

# Modeling Regional House Prices\*

Bram van Dijk

Philip Hans Franses

Richard Paap

Dick van Dijk

*Econometric Institute  
Tinbergen Institute  
Erasmus University Rotterdam*

ECONOMETRIC INSTITUTE REPORT EI 2007-55

## Abstract

We develop a parsimonious panel model for quarterly regional house prices, for which both the cross-section and the time series dimension is large. The model allows for stochastic trends, cointegration, cross-equation correlations and, most importantly, latent-class clustering of regions. Class membership is fully data-driven and based on (i) average growth rates of house prices, (ii) the propagation of shocks to house prices across regions, also known as the ripple effect, and (iii) the relationship of house prices with economic growth and other variables. Applying the model to quarterly data for the Netherlands, we find convincing evidence for the existence of two distinct clusters of regions, with pronounced differences in house price dynamics.

**Keywords:** cross-section dependence, cointegration, ripple effect

**JEL Classification:** C21, C23, C53

---

\*We thank Rene Segers for assistance with the graphs. The corresponding author is Bram van Dijk, Tinbergen Institute, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands, e-mail: [avandijk@few.eur.nl](mailto:avandijk@few.eur.nl), phone: +31-10-4088943, fax: +31-10-4089162. Part of the calculations are done in Ox 4.00 (Doornik, 2002).

# 1 Introduction

Real estate prices in many countries have experienced a dramatic boom in recent years (IMF, 2004). At the same time, the extent of the price increase appears to vary substantially across different regions within a given country. In the Netherlands, for example, it is commonly believed that house prices in *Amsterdam* and the densely populated western part of the country increased far more than prices in the smaller cities and rural areas in the east. As house prices are typically available per region or city, we may analyze these data at such a disaggregate level, to examine whether indeed regions or cities behave differently, perhaps in terms of trends, but also in terms of response to outside economic shocks. In this paper we develop a time series model that suits this purpose.

Most regional house prices have the following properties. First, they tend to display a trend, and historical price patterns suggest that this trend probably is not deterministic but stochastic. In particular, house prices show ‘bubble’-type behavior, where periods of sharp increases of the price level suddenly end with a sharp drop followed by a prolonged period of low price levels, suggesting that trends are unlikely to be deterministic. Second, for different regions within a country these stochastic trends should somehow be linked. It is not plausible that prices in different regions would diverge indefinitely or that certain regions would not respond to common macroeconomic shocks. So, a model for regional house prices should allow for some form of common trends. Third, it can be expected that adjacent regions show similar price patterns, although this may also be the case for regions far apart geographically but with similar economic or demographic characteristics. Hence, a suitable model should allow for similarities in the dynamic behavior of house prices across regions. Put differently, the model should allow for groups or clusters of regions, where house price dynamics in regions within a given cluster are the same, while they can be different across clusters. Preferably, such a model should not require ex-ante or exogenous assignment of regions to specific clusters. In fact it would be best if the data themselves were allowed to indicate if clusters exist and if so, which regions belong to which cluster.

In this paper we extend the latent-class panel time series model introduced by Paap *et al.* (2005) to capture these different properties of regional house prices. The key feature of this model is that the clustering of regions is purely data-driven, where cluster membership is based on three characteristics corresponding to three specific research questions we want to address. The first question is whether prices in all regions have the same average growth rate. Note that a common trend specification across the regions entails that their growth rates must be somehow compatible, but it still leaves open the possibility that house prices in some regions grow faster than in others. The second question concerns the so-called ripple effect, see, for example, Cameron *et al.* (2006). This refers to the phenomenon that price changes start in one particular region (or cluster of regions) and gradually disseminate to other regions in subsequent time periods. Within our model we can examine the speed at which regions react to price changes in other regions. The third question we consider is the way the house prices in each region react to changes in GDP. We examine whether the house prices just follow the trend of GDP, or whether there is another process underlying the trends in house prices.

We apply our model to house price data for the Netherlands, comprising 76 regions for which we have quarterly data for the period 1985Q1-2005Q4. We find that the 76 regions can be grouped into two clusters. The first cluster consists of the capital *Amsterdam* and of rural areas that are close to larger cities, especially close to the *Randstad* (consisting of *Utrecht*, *Amsterdam*, *Den Haag*, *Rotterdam* and other cities in the area). This cluster has a higher growth rate than the other regions in the second cluster. The regions in the first cluster also react faster to changes in the house prices in *Amsterdam* than regions belonging to the second cluster. We find that both classes react equally fast to changes in GDP. However, the extent to which they react to GDP changes differs, where the house prices in regions in the first cluster are more strongly influenced by changes in GDP.

There are not many studies that describe regional house prices. Cameron *et al.* (2006) build a model from inverse demand equations. They have, however, only a limited number (9) of regions, and their model would not work in our situation where we have many more

(76) regions, as we will describe below. Malpezzi (1999) constructs an error correction model for regional house prices. The parameters of his model are however not allowed to vary across regions. Holly *et al.* (2007) model US house prices at the state level. Their model is ‘fully heterogenous’ in the sense that it has different parameters for each region. In this paper we cover the middle ground, that is, the model parameters are allowed to vary across *groups* of regions but not across each region individually.

Before we propose our latent-class model for a large panel of house price series, we first provide some details on the data in Section 2. We consider two decades of quarterly house prices on 76 regions in the Netherlands. We discuss their trending behavior by performing panel unit root tests and we also show that the growth rates in different regions show strong cross-correlations. Using multidimensional scaling techniques we get a first impression if and how these 76 regions could get clustered. Then, in Section 3, we put forward our model specification, highlighting the underlying data-driven clustering mechanism. In addition, we describe the method used for parameter estimation. In Section 4 we first present our estimation results, and give interpretation to the various outcomes. Next, we take a look at impulse response functions of the house prices with respect to a shock in GDP, in the house prices in one leading region and in the interest rate. In Section 5 we conclude with some limitations and we outline topics for further research.

## 2 Data

The Dutch real estate agent association [NVM] publishes quarterly data on house prices for  $N = 76$  regions in the Netherlands. Our dataset covers the sample period 1985Q1-2005Q4 ( $T = 84$  quarters). Hence, we have a panel database where both the cross-section dimension  $N$  and the time dimension  $T$  are fairly large.

The way the country is divided into 76 regions is determined by the NVM. Macroeconomic data, such as output and inflation, are not available for this particular specification of regions. Other (macro) variables that we use in our model are therefore measured at

the country level. In particular, this concerns the interest rate (obtained from the Dutch Central Bank) and quarterly real GDP (from Statistics Netherlands). The GDP series is available until 2005Q2. We obtain real house prices by deflating with the consumer price index [CPI] (from Statistics Netherlands). In addition, we seasonally adjust the real GDP series using the Census X-12 algorithm (available in EViews 5.1). We denote the real house price in region  $i$  at time  $t$  as  $p_{i,t}$ , and  $y_t$  to denote real GDP.

Figure 1 shows time series of  $\log(p_{i,t})$  for three specific regions: *Noordwest-Friesland*, which usually is the least expensive region, *Bunnik/Zeist*, which usually is the most expensive region, and *Amsterdam*, which is in between. On top we also plot  $\log(y_t)$  (scaled to limit the size of the vertical axis in the graph). Comparing the graphs in Figure 1 suggests that real house prices increase slightly faster than real GDP. Prices in *Bunnik/Zeist* and *Amsterdam* show substantial variations in the trend growth rate over time, with alternating periods of steep price increases and of stable or falling prices. Especially the ‘hump’ in the prices around 2000 stands out clearly. This suggests that the trend in the house prices is stochastic rather than deterministic. Furthermore, as the trending behavior of the different price series seems quite similar regional house prices may well be cointegrated.

## 2.1 Unit roots and cointegration

To test whether these visual impressions from Figure 1 can be given more formal statistical support, we perform panel unit root tests on the regional house prices. Two of the most popular tests in the literature are those from Levin *et al.* (2002) [LLC] and Im *et al.* (2003) [IPS], see Breitung and Pesaran (2008). These tests have as null hypothesis the presence of a unit root in all the series in the panel. The alternative hypotheses are different however. Levin *et al.* (2002) assume that the house price dynamics are the same for each region, and therefore the alternative hypothesis is that all regional house prices are stationary. Im *et al.* (2003), however, have as alternative hypothesis that at least one regional house price is stationary. Both these tests assume that there is no cross-

correlation between different series in the panel. In fact, they are not consistent if such a dependency is present, which is quite likely in our case. Alternative tests that do allow for cross-section dependence are available, like the one in Moon and Perron (2004), but these usually rely on asymptotics that require  $T$  to be much larger than  $N$ , while in our case they are about equal.

To meet our data characteristics, we therefore employ the cross-sectionally augmented IPS [CIPS] test, recently developed in Pesaran (2007). This allows for cross-sectional dependence, and is also valid when  $N$  is larger than  $T$ . The idea of the CIPS test is to add the cross-section averages of the lagged levels and first differences to the familiar augmented Dickey-Fuller [ADF] regression equation. If it can be assumed that the cross-correlations are caused by a common factor, then this common factor must also be present in the cross-section averages. Adding these to the ADF equations should then get rid of the common factor in the residuals and thus correct for the presence of cross-correlations.

As the CIPS test is known to have reduced power relative to the IPS and LLC tests in case cross-correlation is not present, we test whether we really should use the CIPS test instead of these simpler tests. For this purpose we use the cross-section dependence [CD] test of Pesaran (2004) and the adjusted LM [ $LM_{adj}$ ] test of Pesaran *et al.* (2008). These tests both use the cross-correlations between the residuals of the individual ADF regressions for the different regions. The CD test takes a simple sum which is scaled such that it has a standard normal distribution under the null hypothesis of no cross-sectional dependence. Therefore, the CD test has little power in the case that there are both positive and negative correlations such that the average is close to zero. The  $LM_{adj}$  test, however, is also valid in this case as it employs the squares of the cross correlations in the construction of the test statistic. However, the  $LM_{adj}$  test is less robust against non-normally distributed error terms and exhibits size distortions, especially when  $N$  is much larger than  $T$ .

Table 1 gives the result of these tests for the panel of quarterly growth rates in house prices  $\Delta \log(p_{i,t})$ , where  $\Delta$  denotes the first-difference filter, and of  $\log(p_{i,t}) - \log(p_{34,t})$ ,

which is the difference of each series with the log house prices in *Amsterdam* (region 34, see Appendix A). The reason for examining the log price differences with respect to *Amsterdam* is that finding these to be stationary, we can conclude that the house prices in each region are cointegrated and have  $(1, -1)$  cointegration relationships. The number of lagged (first) differences is allowed to vary across each (C)ADF equation and is determined by minimizing BIC. Adding a lagged variable means losing one observation, therefore we actually minimize  $\text{BIC}/T$ , see Cameron and Trivedi (2005, pp. 279) or the definition of BIC given in Franses and Paap (2001). Each (C)ADF regression equation contains an intercept and a trend.

From the second column of Table 1 we see that for the first difference of the log of house prices there is substantial cross-sectional dependence, according to both the CD and  $\text{LM}_{adj}$  tests. Next, we see that all three unit root tests reject the presence of a unit root in these growth rate series. Results for the difference between the log price in a region and the log price in *Amsterdam* (region 34) appear in the third column of Table 1. Again, the CD and  $\text{LM}_{adj}$  tests indicate that there is substantial cross-sectional dependence. Next, the LLC and IPS unit root tests do not reject the presence of a unit root, but the CIPS test does. Since the LLC and IPS tests are not valid in case of cross-sectional dependence, we rely on the CIPS test and conclude that the log house prices in each region are cointegrated. Note that the  $(1, -1)$  cointegration relationships suggested by the results in Table 1 are quite plausible. It means that the difference between the log of house prices, or, equivalently the ratio of house prices, in each region is a stationary process. This constrains the long-term growth of house prices in each region to be about the same.

## 2.2 Clusters

Before we turn to our conditional clustering analysis using latent class techniques we consider unconditional clustering based on the correlations of the house price growth rates or of the residuals of the ADF regressions used above. For this purpose, we use multi-

dimensional scaling [MDS], which results in the graphs shown in Figure 2 and 3. Although the graphs in these figures are rather different, they basically lead to the same conclusion that there is just a single cluster. Hence, a clustering of regions based only on the cross-correlations of the regional house prices is not a meaningful possibility. Apparently, we need a more sophisticated clustering method, perhaps based on latent classes, as we will propose in the next section.

### 3 The model

In this section we put forward the specification of the latent-class panel time series model for describing the regional house prices. We first discuss the characteristics of the model, and then we outline the parameter estimation procedure.

#### 3.1 Representation

Our starting point is the latent-class panel time series model developed by Paap *et al.* (2005). The crucial idea behind this model is that the individual time series may be grouped into a limited number of clusters. Within each cluster, a linear model is assumed to describe the dynamic behavior of the time series. The clusters are defined such that the model parameters are the same for all time series within a cluster, but they are different across clusters. Hence, this model covers the middle ground between a pooled regression model, where the model parameters are constrained to be the same for all regions, and a ‘fully heterogenous’ model, where the parameters are allowed to be different for each individual region. Whereas a pooled regression model may be too restrictive, a fully heterogenous model may be too flexible and ignores the possible similarities between regions. Finally, the key feature of the model of Paap *et al.* (2005) is that the number of clusters in the model as well as the allocation of the individual time series to different clusters is purely data-based. This avoids *ex ante*, and necessarily subjective, grouping of regions according to geographical location or economic or demographic characteristics,



for example.

In our model for quarterly growth rates of house prices we allow for more flexibility than was done in Paap *et al.* (2005). As mentioned, there are three research questions we want to answer with our model and each question corresponds to a different parameter that can vary across the latent classes. The first is whether the growth rates of house prices are the same across all regions. We therefore allow the clusters to have a different average growth rate by allowing for a class-specific intercept. To facilitate interpretation, we demean all other variables in the model such that the intercept is equal to the average growth rate of the house prices in the regions in a cluster.

The second question concerns the so-called ripple effect, see Cameron *et al.* (2006), which describes how price changes propagate across regions. To investigate this issue, we consider the parameters of the difference between the log price in region  $i$  relative to *Amsterdam*. As discussed in the previous section, we find that these variables are stationary, such that the log regional prices are cointegrated. As we fix the cointegration relationships at  $(1, -1)$ , this implies that the (log) ratio of the house prices in each pair of regions is stationary. Thus, the long-term growth rates of all regional house prices are constrained to be somewhat similar, but in the short term deviations from this equilibrium are quite possible. The adjustment parameter determines the speed at which a region moves towards the equilibrium. We allow this adjustment parameter to vary across the clusters. House prices in regions with a higher adjustment parameter will respond relatively quickly to deviations within other regions, while those regions with a lower adjustment parameter need more time to react.

As the house prices of a region are  $(1, -1)$  cointegrated with those of each other region, in principle it does not matter which region we choose to put in the cointegration relationship. Here we choose *Amsterdam* in each relationship as a kind of benchmark, also based on the idea that equilibrium deviations might generally start in the capital city *Amsterdam*, and then ripple through to the other regions of the Netherlands.

The third and last question we wish to answer with our model is whether the house

prices in regions follow the trend in real GDP. We add an error correction variable linking regional real house prices and real GDP, where the long-run parameter should be estimated. The adjustment parameter indicates how fast the house prices in a region react to changes in GDP.

Based on the above discussion, we propose the following latent-class panel time series model for regional house prices in the Netherlands

$$\Delta \log(p_{i,t}) = \beta_{0,s_i} + \beta_{1,s_i}[\log(p_{i,t-1}) - \log(p_{34,t-1})] + \beta_{2,s_i}[\log(p_{i,t-1}) + \gamma_{s_i} \log(y_{t-1})] + \eta_{i,t}. \quad (1)$$

The  $\beta$  and  $\gamma$  parameters are class-specific parameters, where the subscript  $s_i = 1, \dots, S$  denotes the latent class which region  $i$  belongs to with  $S$  being the number of latent classes. We denote the probability that a region belongs to latent class  $s$ , the mixing proportions, as  $\pi_s$ . Naturally it must hold that,  $0 < \pi_s < 1$  and that  $\sum_{s=1}^S \pi_s = 1$ .

Even though model (1) includes  $\log(y_{t-1})$ , which is the same for all equations, there may still be some cross-section correlation among the house prices that is not captured. Therefore, following Holly *et al.* (2007), we allow the error term  $\eta_{i,t}$  in (1) to be correlated across regions, but assume that this correlation is due to dependence on certain common factors. To be precise, we consider the specification

$$\eta_{i,t} = \alpha_{1,i} \Delta \log(y_{t-1}) + \alpha_{2,i} \log(I_{t-1}) + \alpha_{3,i} \Delta \log(p_{t-1}) + \varepsilon_{i,t}, \quad (2)$$

where  $I_{t-1}$  denotes the long-term interest rate at time  $t-1$ ,  $p_{t-1}$  denotes the average house price in the Netherlands at time  $t-1$  and where  $\alpha_{k,i}$  for  $k = 1, 2, 3$  are region-specific parameters. The residuals  $\varepsilon_{i,t}$  are now assumed to be independently normally distributed with a region-specific variance  $\sigma_i^2$ .

In the application below, we demean all variables in (1) and (2) and hence the intercepts  $\beta_{0,s}$  in (1) are equal to the average growth rates of the house prices in the latent classes  $s$  for  $s = 1, \dots, S$ .

### 3.2 Estimation

The parameters in our model (1) with (2) can be estimated as outlined in Paap *et al.* (2005), using the EM algorithm of Dempster *et al.* (1977). This makes use of the full data log-likelihood function, that is, the joint density of the house prices and the latent classes  $s_i$ , which we specify in detail below. The EM algorithm is an iterative maximization algorithm, which alternates between two steps until convergence occurs. In the first step (E-step) we compute the expected value of the full data log-likelihood function with respect to the latent classes  $s_i$ ,  $i = 1, \dots, N$ , given the house prices and the current values of the model parameters. In the second step (M-step) we maximize the expected value of the full data log-likelihood function with respect to the model parameters. As the model given the class memberships can be written as a standard linear regression, the M-step amounts to a series of (weighted) regressions. As the EM algorithm maximizes the log-likelihood function, the resulting estimates of the model parameters are equal to the maximum likelihood [ML] estimates. We can therefore compute standard errors of the estimates using the second derivative of the log-likelihood function.

Note that due to the presence of the term  $\beta_{2,s_i}[\log(p_{i,t-1}) + \gamma_{s_i} \log(y_{t-1})]$  the model in (1) is actually nonlinear in the parameters. To deal with this issue, we follow Boswijk (1994) and rewrite the model as

$$\Delta \log(p_{i,t}) = \beta_{0,s_i} + \beta_{1,s_i}[\log(p_{i,t-1}) - \log(p_{34,t-1})] + \beta_{2,s_i} \log(p_{i,t-1}) + \beta_{3,s_i} \log(y_{t-1}) + \eta_{i,t}, \quad (3)$$

where  $\beta_{3,s} = \beta_{2,s} \gamma_s$ . Note that (3) is linear in the parameters, which facilitates estimation. The ML estimate  $\hat{\gamma}_s$  can then be obtained from the ML estimates of  $\beta_{2,s}$  and  $\beta_{3,s}$  as  $\hat{\beta}_{3,s}/\hat{\beta}_{2,s}$ .

The full data likelihood function, that is, the joint density of  $\mathcal{P} = \{\{\Delta \log p_{i,t}\}_{t=1}^T\}_{i=1}^N$  and  $\mathcal{S} = \{s_i\}_{i=1}^N$  is given by

$$l(\mathcal{P}, \mathcal{S}; \theta) = \prod_{i=1}^N \left( \prod_{s=1}^S \left( \pi_s \prod_{t=1}^T \frac{1}{\sigma_i} \phi(\varepsilon_{i,t}^s / \sigma_i) \right)^{I[s_i=s]} \right), \quad (4)$$

where  $\phi(\cdot)$  denotes the probability density function of a standard normal random variable and  $\theta$  is a vector containing all model parameters. The error term at time  $t$  for region  $i$  belonging to cluster  $s$  is defined as

$$\varepsilon_{i,t}^s = \Delta \log p_{i,t} - x'_{i,t} \beta_s - w'_t \alpha_i, \quad (5)$$

where  $x_{i,t}$  is the  $4 \times 1$  vector with the regressors appearing in (3) and  $\beta_s$  contains the corresponding parameters for cluster  $s$ . Similarly,  $w_t$  is the *3times1* vector with common factors in the specification for  $\eta_{i,t}$  in (2), and  $\alpha_i = (\alpha_{1,i}, \alpha_{2,i}, \alpha_{3,i})'$  containing the parameters for region  $i$ .

The expectation of the full data log-likelihood function with respect to  $\mathcal{S}|\mathcal{P}, \theta$  [E-step] is given by

$$\mathcal{L}(\mathcal{P}; \theta) = \sum_{i=1}^N \left( \sum_{s=1}^S \hat{\pi}_{i,s} \left( \ln p_s + \sum_{t=1}^T -\frac{1}{2} \ln \sigma_i^2 - \frac{1}{2} \ln 2\pi - \frac{(\varepsilon_{i,t}^s)^2}{2\sigma_i^2} \right) \right), \quad (6)$$

where  $\hat{\pi}_{i,s}$  denotes the conditional probability that region  $i$  belongs to class  $s$  which is equal to

$$\hat{\pi}_{i,s} = \frac{\pi_s \prod_{t=1}^T \frac{1}{\sigma_i} \phi(\varepsilon_{i,t}^s / \sigma_i)}{\sum_{k=1}^S \pi_k \prod_{t=1}^T \frac{1}{\sigma_i} \phi(\varepsilon_{i,t}^k / \sigma_i)}. \quad (7)$$

In the M-step, we need to maximize (6) with respect to the parameters  $\beta_s$ ,  $\pi_s$ ,  $s = 1, \dots, S$  and  $\alpha_i$ ,  $\sigma_i^2$  for  $i = 1, \dots, N$ . We perform this maximization step sequentially. First, we optimize over  $\beta_s$  keeping the other parameters fixed. This can be done by a simple weighted regression of  $\Delta \log(p_{i,t}) - w'_t \alpha_i$  on  $x_{i,t}$ . The weights are given by  $\sqrt{\hat{\pi}_{i,s}} / \sigma_i$ . Clearly, we want regions with a larger probability of belonging to class  $s$  to have a larger weight in estimating  $\beta_s$ . At the same time, regions with a larger standard deviation of the error term  $\sigma_i$  should get a smaller weight, as their house prices contain relatively more noise and less information about  $\beta_s$ . Each  $\beta_s$ ,  $s = 1, \dots, S$  is estimated in a separate weighted regression.

Second, we optimize the log-likelihood function over  $\alpha_i$  for  $i = 1, \dots, N$ . We do this by regressing  $\sum_{s=1}^S \hat{\pi}_{i,s} [\Delta \log(p_{i,t}) - x_{i,t} \beta_s]$  on  $w_t$ . The dependent variable in this regression is the conditional expectation of  $\eta_{i,t}$ . We perform these regressions for each region separately.

Next, the new estimate of  $\sigma_i^2$  is given by

$$\sigma_i^2 = \frac{1}{T} \sum_{t=1}^T \sum_{s=1}^S \hat{\pi}_{i,s} (\varepsilon_{i,t}^s)^2 \quad (8)$$

for  $i = 1, \dots, N$ . Finally, the mixing proportions are updated by averaging the conditional class membership probabilities, that is,

$$\pi_s = \frac{1}{N} \sum_{i=1}^N \hat{\pi}_{i,s} \quad (9)$$

for  $s = 1, \dots, S$ .

As we maximize over the parameters sequentially in the M-step, we do not reach the optimum of the expected full data log-likelihood function (6) in each iteration of the EM-algorithm. We can repeat the individual update steps until convergence, but this is not necessary. Indeed, Meng and Rubin (1993) have shown that an increase in the full-data log-likelihood function in the M-step is sufficient for the EM algorithm to converge to the maximum of the log-likelihood function.

Determining the appropriate number of latent classes is not straightforward. We cannot use a standard statistical test, due to the Davies (1977) problem of unidentified nuisance parameters under the null hypothesis. The usual approach is using a criterion function balancing the fit and the complexity of the model, where the model fit is measured by the value of the log-likelihood function while the number of model parameters provides a measure of complexity. The most well-known criteria are the Akaike information criterion [AIC] and the Bayesian information criterion [BIC]. Bozdogan (1994) suggests that the AIC should have a penalty factor of 3 instead of 2 in the case of mixture models. Indeed, Andrews and Currim (2003) show that this AIC-3 criterion outperforms other criteria. Bozdogan (1987) modifies the AIC into the so-called consistent Akaike information criterion [CAIC], which is almost equal to BIC. He shows that when the sample size is large the CAIC and BIC criteria perform better than AIC. We will consider all four criteria below.

## 4 Empirical results

In this section we discuss the results of applying our model to the regional house price data for the Netherlands described in Section 2. The effective sample period ranges from 1985Q3 (because we have  $\Delta \log(p_{t-1}) = \log(p_{t-1}) - \log(p_{t-2})$  in our model) to 2005Q2 (because we only have real GDP data until 2005Q2), giving us  $T = 80$  data points in the time series dimension. To obtain a first impression of the extent of similarities across regions, we start by estimating a fully heterogenous model allowing for different parameters for each region. Next, we provide estimation results for the model with a limited number of latent classes. Finally, we consider impulse-response functions for three interesting scenarios to provide further interpretation of the model.

### 4.1 A fully heterogenous model

We first estimate the parameters in a fully heterogenous model, that is, we estimate the model in (1) with (2) allowing for different parameters for each individual region. This essentially is a model with  $S=76$  latent classes, in which case each region forms a separate class.

Figure 4 displays the histograms for the 76 estimated values for each of the parameters  $\beta_j$ ,  $j = 1, 2, 3$ , and  $\gamma$  in (1). The top left panel shows the intercepts, which equal the quarterly growth rates. These are all positive, reflecting the upward trend in the house prices, and range between 0.6% and 1.3% per quarter. The top right panel contains the adjustment parameters for the  $(1, -1)$  cointegration relationship with the house prices in the *Amsterdam* region. We find a few positive values here, which is not as expected as these imply divergence between the prices in those regions and in Amsterdam. Similar results are obtained for the adjustment parameter for the cointegration term with GDP, in the bottom left panel of Figure 4. Finally, the histogram in the bottom right corner shows the parameter  $\gamma$  in the cointegration relationship with GDP, which we expect to be negative as we expect the house prices and GDP to move in the same direction.

We can see from these graphs that some form of aggregation may be useful, as we now

get a wide variety of parameter estimates, with sometimes quite implausible results. At the same time, this variety also suggests that we should perhaps better not restrict the parameters to be the same across all regions. Hence, it may be optimal to allow for a limited number of different clusters.

## 4.2 A model with latent classes

A major issue for successful application of the latent-class panel time series model is of course determining the appropriate number of latent classes. As discussed in Section 3.2, we consider four different information criteria for this purpose. Table 2 shows the values of these criteria for models with one to five and 76 classes. For all criteria, we see that going from a homogenous model (with a single class) to two classes amounts to a relatively large improvement in the balance of model fit and complexity. Both AIC and AIC-3 prefer four classes, while both BIC and CAIC prefer two classes. We choose to focus on the model with two classes, as preferred by the BIC and CAIC criteria, also because some of the (unreported) estimation results with four classes turn out to be difficult to interpret.

The estimation results for the model with two latent classes are given in Table 3. Additionally, Table 4 gives the results for a series of Wald tests which we use to examine whether the parameters for the different classes are significantly different from each other. The estimation results show that the regions in the two latent classes do indeed differ from each other in several important respects. First, the estimated intercepts show that the average growth rates are significantly different at the 5% level, with class 1 having a higher growth rate than class 2.<sup>1</sup> The average growth rate in class 1 is equal to 1.2% per quarter, or 4.9% annually, while the house prices in class 2 grow with 1.0% per quarter, or 4.2% annually.

Second, we find that prices in regions in the high-growth class react faster to changes in the house prices of *Amsterdam* than those in the low-growth class. The interpretation

---

<sup>1</sup>Recall that we demeaned all other variables the model, so the intercepts represent the average growth rates.

is that changes in house prices start in *Amsterdam*, first disseminate to the regions in class 1, and then to regions in class 2. So, this is the ripple effect in the Dutch house prices.

Third, examining the cointegration relationship with GDP, we find that the high-growth class has a larger adjustment parameter, but the difference across the two classes is not significant as can be seen from Table 4. The cointegration relationship itself, however, is significantly different across the classes. For class 1, it is  $(1, -1.51)$ , meaning that in the long run the house prices in the regions in this cluster grow about 50% faster than GDP. In class 2 the cointegration relationship is  $(1, -0.93)$ , which is not significantly different from  $(1, -1)$ . This implies that if GDP increases, so do the house prices in the regions in this cluster, and roughly by the same amount.

The parameters in (2) are region-specific, and full estimation results are not reported to save space. Only 17% of the  $\alpha_{1,i}$  parameters is significant, suggesting that the impact of GDP on the house prices is mostly captured by the cointegration term. The  $\alpha_{2,i}$  parameters are mostly negative, and only two regions have an (insignificant) positive value. Furthermore, for 77% of the regions the  $\alpha_{2,i}$  parameter is significant at the 5% level, indicating that the interest rate indeed influences the house prices in the expected direction. The  $\alpha_{3,i}$  parameters, relating the growth of the house price in a region to growth of the average house price in the Netherlands in the previous quarter, is positive for 84% of the regions, but only significant for about one-third of these regions.

## The latent classes

The parameter estimation results obviously become more interesting if we know which regions belong to each of the two classes. Therefore, we compute the conditional class membership probabilities using (7). The resulting classification of the regions is shown in Figure 5. Regions are colored based on  $\hat{\pi}_{i,1}$ , the probability of belonging to the high-growth class. Regions are colored in five shades of grey. For the regions that are colored in the darkest shade it holds that  $\hat{\pi}_{i,1} > 0.8$ . For regions colored in subsequently lighter



shades of grey it holds that  $0.6 < \hat{\pi}_{i,1} \leq 0.8$ ,  $0.4 < \hat{\pi}_{i,1} \leq 0.6$ ,  $0.2 < \hat{\pi}_{i,1} \leq 0.4$  or  $\hat{\pi}_{i,1} \leq 0.2$ . It can be seen that most regions are either very dark or very light, suggesting that the classification is very clear for most regions. In fact, the average value of  $\max(\hat{\pi}_{i,1}, \hat{\pi}_{i,2})$  is equal to 0.93.

We find that the high-growth class contains mainly rural regions surrounding the big cities in the Netherlands. The main exception is the inclusion of *Amsterdam* itself. Other larger cities included are *Delft* and *Tilburg*. The regions in this class cover parts of *Noord-Brabant*, the *Veluwe* and the south of *Friesland*. Even though the East belongs almost completely to class 1, the larger cities of the East, like *Zwolle*, *Almelo*, *Hengelo*, *Enschede*, *Arnhem* and *Nijmegen* are part of class 2.

Class 2 contains different types of regions. First, it contains many large cities in different parts of the country, like *Breda* and *Groningen*, as well as almost all of the regions in the *Randstad*, the densely populated western part of the country. At the same time some rural regions, like *Zeeland*, *Zuid-Limburg* and regions in the North belong to this class with high probability. Note that these rural regions are not as close to the *Randstad* as most of those in class 1.

A possible explanation for our results is the increased number of commuters that live in the regions belonging to class 1 and who work in the large western cities. If the number of commuters increases, it is likely that they move to regions in cluster 1, as these are still at traveling distance from the *Randstad*. This development has two consequences for the regions in class 1. First, the average income in these regions is likely to increase, as the individuals who move away from the cities are relatively wealthy. The second consequence is an increase of housing quality in these regions, as wealthier people leaving the cities will increase the demand for more luxurious houses.

These potential structural changes within the regions of cluster 1 are consistent with all of our findings. First, the increase in housing quality will result in a larger increase in the average house prices in class 1 as compared to class 2. Our second finding is that house prices in these regions react faster to changes in the house prices in *Amsterdam*. As

a lot of the commuters actually work in the *Amsterdam* region, their decision the move might be influenced by the house prices in *Amsterdam* itself. Our last and most striking finding is that the house prices in class 1 increase roughly 50% faster than GDP. Note however that the increase is not corrected for higher housing quality.

### 4.3 Impulse-response functions

To give further interpretation to our estimation results we compute impulse-response functions for three interesting scenarios, each occurring in the second quarter of 2005. In the first scenario real GDP receives a shock of 10%. In the second scenario real GDP stays the same, but the house price in *Amsterdam* receives an upward shock of 10%. In the third scenario it is the long-term interest rate that receives a shock of 10%. We forecast the house prices for each of the three scenarios and compare with a no-change scenario, for the subsequent three-year period from 2005Q3 until 2008Q2.

In order to compute the impulse responses up to 12 quarters ahead, we also need forecasts for GDP and the interest rate, as these variables also affect house prices, see (1). Here we assume that the interest rate stays the same during the forecast period. In scenario 3, the interest rate is higher, but still assumed to be constant over the whole forecast period. To obtain forecasts for GDP we construct a simple  $AR(q)$  model with intercept for  $\Delta \log y_t$ . We choose  $q$  based on out-of-sample forecasting performance, where we use the last 3 years as a hold-out sample. It turns out that  $q = 8$  gives the best performance.

Figure 6 shows the impulse-response functions of the log house prices with respect to the log of GDP. The  $y$ -axis gives the relative change in house prices between the two scenarios, that is, a value of 0.10 means that the house price is 10% higher than the reference forecast. We show the impulse-response functions for only three regions: *Noordwest Friesland*, *Bunnik/Zeist*, and *Amsterdam*. Of these three, *Amsterdam* belongs to class 1, the high-growth cluster, while the other two regions belong to class 2 with high probability.

We find that the effect of an increase in GDP is initially negative, which is caused by the relatively large negative  $\alpha_{1,i}$  parameter for all three regions. The negative effect is, however, probably not significant as the  $\alpha_{1,i}$  estimates are not significant for these regions. After a while the house prices are higher compared to the reference forecasts, and as expected in *Amsterdam* the prices increase fastest.

Figure 7 shows the impulse-response functions of an increase in the house price in *Amsterdam* with 10%. Naturally, we find that *Amsterdam* has initially a higher price, though the difference with the reference forecast soon diminishes. After three years the impulse responses are about the same for all three regions. This illustrates the effect of the  $(1, -1)$  cointegration relationships between the house prices in each region.

In the last scenario, the log interest rate receives a shock, and increases from 2.06% to 2.27%. We find that the house prices are falling, but the effect is not very large. After three years the house prices are about 2% lower, as compared to the reference forecasts.

## 5 Conclusions

In this paper we developed a latent-class panel time series model for describing several key characteristics of regional house prices in the Netherlands between 1985 and 2005. An important feature of the model is that we cluster the regions in separate classes, where the price dynamics of house prices in regions within the same class are similar, while they are different across the classes. For the 76 regions in the Netherlands we find that two classes are sufficient. The first class contains mainly rural regions close to large cities. The second class contains both the larger cities and some more remote rural regions. The house prices in regions in the first class are characterized by higher average growth rates, a faster response to price changes in the house prices in *Amsterdam*, and stronger reactions to changes in GDP. These findings may be caused by the increased number of commuters. Indeed, the number of people working in the larger cities, but living in the regions of class 1, has increased substantially during our sample period.

Our model allows for the analysis of rather detailed data. To fully exploit its properties

one would want to analyse even further disaggregated data. The collection of such more detailed series is left to further research. Another issue for further research is to make the class probabilities dependent on certain explanatory variables.

## References

- Andrews, R. L. and I. S. Currim (2003), A Comparison of Segment Retention Criteria for Finite Mixture Logit Models, *Journal of Marketing Research*, **40**, 235–243.
- Boswijk, H. P. (1994), Testing for an Unstable Root in Conditional and Structural Error Correction Models, *Journal of Econometrics*, **63**, 37–60.
- Bozdogan, H. (1987), Model Selection and Akaike’s Information Criterion (AIC): The General Theory and its Analytical Extensions, *Psychometrika*, **52**, 345–370.
- Bozdogan, H. (1994), Mixture-Model Cluster Analysis Using Model Selection Criteria and a New Information Measure of Complexity, in H. Bozdogan (ed.), *Proceedings of the First US/Japan Conference on the Frontiers of Statistical Modeling: An Informational Approach*, vol. 2, Kluwer, Boston.
- Breitung, J. and M. H. Pesaran (2008), Unit Roots and Cointegration in Panels, in L. Matyas and P. Sevestre (eds.), *The Econometrics of Panel Data (Third Edition)*, Kluwer, Boston.
- Cameron, A. C. and P. K. Trivedi (2005), *Microeconometrics*, Cambridge University Press, New York.
- Cameron, G., J. Muellbauer, and A. Murphy (2006), Was there a British House Price Bubble? Evidence from a Regional Panel, Economics Series Working Papers 276, University of Oxford, Department of Economics.
- Davies, R. B. (1977), Hypothesis Testing when a Nuisance Parameter is Present Only Under the Alternative, *Biometrika*, **64**, 247–254.
- Dempster, A. P., N. M. Laird, and R. B. Rubin (1977), Maximum Likelihood Estimation from Incomplete Data via the EM Algorithm, *Journal of the Royal Statistical Society, Series B*, **39**, 1–38.

- Doornik, J. A. (2002), *Object-Oriented Matrix Programming using Ox*, 3rd edn., Timberlake Consultants Press, London.
- Franses, P. H. and R. Paap (2001), *Quantitative Models in Marketing Research*, Cambridge University Press, Cambridge.
- Holly, S., M. H. Pesaran, and T. Yamagata (2007), A Spatio-Temporal Model of House Prices in the US, Cambridge Working Papers in Economics No. 654, University of Cambridge, and CESInfo Working Paper Series No. 1826.
- Im, K. S., M. H. Pesaran, and Y. Shin (2003), Testing for Unit Roots in Heterogenous Panels, *Journal of Econometrics*, **115**, 53–74.
- International Monetary Fund (2004), The Global House Price Boom, *World Economic Outlook*, 1491–1517.
- Levin, A., C. F. Lin, and C. J. Chu (2002), Unit Root Tests in Panel Data: Asymptotic and Finite-Sample Asymptotics, *Journal of Econometrics*, **108**, 1–24.
- Malpezzi, S. (1999), A Simple Error Correction Model of House Prices, *Journal of Housing Economics*, **8**, 27–62.
- Meng, X. L. and D. B. Rubin (1993), Maximum Likelihood Estimation via the ECM Algorithm: A General Framework, *Biometrika*, **80**, 267–278.
- Moon, H. R. and B. Perron (2004), Testing for a Unit Root in Panels with Dynamic Factors, *Journal of Econometrics*, **122**, 81–126.
- Paap, R., P. H. Franses, and D. van Dijk (2005), Does Africa Grow Slower Than Asia, Latin America and the Middle East? Evidence from a New Data-Based Classification Method, *Journal of Development Economics*, **77**, 553–570.
- Pesaran, M. H. (2004), General Diagnostic Tests for Cross Section Dependence in Panels, Cambridge Working Papers in Economics No. 435, University of Cambridge, and CESInfo Working Paper Series No. 1229.

Pesaran, M. H. (2007), A Simple Panel Unit Root Test in the Presence of Cross-Section Dependence, *Journal of Applied Econometrics*, **22**, 265–312.

Pesaran, M. H., A. Ullah, and T. Yamagata (2008), A Bias-Adjusted LM Test of Error Cross Section Independence, *The Econometrics Journal*, **forthcoming**.

## A Regions by number

---

1	Noordoost-Groningen	27	Kop v. Noord-Holland	52	Dordrecht
2	Slochteren +s	28	Noord-Kennemerland	53	Gorinchem
3	Grootegast +s	29	West-Friesland	54	Culemborg/Dodewaard
4	Stad Groningen +s	30	Midden-Kennemerland	55	Ede +s
5	Zuidoost-Groningen	31	Waterland	56	Arnhem
6	Noord-Drenthe	32	Zaanstreek	57	Duiven/Westervoort
7	Opsterland	33	Zuid-Kennemerland	58	Elst +s
8	Oost-Friesland	34	Amsterdam	59	Nijmegen
9	Noordwest-Friesland	35	De Bollenstreek	60	Noordoost-Brabant
10	Zuidwest-Friesland	36	Haarlemmermeer	61	Uden +s
11	Zuid-Friesland	37	Almere	62	Oss +s
12	Zuidwest-Drenthe	38	Het Gooi	63	Den Bosch
13	Zuidoost-Drenthe	39	Amersfoort	64	Waalwijk/Drunen
14	Hardenberg +s	40	Barneveld	65	Zeeuwse Eilanden
15	Kop van Overijssel	41	Bunnik/Zeist	66	Zeeuws-Vlaanderen
16	Zwolle +s	42	Utrecht	67	Bergen op Zoom +s
17	Raalte +s	43	Woerden	68	West-Brabant
18	Almelo Tubbergen	44	Alphen	69	Breda
19	Hengelo Enschede	45	Leiden	70	Tilburg/Oirschot
20	Ruurlo Eibergen	46	Den Haag	71	Eindhoven +s
21	Doetinchem +s	47	Gouda	72	Zuidoost-Brabant
22	Zutphen +s	48	Delft +s	73	Noord-Limburg
23	Apeldoorn +s	49	Rotterdam	74	Weert +s
24	Nunspeet +s	50	Westland	75	Roermond +s
25	Lelystad	51	Brielle/Goeree	76	Zuid-Limburg
26	Den Helder/Texel				

---

Note: +s means including surrounding area.






Figure 1: Log house prices for 3 distinct regions, and log GDP.




Figure 2: Multidimensional scaling plot of the regions, based on the correlations of the first differences of the log house prices over the period 1985Q1-2005Q4.




Figure 3: Multidimensional scaling plot of the regions, based on the correlations of the residuals of the ADF regressions for the log house prices over the period 1985Q1-2005Q4.

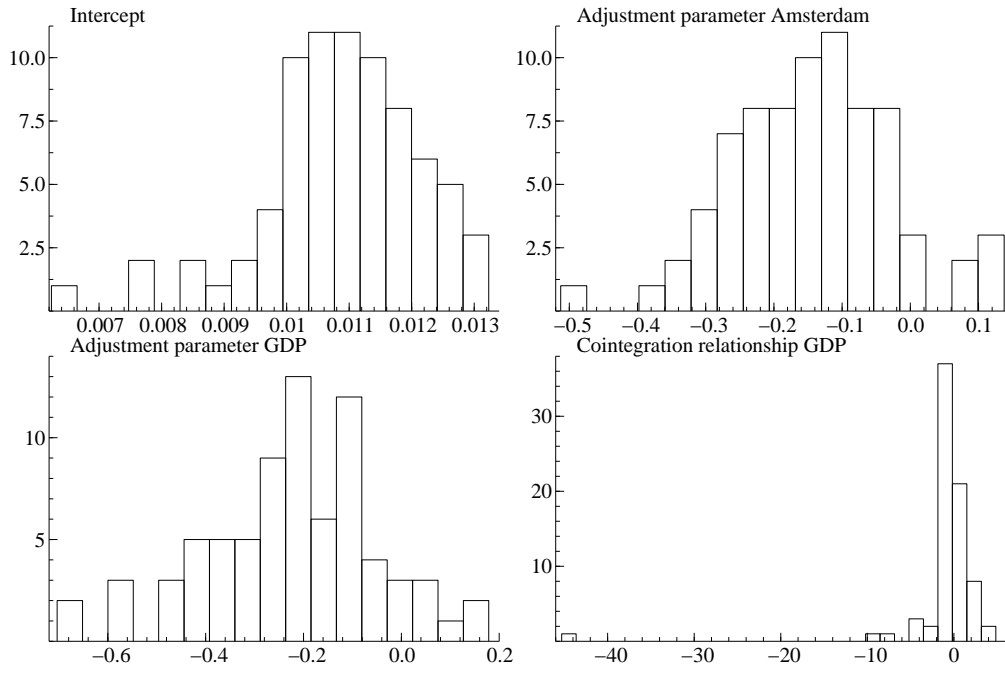


Figure 4: Histograms of the estimated values of the parameters  $\beta_j$ ,  $j = 1, 2, 3$ , and  $\gamma$  in (1) in the fully heterogenous model with 76 classes.




Figure 5: Clustering of regions. Regions with a high probability of belonging to the high-growth class are colored dark, regions with a low probability of belonging in the high-growth class are colored lighter. The numbers inside the regions correspond to the ones in Appendix A.




Figure 6: Impulse-response function of  $\log(p_{i,t})$  with respect to  $\log(y_t)$  for 3 regions.




Figure 7: Impulse-response function of  $\log(p_{i,t})$  with respect to  $\log(p_{34,t})$  for 3 regions.




Figure 8: Impulse-response function of  $\log(p_{i,t})$  with respect to  $\log(I_t)$  for 3 regions.



Table 1: Results of the CD test, the  $LM_{adj}$  test and three different tests for a unit root for two series (boldface numbers indicate rejection of the null hypothesis).

Test	Series	$\Delta[\log(p_{i,t})]$	$\log(p_{i,t}) - \log(p_{34,t})$
CD <sup>a</sup>		<b>92.0</b>	<b>144.2</b>
$LM_{adj}$ <sup>a</sup>		<b>60.4</b>	<b>175.1</b>
LLC <sup>a</sup>		<b>-61.2</b>	2.0
IPS <sup>a</sup>		<b>-55.9</b>	1.9
CIPS <sup>b</sup>		<b>-8.9</b>	<b>-3.5</b>

<sup>a</sup> Test statistic is asymptotically distributed as normal

<sup>b</sup> Tables with critical values for various values for  $N$  and  $T$  are given by Pesaran (2007), in the presence of and intercept and a trend in the CADF equations and for  $N = T = 70$  the critical value at the 95%-level is  $-2.58$ , for  $N = T = 100$  it is  $-2.56$ .

Table 2: Criteria values for different numbers of latent classes (boldface numbers indicate the optimum).

Criterion \ $S$	1	2	3	4	5	76
AIC	-3.985	-4.102	-4.103	<b>-4.104</b>	-4.104	-3.988
AIC3	-3.934	-4.050	-4.050	<b>-4.051</b>	-4.050	-3.875
BIC	-3.645	<b>-3.756</b>	-3.752	-3.748	-3.742	-3.234
CAIC	-3.594	<b>-3.705</b>	-3.699	-3.695	-3.688	-3.121

Table 3: Estimation results for  $S = 2$  latent classes.

Class	Estimate	Standard error
<i>intercept <math>\beta_{0,s}</math></i>		
1	0.012	0.001
2	0.010	0.000
<i>adjustment parameter Amsterdam <math>\beta_{1,s}</math></i>		
1	-0.163	0.015
2	-0.114	0.010
<i>adjustment parameter GDP <math>\beta_{2,s}</math></i>		
1	-0.122	0.010
2	-0.097	0.008
<i>cointegration relationship GDP <math>\gamma_s</math></i>		
1	-1.511	0.076
2	-0.935	0.103
<i>mixing proportions <math>\pi_s</math></i>		
1	0.381	0.071
2	0.619	.

Table 4: Wald tests for equality of the parameters across the two classes in (1).

Restriction	Wald statistic	p-value
$\beta_{0,1} = \beta_{0,2}$	4.01	0.045
$\beta_{1,1} = \beta_{1,2}$	7.13	0.008
$\beta_{2,1} = \beta_{2,2}$	3.37	0.066
$\gamma_1 = \gamma_2$	18.54	0.000