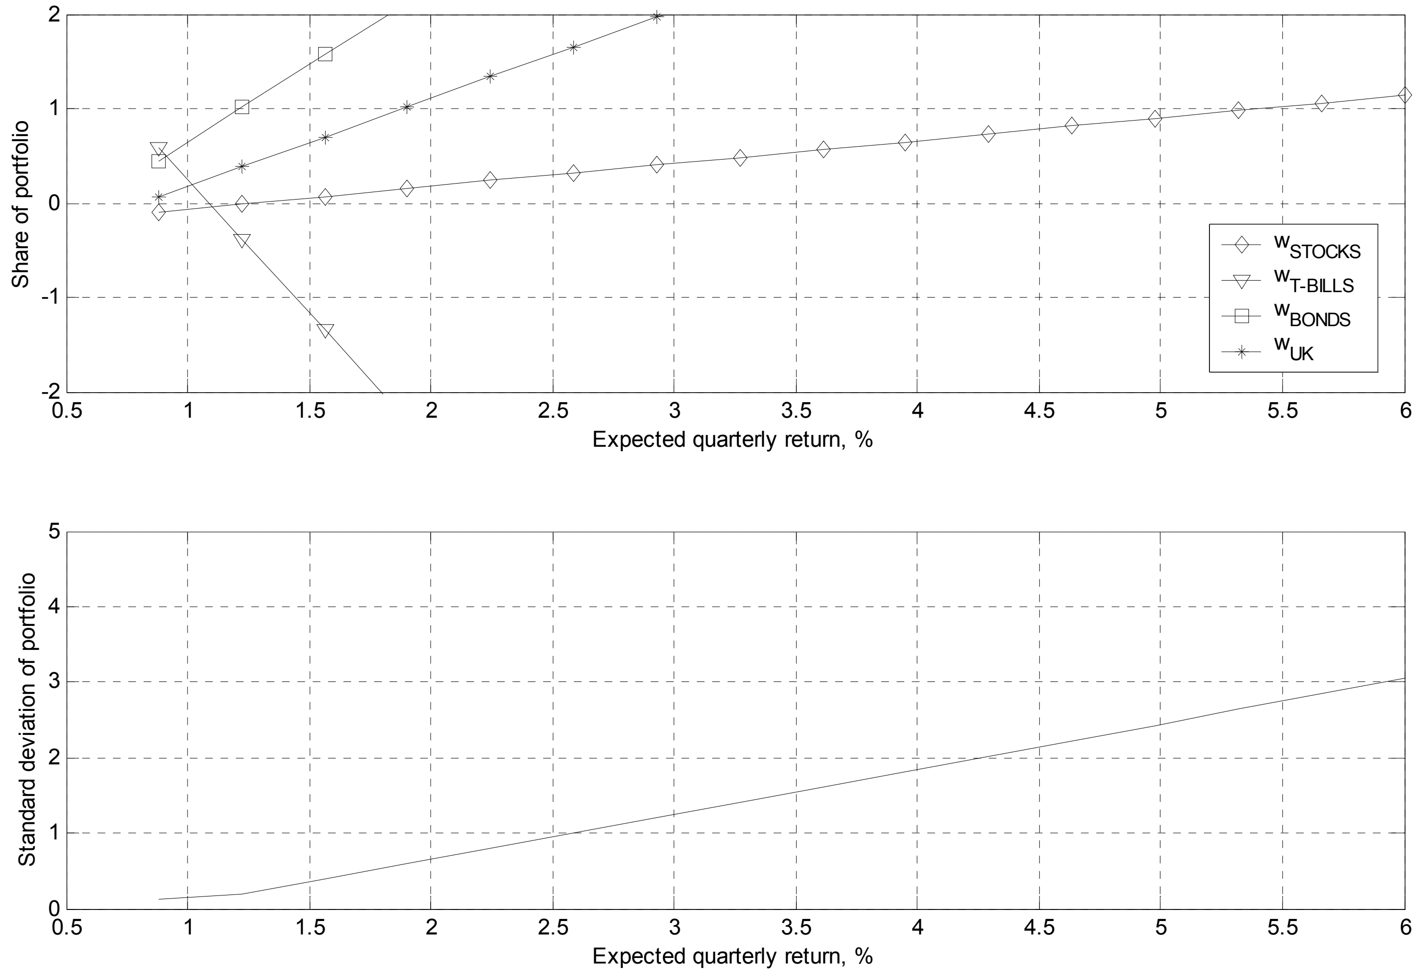

Figure 4. Unrestricted Portfolio Weights, Stocks, *t*-bills, bonds, UK housing index

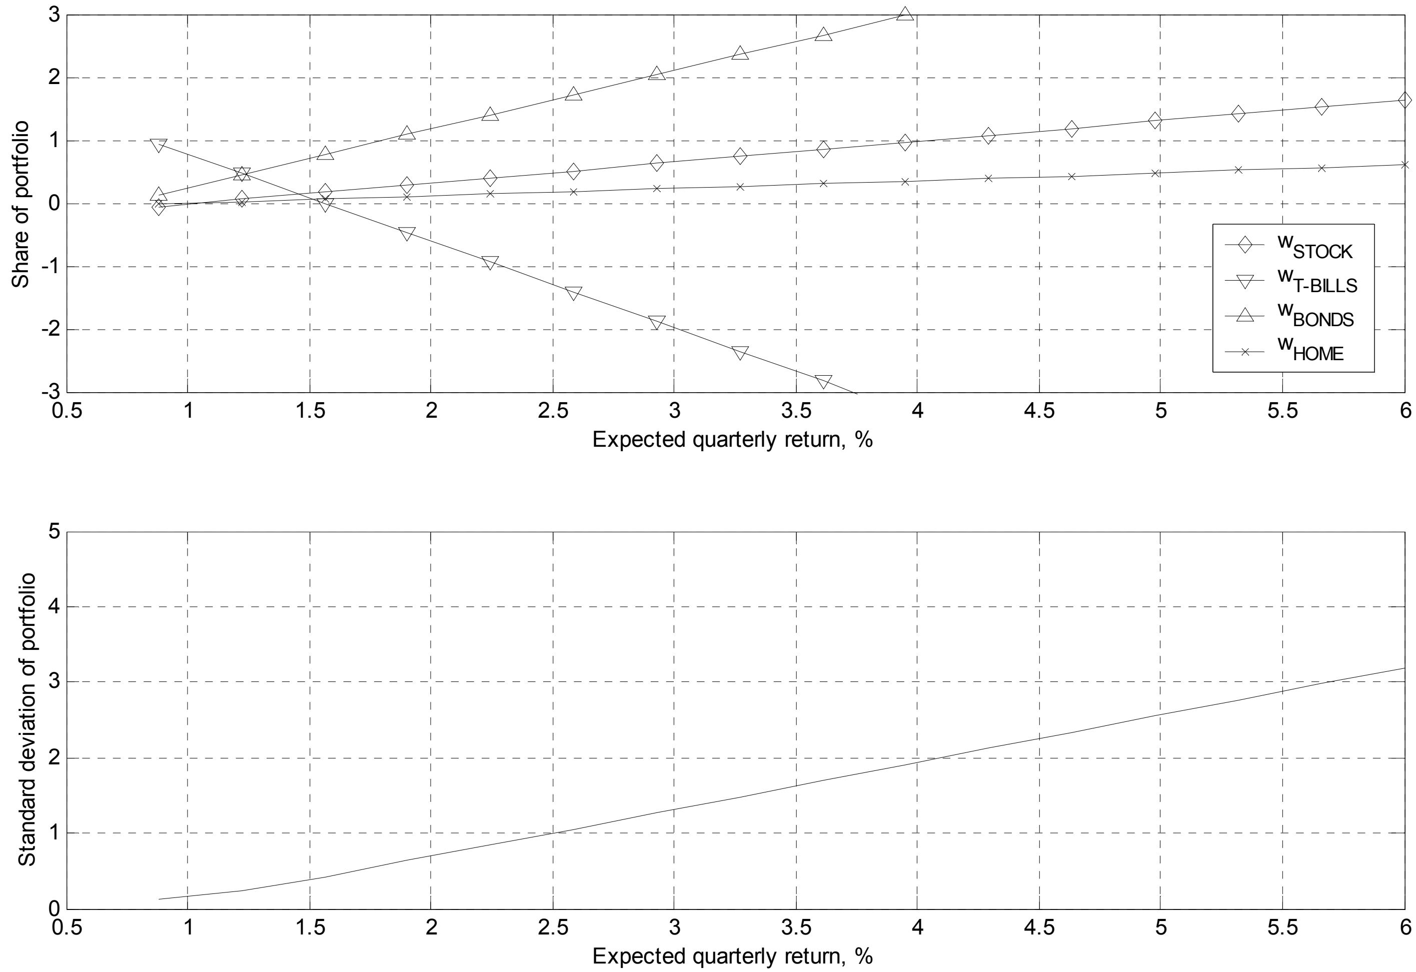

Figure 5. Unrestricted Portfolio Weights, Stocks, *t*-bills, Bonds, London Home

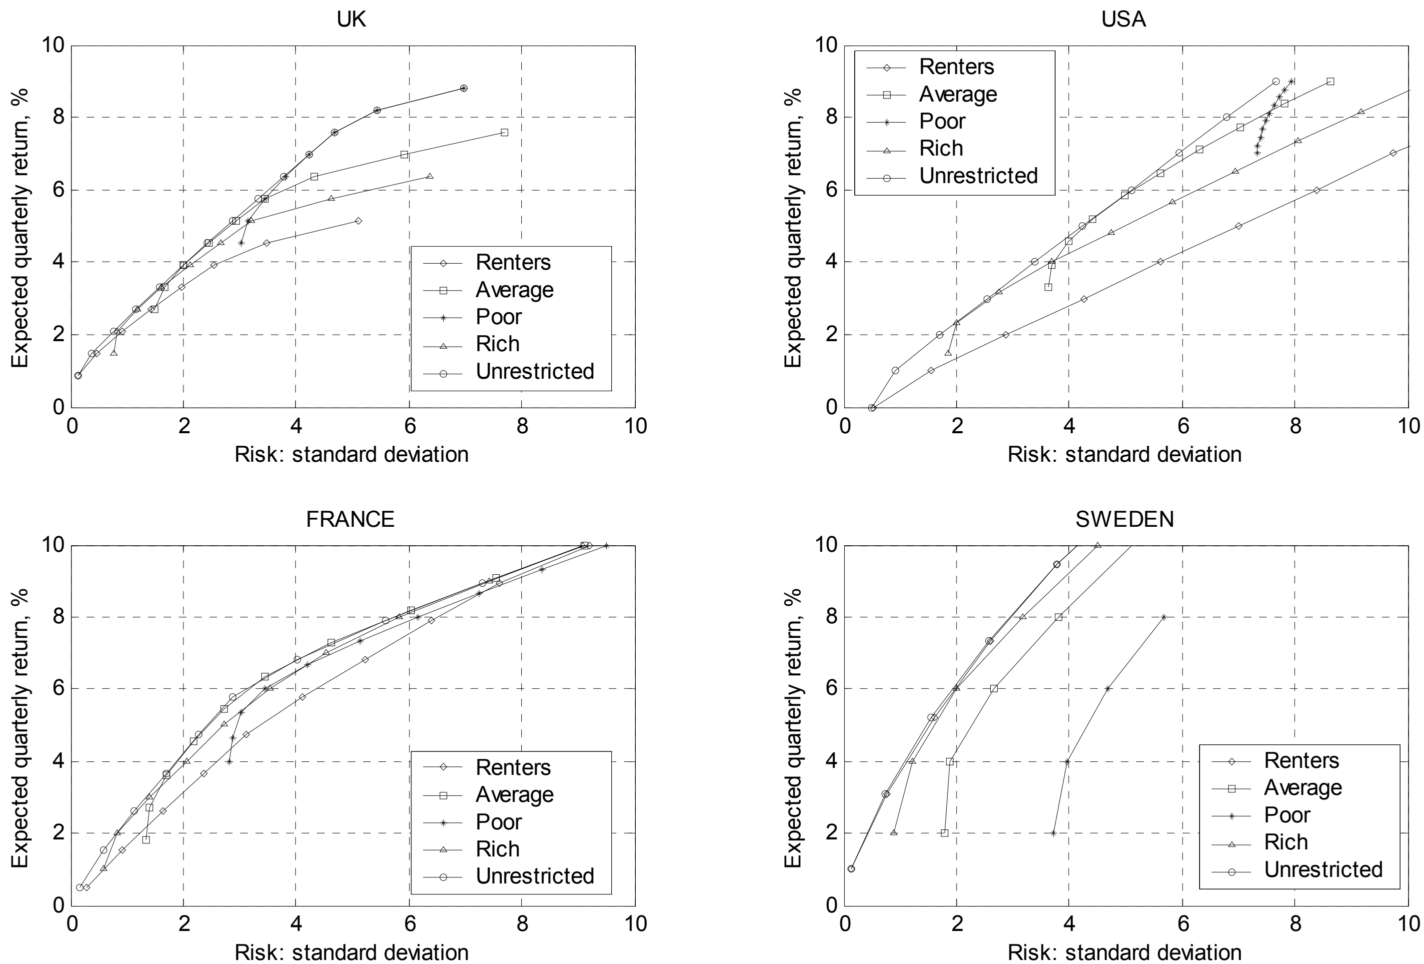

Figure 6. Constrained Mean-Variance Frontiers, Country Comparison

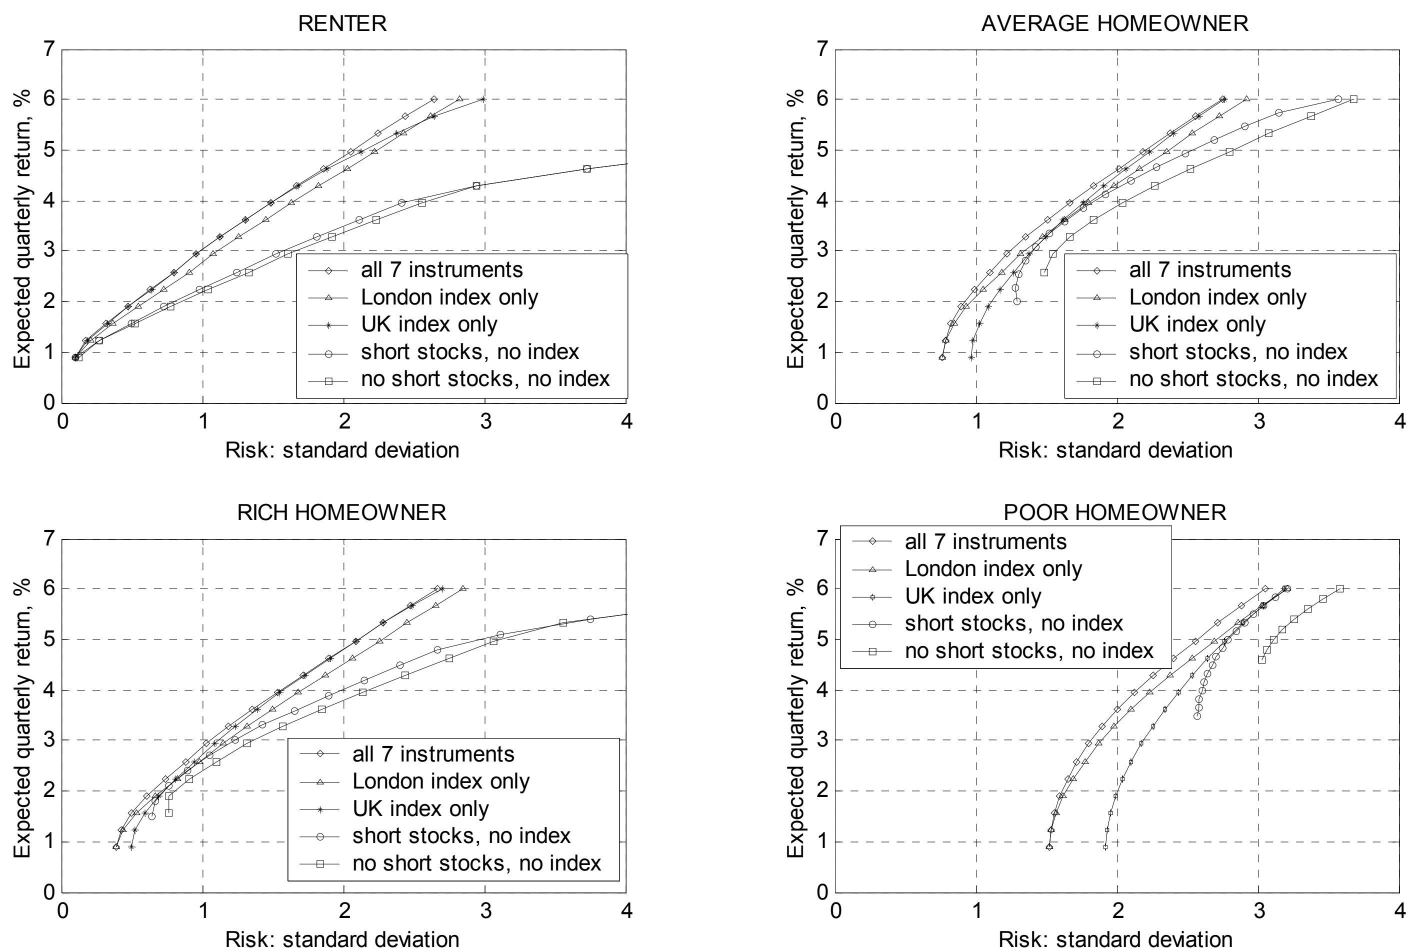

Figure 7. Constrained Mean-Variance Frontiers, Various Assets