

# The Hierarchical Repeat Sales Model for Granular Commercial Real Estate and Residential Price Indices

Marc K. Francke<sup>1,2</sup> · Alex van de Minne<sup>3</sup>

Published online: 29 August 2017  
© Springer Science+Business Media, LLC 2017

**Abstract** This paper concerns the estimation of granular property price indices in commercial real estate and residential markets. We specify and apply a repeat sales model with multiple stochastic log price trends having a hierarchical additive structure: One common log price trend and cluster specific log price trends in deviation from the common trend. Moreover, we assume that the error terms potentially have a heavy tailed ( $t$ ) distribution to effectively deal with outliers. We apply the hierarchical repeat sales model on commercial properties in the Philadelphia/Baltimore region and on residential properties in a small part of Amsterdam. The results show that the hierarchical repeat sales model provides reliable indices on a very detailed level based on a small number of observations. The estimated degrees of freedom for the  $t$ -distribution is small, largely rejecting the commonly made assumption of normality of the error term.

**Keywords** Commercial real estate · Residential real estate · Property price indices · Thin markets · Bayesian inference · Kalman filter · State space models

---

✉ Marc K. Francke  
m.k.francke@uva.nl  
Alex van de Minne  
avdminne@mit.edu

<sup>1</sup> Amsterdam Business School, Plantage Muidergrecht 12, 1018 TV Amsterdam, The Netherlands

<sup>2</sup> Ortec Finance, Naritaweg 51, 1043 BP Amsterdam, The Netherlands

<sup>3</sup> MIT Center for Real Estate, 105 Massachusetts Ave, Cambridge, MA 02139, USA

## Introduction

Real estate is the largest single class of investment in the world and the biggest source of household wealth in most industrialized countries. Estimating the risk and return of real estate is therefore important for underwriters, investors, policy makers and households alike.

Increasing attention has been paid to track real estate market movements on a detailed level (Constantinescu and Francke 2013). Granular indices give a better representation of the space and asset market the properties are traded in and are thus more informative. Also, having more granular indices enables tradable derivatives based on these indices (Bokhari and Geltner 2012; Deng and Quigley 2008). Estimating granular indices is challenging, because real properties are infrequently traded and are – especially in the case of commercial real estate or single-family housing – very heterogeneous (Deng et al. 2012, 2014). Geltner and Ling (2006) discuss the general trade-off that arises because of this: more granular indices are more useful, but *ceteris paribus* also more noisy, so less reliable and as a consequence, less useful.

In order to address this trade-off, we introduce a Hierarchical Repeat Sales Model (HRS). The HRS is a repeat sales model with multiple stochastic log price trends having a hierarchical additive structure: one common log price trend and cluster specific log price trends in deviation from the common trend. Both the number of clusters and the number of entries within a cluster are flexible. For example, consider the clusters property type and region. The log price trend for property type A in region B is the sum of the common log price trend, the property type A log price trend and the region B log price trend. The specification of the stochastic trends is flexible, for example a local linear trend model for the common trend and random walks for the cluster specific trends. The stochastic trends replace the commonly used fixed coefficients of time dummy variables.

The HRS basically consists of two sets of equations. The first equation is the measurement equation which relates the individual property log returns to changes in log price levels. The second set of equations is the transition equation which specifies the structure of the common and cluster trends. The hierarchical trend structure has successfully been applied in hedonic price models, see Francke and De Vos (2000), and Francke and Vos (2004), but has never been applied in repeat sales models. The HRS can be interpreted as a generalization of the structural time series repeat sales model (Francke 2010), in which only a common stochastic trend has been specified.

In order to demonstrate the strength of the hierarchical repeat sales model, we apply it in two different markets. The first application is commercial real estate (office, retail, industrial and multi-family) in the Philadelphia/Baltimore region in the United States, and the second application is owner-occupied housing (single-family housing and apartments) in a small area in Amsterdam in the Netherlands. We are able to estimate reliable quarterly price indices over a period of 15 years using less than 100 transaction pairs per index on average, about 3 transactions per quarter.

As a robustness check we also estimate our main model on (sub)markets with plenty of observations in itself first; multi-family in Manhattan (New York) and Los Angeles. Subsequently we (uniformly) sample without replacement 25% of the data, which makes the data scarce again. Our results show that omitting 75% of the data does not affect our indices considerably.

In this paper we make multiple contributions to the field of index construction. Firstly, we imply a hierarchical trend structure to the repeat sales model. The main benefit of using a hierarchical structure is that an aggregate dataset can be used to estimate indices on a more granular level, and thus use more information efficiently. Hierarchical trends have so far not been applied in repeat sales models, because the estimation is technically much more challenging compared to hedonic price models.

Secondly, we show that it is more appropriate to assume a  $t$ -distribution for the error term, instead of the more generally used normal distribution. This holds in specific for small samples, as is the focus in our application. The *estimated* degrees of freedom of the  $t$ -distribution are mostly between 2.5 and 3, very clearly rejecting the normal distribution, and the model fits the data better by a large margin.

Thirdly, our results reveal that – indeed – there are large differences between the different real estate markets, even though (spatially) they are relatively close by. For example, Baltimore had consistently more price appreciation during our analyzed period than Philadelphia. Also, the results from Amsterdam show that single-family homes had a higher risk-return profile compared to apartments, which is consistent with earlier findings.

Finally, to the best of our knowledge, this paper is the first real estate application which is estimated using a Markov Chained Monte Carlo algorithm. In our case we used Gibbs sampling.

The remainder of this paper is structured as follows. “[Model Specification and Estimation](#)” gives the evolution of repeat sales models in markets with small number of observations and our proposed methodology. “[Application](#)” first provides the data and some descriptive statistics. Next the results of the models are presented. Finally “[Conclusions](#)” concludes.

## Model Specification and Estimation

### The Repeat Sales Model in Small Samples

The existing literature on real estate index construction has been developed initially around the hedonic (Rosen 1974) and the repeat sales methodology (Bailey et al. 1963). The main benefit of the repeat sales over the hedonic model, is that the first is less affected by specification errors<sup>1</sup> and characteristics that are not observed in the data. This issue is exacerbated in markets with heterogeneous goods, such as commercial real estate or single-family housing.<sup>2</sup> On a negative side, only selling

---

<sup>1</sup>Although one important implicit assumption in the repeat sales model is that the characteristics and their impact on prices do not change over time. This assumption is obviously not correct for the age of the property (Harding et al. 2007). In the repeat sales model the effect of age is embedded into the model's estimates of the price levels, see Cannaday et al. (2005).

<sup>2</sup>For example, in an earlier version of this research we also fitted a hedonic price model with hierarchical trends on data for commercial real estate (including single sales), using characteristics like structure size, land size and age of the structure. The model contained property type and location specific cluster trends. The root mean squared errors for the different analyzed regions were in the range of 0.5 to 0.6.

prices of properties sold more than once are included in a repeat sales model: all single-sales are left out. The issue that (repeat) sales are a non-random selection of the entire property stock (sample selection bias) is addressed by Gatzlaff and Haurin (1997), Hwang and Quigley (2004), and Chinloy et al. (2013). In this article, however, we do not address the problem of sample selection bias.

The hedonic price model with time dummy variables (pooled regression) is given by

$$p_{it} = \mu_{it} + X_{it}\beta + Z_{it}\gamma + \varepsilon_{it}, \quad (1)$$

where  $p_{it}$  is the log transaction price of property  $i$ ,  $i = 1, \dots, n_t$  at time  $t$ ,  $t = 1, \dots, T$ ,  $X_{it}$  are observed characteristics and  $Z_{it}$  are unobserved characteristics, with corresponding parameters  $\beta$  and  $\gamma$ .  $\mu_t$  is a time varying constant, capturing the log price movement over time. The error term  $\varepsilon_{it}$  is typically assumed to be independently normally distributed with mean zero and variance  $\sigma_\varepsilon^2$ .

Under the assumptions that both the observed and unobserved variables  $X$  and  $Z$  do not change over time, we can take ‘first differences’ for pairs of sales. The repeat sales model is now provided by

$$\tilde{p}_i = p_{it} - p_{is} = \mu_t - \mu_s + \varepsilon_{it} - \varepsilon_{is} = \Delta\mu_{s+1} + \dots + \Delta\mu_t + \varepsilon_{it} - \varepsilon_{is}, \quad (2)$$

where  $s$  represents the time of buy (as opposed to time of sale  $t$ , with  $s < t$ ), and  $\mu_1 = 0$  for identification. Please note that we follow the broad strand of literature by assuming *pair* fixed effects, instead of a *property* fixed effects, leading to zero correlation between pairs of sales of the same property. Traditionally, the specification of the time effect is simply a dummy variable approach with fixed parameters  $\mu_t$ .

Equation 2 can alternatively be expressed as

$$\begin{pmatrix} p_{is} \\ p_{it} \end{pmatrix} = \begin{pmatrix} \mu_s \\ \mu_t \end{pmatrix} + \begin{pmatrix} 1 \\ 1 \end{pmatrix} \delta_i + \begin{pmatrix} \varepsilon_{is} \\ \varepsilon_{it} \end{pmatrix}, \quad (3)$$

where  $\delta_i$  is a pair fixed effect.<sup>3</sup> Conditional on  $\delta_i$ , the estimate of  $\mu_t$  is the average selling price at time  $t$ . This means that the estimate of  $\mu_t$  does not depend on preceding and subsequent periods. However, the estimate of  $\mu_t$  is sensitive to transaction price noise, in particular in small samples when the number of transactions per period is low. This happens, for example, with local price indices, short time periods, and/or in case of severe outliers, when the transaction price differs from its true market value by a large amount. The resulting price indices may then become very volatile (Francke 2010).

Multiple approaches have been proposed in order circumvent volatile indices in small samples. One method is to ‘simply’ smooth the estimated index  $\hat{\mu}_t$  ex post. Examples can be found in Cleveland (1979) and Clapp (2004), who both use a form of the local polynomial regression to smooth the index. The main drawback of this

<sup>3</sup>Please note that when we take the difference of the bottom and the top row of Eq. 3, the pair fixed effect  $\delta_i$  cancels out, so we lose one observation for one unknown parameter. Estimation results for  $\mu_t$  based on Eqs. 2 and 3 will therefore be identical.

two-step procedure is that it does not take into account the uncertainty in the estimates of  $\mu_t$ . One can also replace the dummy variables by a deterministic trend, like cubic or piecewise linear splines. Thorsnes and Reifel (2007) for example, use the Fourier form approach, to estimate a hedonic model. Bokhari and Geltner (2012) use a two-stage data frequency conversion procedure in a repeat sales model, by first estimating lower-frequency indices (using the Case and Shiller methodology) staggered in time, and then applying a generalized inverse estimator to convert from lower to higher frequency return series. Recently authors have also created ‘artificial pairs’ (thus increasing the number of observations) by matching transactions based on their hedonic characteristics, see for example McMillen (2012) and Guo et al. (2014).

Goetzmann (1992) introduced into the real estate literature what is perhaps the major approach to date for addressing small-sample problems in property price indices, namely, the use of Bayesian inference within repeat sales regression models. Goetzmann (1992) specifies the prior distribution for the periodic returns in Eq. 2 to be normally distributed,  $\Delta\mu_t \sim N(\kappa, \sigma_\eta^2)$ , implying a random walk with drift for the log price index  $\mu_t$ :  $\mu_{t+1} = \mu_t + \kappa + \eta_t$ . The variance parameters of the model (the signal  $\sigma_\eta^2$  and noise  $\sigma_\varepsilon^2$ ) are estimated in a first step and ‘plugged in’ to the second step of the Bayesian procedure, which can lead to biased estimates of the variances.<sup>4</sup>

Francke (2010) generalized this model by assuming that the trend component follows a local linear trend, and by providing the loglikelihood function  $\ell(\tilde{p}; \sigma_\eta^2, \sigma_\varepsilon^2)$  to estimate the signal  $\sigma_\eta^2$  and noise  $\sigma_\varepsilon^2$  directly (for example by maximum likelihood), avoiding the somewhat ad hoc two-step procedure proposed by Goetzmann (1992). The local linear trend repeat sales model is given by

$$p_{it} - p_{is} = \mu_t - \mu_s + \varepsilon_{it} - \varepsilon_{is}, \quad \varepsilon_{it} \sim N(0, \sigma_\varepsilon^2), \quad (4)$$

$$\mu_{t+1} = \mu_t + \kappa_t + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2), \quad (5)$$

$$\kappa_{t+1} = \kappa_t + \zeta_t, \quad \vartheta_t \sim N(0, \sigma_\zeta^2). \quad (6)$$

Note that the local linear trend is very flexible and includes different specifications, like random walk ( $\sigma_\zeta^2 = 0$  and  $\kappa_1 = 0$ ), random walk with drift ( $\sigma_\zeta^2 = 0$ ) (Goetzmann 1992) and smoothed trend ( $\sigma_\eta^2 = 0$ ), see Harvey (1989) for a description of the different functional forms.

The local linear trend repeat sales model is an example of a structural time series model. A structural time series model is a model in which the trend, error terms, plus other relevant components, are modeled explicitly. In contrast to the dummy variable approach, the structural time series model enables the prediction of the price level based on preceding and subsequent information. This means that even for particular time periods where no observations are available, an estimate of the price level can

<sup>4</sup>In matrix notation, Eq. 2 can alternatively be expressed as  $\tilde{p} = X\Delta\mu + \tilde{\varepsilon}$ , where  $X$  is a matrix with row vectors  $X_i$ , having ones for the columns representing time  $s + 1$  up to  $t$ , and zeros elsewhere. The error term is distributed as  $\tilde{\varepsilon} \sim N(0, \Omega)$ , where the variance-covariance matrix is  $\Omega = 2\sigma_\varepsilon^2 I$  (or an alternative specification depending on the time between sales). The prior for  $\Delta\mu$  is provided by  $\Delta\mu \sim N(\kappa j, \sigma_\eta^2 I)$ , where  $j$  and  $I$  are a vector of ones and the identity matrix with appropriate dimensions. Conditional on the variance parameters ( $\sigma_\eta^2, \sigma_\varepsilon^2$ ), the posterior for  $\Delta\mu$  is given by  $\Delta\mu | (\tilde{p}, \sigma_\eta^2, \sigma_\varepsilon^2) \sim N(m^*, W^*)$ , where  $W^* = (X'\Omega^{-1}X + \frac{1}{\sigma_\eta^2}(I - \frac{1}{T-1}jj'))^{-1}$ , and  $m^* = W^*X'\Omega^{-1}\tilde{p}$ .

be provided. State space models have also been successfully used to estimate hedonic price indices, see Schwann (1998), Francke and De Vos (2000) and Francke and Vos (2004). We built on the work of Goetzmann (1992) and Francke (2010) by estimating the repeat sales model in a structural time series framework.

### The Hierarchical Repeat Sales Model

A way to extend the repeat sales model is to imply a hierarchical trend structure. The main benefit of using a hierarchical structure is that an aggregate dataset can be used to estimate indices on a more granular level, and thus use more information efficiently. This was done for the hedonic price model in Francke and De Vos (2000) and Francke and Vos (2004). In this model  $\mu_t$  represents a common trend, while we add cluster specific trends. The Hierarchical Repeat Sales Model (HRS) is specified by

$$\lambda_{t+1}^j = \lambda_t^j + \zeta_t^j, \quad \zeta_t^j \sim N(0, \sigma_{\zeta_j}^2 I_{n_j}), \quad (7)$$

$$p_{it} - p_{is} = \mu_t - \mu_s + \sum_{j=1}^k d_i^j (\lambda_t^j - \lambda_s^j) + \varepsilon_{it} - \varepsilon_{is}, \quad \varepsilon_{it} \sim N(0, \sigma_\varepsilon^2). \quad (8)$$

The vectors of cluster trends  $\lambda_t^j, j = 1, \dots, k$ , are modeled as (random walk) deviations from the common trend  $\mu_t$ , where  $k$  is the number of clusters. The common trend is specified as a local linear trend model, given by Eqs. 5–6.<sup>5</sup> The number of elements in cluster  $j$  is  $n_j$ , so vector  $\lambda_t^j$  has dimension  $n_j$ . The selection row vector  $d_i^j$  has dimension  $n_j$  and consists of zeros and a one to select the appropriate element of cluster  $j$  for observation  $i$ .  $I_{n_j}$  is an identity matrix with dimension  $n_j$ . We impose  $\mu_1 = \lambda_{1,1}^j = 0, j = 1, \dots, k, l = 1, \dots, n_j$  for identification reasons. The total number of different price trends that can be calculated from the HRS is  $\prod_{j=1}^k n_j$ .

In our empirical application we will use two clusters, one for property types  $\lambda_t^1$  and one for locations  $\lambda_t^2$ . For example, for the US market we have 4 property types – Multi-family, Industrial, Office, and Retail – and 3 locations – Baltimore, Philadelphia and Rest. In total  $4 \times 3 = 12$  price indices are calculated by HRS. The index for offices (Off) in Philadelphia (Phi) can be calculated as the sum of the common trend, an office specific trend and a Philadelphia specific trend,  $\mu_t + \lambda_{\text{Phi},t}^1 + \lambda_{\text{Off},t}^2$ .

The Hierarchical Repeat Sales (HRS) has multiple benefits over the structural time series repeat sales model (STRS) by Francke (2010). First of all, it uses more information for each trend, resulting in being able to estimate more reliable indices in markets with even smaller numbers of observations. In the extreme, with no observations at all in a specific sub-market (combinations of elements from the cluster trends), it is still possible to produce an index based on the common trend and the trends for location and property type. Secondly, indices estimated by the HRS are less likely to ‘break’. A break is defined as an index which ends prematurely, due to

<sup>5</sup>Alternative specifications are also possible.

**Table 1** Summary of the models

Model	Structural Time Series	t-dist.	Use all data	Element specific variance of $\zeta_t^j$	Element specific variance of $\varepsilon_t$	Equations
STRS	✓	✓		NA	NA	Eqs. 4–6
HRS FEn	✓		✓			Eqs. 5–8
HRS FE	✓	✓	✓			Eqs. 5–8
HRS VE	✓	✓	✓	✓		Eqs. 5–8
HRS VEN	✓	✓	✓	✓	✓	Eqs. 5–8

missing data at the *end* (or *start*) of the sample. Even though the STRS can handle missing observations for earlier time periods, no price level will be provided in the final (initial) period(s) if corresponding data is missing.<sup>6</sup>

### More Flexible Specifications of the Hierarchical Repeat Sales Model

Note that the variance of  $\zeta_t^j$  in Eq. 7 is assumed to be equal for all elements within a cluster. Another assumption is that the variance of the noise term  $\varepsilon_{it}$  in Eq. 8 does not vary over combinations of clusters. In this subsection we will relax both assumptions: all elements are allowed to have their own variances.

Another generalization is that we replace the normality assumption of the error term  $\varepsilon$  by a potentially heavy-tailed distribution, the *t*-distribution  $\varepsilon_{it} \sim t_\nu(0, \sigma_\varepsilon^2)$ , where  $\nu$  is the degrees of freedom. The *t*-distribution has fatter tails compared to the normal distribution, and can better deal with outliers. For large values of  $\nu$  ( $\nu \rightarrow \infty$ ) the *t*-distribution coincides with the normal distribution. If  $\nu = 2$  the error distribution collapses to a Cauchy distribution. Only for  $\nu > 2$  the second moment exists. The degrees of freedom will be treated as a parameter to be estimated from the data.

Table 1 provides an overview of the different models that will be applied in the empirical section. The local linear trend methodology as proposed by Francke (2010) is denoted STRS. The Hierarchical Repeat Sales model is denoted HRS. The suffix FE (VE) is added when all elements in vector  $\zeta_t^j$  have the same (different) variance(s). The VEN model has in addition different variances of the error term  $\varepsilon_t$  for combinations of clusters. For all models, except for the HRS FEn model, we assume a *t*-distribution for the error term  $\varepsilon_t$ .

### Estimation

Hierarchical trends have not previously been applied to repeat sales models, mainly because of the more complex estimation procedure compared to hedonic price models with a hierarchical trend structure. In this section we provide two different estimation methods for the hierarchical repeat sales model. The first is by using

<sup>6</sup>Note that if there are no observations at all for any of the trends at the end of the sample, the HRS will still ‘break’.

state space methods (Kalman filter) and the second by Markov chain Monte Carlo techniques (Gibbs sampling).

For the Kalman filter we cannot use the ‘first difference’ specification (8), because the individual return data depend on the difference of the state vector in two moments in time, with varying time spans. The Kalman filter assumes that the state vector is a Markov chain. Therefore we have to use the alternative specification including pair or property fixed effects  $\delta$ , such that the measurement equation does not depend on preceding states,

$$p_t = D_t^\delta \delta + \mu_t + \sum_{j=1}^k D^j \lambda_t^j + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma_\varepsilon^2 I), \quad (9)$$

where  $D_t^\delta$  is a selection matrix for the pair or property fixed effects, and the state vector  $\alpha_t$  is defined by  $\alpha_t = (\mu_t, \kappa_t, \lambda_t^1, \dots, \lambda_t^k)'$ , see Eqs. 5–7. Define  $\psi$  as the vector of all variance parameters in the state and measurement equations.

The downside of this specification is that we have to deal with a (large amount) of pair or property fixed effects  $\delta_i$ . This can be done efficiently by the *augmented* Kalman filter (De Jong 1991). This filter gives estimates of the state vector  $\alpha_t$  conditional on all information up to time  $t$  and  $\psi$ .<sup>7</sup> Moreover, the augmented Kalman filter produces the loglikelihood  $\ell(p; \psi)$ . The loglikelihood can be used to obtain maximum likelihood estimates of  $\psi$  or it can be used for Bayesian inference. Estimates of the state vector conditional on the full sample and  $\psi$  are computed by the Kalman smoother.

The standard (and augmented) Kalman filter assumes that the error terms in both the measurement and transition equations are normally distributed. Shephard (1994) and Durbin and Koopman (2012)[for details see Chapter 10.8.3] provide efficient ways to estimate state space models with  $t$ -distributions for the error terms in the measurement and transition equations. For example, if the error term in the measurement equation  $\varepsilon_{it}$  has a  $t$ -distribution with  $\nu (> 2)$  degrees of freedom and variance  $\sigma_\varepsilon^2$ , then it has the representation  $\varepsilon_{it} = (\nu - 2)^{1/2} \varepsilon_{it}^* / c_{it}^{1/2}$ , where  $1/2 c_{it} \sim \text{Gamma}(\nu/2, \nu/2)$  and  $\varepsilon_{it}^* \sim N(0, \sigma_\varepsilon^2)$  are independent. Conditional on  $c_{it}$  the model is linear Gaussian and the standard (augmented) Kalman filter can be applied. Markov chain Monte Carlo methods can be used to integrate  $c$  out. An identical procedure can be used for  $t$ -distributed error terms in the transition equation.

In the second estimation method we don’t apply state space techniques, and use directly the model formulated in ‘first differences’ given by Eqs. 5–8, avoiding the inclusion of the property or paired fixed effects  $\delta$ . The second estimation method is

<sup>7</sup>The augmented Kalman filter requires first running the standard Kalman filter assuming  $\delta = 0$ , and subsequently rerunning a part of the Kalman filter for each column of matrix  $D_t^\delta$ , where this column is treated as observation vector (Durbin and Koopman 2012, see Chapter 5.7). Note that in case of pair fixed effects each column of  $D_t^\delta$  has only 2 observations over the total time period, the rest has missing values, so the computations are really fast. A computational restriction is that the augmented Kalman filter requires taking the inverse of a symmetric matrix with dimension the number of pair or property fixed effects. Estimation code has been written in Ox, using the package SsfPack (Koopman et al. 2008), allowing for a maximum of 25,000 property or pair fixed effects.



full Bayesian inference, where non-informative priors are specified for all variance parameters as  $\sigma^{-2} \sim \text{Gamma}(0.001; 0.001)$ . It is assumed that the degrees of freedom  $\nu$  in the  $t$ -distribution has an exponential distribution:  $\nu \sim \text{Exp}(1/3)$ . Although a full Bayesian analysis can be based on the loglikelihood function produced by the Kalman filter and the specified prior distributions for the variance parameters and the degrees of freedom, we used the program WinBUGS (Bayesian Analysis Using Gibbs Sampling), see Lunn et al. (2000). The main reason is the flexibility of the program: It only requires the model specification (in the ‘first difference’ representation (8), so without the pair or property fixed effects) and the specification of the priors. Appendix C provides an excerpt of the actual WinBUGS code used (Sturtz et al. 2005). All presented estimation results are from WinBUGS.

## Application

### Data and Descriptive Statistics

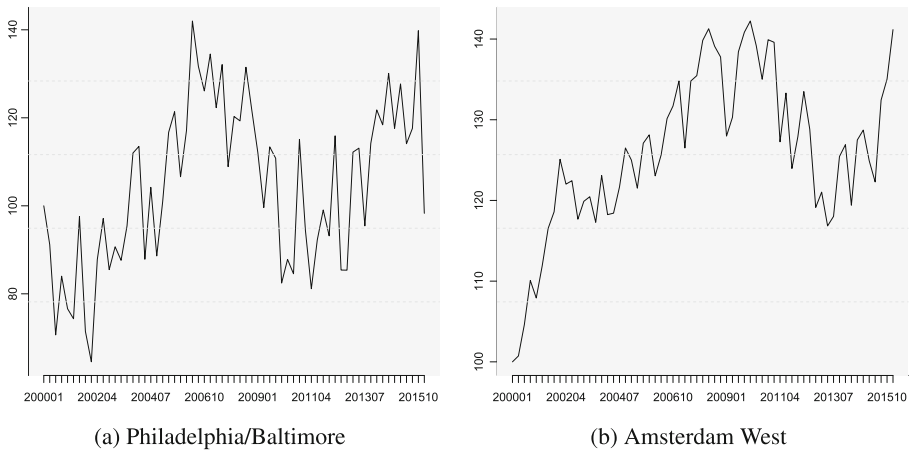
The Kadaster (Dutch Land Registry Office) provided the transaction data of individual houses in the Netherlands. The Kadaster is the official institution responsible for the registration of real estate properties. The Kadaster together with Statistics Netherlands construct an index based on the Sales Price Appraisal Ratio (SPAR) method (Bourassa et al. 2006). Price indices are provided on a national level as well as on regional (including Amsterdam) and house type levels. More details on the index construction method and the database can be found in De Vries et al. (2009). In this research we limit ourselves to a small area in Amsterdam-West, consisting of only five four-digit zip code areas. This area is of interest for us as it has a good mix of single-family housing and apartments (whereas Amsterdam is usually dominated by apartments). See Appendix A for more details on the exact location in Amsterdam.

Real Capital Analytics (RCA) provided us with repeat sales data of commercial real estate for the Philadelphia/Baltimore region, as defined by RCA. The Philadelphia/Baltimore region consists of the entire Delaware Valley and Baltimore metro region. It thus spans parts of the states: Pennsylvania, Maryland, Delaware and New Jersey.<sup>8</sup> RCA publishes a repeat sales index on a quarterly basis for the entire Philadelphia/Baltimore region, based on the frequency conversion methodology described by Bokhari and Geltner (2012).

In both cases, the total number of pairs is low: 915 pairs for our 5 zip codes in Amsterdam-West and 795 pairs for the combined Philadelphia/Baltimore region, between 2000 and the first quarter of 2016. Since we are interested in a quarterly index, we end up with an average of 28 transactions per quarter in Amsterdam-West

---

<sup>8</sup>Examples of cities in this region (apart from Baltimore and Philadelphia) are: Columbia, Cherry Hill, Exton, Jessup, Wilmington, Annapolis, Bensalem, Hanover, West Chester, Elkrigde, King of Prussia and Marlton.



**Fig. 1** Bailey et al. (1963) indices for both regions

and 24 transactions per quarter in Philadelphia/Baltimore. The resulting ‘standard’ Bailey et al. (1963) repeat sales quarterly index for both regions is shown in Fig. 1.

Both indices are very volatile (noisy) and other estimation techniques would probably fit the data better. Both indices show high price appreciation before the crisis, followed by a subsequent price drop and subsequent recovery. The price drop in Amsterdam is observed a few years after the start of the crisis. The most common explanation is the high mortgage debt levels of Dutch households and subsequent negative equity, which resulted in a big drop in transaction volume, but less in prices (Genesove and Mayer 1997).<sup>9</sup>

In our research we are interested in even more granular indices. For example, it is expected that the Philadelphia market is different (at least to a certain extent), than the Baltimore market. Also, offices, retail, industrial and multi-family housing complexes have their own demand and supply characteristics. The same goes for ‘the’ housing market in Amsterdam-West. Housing in every zip code, and per property type (apartment and single-family housing) can have a different price appreciation, volatility, etc. Obviously, the number of pairs observed is even lower per sub-market. Tables 2 and 3 give the number of pairs per sub-market for both Amsterdam and Philadelphia/Baltimore respectively.<sup>10</sup>

We end up with an average of only 90 pairs per market, or 3 transactions on average per quarter for both markets. However, note that the HRS uses information of the

<sup>9</sup>The allowed max Loan-to-Value, has been over 100% for a long period of time in the Netherlands. Combined with mortgage interest rate deductibility, this resulted in average LTVs at origination of over 1 (Brueckner 1994) and a third of all households who are ‘under water’ as of 2013.

<sup>10</sup>Note that we are mainly interested in the differences between Philadelphia and Baltimore for the commercial indices. We ‘naively’ grouped all other cities into one ‘rest’ category. However, in reality, some of the ‘rest’ cities might be considered to be part part of either Philadelphia or Baltimore. Examples are the cities of Cherry Hill and Marlton who are in the agglomerate of Philadelphia. Still, the ‘best’ subdivision of geographical areas is out of the scope of this paper and this way we keep sufficient observations in the ‘rest’ category.

**Table 2** Number of pairs in Amsterdam-West

Zip code	Single-family	Apartment	Total
1063	61	95	156
1064	75	137	212
1065	27	37	64
1066	148	124	272
1068	5	206	211
Total	316	599	915

*entire* region and *all* property types, but still produced the individual indices. In other words, the number of pairs used by the HRS is still 795 for Philadelphia/Baltimore and 915 for Amsterdam-West. This is inherently different from the ‘standard’ Bailey et al. (1963), Case and Shiller (1987) and ‘Bayesian’ Goetzmann (1992) and Francke (2010) repeat sales methodologies.

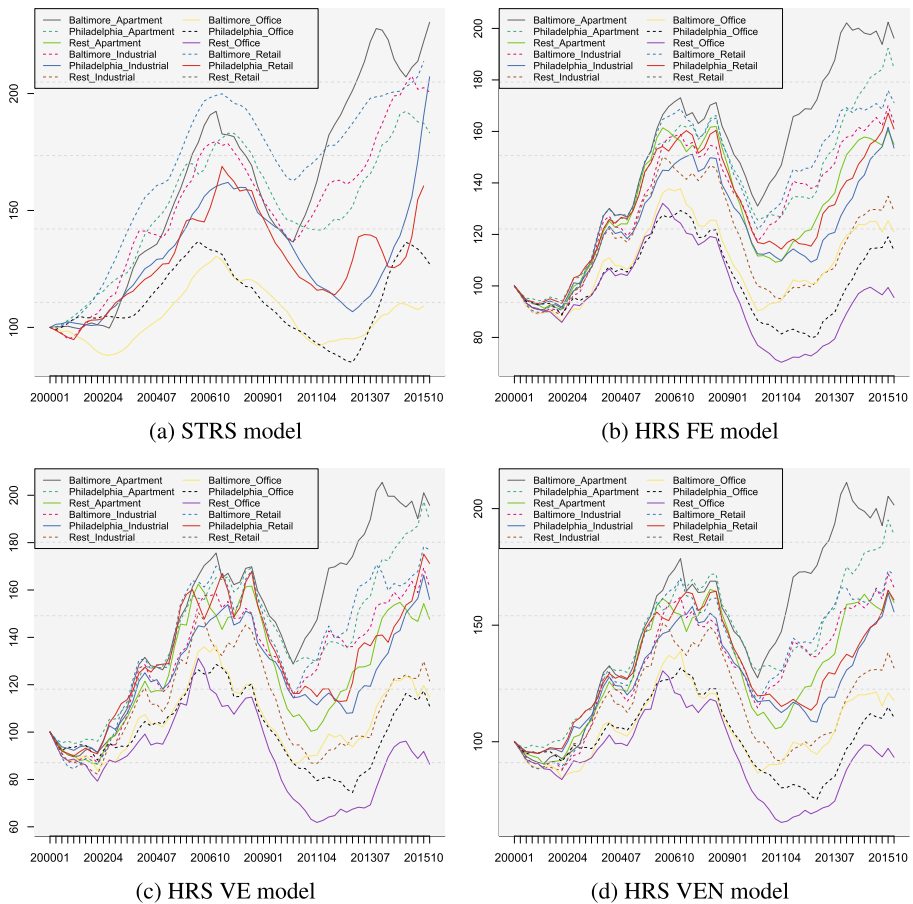
### Comparison of the Index Methodologies

In this Section the results of 5 different methodologies are compared, see Table 1 for an overview. Figures 2 and 3 show the price indices for the different combinations of location and property type in Philadelphia/Baltimore and Amsterdam-West respectively. Note that we suppress the indices for the ‘rest’ area in Philadelphia/Baltimore to conserve space, however the indices are available upon request. The indices of the HRS FEn model for both Philadelphia/Baltimore and Amsterdam-West are presented in Appendix B. Table 4 shows the estimation results of the variance parameters and the degrees of freedom  $\nu$  for different data sets and models. Table 5 provides some descriptive statistics of the (log) returns of the indices: mean, absolute mean, standard deviation, autocorrelation and correlation between sub-indices.

For the Philadelphia/Baltimore region it is evident (Fig. 2 and Tables 4 and 5) that the STRS model results in more smooth indices. The standard deviation of the quarterly returns is 2.3% on average for the STRS model compared to 3.5% for the HRS type models (see Table 5). An obvious reason for the smoothness in the STRS indices is the following: unlike the HRS index, the STRS index is based only on observations in the specific cluster. When very little data is available in a specific cluster, the STRS model just interpolates (according to the local linear trend structure) for the periods without observations, resulting in relatively smooth indices. This ‘over-smoothing’ is

**Table 3** Number of pairs Philadelphia/Baltimore

Metro	Multi-family	Industrial	Office	Retail	Total
Baltimore	99	111	60	44	314
Philadelphia	77	88	122	45	332
Rest	37	44	42	26	149
Total	213	243	224	115	795

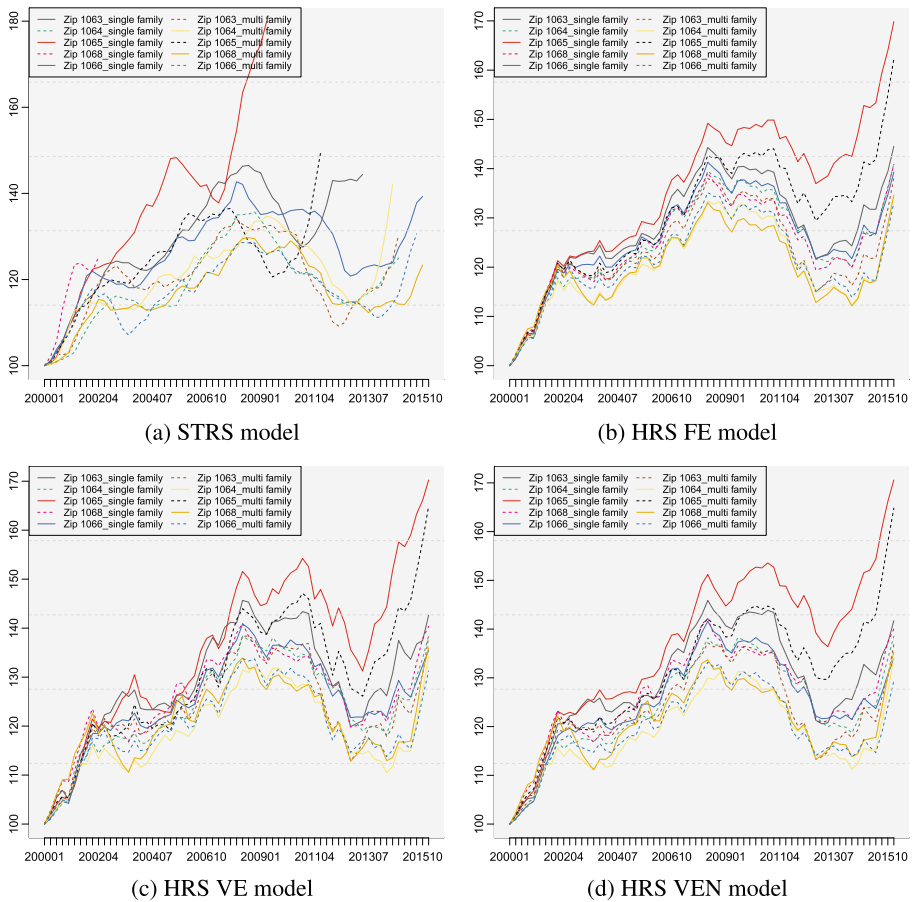


**Fig. 2** Price indices for the Philadelphia/Baltimore metro area

not evident in Amsterdam-West. However, it is difficult to compare the STRS models with the results from the HRS models, as so many indices ‘break’ using the first methodology. Finally, the timing of the STRS models seem off at certain points in time. In both Philadelphia/Baltimore and Amsterdam-West, the big downturn is a full year earlier than one would expect on average.

One issue with the HRS FE model is that the sub-indices seem have a very high correlation, see Table 5. In Amsterdam-West, the correlation of the (log) index returns is even close to 1. The reason for this behavior is that when subclusters have only a few observations, the cluster indices from HRS FE model will closely follow the common trend. The correlation drops considerably when allowing for a more flexible structure (VE and VEN models).

The HRS VE and VEN results reveal that in Philadelphia/Baltimore offices and retail are the most volatile markets, and multi-family the least. Within Amsterdam-West, the largest differences are between the zip codes and not so much between the property types (single family versus apartments).



**Fig. 3** Price indices for Amsterdam-West

The estimates of the scale parameter  $\sigma_\varepsilon$  for the individual STRS models coincide largely with those of the HRS VEN model. This holds for both Philadelphia/Baltimore and Amsterdam-West, see Fig. 4. On average  $\hat{\sigma}_\varepsilon$  is larger in the Philadelphia market (compared to Baltimore) and for offices and retail. There is a large difference in noise between property types in Amsterdam-West. In all cases single family houses have a larger  $\hat{\sigma}_\varepsilon$  compared to apartments. This is not surprising as single family housing is more heterogeneous than apartments in general.

Interestingly, the difference between trends from the model with a normal distribution for the error term in the measurement Equation (HRS FEn), and its  $t$ -distributed counterpart (HRS FE), is quite large. For example, the Philadelphia office index ends at 120 (100) for the  $t$ -distributed (normally distributed) version. The multi-family index in Baltimore ends at 200 (220) for the  $t$ -distributed (normally distributed) version. In general the results clearly show that the errors in the measurement equation follow a  $t$ -distribution. In almost all cases the estimated degrees of freedom ( $\nu$ ) are

**Table 4** Estimation results from repeat sales models over different data samples

	STRS	HRS FE	HRS VE	HRS VEN	HRS FEn
1. Philadelphia/Baltimore					
DIC		798	773	753	968
$\nu$	3.702 <sup>a</sup>	2.647	2.583	3.162	
$\sigma_\varepsilon$	0.298 <sup>a</sup>	0.266	0.259	0.292 <sup>b</sup>	0.445
$\sigma_\eta$	0.058 <sup>a</sup>	0.039	0.039	0.037	0.041
$\sigma_\zeta$	0.025 <sup>a</sup>	0.021	0.021	0.020	0.021
$\sigma_{\zeta\text{-Location}}$		0.035	0.050 <sup>b</sup>	0.048 <sup>b</sup>	0.031
$\sigma_{\zeta\text{-Property Type}}$		0.027	0.047 <sup>b</sup>	0.041 <sup>b</sup>	0.031
2. Amsterdam-West					
DIC		-1,565	-1,598	-1,503	-1,300
$\nu$	6.266 <sup>a</sup>	2.709	2.722	5.741	
$\sigma_\varepsilon$	0.142 <sup>a</sup>	0.071	0.070	0.133 <sup>b</sup>	0.120
$\sigma_\eta$	0.042 <sup>a</sup>	0.020	0.020	0.019	0.022
$\sigma_\zeta$	0.026 <sup>a</sup>	0.015	0.015	0.015	0.016
$\sigma_{\zeta\text{-Location}}$		0.010	0.023 <sup>b</sup>	0.022 <sup>b</sup>	0.013
$\sigma_{\zeta\text{-Property Type}}$		0.012	0.017 <sup>b</sup>	0.017 <sup>b</sup>	0.014

*STRS* structural time series model, *HRS* hierarchical repeat sales model, *FE* fixed variances within a cluster, *VE* varying variances within a cluster, *VEN* both varying variances within a cluster and varying noise per combination of clusters. Lower case **n** represents a model with a normal distribution for the error in the measurement Equation. *DIC* stands for Deviance Information Criterion and  $\nu$  are the estimated degrees of freedom in the *t*-distribution

<sup>a</sup>These are averages of the results for the STRS model over all the different subsets of data. See Fig. 4 for the individual estimates

<sup>b</sup>These numbers represent the average of the estimates within the model. See Fig. 4 for the individual estimates

between 2.5 and 3. Only in Amsterdam-West the error terms in the STRS and HRS VEN models are somewhat closer to normality.

Finally note that the Deviance Information Criterion (DIC) (Spiegelhalter et al. 2002) is vastly superior for the HRS FE with a *t*-distribution over the normal distribution for both Philadelphia/Baltimore and Amsterdam-West. The DIC is lowest for the HRS VEN model in Philadelphia/Baltimore and for the HRS VE model in Amsterdam-West.<sup>11</sup>

<sup>11</sup>As the STRS model is estimated over different subsamples we are not able to make a comparable DIC statistic.

**Table 5** (log) Return statistics

	STRS	HRS FE	HRS VE	HRS VEN	HRS FEn
1. Philadelphia/Baltimore					
mean	0.008	0.006	0.006	0.006	0.005
mean	0.020	0.027	0.033	0.030	0.024
sd.dev	0.023	0.032	0.040	0.036	0.029
ACF.1	0.674	0.557	0.458	0.485	0.635
ACF.6	0.110	0.129	0.073	0.116	0.106
Correlation	0.331	0.886	0.732	0.773	0.919
2. Amsterdam-West					
mean	0.008	0.006	0.006	0.005	0.006
mean	0.014	0.015	0.017	0.015	0.016
sd.dev	0.016	0.018	0.020	0.018	0.018
ACF.1	0.565	0.331	0.259	0.417	0.479
ACF.6	-0.056	0.059	0.017	0.057	0.117
Correlation	NA <sup>a</sup>	0.960	0.797	0.863	0.938

All numbers are based on the mean statistics *per* index. For example, in Philadelphia/Baltimore we estimate 12 indices per model. We first calculate the mean statistic per index. Subsequently, the number presented in the table is the average of those means

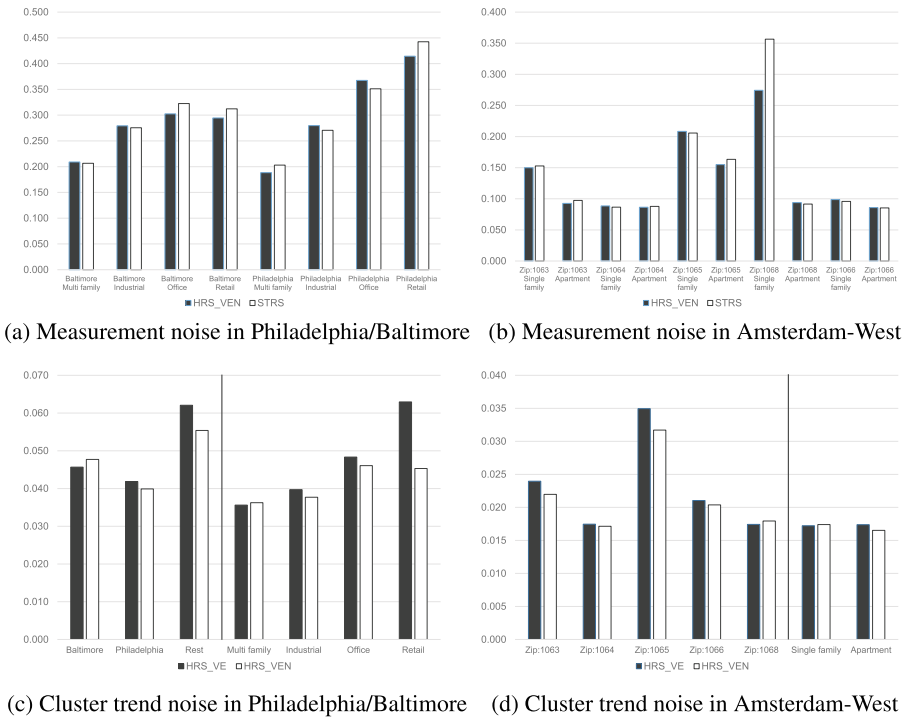
<sup>a</sup>Most indices in this group 'break', preventing us from calculating meaningful correlations

## Robustness

In this Section we will estimate our model (the HRS VEN) on large (sub)markets with plenty of observations. Subsequently we (uniformly) sample without replacement 25% of the data, which makes the data scarce again. This random sampling is done three times. The idea is to compare the different indices and assess to what extent the indices based on the 'scarce' dataset follow the index on the full sample. We also include the standard Bailey et al. (1963) repeat sales index (using the full sample of the submarket) for comparison.

More specifically, we are interested in a multi-family index for both Manhattan (739 observations) and Los Angeles proper (1,388 observations) using - again - the RCA data.<sup>12</sup> The multi-family properties are placed in a property type cluster with retail, industrial and offices for both indices. The location cluster for the Manhattan index consists of the (other) New York City boroughs and the rest of the New York metro area (according to the RCA definition, which closely follows the official MSA of New York). Inland Empire and Orange County are the other regions in the location

<sup>12</sup>These submarkets are selected as they have the highest number of counts in the RCA data.



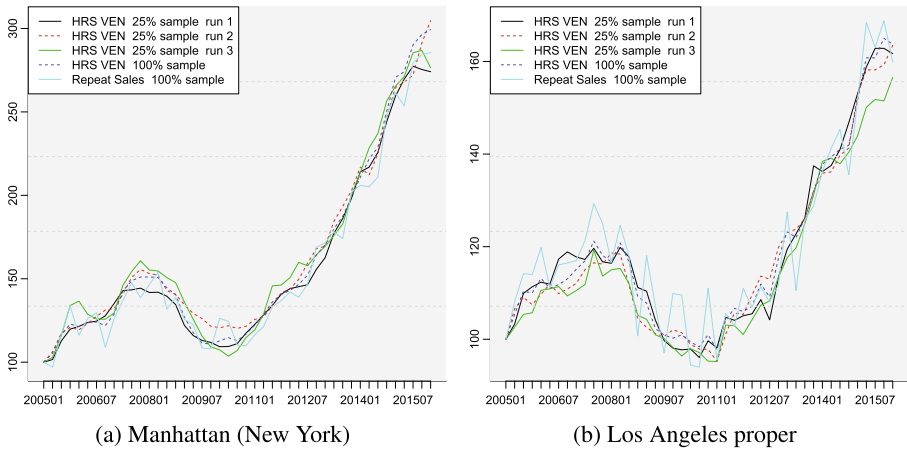
**Fig. 4** Estimates of scale parameters in the different repeat sales models

cluster for the LA index. The total number of observations, after adding the other property types and locations, is 3,417 (4,724) for the New York index (Los Angeles index).

Note that we did the sampling on the ‘full’ markets (all locations and property types). This makes the underlying data of the different runs even more random. For example, one of the New York runs only had 100 Manhattan multi-family observations, whereas in another run the number of observations for said submarket was 195. The exact number of observations, the descriptive statistics of these markets and the estimation results (of the hyperparameters) are omitted for brevity (however are available upon request). The resulting indices are displayed in Fig. 5.

First, Fig. 5 reinforces earlier findings in the field of structural times series in that the HRS is less volatile and more positively autocorrelated than the ‘standard’ Bailey et al. (1963) repeat sales index. Counterintuitive to prior expectation is that the ‘standard’ repeat sales index for Los Angeles is (a lot) more volatile than the New York index, even though there are twice as many observations. This might be caused differences in heterogeneity of the properties between the markets. More interestingly however, is that omitting 75% of the data does not affect the index considerably.





**Fig. 5** Multi-family indices on sub-samples

Arguably, only the third run of the Los Angeles index goes a bit awry in the final year or two.<sup>13</sup>

## Conclusions

In this paper we focus on computing property price indices in small samples by hierarchical repeat sales (HRS) models. The HRS model can be seen as an extension of the structural time series repeat sales model (STRS) by Francke (2010). The STRS model assumes one common trend for the total market, modeled by a local linear trend model, whereas the HRS model assumes both a common trend, and cluster trends in deviation from the common trend (modeled by random walks). The clusters can for example represent locations and property types. Both the STRS and HRS model specify the trends as a stochastic time series models, replacing fixed time dummy variables in the original repeat sales model by Bailey et al. (1963).

Another modification to commonly used versions of repeat sales models is that we allow for a heavy-tailed error distribution to better deal with potential outliers in the data. We specify the error term by a  $t$ -distribution with  $\nu$  degrees of freedom. The degrees of freedom is not set a priori, so will be estimated from the data.

We provide estimation methods that efficiently deal with the complex error structure induced by the stochastic trends and the  $t$ -distribution. The first estimation method uses the augmented Kalman filter and the second by Markov chain Monte Carlo techniques (Gibbs sampling).

<sup>13</sup>Although it is possible that future new pairs will revise the index to be more in line (Clapp and Giaccotto 1999; Deng and Quigley 2008).

We apply the hierarchical repeat sales model on two different markets: commercial property in the Philadelphia/Baltimore region in the United States and residential property in Amsterdam-West in the Netherlands. In both markets we have on average about 30 transactions per quarter. Unlike the standard repeat sales model the STRS model is able to produce reliable indices on the total market level. The HRS goes a step further: it subdivides the market in sub-markets based on location and property type, resulting in 10 (Amsterdam-West) and 12 (Philadelphia/Baltimore) sub-indices. In specific, versions of the HRS model where the noise components have a  $t$ -distribution, and the variance parameter of the cluster trend noise varies over the different elements in the cluster, perform well. The estimated degrees of freedom are small, rejecting the commonly made assumption of normality of the error terms.

For robustness we also estimate our main model on (sub)markets with approximately 1,000 observations each, namely multi-family properties in Manhattan (New York) and in Los Angeles proper. The number of observations are plentiful and should therefore result in reliable indices. The two clusters are - again - completed by (1) other property types and (2) other (sub)locations within the metro area. Subsequently we (uniformly) sample without replacement 25% of the data, and re-estimate the model. This way we can assess how accurate our index methodology is in a scarce data environment by comparing it to the 'true' index (volatility and mean returns-wise for example). Our robustness check shows that omitting 75% of the data does not affect our indices considerably.

**Acknowledgements** The authors greatly appreciated the comments given during the JREFE special conference in Cambridge (UK) in 2016. Special gratitude goes to Yenghong Deng and John Clapp. We would also like to thank the people at Real Capital Analytics for providing us the data, and commenting on the model in an earlier stage. Especially the comments by Bob White, Elizabeth Szep, Willem Vlaming and Jim Sempere of RCA were helpful and insightful. Finally, we would like to thank the anonymous referee for his comments on this paper.

## Appendix A: Map of Amsterdam West

In this research we use repeat sales data on a region in the west of Amsterdam for the zip codes: 1063, 1064, 1065, 1066 and 1068. Said zip codes are found in the red rectangular box in Fig. 6 and falls under a larger area called the 'Westelijke Tuinsteden'. In the middle the Sloterplass (lake) is observed. As of 2013, approximately 70,000 people populated these zip codes.

Most of the housing was built during the 1950s and 1960s as part of a bigger urban plan designed in the 1930s, loosely following the 'garden city' principles. The Westelijke Tuinsteden is also known for the big urban renewal program, executed between 1999 and 2015, when in total more than 10,000 dwellings were demolished and approximately 25,000 new dwellings were constructed. For your orientation, zip codes: 1011 and 1012 include the historic center of the city and zip codes: 1015 through 1018 include the famous canals.

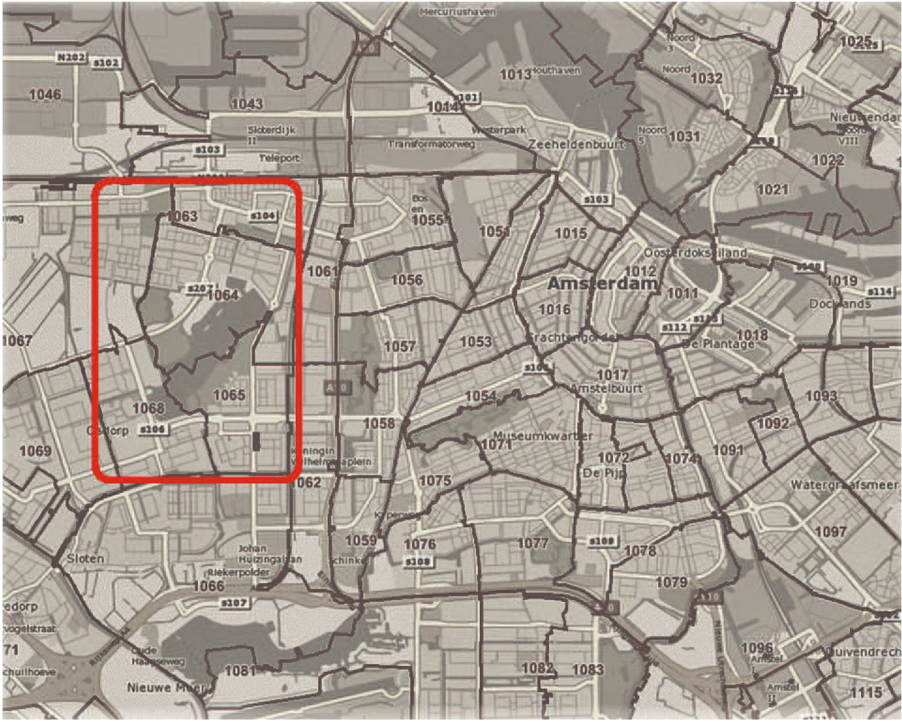


Fig. 6 Map of Amsterdam West

Appendix B: Indices for the HRS FEn model

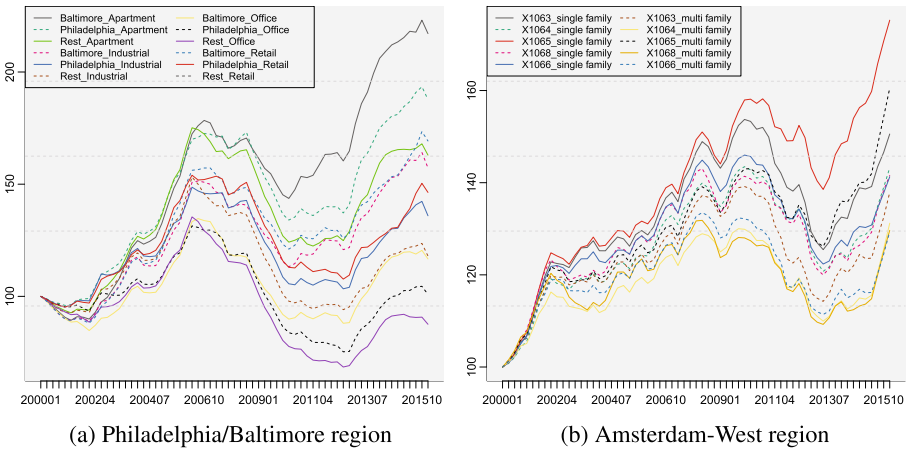


Fig. 7 HRS FE model with normal distribution

## Appendix C: Excerpt of the WinBUGS code for the HRS FE model

```
#####
#                               MODEL
#####
model
for(i in 1:N){
  y.Var[i] ~ dt(mu[i],tau,df)
  mu[i] <- (beta[Index.Sell[i]] - beta[Index.Buy[i]]) +
  (alpha[AG[i],Index.Sell[i]] - alpha[AG[i],Index.Buy[i]]) +
  (delta[PT[i],Index.Sell[i]] - delta[PT[i],Index.Buy[i]])}

#####
#                               TRENDS
#####
for(t in 2:Nt){
  beta[t] ~ dnorm(mu.beta[t],tau.eta)
  mu.beta[t] <- beta[t-1] + kappa[t-1]
  kappa[t] ~ dnorm(kappa[t-1],tau.kappa)}

for(p in 1:Na){
  for(t in 2:Nt){ alpha[p,t] ~ dnorm(alpha[p,t-1],tau.alpha)}}
for(p in 1:Np){
  for(t in 2:Nt){ delta[p,t] ~ dnorm(delta[p,t-1],tau.delta)}}

#####
#                               PRIORS
#####
beta[1] <- 0
kappa[1] <- 0
for(p in 1:Na){ alpha[p,1] <- 0}
for(p in 1:Np){ delta[p,1] <- 0}

tau.alpha ~ dlnorm(5,1)
tau.delta ~ dlnorm(5,1)
tau.eta ~ dgamma(0.001,0.001)
tau.kappa ~ dgamma(0.001,0.001)
tau ~ dgamma(0.1,0.1)
df ~ dexp(adf)
adf <- 1/3

#####
```

Here  $N$  is the total number of observations,  $y.Var$  is the log price differences between buy and sell ( $p_{it} - p_{is}$ ).  $Na$  are the amount of location categories  $AG$ ,  $Np$  are the amount of property type categories  $PT$  and  $Nt$  are the number of periods.  $Index.Sell$  is the time period on a scale of  $1, \dots, Nt$  on which the property was sold and  $Index.Buy$  when the property was bought.

## References

- Bailey, M.J., Muth, R.F., & Nourse, H.O. (1963). A regression method for real estate price index construction. *Journal of the American Statistical Association*, 58(304), 933–942.
- Bokhari, S., & Geltner, D. (2012). Estimating real estate price movements for high frequency tradable indexes in a scarce data environment. *The Journal of Real Estate Finance and Economics*, 45(2), 522–543.
- Bourassa, S.C., Hoesli, M., & Sun, J. (2006). A simple alternative house price index model. *Journal of Housing Economics*, 15, 80–97.
- Bruelckner, J. (1994). The demand for mortgage debt: some basic results. *Journal of Housing Economics*, 3(4), 251–262.
- Cannaday, R.E., Munneke, H.J., & Yang, T.L. (2005). A multivariate repeat-sales model for estimating house price indices. *Journal of Urban Economics*, 57, 320–342.
- Case, K.E., & Shiller, R.J. (1987). Prices of single family homes since 1970: new indexes for four cities. *New England Economic Review* 45–56.
- Chinloy, P., Hardin III, W.G., & Wu, Z. (2013). Transaction frequency and commercial property. *Journal of Real Estate Finance and Economics*, 47, 640–658.
- Clapp, J.M. (2004). A semiparametric method for estimating local house price indices. *Real Estate Economics*, 32, 127–160.
- Clapp, J.M., & Giaccotto, C. (1999). Revisions in repeat-sales price indexes: here today, gone tomorrow. *Real Estate Economics*, 27, 79–104.
- Cleveland, W.S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74, 829–836.
- Constantinescu, M., & Francke, M.K. (2013). The historical development of the Swiss rental market—a new price index. *Journal of Housing Economics*, 22(2), 135–145.
- De Jong, P. (1991). The diffuse Kalman filter. *The Annals of Statistics*, 2, 1073–1083.
- De Vries, P., de Haan, J., van der Wal, E., & Mariën, G. (2009). A house price index based on the SPAR method. *Journal of Housing Economics*, 18(3), 214–223.
- Deng, Y., & Quigley, J.M. (2008). Index revision, house price risk, and the market for house price derivatives. *The Journal of Real Estate Finance and Economics*, 37(3), 191–209.
- Deng, Y., McMillen, D.P., & Sing, T.F. (2012). Private residential price indices in Singapore: a matching approach. *Regional Science and Urban Economics*, 42(3), 485–494.
- Deng, Y., McMillen, D.P., & Sing, T.F. (2014). Matching indices for thinly-traded commercial real estate in Singapore. *Regional Science and Urban Economics*, 47, 86–98.
- Durbin, J., & Koopman, S.J. (2012). *Time series analysis by state space methods*. Oxford University Press.
- Francke, M.K. (2010). Repeat sales index for thin markets: a structural time series approach. *Journal of Real Estate Finance and Economics*, 41, 24–52.
- Francke, M.K., & De Vos, A.F. (2000). Efficient computation of hierarchical trends. *Journal of Business and Economic Statistics*, 18, 51–57.
- Francke, M.K., & Vos, G.A. (2004). The hierarchical trend model for property valuation and local price indices. *Journal of Real Estate Finance and Economics*, 28, 179–208.
- Gatzlaff, D.H., & Haurin, D.R. (1997). Sample selection bias and repeat-sales index estimates. *Journal of Real Estate Finance and Economics*, 14, 33–50.
- Geltner, D.M., & Ling, D. (2006). Considerations in the design and construction of investment real estate research indices. *Journal of Real Estate Research*, 28(4), 411–444.
- Genesove, D., & Mayer, C. (1997). Equity and time to sale in the real estate market. *American Economic Review*, 87, 255–269.
- Goetzmann, W.N. (1992). The accuracy of real estate indices: repeats sale estimators. *Journal of Real Estate Finance and Economics*, 5, 5–53.
- Guo, X., Zheng, S., Geltner, D.M., & Liu, H. (2014). A new approach for constructing home price indices: the pseudo repeat sales model and its application in China. *Journal of Housing Economics*, 25, 20–38.
- Harding, J.P., Rosenthal, S.S., & Sirmans, C.F. (2007). Depreciation of housing capital, maintenance, and house price inflation: estimates from a repeat sales model. *Journal of Urban Economics*, 61, 193–217.
- Harvey, A. (1989). *Forecasting structural time series models and the kalman filter*. Cambridge: Cambridge University Press.
- Hwang, M., & Quigley, J.M. (2004). Selectivity, quality adjustment and mean reversion in the measurement of house values. *Journal of Real Estate Finance and Economics*, 28, 161–178.

- Koopman, S.J., Shephard, N., & Doornik, J.A. (2008). *Statistical algorithms for models in state space form: Ssfpack 3.0*. Timberlake Consultants Press.
- Lunn, D.J., Thomas, A., Best, N., & Spiegelhalter, D. (2000). WinBUGS – a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing*, *10*, 325–337.
- McMillen, D.P. (2012). Repeat sales as a matching estimator. *Real Estate Economics*, *40*(4), 745–773.
- Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *The Journal of Political Economy*, *82*, 34–55.
- Schwann, G.M. (1998). A real estate price index for thin markets. *Journal of Real Estate Finance and Economics*, *16*, 269–287.
- Shephard, N. (1994). Partial non-gaussian state space. *Biometrika*, *81*, 115–131.
- Spiegelhalter, D.J., Best, N.G., Carlin, B.P., & van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *64*(4), 583–639.
- Sturtz, S., Ligges, U., & Gelman, A.E. (2005). R2WinBUGS: a package for running WinBUGS from R. *Journal of Statistical software*, *12*(3), 1–16.
- Thorsnes, P., & Reifel, J.W. (2007). Tiebout dynamics: neighborhood response to a central-city/suburban house-price differential. *Journal of Regional Science*, *47*, 693–719.