2014 V42 3: pp. 662–689
DOI: 10.1111/1540-6229.12023
REAL ESTATE
ECONOMICS

Inflation Protection from Homeownership: Long-Run Evidence, 1814–2008

Dirk Brounen,* Piet Eichholtz,** Stefan Staetmans***
and Marcel Theebe****

This article examines the inflation hedging capacity of the private home. We employ unique long-term data for inflation, house price dynamics and rents for Amsterdam dating back to 1814, allowing us to study total housing returns in different inflation regimes and for varying investment horizons. Our Amsterdam data show that homeownership's protection against actual and expected inflation increases with the investment horizon. This increase is especially strong for horizons up to 10 years. Inflation protection from housing is stronger when inflation is persistent, and the hedging capacities of housing regarding unexpected inflation are weak.

Inflation is among the key risk sources in investment, especially in the long run. Starting with the seminal theoretical work of Fisher (1930), inflation protection has been under continuous attention from researchers. Contrary to the Fisher model's prediction, the empirical evidence suggests that stocks and bonds offer poor protection against inflation, both in the United States[1] and abroad.[2]

Unlike stocks and bonds, housing has hardly been studied from an inflation risk perspective, even though housing generally has a dominant position in the household wealth portfolio. Also, the few available housing studies solely look

*Department of Economics, Tilburg University 5000, LE, Tilburg, The Netherlands or d.brounen@uvt.nl.
**Finance Department, Maastricht University 6200, MD, Maastricht, The Netherlands or p.eichholtz@maastrichtuniversity.nl.
*** Finance Department, Maastricht University 6200, MD, Maastricht, The Netherlands or s.straetmans@maastrichtuniversity.nl.
****Finance Department, University of Amsterdam 1018, TV, Amsterdam, The Netherlands or m.a.j.theebe@uva.nl.

[1]See, for example, Bodie (1976), Nelson (1976), Fama (1975, 1981 and 1990), Fama and Schwert (1977), Feldstein (1980), Geske and Roll (1983), Jaffe and Mandelker (1976) and Lee (1992).

[2]See Solnik (1983) and Gultekin (1983).

at house prices, ignoring rents,[3] and focus on short time periods.[4] However, studying the long-term impact of inflation risk seems the relevant thing to do given the housing owners' relatively long investment horizons. Investigating the inflation-hedge capacity of housing for different horizons (and especially for the longer term) therefore constitutes the main contribution of this study. Moreover, instead of using house prices only, we look at total housing returns by combining a long-term house price index and a long-term market rent index. We use a unique dataset of 195 years of rents, house prices and inflation for Amsterdam, covering a wide variety of market circumstances and inflation regimes.

Studying the long-term hedge properties of housing against inflation risk has become increasingly relevant because of the increased weight of housing in household portfolios. In all Western economies besides Germany and Switzerland, well over half of all households are homeowners.[5] In Europe, housing accounts for 40–60% of total household wealth, opposed to 19% for the average household in the United States. Homeownership is also more widespread than ownership of financial assets like stocks and mutual funds. For example, more U.S. households own a home than that they hold stocks (68% vs. 52% in 2001), and homeownership is more evenly spread across income deciles (Belsky and Prakken 2006). Sinai and Souleles (2005) show convincingly that rent risk is an important motivation for homeownership, and it is likely that inflation and rent changes are linked. However, empirical evidence for the (long-term) relation between total housing returns and inflation is still lacking.

A crucial parameter when investigating the hedging capacity of physical or financial assets is the relevant investment horizon. A typical home-owning household is exposed to the housing market for several decades. Homeowners typically inhabit the same house for a period of 12 years on average. Besides

[3]If changes in the flow of housing services are perfectly correlated with house price changes, or if this service flow is constant, this is not a problem. However, it is doubtful whether that is the case.

[4]Fama and Schwert (1977), Gyourko and Linneman (1988), Bond and Seiler (1998) and Sing and Liow (2000) conclude that house prices provide a hedge against both expected and unexpected inflation, and they use housing returns measured at an annual, quarterly or even monthly frequency. Rubens, Bond and Webb (1989) find residential properties to protect only against unexpected inflation, while Stevenson (2000) finds little evidence for an effective hedge against either expected or unexpected inflation. Anari and Kolari (2002) conclude that U.S. house prices have provided a stable inflation hedge for the period 1968–2000.

[5]See, for a more detailed discussion on international homeownership, Brounen, Neuteboom and Van Dijkhuizen (2006).

that, mortgage contracts last for up to 30 years, and home equity often serves as an implicit pension insurance. These issues suggest that the investment horizon for homeowners is long, further illustrating the need for a study at long and different time horizons.

Unlike for housing, long-horizon inflation hedging potential has been investigated for stocks. Boudoukh and Richardson (1993) analyze U.S. and U.K. stock markets using annual stock price and inflation data from 1802 to 1990 and find that stocks provide a hedge against inflation risk at a five-year horizon, while this is not the case for shorter horizons. Schotman and Schweitzer (2000) determine inflation hedge ratios for different investment horizons and show that stock investors can have positive hedge ratios even if returns are negatively correlated with unexpected inflation shocks and only moderately positively related to expected inflation. Both studies suggest that the degree of inflation protection for long horizons may well be different from the short-term protection.

To investigate whether and to what extent housing investments hedge against inflation over long and varying investment horizons, we will apply the above-mentioned frameworks of Boudoukh and Richardson (1993) and Schotman and Schweitzer (2000) to total housing returns in Amsterdam and inflation between 1814 and 2008.

Anticipating our results, we find that home ownership provides long-run protection against actual and expected inflation, but not against unexpected inflation. Investment horizons matter in this regard: especially between 1 and 10 years, the inflation hedge ratio increases strongly with the horizon. For longer investment horizons, this increase slows down. The hedging potential of houses against inflation also seems to increase with the persistence of inflation. During periods when inflation was not found to be persistent (like in the nineteenth century), we no longer find a positive relation between housing returns and inflation, and the inflation hedge ratio is substantially lower. So, when inflation gets more persistent and investment horizons get longer, our results get closer to the Fisher hypothesis.

The remainder of this article is organized as follows. The Empirical Approaches section discusses the two methodological approaches applied in our subsequent analysis. After discussing the data in the Data section, the Empirical Results section synthesizes the empirical results on long-term inflation hedge capacity of home ownership, both for the 195-year sample period and for two economically relevant subperiods. Finally, the last section summarizes and presents our main conclusions.

Empirical Approaches

In this section, we shortly review regression-based approaches as well as the Schotman and Schweitzer hedge ratio approach toward evaluating the inflation protection potential of housing.

Regression-Based Approaches

The most commonly used version of the Fisher equation states that expected nominal rates of return on risky assets move one-to-one with expected inflation, see, *e.g.*, Boudoukh and Richardson (1993, p. 1347). A straightforward intuitive test of this relationship is to investigate the relation between *actual* risky asset returns and *actual* inflation:

$$R_{t+1} = c_1 + \gamma_1 \pi_{t+1} + \nu_{t+1}, \tag{1}$$

where $R_{t+1} = \ln(\frac{P_{t+1}+HR_{t+1}}{P_t})$ denotes the continuously compounded annual total housing return based on the house prices (*P*) and cash flows in the form of housing rents (*HR*). Accordingly, $\pi_{t+1} = \ln(\frac{CPI_{t+1}}{CPI_t})$reflects the continuously compounded annual inflation rates based on the (log) changes in the consumer price index (CPI).

As homeowners typically live several years in their asset before selling it, we extend Equation (1) by varying the investment horizon from one to *k* years:

$$R_{t+k}^{(k)} = c_k + \gamma_k \pi_{t+k}^{(k)} + \nu_{t+k}^{(k)}, \tag{2}$$

and where $R_{t+k}^{(\mathrm{k})}$, $\pi_{t+k}^{(k)}$ and $\upsilon_{t+k}^{(k)}$ refer to cumulative (multiperiod) total housing returns, inflation rates and disturbance terms from time t to time $t + k$, respectively.[6]

OLS regressions based on (1)–(2) enable one to assess the degree of protection against actual inflation for different investment horizons. However, when using overlapping time windows as in (2), standard errors might be distorted by serially correlated residuals. We therefore employed the Newey–West (1987) method to produce heteroskedasticity and autocorrelation consistent estimates, and we use this method for all regressions throughout the article.

Next to OLS regressions of actual returns on actual inflation, we also want to quantify the relation between actual returns and expected inflation as this

[6]Continuous compounding implies that multiperiod total housing returns, inflation rates and disturbance terms are equal to the sum of their single-period counterparts. For total housing returns, for example, the additivity property boils down to $R_{t+k}^{(k)} = R_{t+1} + R_{t+2} + \cdots + R_{t+k}$. The same additivity property holds for single-period versus multiperiod inflation rates and disturbance terms in Equations (1) and (2).

comes closer to what the Fisher hypothesis is actually about. Boudoukh and Richardson (1993) argue that regressing actual returns on actual inflation like in (1) and (2) at best produces a biased estimate of the hedge potential against *expected* inflation due to the fact that actual inflation equals expected inflation plus measurement error. To solve this errors-in-variables problem, the same authors suggest an IV-GMM estimator that we also apply on our dataset. The advantage of the IV approach is that we do not need to specify an explicit model for expected inflation.

Following Boudoukh and Richardson (1993), we rewrite the multiperiod regression model (2) in the following IV setting:

$$E\left[\left(R_{t+k}^{(k)} - c_k - \gamma_k \pi_{t+k}^{(k)}\right) \otimes Z_{kt}\right] = 0, \quad \forall k \tag{3}$$

where Z_{kt} refer to the instrument set for the multiperiod regression. For the single-period regression, this set of instruments reduces to a constant and the lagged annual inflation rate. Under standard IV assumptions, $\hat{\gamma}_k^{IV}$ will provide consistent estimates of γ_k^{IV} in Equation (3). Instruments have to be selected that are believed to be correlated with expected inflation but uncorrelated with unexpected inflation. Boudoukh and Richardson (1993) selected lagged values of short-term and long-term interest rates and inflation rates as instruments. Due to data limitations, we are limited to using lagged inflation rates as instruments. We estimate equations for the k-year time horizon by using the following sets of instruments:

$$Z_{kt} = \left(1, \sum_{i=1}^{k} \pi_{t-k+i}\right),$$

where Z_{kt} refers to the instrument set for the single-period regression as well as the multiperiod regression.[7] Notice that there is a long tradition in the empirical finance literature on the Fisher equation to use past inflation rates as predictors of future inflation. Fama and Gibbons (1984) compare different forecasting methods for inflation, whereas Nelson (1976) provides an explicit test of the Fisher equation using past inflation rates.

The Schotman and Schweitzer Approach

In order to perform some robustness analysis of the regression-based results, we also perform an alternative approach toward assessing the long-run inflation protection potential of housing investments based on hedge ratios, see

[7]Notice that the summation of single-period inflation rates in the instrument set is the lagged multiperiod inflation rate. For the single-period regressions, the instrument boils down to the lagged annual inflation rate.

Schotman and Schweitzer (2000). The main advantage of this alternative approach is that one only needs single-period parameter estimates using annual (nonoverlapping) data on total housing returns and inflation. In this way, one can avoid the econometric problems that are inherent to long-horizon regressions.

When investing in financial assets, a hedge ratio expresses the number of hedging vehicles required to offset the risk of an unprotected position. Similarly, if a residential property is used as a vehicle to hedge against inflation, one could determine the amount of housing required to protect against inflation by means of a hedge ratio. Schotman and Schweitzer derive a long-term hedge ratio by combining mean-variance portfolio optimization for a stock/bond portfolio with simple driving equations for stock returns and inflation. Without loss of generality, we make the same model assumptions for housing as a risky security. Annual total housing returns are assumed to be driven by

$$R_{t+1} = c + \beta E_t\left[\pi_{t+1}\right] + \phi\left(\pi_{t+1} - E_t\left[\pi_{t+1}\right]\right) + \varepsilon_{t+1}, \tag{4}$$

where R_{t+1} and π_{t+1} are defined in the same way as in Equation (1). Notice that this is the well-known Fama and Schwert (1977) regression, which augments the Fisher relation with unexpected inflation $\eta_{t+1} \equiv \pi_{t+1} - E_t\left[\pi_{t+1}\right]$. Housing returns fluctuate around a constant c and the fluctuations depend on expected inflation $E_t\left[\pi_{t+1}\right]$, unexpected inflation $\eta_{t+1} \equiv \pi_{t+1} - E_t\left[\pi_{t+1}\right]$ (with standard deviation σ_η), and a mean zero disturbance term ε_{t+1} (with standard deviation σ_ε). Unexpected inflation is interpretable as the actual inflation risk. The Fama–Schwert regression basically reflects a "horse race" between the impact of expected and unexpected inflation on risky security returns.

Next, actual inflation is assumed to be generated by an AR(1) time series model:

$$\pi_{t+1} - \mu = \alpha\left(\pi_t - \mu\right) + \eta_{t+1}, \tag{5}$$

with $|\alpha| < 1$ reflecting the first-order serial correlation in inflation. This is the simplest time series model one can think of that reflects the well-documented stylized fact that inflation series are persistent but stationary and predictable. In (5), the current inflation depends on the long-term inflation average μ, the deviation of inflation in the previous period from the long-term mean and an independent shock η_{t+1}. The persistence parameter α reflects how fast actual inflation π_t returns to its long-run average μ in case of a deviation. The disturbance term η_{t+1} is identical to the unexpected inflation component in Equation (4).

Finally, using these two driving equations for housing returns and inflation, we can evaluate the multiperiod hedge potential of housing by deriving the demand for housing in a mean-variance framework. We consider a portfolio where a fraction w is invested in housing, and $(1 - w)$ is invested in a nominally risk-free discount bond of maturity k. The bond's maturity equals the horizon of the investor (Schotman and Schweitzer 2000). The housing demand that maximizes the investor's objective function over the portfolio constraint consists of two parts: a first term that represents the demand for housing resulting from the housing risk premium and a second term that reflects the hedging demand for housing depending on the covariance between nominal multiperiod housing returns with multiperiod inflation. The hedging demand (or hedge ratio) is[8]

$$\Delta^{(k)} \equiv \frac{\operatorname{cov}\left[R_{t+k}^{(k)}, \pi_{t+k}^{(k)}\right]}{\operatorname{var}\left[R_{t+k}^{(k)}\right]} = \frac{(k\phi + \psi_{1k}(\phi + \beta) + \psi_{2k}\beta)\,\sigma_\eta^2}{k\sigma_\varepsilon^2 + \left(k\phi^2 + 2\psi_{1k}\phi\beta + \psi_{2k}\beta^2\right)\sigma_\eta^2}, \tag{6}$$

with

$$\psi_{1k} = \frac{\alpha}{1-\alpha}\left(k - \frac{1-\alpha^k}{1-\alpha}\right)$$

$$\psi_{2k} = \frac{\alpha^2}{(1-\alpha)^2}\left(k - 2\frac{1-\alpha^k}{1-\alpha} + \frac{1-\alpha^{2k}}{1-\alpha^2}\right).$$

Clearly, the higher this (scaled) covariance term between long-horizon housing returns and inflation, the better housing protects against inflation that implies a higher hedge ratio and a higher investment weight w in housing. But what determines the level of the hedge ratio and its changes over time? The formula above shows that it is a nonlinear function of the inflation persistence parameter α, the expected inflation hedge parameter β, the unexpected inflation parameter φ and the variances of the error terms of the driving equations (4)–(5). For nonoverlapping annual data ($k = 1$), it can be easily shown that the single-period hedge ratio $\Delta^{(1)}$ only depends on the coefficient for unexpected inflation φ but not on the Fisher coefficient β or the inflation persistence parameter α. For longer time horizons ($k > 1$), however, the Fisher coefficient β and the inflation persistence α also start to influence the hedge ratio. In fact, Schotman and Schweitzer (2000) graphically show that their influence dominates the short-term inflation risk effect φ, provided that α and β are sufficiently big. More specifically, they show that $\Delta^{(k)}$ increases in α (given β, φ and k) and in β (given α, φ and k). Also, $\Delta^{(k)}$ increases in k provided that α and β are not too small. These differing horizon effects induced by the parameters φ, β and α can be understood by the fact that unexpected inflation shocks η_{t+1} both have static and dynamic effects on housing returns and inflation via (4)–(5). On the one hand, Equation (4) shows an immediate static impact of unexpected inflation

[8] See Schotman and Schweitzer (2000) for the formal derivation.

shocks on contemporaneous housing returns via the parameter φ. However, the persistence of inflation, as reflected by Equation (5), implies that current inflationary shocks should also exhibit a dynamic impact on all future inflation rates. As a result, the inflationary expectations in (4) need to be updated using the news of the unexpected inflationary shock. This updating, in turn, impacts future housing returns via Equation (4). The higher the inflation persistence, the more the expectation revisions accumulate over the investment horizon k.

Summarizing the model predicts that the hedging properties of housing improve with the horizon k when inflation is persistent (high value of α) and when there is at least a partial feedback from expected inflation to nominal housing returns (high value of β). Moreover, the hedging properties also increase for given time horizons k when the inflation persistence (or the Fisher effect) increases.

As our dataset comprises nearly two centuries, the presence of structural breaks in the data is not unlikely. As this will hamper the interpretation of regression outcomes, we test for structural breaks by applying a regime shift model, which cuts our total sample into two distinct eras. These differ with respect to their moving average and volatility of inflation. As the periods 1814–1914 and 1915–2008 appear to be very different in data terms, we will perform all our subsequent analyses both for the full sample as well as for the two separate subperiods.

The Data

To investigate the ability of home ownership to protect against inflation risk, we start with an index of consumer prices, based on a careful combination of sources. To get a complete picture of the total return derived from home ownership, we analyze a total return series, including capital appreciation and (imputed) rental income. For each of the indices used in this article, we discuss sources, construction and time series behavior below.

Consumer Prices

The first long-run data series we need concerns consumer prices. No single index exists that reflects all the 195 years we study. We therefore use different sources to construct one. Van Zanden (2005) is the source for the development of the general consumer price level until 1850. This index is based on a basket of consumer goods, including rye bread, beer, butter, meat, potatoes, peas, different types of fish and various textiles. This basket changes with the broad societal use of the products. For the period between 1850 and 1900, we employ Van Riel (2006), who uses a similar basket of goods and adds housing rental expenses. From 1900 onwards, we use the CPI calculated by the Dutch

Table 1 ■ Characteristics of nominal changes in house prices, rents and CPI.

Time Period	Inflation		Clean Price Returns		Rent Changes	
	Mean (%)	Std. (%)	Mean (%)	Std. (%)	Mean (%)	Std. (%)
Full Sample	1.65	7.86	2.52	19.45	2.09	4.22
1814–1914	0.38	5.91	1.00	14.46	0.22	4.06
1915–2008	3.00	9.36	4.12	23.61	4.08	3.42
1814–1850	−0.00	8.34	0.91	18.01	0.34	4.34
1851–1914	0.84	3.95	1.05	12.18	0.16	3.93
1915–1939	−0.00	14.98	0.62	34.36	2.48	3.21
1940–1980	5.51	7.30	6.21	23.02	5.17	3.99
1981–2008	2.38	1.55	4.20	8.15	3.90	1.81

Note: This table presents our data regarding the Dutch inflation rates, nominal house price changes and rent dynamics for the full sample period 1814–2008. Inflation numbers are derived from general price indices that cover the Dutch market over different periods. Until 1850, these data have been compiled by Van Zanden (2005), for the period between 1850 and 1900, we employ Van Riel (2006), and from 1900 onwards, our inflation data originate from the Dutch Bureau of Statistics. The house price data in our sample are based on housing transaction data for the Herengracht, which are compiled into a repeated-measures regression index. After 1973, house price series are taken from the Dutch Association of Real Estate Agents. Finally, the market rent series comes from three sources; until 1850 data are taken from Eichholtz, Staetmans and Theebe (2012), complemented by data from Van Riel (2006) and the Dutch Bureau of Statistics for the later years.

Central Bureau of Statistics (CBS), which is based on a much broader (and also changing) basket of commonly used consumer goods.

Descriptive statistics for the resulting annual inflation series are presented in the left panel of Table 1. We provide the development of the annual inflation rate graphically in Figure 1. The annual inflation rate for the complete period is 1.65%, with a 7.86% standard deviation. However, the average inflation rate differed strongly in different periods. We have divided our 195-year sample period into five subperiods, based on inflation regime and economic circumstances. The first period runs from 1814 through 1850. The latter year is generally regarded as an important turning point for the Dutch economy, which was very slow to recover from the Napoleonic era and the collapse of the Dutch Republic and started industrializing around that time.[9] The stagnant shape of the Dutch economy is reflected in the development of the general price level, with an average annual inflation rate close to zero, but the high standard

[9] See, for example, De Jonge (1968) and Wintle (2000).

Figure 1 ■ Inflation, 1814–2008.

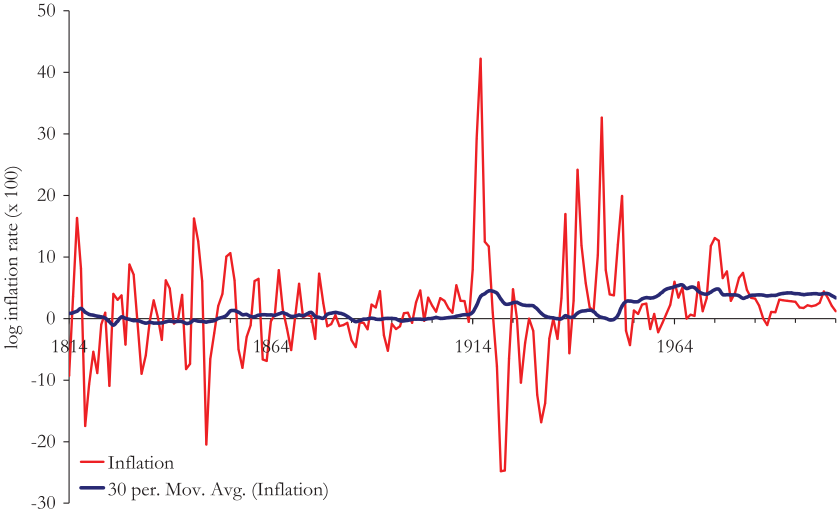

Note: This figure graphically presents the development of inflation in the Dutch economy since 1814. Inflation numbers are derived from general price indices for different periods. Until 1850, these data have been compiled by Van Zanden (2005), for the period between 1850 and 1900 we employ Van Riel (2006) and from 1900 onwards our inflation data originate from the Dutch Central Bureau of Statistics.

deviation and the graph show that the price level was not at all stable in that period.

The next period runs from 1851 through 1914, when the gold standard was—temporarily—abolished for the first time. This period had a considerably higher average annual price growth, but at 0.84%, it is still quite low by today's standards. The period from 1915 until 1939 had a very low average inflation rate, but this was a result of very high inflation during the First World War and severe deflation in the years directly after that. This severe price reversal was the result of the abolishment and subsequent reintroduction of the gold standard, while the deflation in the 1930s coincides with the depression. These developments are reflected in an unprecedented standard deviation of 14.98%. The graph confirms the very high volatility of consumer price in that period. As the graph shows, the persistence and level of inflation during the period from 1940 to 1980 is a historic exception, given the annual average level of 5.51%. The standard deviation is in line with the long-term average. The last subperiod we distinguish runs from 1981—when monetary policy became structurally and institutionally focused on inflation control—to 2006. For this period, the annual average inflation rate is 2.38%, with a record low standard deviation of 1.55%.

House Prices

Regarding house prices for the period 1814–1965, we use an index based on the same housing transactions data as Eichholtz' (1997) biennial Herengracht index. This series is based on transactions of the houses on the Herengracht, the main canal in Amsterdam, and it is constructed using repeated-measures regression. We construct an index for the full time period for which transaction price is available: 1650 through 1973. Sparse data in the first and last years of the Herengracht transaction price sample period forced Eichholtz to estimate a biennial index. We only use the index estimates for 1814–1965, allowing us to construct an annual index. For the remaining period from 1973 through 2008, the house price series are based on a local index constructed by the Dutch Association of Real Estate Agents (NVM) and the Central Bureau of Statistics (CBS). This index is a mean index of Amsterdam house transaction prices. The middle panel of Table 1 provides descriptive statistics for annual changes in house prices. The table shows that nominal house price increases were far lower during our nineteenth-century subsample than in the more recent subperiod, *i.e.*, an annual average of 1% against 4.12% for the twentieth century. It also shows that nominal house price growth was especially strong in the period between 1940 and 1980.

The house price index is provided graphically in Figure 2, both in nominal terms and in real terms, on a logarithmic scale. The graph shows that nominal Amsterdam house prices only grew substantially in the twentieth century. Real rents show far less growth. The gap between the nominal and real price series is very small between 1814 and approximately 1900. In the years preceding the First World War, nominal and real house price began to diverge, and this divergence got especially strong from 1914 onwards, only to be reversed in the decade after the war. From the end of the Second World War, the divergence between nominal and real house prices became structural. In nominal terms, the house price index used in this article increased from 100 in 1814 to 11,669 in 2008. In real terms, the index value is 510 in 2008. So, a very large part of the observed price changes was caused by inflation.

Rents

The second component of the total return to housing is the rental income. We do not have rental income for the houses on the Herengracht, since these were mostly owner-occupied. The rental income series therefore has to be based on market rents, which are subsequently related to house prices to get

Figure 2 ■ Nominal and real house prices, 1814–2008.

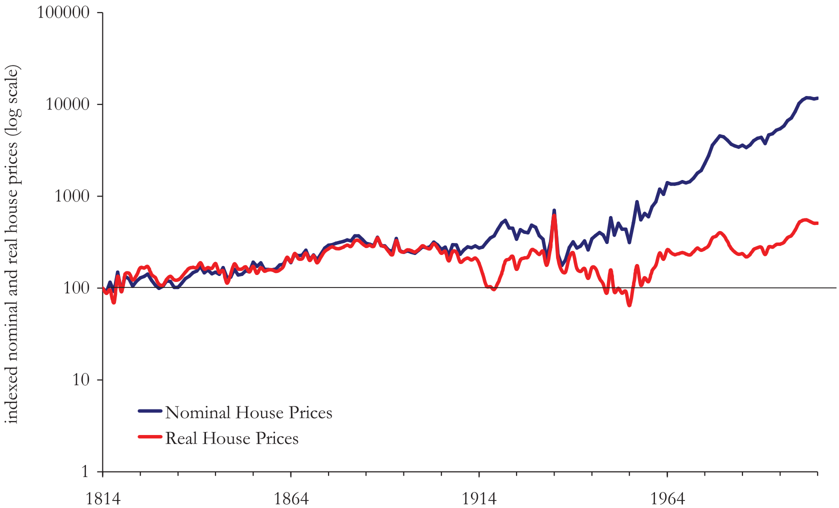

Note: This figure presents the house price dynamics both in real and nominal terms for the period 1814–2008 based on the Amsterdam Herengracht index. Inflation numbers are derived from general price indices that cover the Dutch market over different periods. Until 1850, these data have been compiled by Van Zanden (2005), for the period between 1850 and 1900 we employ Van Riel (2006) and from 1900 onward, our inflation data originate from the Dutch Central Bureau of Statistics. The house price data in our sample are based on housing transactions data for the Herengracht, which are compiled into an index using repeated-measures regressions. After 1973, house price series are taken from the Dutch Association of Real Estate Agents.

an estimate of rental income. The market rent series comes from different sources. For the period 1814–1850, we use a repeated rent index developed by Eichholtz, Straetmans and Theebe (2012). That index is constructed from a sample of 1,055 dwellings owned by the precursors of today's institutional investors: orphanages, hospitals and poor-relief boards. Only cash flows of new rental contracts were used to create the market rent index, which was based on repeated-measures regression. The resulting index is truly market-based as government interference in the rental housing market was nonexistent before 1914.

For the period after 1850, the same rental data are not available, and we therefore have to turn to other sources to extend the market rent index to 2008. As far as we know, no rental data are available for Amsterdam for the complete 1850–2008 period, so we are forced to use national data instead. Van Riel (2006) has estimated market rent developments for the period 1850–1913. He has done that by estimating the gross rental value of houses on the basis of the wealth tax that was levied on the estimated (imputed) rental income of Dutch houses, which was in its turn based on the actual rents of comparable houses in the vicinity. Tax collectors kept up with the annual development of rents (Horlings 1995). Van Riel's rent index only involves residential properties and includes maintenance and repair costs.

The Dutch Central Bureau of Statistics' (CBS) collection of national rent data started in 1914. We use five different CBS sources to construct a rent index for the period 1914–2008. Especially in the first decades of the twentieth century, a large number of different sources had to be used to collect the appropriate data. For the decades after the Van Riel index, rent data are from CBS (1939, 1948), which have been calculated in the same manner as Van Riel. From 1937 onwards, the CBS has published an index of the national rent development. To show the exact movement of the rent, the CBS rebases the index every five years. Because this reestimation procedure can create flattening of the index in the first years, we go back to the earliest source possible (CBS 1959) to provide the rent data of this period. It was fortunate that the CBS provided a long time series on the rent development in the Netherlands. For the first decades after 1937, the CBS continued to use the same method for index calculation (CBS 1949), but they replaced that by a panel survey in 1973 (CBS 2002). The panel survey method takes into account quality changes so that the rent index mirrors the pure rent development. This change in index methodology may affect the index' behavior, but the CBS (1999) has compared the two index methods and concluded that the indices differ by a stable factor. So, the two different concepts can be used in one continuous rent index.

To conclude, we have an index covering nearly two centuries, but it has two drawbacks. The first is that it is based on three different construction methods: repeated observed rent regression (1814–1850), measurement of imputed rent by looking at comparable market rents (1851–1973) and survey techniques (1974–2006). The second drawback is that the index is partly based on data for Amsterdam (1814–1850) and partly on data for the Netherlands as a whole (1851–2006).

Concerning the first issue, we already quoted the CBS (1999) study, indicating that the second and third index construction methods yield equivalent results. By comparing the statistical properties of the rent index until 1850 with that from 1851 onwards, we can judge the correspondence between the first method and the other two.

The second issue, the fact that we look at Amsterdam rents between 1814 and 1850 and national rents afterwards, is only a problem if Amsterdam's rents have behaved very differently than national rents. We are fortunate to have two periods for which rent developments in Amsterdam are available. One CBS study (1938) provides rent data for 1921–1936, and the correlation between this index and the national rent index is 0.98. Furthermore, an overview of the rent development of period 1995–2006 was provided by the Department of Housing of the city of Amsterdam (2006). Similar to the CBS (1938) study, they conducted a panel survey to measure the rent development of all rents in Amsterdam. The correlation between this rent index and the national rent index of the CBS was 0.99. With caution due to the short time periods, it can be concluded that the Amsterdam rent development and the national rent development display similar patterns.

The right panel of Table 1 includes descriptive statistics of annual changes in the market rent index. Between 1814 and 2008, market rents have grown 2.09% per year, on average. As we saw for house prices, the rent increase was considerably higher during the second subperiod covering most of the twentieth century (4.08% per annum) than it had been in the preceding century (0.22% per annum). The standard deviation of the rent changes is rather constant across time. For the full sample period, we observe a standard deviation of 4.22%, which is very similar to the values we find for the two regime periods before and after 1914 (4.06% vs. 3.42%). Regarding the five historic subperiods, our results show that in line with inflation rents increased most fiercely since World War II. Figure 3 provides graphic representations of the real and nominal rent indices.

The Total Return Index

The last step in creating a total return index for housing involves the link between the price and the rent index. The capital component of the annual total return is, of course, the annual percentage change in the price index. Next, we derive the income component of total returns by reconstructing the dynamics of Amsterdam rent levels, which were capitalized on house values. This is performed using various capitalization rates, which were retrieved from market reports of the last 20 years. By rescaling these rental returns into the total return index, we obtained a series, which was then again compared to the total return series that are available since 1978 for the Dutch housing market. This procedure resulted in a series that has a 0.92 correlation with an appraisal-based (but unsmoothed) total housing return index for housing since 1978.[10]

[10]Besides the analysis in total returns, we have done the full analysis for capital returns only, and the results do not materially differ.

Figure 3 ■ Nominal and real rent indices, 1814–2008.

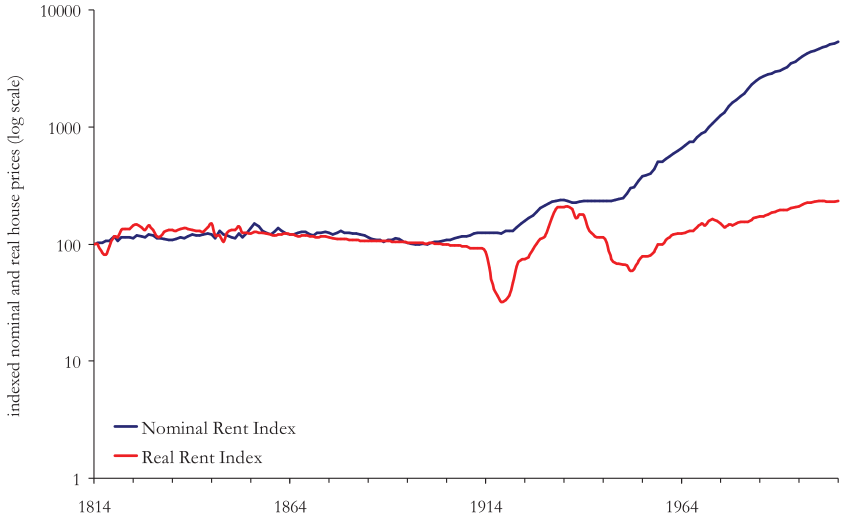

Note: This figure presents the development of both the nominal and real rent index for the Herengracht for the period 1814–2008. Inflation numbers are derived from general price indices that cover the Dutch market over different periods. Until 1850, these data have been compiled by Van Zanden (2005), for the period between 1850 and 1900 we employ Van Riel (2006) and from 1900 onward, our inflation data originate from the Dutch Central Bureau of Statistics. The market rent series comes from three sources; until 1850 data are taken from Eichholtz, Straetmans and Theebe (2012), complemented by data from Van Riel (2006) and the Dutch Central Bureau of Statistics for the later years.

Before proceeding with the empirical analysis on the hedging properties of housing investments, we shortly investigate the stationarity properties of total housing returns and inflation, the two series of interest in the empirical section. Stationarity is crucial because the validity of the empirical approaches to be implemented hinges upon the stationarity assumption.[11] We investigate stationarity of the mentioned series by doing Augmented Dickey-Fuller (ADF) unit root tests and by calculating serial correlations for different lag lengths. Estimation and testing results are summarized in Table 2. To stay consistent with Table 1, we both report full sample results and results for the two main subsamples (taking 1914 as breakpoint) because inflation seems to behave differently across these two subsamples. The upper panel contains the outcomes of the unit root tests, whereas the lower panel reports the serial correlations

[11] Regressions like (1), (2), (4) and (5) render spurious outcomes if regressors and regressands contain unit roots.

Table 2 ■ Unit root tests and serial correlations for housing returns and inflation.

	Full Sample		1814–1914		1915–2008	
	R_t	π_t	R_t	π_t	R_t	π_t
Panel A: Unit Root Tests (Augmented Dickey-Fuller) for the Null Hypothesis of a Unit Root						
ADF	-13.11^{***}	-8.50^{***}	-17.26^{***}	-8.64^{***}	-9.47^{***}	-5.25^{***}
CV(10%)	−2.57		−2.58		−2.58	
CV(5%)	−2.88		−2.89		−2.89	
CV(1%)	−3.46		−3.50		−3.50	
p-value	0.000	0.000	0.000	0.000	0.000	0.000
Lags	1	1	0	1	1	0
Panel B: Serial Correlations in Housing Returns and Inflation						
First Order	-0.356^{***}	0.650^{***}	-0.479^{***}	0.332^{***}	-0.328^{**}	0.757^{***}
	(0.099)	(0.096)	(0.125)	(0.099)	(0.137)	(0.153)
Second Order	−0.223	-0.238^{***}	−0.066	-0.227^{**}	−0.298	-0.319^{**}
	(0.166)	(0.077)	(0.114)	(0.103)	(0.196)	(0.130)
Third Order	−0.127	0.067	−0.025	−0.108	−0.168	0.151
	(0.091)	(0.103)	(0.096)	(0.167)	(0.091)	(0.135)
Fourth Order	−0.038	0.027	−0.119	−0.043	−0.036	−0.022
	(0.058)	(0.106)	(0.091)	(0.105)	(0.088)	(0.131)

Note: This table presents serial correlations for different lag lengths and Augmented Dickey–Fuller (ADF) unit root tests for annual total housing returns and inflation. Unit root testing outcomes are reported in panel A. ADF regressions are run with an intercept but without a trend term. The ADF test statistic is reported in the first row of panel A with corresponding ADF critical values just below. Critical values are the same for a given (sub) sample size. (McKinnon) p-values are reported below the critical values. The number of significant ADF regression lags is reported in the last row and determined using the Schwarz information criterion. The lower panel B reports serial correlation estimates up to four lags. Given the serial correlation and heteroskedasticity present in annual housing return and inflation series, we complement the serial correlation estimates with Newey–West HAC standard errors (in parentheses below each estimate). Regression and testing outcomes are reported for the full sample and two subsamples. Statistically significant serial correlation estimates and rejections of the unit root hypothesis are marked with *, ** and *** to indicate statistical significance at 90%, 95% and 99%, respectively.

for an AR(4) linear time series model. The ADF unit root testing regressions include an intercept but no trend because the two time series seem to fluctuate around a nonzero mean but do not seem to exhibit trending behavior. Upon comparing the ADF test statistic with the ADF critical values, we find a rejection of the unit root null hypothesis at the 1% level for both series and for both the full sample and the two subsamples. This is further confirmed by the extremely low p-values, and one can safely conclude that both total housing returns and inflation are stationary regardless of the considered sample.[12] The lower panel of Table 2 shows that higher order serial correlation estimates are statistically significant up to second order but we do not find evidence of higher order dynamics. The results show that inflation—although stationary—is quite persistent and that the persistence after 1915 is much higher than in the first part of the sample. We argued earlier in the methodology section that the inflation persistence is likely to have important implications for the long-term hedge potential of housing investments against inflation.

Finally, we tested for the existence of Granger causality between total housing returns and inflation. The Granger causality tests are again performed for the full sample and the two earlier considered subsamples and for different lag lengths ($p = 1, 2, 3$).

The Granger causality testing outcomes in Table 3 provide some empirical evidence that inflation Granger causes total housing returns over the full sample. However, Granger causality from inflation to housing returns disappears for the considered subsamples. Also, we do not find any evidence of reverse causality from housing returns to inflation for any sample period. These outcomes give credence to the Schotman and Schweitzer approach, so we proceed to implement it in the empirical section.[13]

Empirical Results

The regression-based OLS and IV-GMM outcomes for the Fisher equation are presented and discussed below, and the subsequent subsection contains a discussion of the outcomes of our inflation hedge ratio analysis.

Single-Period and Multiperiod Regression Results

We start with discussing results of the simple benchmark regression (1) that relates *ex post* housing returns with contemporaneous inflation for nonoverlapping annual data. Table 3 contains full sample as well as subsample OLS results in order to assess potential instabilities in the hedging potential over time. Standard errors are adjusted for serial correlation and heteroskedasticity

[12]Notice also that the ADF testing regression is only augmented with one lagged term in half of the cases, which implies that there is not much residual serial correlation in the ADF regression.

[13]We also performed a Johansen cointegration test for the total housing return index and the CPI but we were unable to find could a statistically significant long-term cointegration relation.

Table 3 ■ Granger causality tests for total housing returns and inflation.

	Full Sample		1814–1914		1915–2008	
VAR(1)						
$H_0 : \pi \rightarrow R$	$F = 7.23^{***}$ ($p = 0.01$)	–	$F = 4.15$ ($p = 0.04$)	–	$F = 2.48$ ($p = 0.12$)	–
$H_0 : R \rightarrow \pi$	–	$F = 0.3$ ($p = 0.56$)	–	$F = 0.01$ ($p = 0.93$)	–	$F = 0.72$ ($p = 0.40$)
VAR(2)						
$H_0 : \pi \rightarrow R$	$F = 3.96^{**}$ ($p = 0.020$)	–	$F = 1.95$ ($p = 0.15$)	–	$F = 1.40$ ($p = 0.25$)	–
$H_0 : R \rightarrow \pi$	–	$F = 0.10$ ($p = 0.90$)	–	$F = 0.51$ ($p = 0.60$)	–	$F = 0.43$ ($p = 0.65$)
VAR(3)						
$H_0 : \pi \rightarrow R$	$F = 3.42^{**}$ ($p = 0.018$)	–	$F = 2.10$ ($p = 0.10$)	–	$F = 1.47$ ($p = 0.23$)	–
$H_0 : R \rightarrow \pi$	–	$F = 0.49$ ($p = 0.69$)	–	$F = 0.17$ ($p = 0.91$)	–	$F = 0.42$ ($p = 0.74$)

Note: This table reports Granger causality tests (F-test statistics and corresponding p-values) for (log) annual total housing returns (R) and inflation (π). We test for Granger causality from inflation to housing returns and vice versa. The F-test is performed for various lags p of a corresponding bivariate VAR model (p = 1, 2, 3) of housing returns and inflation. Testing outcomes are reported for the full sample and two subsamples. Rejections of the null hypothesis of the absence of Granger causality are marked with *, ** and *** and are statistically significant at 90%, 95% and 99%, respectively.

by applying the Newey–West algorithm (1987). Moreover, we run the regression both for house price returns (excluding rents) as well as for total housing returns in order to check whether cash flows from housing influence its potential as an inflation hedge.

The full sample regression coefficients of capital returns and total housing returns on annual inflation equal 0.240 and 0.238, respectively. Moreover, they are both significant at the 5% level. Notice, however, that the overall goodness of fit is quite poor and the slope coefficient γ is very unstable across subsamples.[14] The post-1915 slope parameter is notably higher than in the first part of the sample. As will be argued later on when calculating hedge ratios, the higher slope values in the latter part of the sample coincide with a higher inflation persistence during that era. The table further shows that the Fisher hypothesis of a unit slope $\gamma = 1$ can be rejected for all cases.

[14]Chow breakpoint tests reject the null hypothesis of time constant intercept and slope parameters.

Last but not least, notice that the slope coefficients for capital gains regressions (upper panel) versus total return regressions (lower panel) hardly differ. In other words, additional cash flows arising from renting a house do not seem to have much influence on the hedge potential of houses against actual inflation. We do see, however, a significantly higher constant term when using total housing returns rather than just price changes: 5.8% for the complete sample, and 6.8% for most of the twentieth century.[15]

Next, we evaluate the inflation hedging potential of housing for longer time horizons up to five years by running simple OLS regressions of the long-horizon regression (2). Standard errors are again adjusted for serial correlation and heteroskedasticity by applying the Newey–West algorithm (1987).[16] Table 5 reports OLS regression results of k-year housing returns on k-year actual inflation rates ($k = 1, \ldots, 5$).

For the full sample, the Fisher slopes γ_k are significantly positive for all horizons and the coefficients as well as the goodness-of-fit measure (R^2) have a tendency to rise with the time horizon k. In other words, over longer holding periods, housing appears to have a more positive relation with inflation; the long-run coefficient over the full sample equals 0.374 as compared to 0.238 for one year. However, upon considering subsample results, slopes and goodness of fit results are only rising with the time horizon in the second part of the sample after 1915. Also, just as for the single-period regression results, the Fisher effect seems stronger for the post-1915 sample irrespective of the investment horizon. Upon testing whether slopes γ_k are equal to 1 (as implied by the Fisher hypothesis), one finds that slopes are all significantly below 1.

Next to the contemporaneous regressions of actual returns on actual inflation, we also quantify the relation between actual returns and expected inflation as this comes closer to what the Fisher hypothesis is actually about. Boudoukh and Richardson (1993) propose an Instrumental Variables – Generalized Method of Moments (IV-GMM) estimation procedure, which we apply to our data) for more details.

[15]All regressions have also been performed for the full sample excluding the First- and Second-World-War years (because housing markets were essentially put on hold). The results of this robustness check show no significant deviations from the outcomes presented above. All these complementary results were omitted for sake of space considerations but are available upon request from the authors. The remainder of this study presents results for total housing returns only. Results for house price returns are also available upon request.

[16]The serial correlation problem is further aggravated by overlapping data arising from the use of multiperiod returns and inflation rates.

Table 4 ■ OLS estimates for protection against actual inflation: single-period results.

Time Period	c_1	γ_1	R^2
(log) Price Returns			
Full Sample (1814−2008)	0.021**	0.240^{**,a}	0.009
	(0.009)	(0.136)	
1814−1914	0.010***	0.021^{a}	0.000
	(0.008)	(0.201)	
1915−2008	0.033***	0.294^{a}	0.014
	(0.019)	(0.186)	
(log) Total Returns			
Full Sample (1814−2008)	0.058***	0.238^{**,a}	0.010
	(0.009)	(0.134)	
1814−1914	0.048***	0.015^{a}	0.000
	(0.009)	(0.205)	
1915−2008	0.068***	0.297^{a}	0.014
	(0.018)	(0.182)	

Note: The table contains OLS estimates of Equation (1) that regresses housing returns on actual inflation. The housing returns either refer to (log) price returns or capital gains (upper panel) or to (log) total housing returns (including rents). The data are nonoverlapping (frequency and time horizon coincide). Given the serial correlation and heteroskedasticity present in annual housing return and inflation series, we complement the OLS estimates with Newey–West HAC standard errors (in parentheses below each estimate). Regression outcomes are reported for the full sample and two subsamples. Coefficients marked with *, ** and *** are statistically significant at 90%, 95% and 99%, respectively. For slope estimates marked with subscript "a," the Fisher hypothesis that the slope equals 1 is rejected at the 5% level.

Results of the IV-GMM estimation approach are reported in Table 6. We see that most slope estimates are higher than in Table 5 for the same time horizons. Moreover, and in line with the previous results, the slope estimates rise with the investment horizon except for the first subsample. Also, and conforming with the previous tables, the post-1915 Fisher slope nearly always exceeds its pre-1915 counterpart regardless of the time horizon. However, the slope estimates are often insignificantly different from zero and one due to relatively high standard errors. Stated otherwise, based upon the point estimates, housing seems to offer better protection against expected inflation than against actual inflation (and the downward bias in the OLS slopes seems to be removed by using IV estimators) but the statistical significance of the outcomes is rather weak.

Notice that the majority of the empirical findings reported in Tables 4, 5 and 6 are in line with Boudoukh and Richardson's (1993) findings for stocks: their OLS slopes for stock returns on actual inflation also tend to rise with the investment

Table 5 ■ OLS estimates for protection against actual inflation: multiperiod results.

	Full Sample			1814–1914			1915–2008		
k	c_k	γ_k	R^2	c_k	γ_k	R^2	c_k	γ_k	R^2
1	0.058***	0.238^{*,a}	0.01	0.048***	0.016^{a}	0.00	0.068***	0.297^{a}	0.01
	(0.009)	(0.134)	0	(0.009)	(0.205)	0	(0.018)	(0.182)	4
2	0.115***	0.290^{***,a}	0.031	0.096***	0.190^{a}	0.016	0.135***	0.284^{*,a}	0.027
	(0.016)	(0.111)		(0.016)	(0.144)		(0.030)	(0.157)	
3	0.171***	0.300^{***,a}	0.048	0.143***	0.142^{a}	0.009	0.203***	0.290^{*,a}	0.046
	(0.021)	(0.114)		(0.022)	(0.170)		(0.039)	(0.154)	
4	0.226***	0.338^{***,a}	0.078	0.192***	0.101^{a}	0.004	0.265***	0.327^{**,a}	0.082
	(0.027)	(0.115)		(0.028)	(0.209)		(0.047)	(0.146)	
5	0.278***	0.374^{***,a}	0.111	0.237***	0.027^{a}	0.000	0.325***	0.362^{**,a}	0.119
	(0.032)	(0.115)		(0.032)	(0.224)		(0.053)	(0.139)	

Note: This table presents OLS estimates of Equation (2) that regresses (log) total housing returns on actual inflation for time horizons k ranging from one to five years. The data are overlapping. The OLS estimates are complemented with Newey–West HAC standard errors (in parentheses below each estimate). Regression outcomes are reported for the full sample and for two subsamples. Coefficients marked with *, ** and *** are statistically significant at 90%, 95% and 99%, respectively. For slope estimates marked with subscript "a," the Fisher hypothesis that the slope equals 1 is rejected at the 5% level.

Table 6 ■ IV-GMM estimates for protection against expected inflation.

	Full Sample		1814–1914		1915–2008	
k	c_k	γ_k	c_k	γ_k	c_k	γ_k
1	0.051****	0.706***	0.044***	1.394	0.060***	0.540**
	(0.009)	(0.249)	(0.012)	(0.906)	(0.018)	(0.261)
2	0.104***	0.642	0.094***	0.327	0.121***	0.532
	(0.022)	(0.512)	(0.016)	(0.811)	(0.044)	(0.508)
3	0.115**	1.440	0.138***	0.515	0.116	1.248
	(0.045)	(0.872)	(0.023)	(0.607)	(0.076)	(0.691)
4	0.139	1.528	0.162***	1.371	0.101	1.674
	(0.104)	(1.620)	(0.034)	(1.075)	(0.227)	(1.937)
5	0.177	1.470	0.233***	0.402	0.120	1.711
	(0.159)	(1.918)	(0.032)	(0.744)	(0.508)	(3.448)

Note: This table presents IV-GMM estimates of Equation (2) that regresses (log) total housing returns on inflation for time horizons k ranging from one to five years. The data are overlapping. The set of instruments consists of a constant and the lagged inflation rate. The point estimates are complemented with Newey–West HAC standard errors (in parentheses below each estimate). Regression outcomes are reported for the full sample and for two subsamples. Coefficients marked with *, ** and *** are statistically significant at 90%, 95% and 99%, respectively. For slope estimates marked with subscript "a," the Fisher hypothesis that the slope equals 1 is rejected at the 5% level. All regressions have also been performed on the full sample excluding the First- and Second-World-War years (because housing markets were put on hold). The results of this robustness check show no significant deviations from the outcomes presented below.

horizon but stay significantly below 1. Moreover, all their IV-GMM slopes for stocks versus expected inflation exceed the OLS counterparts and are also rising with time horizons. However, the significance of these outcomes depends on the used instruments and considered subsample periods. Finally, they cannot reject the Fisher hypothesis for the IV-GMM estimation in a majority of the cases.

One important difference between housing and stocks as concerns their hedging properties seems that inflation protection for housing is already present at short horizons of one year, albeit weak, whereas stocks are often found to be unrelated (or even negatively correlated) with actual and expected inflation over this time horizon. As a result, the horizon effects—in terms of slope increases against actual and expected inflation for rising horizons k—appear less spectacular than in the Boudoukh–Richardson tables because the yearly housing slope values are already considerably above zero prior to lengthening the investor's holding period.

Multiperiod Hedge Ratio Analysis

As a second test for the importance of the investment horizon when it comes to inflation protection from homeownership, we calculate hedge ratios as discussed by Schotman and Schweitzer (2000). If a residential property is used as a vehicle to hedge inflation risk, one could determine the amount of housing required to protect against inflation by means of a hedge ratio. As explained earlier, important inputs for this k-period ratio are the *single-period* estimates for the inflation persistence, α, the hedge parameters against expected and unexpected inflation, β and φ, and the variances of the residuals ε and η from both the Fama–Schwert regression (4) and the inflation regression (5), respectively. Table 7 (upper panel) reports results for the Fama–Schwert regression parameters; the lower panel contains the k-period hedge ratios that are determined on the basis of the Fama–Schwert parameters imputed in Equation (6). In line with the previous analyses, we again distinguish between full sample results and subsample results.

The upper panel results show that inflation is quite persistent but also unstable over time, *i.e.*, the persistence after 1915 is nearly double the level of the first part of the sample. Given the important role of the inflation persistence parameter in the hedge ratio formula (6), this observation justifies the sample split. The fitted values and residuals of the AR(1) inflation regressions are used as proxies for expected and unexpected inflation in the Fama–Schwert regression (4). The Fama–Schwert outcomes suggest that housing provides a better hedge against expected inflation than against unexpected inflation regardless of the considered sample period ($\beta > \phi$). For the full sample period of 1814–2008, the coefficient for expected inflation equals a significant 0.706, while unexpected inflation is associated with an (insignificant) value of only 0.065. For both subperiods, we find that expected inflation relates more strongly to our home ownership returns than unexpected inflation, albeit only in a statistically significant way for the second subsample. The hedge parameter against unexpected inflation is always close to zero (and even negative in the first subsample).

The lower panel of Table 7 contains full sample and subsample hedge ratios $\Delta^{(k)}$for varying time horizons k. We also assess the statistical significance of the multiperiod hedge ratio$\Delta^{(k)}$ by block bootstrapping two-sided confidence intervals.[17] The multiperiod hedge ratios are statistically significant at the 5%

[17]Optimal block lengths for bootstraps on serially dependent data can be selected in various ways, see Hall, Horowitz and Jing (1995) or Politis and White (2004) for examples of two more advanced algorithms. However, we found that exogenous calibration of block length to reproduce bootstrapped inflation persistence parameters and Fama–Schwert regression parameters that come close to their real counterparts often works better than more involved block selection algorithms.

Table 7 ■ Single-period Fama–Schwert regressions and long-term hedge ratios.

	Full Sample	1814–1914	1915–2008
Panel A: Parameter Inputs for Hedge Ratio: Estimates			
α	0.522***	0.294***	0.597***
	(0.071)	(0.068)	(0.100)
β	0.706***, [a]	1.394	0.540*, [a]
	(0.248)	(0.971)	(0.276)
Φ	0.065	−0.104	0.162
	(0.167)	(0.238)	(0.233)
R^2	0.024	0.030	0.020
σ_η	0.067	0.057	0.075
σ_ε	0.189	0.141	0.229
Panel B: Estimated Hedge Ratios			
$k = 1$	0.008	−0.017	0.017
$k = 5$	0.117**	0.069	0.122
$k = 10$	0.159**	0.087	0.174
$k = 15$	0.174**	0.093	0.193
$k = 20$	0.181**	0.096	0.203
$k = 25$	0.186**	0.098	0.209
$k = 30$	0.189**	0.099	0.213

Note: The upper panel A of the table presents OLS estimates of the Fama–Schwert regression equation (4) that regresses (log) total housing returns on proxies for expected inflation and unexpected inflation. The proxies are constructed by assuming that inflation is driven by the stationary AR(1) process in Equation (5) and that expectations are formed accordingly. The data are annual and thus nonoverlapping. The OLS estimates are complemented with Newey–West HAC standard errors (in parentheses below each estimate). Regression outcomes are reported for the full sample and for two subsamples. The lower panel reports k-period hedge ratios that use the upper panel parameters as input according to formula (6). The 90%, 95% and 99% confidence intervals for the hedge ratios are block bootstrapped (not shown in the table). Block lengths are chosen such as to mimic the parameter estimates in the upper panel. This leads to block lengths of 30, 20 and 30 for the full sample and the two subsamples, respectively. Coefficients and hedge ratios marked with *, ** and *** are statistically significant at 90%, 95% and 99%, respectively. For slope estimates marked with subscript "a," the Fisher hypothesis that the slope equals 1 is rejected at the 5% level.

level over the full sample period but statistical significance disappears for the two subsamples. This is caused by subsample estimation risk which, combined with a lower number of observations, results in a lack of power for testing whether the hedge ratios equal zero or not.

A graphical representation of the three hedge ratio curves is provided in Figure 4. Remember that the hedge ratio is the fraction of total housing

Figure 4 ■ Hedge ratios for full period and for subperiods.

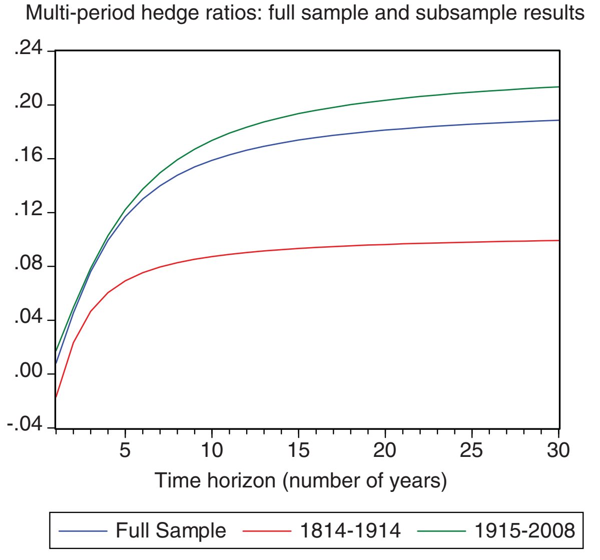

Note: This figure graphically represents the multiperiod hedge ratios for the full sample and for the two subsamples of Table 7. The hedge ratios at different investment horizons are estimated for the periods 1814–1914 and 1915–2008. The ratio is computed using Equation (6) in the article.

investment that is driven by an inflation hedging motive. Thus, a higher hedge ratio implies a higher hedging demand as part of total housing demand. The hedge ratio can also be interpreted as a scaled covariance between (long horizon) housing returns and inflation, which immediately makes clear why hedging demand rises with the hedge ratio: if the covariance with inflation rises, housing becomes a better hedge against inflation and investors would like to have more of it in their portfolios. Table 7 as well as Figure 4 reveal that hedge ratios rise with the investment horizon k for all three considered samples. Moreover, they seem concave functions of k: especially between 1 and 10 years, the hedge ratio increases fast; for horizons in excess of 10 years, the ratio increases only slightly. The table and figure also confirm our earlier prediction that the hedging potential of housing increases when inflation becomes persistent: the post-1915 hedge ratios clearly dominate their pre-1915 counterparts and also rise more quickly with k. This should not surprise given the higher inflation persistence in the second subsample. Also, the post-1915 hedge ratio dominance does not seem to be mitigated by the lower post-1915 value of the Fisher parameter β. Summarizing, hedge ratio estimates suggest that housing offers better protection for long holding periods up to 30 years as compared to short holding periods but the results have to be interpreted with care given the relatively large estimation risk.

Conclusions

Analyzing 195 years of Amsterdam housing returns has provided a number of interesting insights into housing's ability to protect against inflation risks. We add to the literature in three ways. First, we look at housing markets in the very long run, which enables us to investigate the inflation hedge issue for subsamples corresponding with greatly varying market circumstances and inflation regimes. Second, homeowners typically exhibit long time horizons, and the long-term time series offer us the opportunity to analyze the hedging potential over these time horizons. Lastly, the fact that we can combine house prices and rents implies that we can construct a total return index to housing, painting a complete picture of the return a homeowner can expect.

Our long-horizon regression results show that owning a house offers inflation protection in the longer run, both against actual as well as expected inflation. The long-horizon results for expected inflation are more in line with the Fisher hypothesis than single-period regressions using expected inflation or long-horizon regressions against actual inflation. Judged by results obtained from an alternative methodology—the hedge ratio—the effect of the time horizon is especially relevant up to horizons of approximately 10 years, after which the hedge ratio flattens. Interestingly, we also find a positive relation between the persistence of inflation and the long-term inflation hedge potential of housing: both long-term hedge ratios as well as multiperiod regression coefficients against expected and actual inflation increase during persistent inflation regimes. Housing offers much less protection against inflation risk when inflation is not persistent.

The authors thank Shaun Bond, Sylvester Eijffinger, John Glascock, Frank de Jong, Jianping Mei, John Quigley, Peter Schotman, Marion Steele, Coen Teulings, Marno Verbeek, Christian Wolff and participants of seminars at Netspar, the University of Tilburg, the University of Maastricht, the RSM Erasmus University, The University of Cambridge and participants to the 2000 Allied Social Science Associations conference in Boston for their valuable feedback. Jan Luiten van Zanden is thanked for providing his historic inflation series, and Victor Burger for his help with the rent data.

References

Anari, A. and J. Kolari. 2002. House Prices and Inflation. *Real Estate Economics* 30: 67–84.

Belsky, E. and J. Prakken. 2006. Housing Wealth Effects: Housing's Impact on Wealth Accumulation, Wealth Distribution and Consumer Spending. *Joint Center for Housing Studies Report* W04-13, Harvard University.

Bodie, Z. 1976. Common Stocks as a Hedge against Inflation. *Journal of Finance* 31: 459–470.

Bond, M.T. and M.J. Seiler. 1998. Real Estate Returns and Inflation: An Added Variable Approach. *Journal of Real Estate Research* 15: 327–338.

Boudoukh, J. and M. Richardson. 1993. Stock Returns and Inflation: A Long Horizon Perspective. *American Economic Review* 83: 1346–1355.

Brounen, D., P. Neuteboom and A. van Dijkhuizen. 2006. House Prices and Affordability: A First and Second Look across Countries. *DNB Working Paper Series* No. 83.

Central Bureau of Statistics. 1938. De Woningbouw in Amsterdam [Housing Construction in Amsterdam]. *Nederlandse Conjuctuur 9. Afzonderlijke Bedrijfstakken*. CBS: Voorburg, The Netherlands.

——. 1939. De Nederlandse Volkshuishouding, 1914–1918 [The Dutch National Economy, 1914–1918]. *Nederlandse Conjuctuur Speciale Onderzoekingen 3. Een Economisch-Statistische Schets*. CBS: Voorburg, The Netherlands.

——. 1948. Het Nationale Inkomen van Nederland, 1921–1939 [The Dutch National Income, 1921–1939]. *Monografieën van de Nederlandse Conjunctuur*, Vol. 7. CBS: Voorburg, The Netherlands.

——. 1949. Een Indexcijfer van de Consumptieve Uitgaven van Verbruikers [An Indexation of the Consumptive Expenditure]. *Statistische en Econometrische Onderzoekingen*, Vol. 4. CBS: Voorburg, The Netherlands.

——. 1959. *Zestig Jaar Statistiek in Tijdreeksen* [Sixty Years Statistics in Time Series]. CBS: Voorburg, The Netherlands.

——. 1999. *De Raming van Diensten uit Eigenwoningbezit* [An Estimation of the Services from Homeownership]. *Sector Nationale Rekeningen*. CBS: Voorburg, The Netherlands.

——. 2002. *Maandstatistiek van de Prijzen* 27(1) [Monthly Price Statistics]. CBS: Voorburg, The Netherlands.

De Jonge, J.A. 1968. *De Industrialisatie in Nederland Tussen 1850 en 1914 [Industrialization in the Netherlands: 1850 – 1914].* Amsterdam.

Department of Housing of the City of Amsterdam. 2006. *Wonen in Amsterdam 6 [Housing in Amsterdam].* City of Amsterdam: Amsterdam, The Netherlands.

Eichholtz, P. 1997. A Long Run House Price Index: The Herengracht Index. 1628–1973. *Real Estate Economics* 25: 175–192.

Eichholtz, P., S. Straetmans and M.A.J. Theebe. 2012. The Amsterdam Rent Index: The Housing Market and the Economy 1550–1850. *Journal of Housing Economics* 21: 269–282.

Fama, E.F. 1975. Short-term Interest Rates as Predictors of Inflation. *American Economic Review* 6: 269–282.

——. 1981. Stock Returns, Real Activity, Inflation, and Money. *American Economic Review* 71: 545–565.

——. 1990. Stock Returns, Expected Returns, and Real Activity. *Journal of Finance* 45: 1089–1108.

Fama, E.F. and M.R. Gibbons. 1984. A Comparison of Inflation Forecasts. *Journal of Monetary Economics* 13: 327–348.

Fama, E.F. and G.W. Schwert. 1977. Asset Returns and Inflation. *Journal of Financial Economics* 5: 115–146.

Feldstein, M. 1980. Inflation and the Stock Market. *American Economic Review* 70: 839–847.

Fisher, I. 1930. *The Theory of Interest*. MacMillan: New York.

Geske, R. and R. Roll. 1983. The Fiscal and Monetary Linkage between Stock Returns and Inflation. *Journal of Finance* 38: 1–33.

Gultekin, N.B. 1983. Stock Market Returns and Inflation: Evidence from Other Countries. *Journal of Finance* 38: 49–65.

Gyourko, J. and P. Linneman. 1988. Owner-Occupied Homes, Income-Producing Properties, and REITs as Inflation Hedges: Empirical Findings. *Journal of Real Estate Finance and Economics* 4: 347–372.

Hall, P., J. Horowitz and P. Jing. 1995. On Blocking Rules for the Bootstrap with Dependent Data. *Biometrika* 82: 561–574.

Horlings, E. 1995. *The Economic Development of the Dutch Service Sector, 1800–1850*. University of Utrecht: Utrecht.

Jaffe, J.J. and G. Mandelker. 1976. The 'Fisher Effect' for Risky Assets: An Empirical Investigation. *Journal of Finance* 31: 447–458.

Lee, B. 1992. Causal Relations among Stock Returns, Interest Rates, Real Activity, and Inflation. *Journal of Finance* 47: 1591–1603.

Nelson, C.R. 1976. Inflation and Rates of Return on Common Stock. *Journal of Finance* 31: 471–483.

Newey, W.K. and K.D. West. 1987. A Simple, Postive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55: 703–708.

Politis, D.N. and H. White. 2004. Automatic Block-Length Selection for the Dependent Bootstrap. *Econometric Reviews* 23(1): 53–70.

Rubens, J.H., M.T. Bond and J.R. Webb. 1989. The Inflation-Hedging Effectiveness of Real Estate. *Journal of Real Estate Research* 4: 45–55.

Schotman, P.C. and M. Schweitzer. 2000. Horizon Sensitivity of the Inflation Hedge of Stocks. *Journal of Empirical Finance* 7: 301–315.

Sinai, T. and N.S. Souleles. 2005. Owner-Occupied Housing as a Hedge against Rent Risk. *Quarterly Journal of Economics* 120: 763–789.

Sing, T.F. and S.H.Y. Liow. 2000. The Inflation-Hedging Characteristics of Real Estate and Financial Assets in Singapore. *Journal of Real Estate Portfolio Management* 6: 373–386.

Solnik, B. 1983. The Relation between Stock Prices and Inflationary Expectations: The International Evidence. *Journal of Finance* 38: 35–48.

Stevenson, S. 2000. A Long-Term Analysis of Regional Housing Markets and Inflation. *Journal of Housing Economics* 9: 24–39.

Van Riel, A. 2006. *Prices and Economic Development in the Netherlands, 1800–1913; Markets, Institutions and Policy Restraints*. Universiteit Utrecht: Utrecht.

Van Zanden, J.L. 2005. What Happened to the Standard of Living before the Industrial Revolution? New Evidence from the Western Part of the Netherlands. R.C. Allen, T. Bengtsson and M. Dribe, editors. *Living Standards in the Past*. Oxford University Press: Oxford. 173–195.

Wintle, M. 2000. *An Economic and Social History of the Netherlands, 1800–1920*. Cambridge University Press: United Kingdom.