

School of Finance



University of St.Gallen

THE ROLE OF SPATIAL AND TEMPORAL STRUCTURE FOR RESIDENTIAL RENT PREDICTIONS

ROLAND FÜSS

JAN KOLLER

WORKING PAPERS ON FINANCE NO. 2015/23

SWISS INSTITUTE OF BANKING AND FINANCE (S/BF – HSG)

NOVEMBER 2015



The Role of Spatial and Temporal Structure for Residential Rent Predictions

Roland Füss* Jan Koller[†]

November 18, 2015

Abstract

This paper examines the predictive power of five linear hedonic pricing models for the residential market with varying complexity in their spatial and temporal structure. In contrast to similar studies, we extend the out-of-sample forecast evaluation to one-day-ahead predictions with a rolling estimation window, which is a reasonable setting for many practical applications. We can show that in-sample fit and cross-validation prediction accuracy improve significantly when we account for spatial heterogeneity. In particular, for one-day-ahead forecasts, the spatiotemporal autoregressive (STAR) model demonstrates its superiority compared to model specifications with alternating spatial and temporal heterogeneity and dependence structures. In addition, sub-market fixed-effects, constructed on the basis of statistical TREE methods, further improve the results of predefined local rental markets.

JEL classification: *C1, C2, R3.*

Keywords: Classification and regression tree (CART) technique; forecast evaluation; hedonic pricing model; rental prices; spatiotemporal autoregressive (STAR) model.

* *Corresponding author:* Swiss Institute of Banking and Finance (s/bf), University of St.Gallen, Rosenbergstrasse 52, 9000 St. Gallen, Switzerland and Research Associate at the Centre for European Economic Research (ZEW), Mannheim, Germany; Phone: +41 (0)71 224-7055, Fax: +41 (0)71 224-7088; Email: roland.fuess@unisg.ch

[†]Swiss Institute of Banking and Finance (s/bf), University of St.Gallen, Rosenbergstrasse 52, 9000 St. Gallen, Switzerland, Phone: +41 (0)71 224-7059, Fax: +41 (0)71 224-7088, Email: jan.koller@c-alm.ch.

1 Introduction

Hedonic pricing models are by far the most extensively applied approach to predicting rental and residential sales prices. In their most basic form, they model real estate prices as a linear function of the properties' attributes. Recent approaches have focused on the temporal and have paid particular attention to the spatial structure. In this paper, we compare the predictive power of hedonic models with a varying complexity of the spatial and temporal model structure. We focus mainly on the comparison between in-sample model fit as well as cross-validation prediction and one-day-ahead forecasting accuracy. We stress the fact that one-day-ahead forecasting, without ex-ante information, is the very nature of many practical applications but still has gained little attention in the literature.

The accurate prediction of rental prices is essential for participants in the real estate market such as investors, regulators, and policy makers. For instance, rent forecasts play a crucial role in property valuation adopted in discounted cash flow models. In addition, when imputed market rents for owner-occupied dwellings are subject to income tax, as they are in Switzerland, local rent predictions serve as a tax base. A profound knowledge of the structure and development of rents in local and national housing markets is also of importance for public housing policy.

We restrict the analysis to fully parametric and linear models in order to keep the results interpretable for practical purposes. Starting with a simple *baseline* model including only physical attributes of the dwellings, we incrementally add elements, accounting for spatial and temporal effects. We show that the consideration of a few spatial and temporal components can increase forecast accuracy substantially. Likewise, significant spatial and temporal effects underpin the inconsistency of hedonic coefficients in the absence of these components, since the assumption of *i.i.d.* errors of the *baseline* model is violated. We demonstrate the empirical bias resulting from spatial heterogeneity, which in general is an issue for causal inference. Thus, we aim at improving the prediction accuracy for rental prices by accounting for both the space and time dimension in our

model specifications.

Our paper contributes to the literature in many ways: First, we analyze the importance of the temporal and spatial structure of a model for residential markets using an exclusive rental apartment data set for the canton Zurich, Switzerland. In countries where the majority of households are tenants, the higher turnover in the rental market reduces the estimation error due to the rich data availability. Thus, the Swiss housing market with its low homeownership rate serves as an ideal test ground to evaluate the predictive power of different hedonic rental price models.¹ Most notably, the models include a spatiotemporal autoregressive (*STAR*) specification as suggested in Pace, Barry, Clapp, and Rodriguez (1998).² Other contributions of the literature of rental price modeling include Sirmans and Benjamin (1991) and Valente, Wu, Gelfand, and Sirmans (2005). As in our study, the latter predict apartment rents using spatial econometric techniques. Unlike our approach, they model an explicit spatial process, whereas our *STAR* model makes strong structural assumptions about the spatiotemporal dependencies.

Second, through the choice of our model specifications, we are able to determine the *marginal effects* of increasing complexity concerning the spatial and temporal structure. Importantly and often underrated in similar studies, we pay particular attention to the different prediction approaches. In the forecasting application with a rolling estimation window, we account for the temporal heterogeneity in all models to the same degree. We then incrementally improve the *baseline* model by gradually accounting for spatial heterogeneity as well as temporal and spatial dependence. Finally, we account for spatial heterogeneity by constructing rental sub-markets using a classification and regression tree (*CART*) method. To our knowledge, only Clapp and Wang (2006) use *CART* methodology for sub-market construction in the context of housing price *sales* data.

¹On an international scale, Switzerland has one of the lowest homeownership rates with approx. 44% according to Eurostat (2015).

²See also Pace, Barry, Gilley, and Sirmans (2000) for an application of the *STAR* model to housing prices in Baton Rouge, Louisiana.

Our empirical results show that including sub-market dummy variables constituted by ZIP codes significantly improves prediction accuracy. However, while this approach systematically ignores the spatial structure, sub-market construction based on regression tree technique displays a superior estimation strategy. The forecast evaluation shows that augmenting the hedonic model by spatially lagged variables is particularly successful in one-day-ahead forecasting. These results highlight the STAR model’s superiority compared to other specifications and emphasize the importance of local dynamics in one-day-ahead forecasting.

The rest of the paper is organized as follows: In Section 2, we provide an overview of theoretical and empirical approaches for modeling the temporal and spatial structure in hedonic pricing models. Section 3 introduces the research design and the models that are compared for prediction purposes. The empirical predictive power of these models is evaluated in Section 4. Section 5 summarizes our results and gives concluding remarks.

2 Space-Time Modeling

Hedonic pricing models have been the workhorse in the housing literature for decades. Since the seminal work by Rosen (1974), the capitalization of dwelling amenities in mainly linear hedonic functions has been studied in a large strand of literature.³ One of the most important distinguishing features of hedonic housing price models is an underlying assumption about the spatial and temporal structure. Moreover, hedonic pricing models take a wide range of functional forms. In this section, we demonstrate the range of spatial and temporal effects in empirical housing applications, which serve as individual components in our model comparison and evaluation. Concerning the functional form, the potential model complexity has largely been influenced by advances in computer technology. In particular, semi- and non-

³Instead of using house prices, the present study is one of the few to apply hedonic pricing techniques to rental price data. For reviews of the earlier literature see Bourassa, Hoesli, and Peng (2003), Malpezzi (2003) and Sirmans, Macpherson, and Zietz (2005).

parametric approaches have gained much attention. For an overview of these models, see Anglin and Gencay (1996) and McCluskey, McCord, Davis, Haran, and McIlhatton (2013), as well as Martins-Filho and Bin (2005) for the application of artificial intelligence methods in the field of real estate.

Despite the successful application of non-linear and non-parametric methods, we restrict our analysis to linear and parametric methods for two reasons: First, the focus of this study is on the marginal effects of the spatial and temporal components. Comparisons between models with different functional forms would unnecessarily dilute this intention. Second, semi- and non-parametric approaches often yield results that are difficult to interpret. For the purpose of price prediction, especially in practical applications, it is therefore questionable whether the parsimony of time- and space-discrete models should be dismissed in favor of more complex approaches. Restricting our focus to linear models throughout this study, we use the following general pricing model:

$$\begin{aligned}
p_{t,s} = & \alpha(t, s) + \beta X_{t,s} \\
& + \phi^S X_{t,s-\sigma} + \phi^T X_{t-\tau,s} + \phi^{ST} X_{t-\tau,s-\sigma} \\
& + \psi^S p_{t,s-\sigma} + \psi^T p_{t-\tau,s} + \psi^{ST} p_{t-\tau,s-\sigma} + \varepsilon_{t,s}
\end{aligned}$$

with: $\alpha(t, s)$ = temporal and spatial heterogeneous intercept,

β = hedonic prices of physical attributes,

ϕ^T, ψ^T = temporal dependence,

ϕ^S, ψ^S = spatial dependence, and

ϕ^{ST}, ψ^{ST} = spatiotemporal cross-effects,

(1)

where subscripts t and s refer to time and space, respectively. The rental price is denoted by $p_{t,s}$, i.e., it depends on time and location. The vector $X_{t,s}$ comprises the housing-specific characteristics. Note that the intercept α is potentially time- and location-specific, i.e., it accounts for *heterogeneity* along these dimensions. The coefficients ϕ and ψ measure

the effects of temporal and spatial *dependencies*. In particular, these are the coefficients of spatially and temporally lagged variables, indicated by subscripts $t - \tau$ and $s - \sigma$, respectively. The way to model the temporal and spatial effects is to a large extent determined by the underlying data structure. In the next section, we therefore present the most widely applied approaches as well as combinations of them.⁴

2.1 Aggregation of Space and Time

Housing naturally has a specific location, while observations of price offers or transactions naturally happen at specific points in time. The resulting observations are therefore principally continuous in time and space. For this reason, real estate prices should not be treated as cross-sectional, time-series or panel data, as Dube, Legros, and Thanos (2014) point out. The aggregation along either of these dimensions is problematic and a subject of current research.⁵ In addition, the way space and time are aggregated (or not) largely determines the set of feasible models and estimation methods. The models in this paper exhibit both *continuous* and *discrete* elements. In the *continuous* case, the location s is a pair of values containing longitude and latitude coordinates:

$$s = \{lat, lon\}. \quad (2)$$

In the *discrete* case, the Euclidian space is partitioned into K aggregate sub-markets implicitly defined by

$$k : I(s \in R_k), \quad (3)$$

where $k = 1 \dots K$ and $I(.)$ is the indicator function. Similarly, the time dimension can either be divided into discrete partitions (e.g., in years) or treated as continuous value. The left panel of Figure 1 illustrates the aggregation of time and space into discrete partitions by

⁴For instance, Liu (2013) uses time and regional dummies combined with a spatiotemporal autoregressive (STAR) model.

⁵See Dube, Legros, and Thanos (2014) for a detailed discussion of these issues.

including simple time and location dummies to account for potential *heterogeneity*. In the spatiotemporal specification, as illustrated in the right panel, spatial and temporal *dependency* is modeled. In particular, the current observation (dark shaded dot) may depend on temporally lagged ($t1$, $t2$, and $t3$) or spatially lagged ($s1$, $s2$, $s3$, and $s4$) observations.

[INSERT FIGURE 1 HERE]

2.2 Spatial Dimension

It is widely accepted that real estate transaction prices not only depend on physical attributes, but also on the local “market conditions.” There are two ways to interpret local market conditions: The first is to regard them as an unobserved effect that constitutes a locally homogeneous environment by differentiating the local from the global hedonic pricing coefficients. This is referred to as *spatial heterogeneity*. The second interpretation of local market conditions is traced back to a contagion effect between transactions. In particular, the occurrence of market transactions affects the price of (spatially lagged) observations. This is referred to as *spatial dependence*. Can (1990) refers to spatial heterogeneity as neighborhood effects and to spatial dependence as spillover effects.⁶ Although the theoretical distinction between these two effects has been widely accepted, the way these effects are incorporated into econometric models is a subject of ongoing discussion. Since spatial dependence leads to empirical spatial heterogeneity, it is hard to identify the true data generating process (*DGP*). Therefore, the choice of the model structure is, to a certain degree, arbitrary. In most applications, the choice is simply driven by data availability.⁷

Spatial Heterogeneity. As it is for panel data, allowing for fixed-effects is the simplest and most obvious way to account for potential cross-sectional dependence. In

⁶See also Can (1992) for a discussion of this distinction in the spatial dimension.

⁷Anselin (2010), who provides an extensive review of the evolution of spatial econometrics over the last three decades, illustrates the trade-offs emerging from these issues.

spatial data, cross-sectional dependence often exhibits a spatial structure. If this is indeed the result of the spatial heterogeneous hedonic pricing function, including location-specific dummy variables in the case of a discrete space is a promising way to control for this type of heterogeneity. Several studies have found evidence for local differences in marginal effects of property characteristics (see, e.g., Goodman and Thibodeau (1998)). Pace, Barry, Clapp, and Rodriguez (1998) argue that the use of indicator variables is the simplest way to control for spatial and temporal dependencies. However, they also argue that this strategy is only feasible with a limited number of dependencies, since each indicator is related to a coefficient, which for large N induces an incidental parameter problem.

In this context, the definition of sub-markets is given particular attention and has been addressed by several studies (see, e.g., MacLennan (1977), Bourassa, Hamelink, Hoesli, and MacGregor (1999), and Michaels and Smith (1990)). The problem with administrative boundaries, e.g., based on MSA or ZIP code definitions, is that such sub-markets do not follow the spatial heterogeneity structure. Typically, housing data comes with information on the spatial position of the observations (e.g., ZIP codes or school districts). These predefined sub-markets have widely been used as dummy variables in hedonic regression equations. In the last two decades, several approaches have been examined for housing market segmentation based on statistical methods. Bourassa, Hoesli, and Peng (2003) use principal component analysis to identify orthogonal factors of the dwelling properties. On the basis of these factors, they utilize cluster analysis to identify homogeneous sub-markets. Goodman and Thibodeau (2003) apply a hierarchical method for sub-market construction. In this paper, we follow the approach suggested by Clapp and Wang (2006) and use the regression tree approach.

Goodman (1981) identifies three requirements that a concise sub-market definition must fulfill: homogeneity, parsimony, and contiguity.⁸ Indeed, finding contiguous homo-

⁸See also Clapp and Wang (2006) for theoretical requirements of sub-market construction.

geneous sub-markets is supposed to solve the spatial heterogeneity problem. Comparisons in terms of prediction accuracy show that constructed sub-markets are not superior to *a priori* sub-market definitions. In an analysis of residential property transactions in Auckland, Bourassa, Hoesli, and Peng (2003) find a better performance of sub-markets as defined by professional appraisers. They compare the predictive power of models with sub-markets based on a principal component analysis with a predefined set of sub-markets. A main result of their analysis is that sub-market definitions that disregard spatial contiguity are worse for practical purposes. Goodman and Thibodeau (2003) derive housing sub-markets using a data set for metropolitan Dallas containing 28,561 transactions of single-family houses. Using a segmentation approach based on a hierarchical model, they find that predictions based on smaller sub-markets outperform pooled estimates. However, the predictive performance of ZIP code pre-defined sub-markets is comparatively good. The literature on the construction and analysis of segments in the *rental* market is sparse. Of the few contributors, Des Rosiers and Theriault (1996) analyze rental property data in the Quebec area and identify five significantly different market segments.⁹

Spatial Dependence. Most recent studies have focused on the modeling of spatial dependence of house prices by applying new techniques from the field of spatial statistics to housing data (see, e.g., Dubin (1998), Can (1992), Pace and Gilley (1997), and Basu and Thibodeau (1998)). Similarly, standard lattice models such as spatial autoregressive (SAR) and conditional autoregressive (CAR) models have been widely applied in the real estate literature.¹⁰ A comparison and summary of the two different specifications of spatial dependence are given in Dubin (1998). A taxonomy of recent applications including spatial dependencies is provided in Bourassa, Cantoni, and Hoesli (2010), who differentiate between models with location dummies, lattice models of spatial dependence, and geostatistical models. Geostatistical models are characterized by a continuous spatial

⁹Concerning heterogeneity in a continuous sense, *soft boundary* methods, originally developed in the field of geostatistics, are available. In this study, we restrict our focus to *hard boundaries*. See Bourassa, Cantoni, and Hoesli (2007) for a more detailed discussion of geostatistical models with continuous spatial domains and lattice models.

¹⁰See, e.g., Anselin (2010) for an overview.

dimension. Using a data set of U.S. house price transactions, the authors find a good performance of geostatistical models. While their methodological approach is closely related to this study, we additionally test models with temporal dependence.

2.3 Temporal Dimension

While there has been a significant amount of research on spatial effects in the last three decades, the temporal dimension has gained very little attention. A few studies have used aggregate panel data to determine the fundamental drivers of housing prices (see, e.g., Hort (1998), Harter-Dreiman (2004), and Adams and Füss (2010)). In many hedonic pricing applications, the time dimension is entirely neglected since transactions take place within a relatively short time window (i.e., within one year). This special case of temporal aggregation shrinks the time dimension to a single point (see Sub-section 2.1).¹¹ In the following subsection, we discuss these remaining issues, i.e., temporal heterogeneity and dependence, which are conceptually similar to the spatial dimension.¹²

Temporal Heterogeneity. Temporal heterogeneity is the variation of hedonic parameters over time. In applications with a relatively long time period, temporal heterogeneity seems almost natural. For the purpose of hedonic price index construction, the time dimension is even the major focal point.¹³ Two estimation methods for index construction are widely applied: incorporating time dummies in the regression or estimating the hedonic equation separately for each period. In the first case, parameters in β are assumed constant over time and only the intercept parameter $\alpha(t, s)$ is allowed to vary over time. In the second case, the assumption of time-invariant hedonic slope parameters β is

¹¹For instance, Bourassa, Cantoni, and Hoesli (2010) treat the data as cross-sectional. By neglecting time effects, the implicit assumption is made that observations are simultaneous, which is generally not a realistic assumption. Still, if the time window of the presumable cross-sectional data is short, planned but not yet realized transactions or price offers may still have an effect on previous transactions.

¹²One important exception is that spatial econometrics allows for feedback loops due to the connectivity induced by the weighting matrix. The size of spillover and feedback effects depends on the estimated spatial lag as well as the strength of the spatial weights.

¹³See Diewert, de Haan, and Hendriks (2015) for a general discussion of residential real estate index construction and Diewert, Saeed, and Silver (2009) for a specific discussion of the time dimension in indexes. For residential real estate index construction in Switzerland, see Fahrlander (2008).

relaxed. Note that the rolling window approach in our out-of-sample predictions allows us to estimate β on a day-to-day basis. In the one-day-ahead forecasting approach, we therefore account, to a certain degree, for temporal heterogeneity in all models, similarly to Munneke and Slade (2001) and Liu (2013). A crucial distinction between the cross-validation and the one-day-ahead forecasting results is the incremental improvement by accounting for temporal heterogeneity.

Temporal Dependence. For discrete-time data, dynamic panel models have been the first choice in the literature on housing markets. However, these models are used if data comes in an aggregated form. This is not the case for individual housing transaction or offer price data. In the case of non-aggregated data, the continuous-time information naturally often comes with continuous-space data and is captured by spatiotemporal models, which typically include autoregressive components. The temporal dimension has only recently been addressed by a few studies. Particularly, Pace, Barry, Clapp, and Rodriguez (1998) formulate a model incorporating both the spatial and temporal dependencies in a concise methodological framework.¹⁴ As in our study, Liu (2013) applies the STAR model based on Pace, Barry, Clapp, and Rodriguez (1998) to a Dutch data set of housing transactions. The main reason for the increased application of spatiotemporal models has been the more structured way of modeling spatial patterns.

In addition, the temporal dependence structure reveals potential for forecasting applications. In particular, in many practical applications, only historical information is available for pricing a dwelling. In this case, time dummy methods are limited. This favorable aspect of the *STAR* model has not yet been stressed adequately in the literature, and thus is a main focus in the one-day-ahead forecasting application of this study.¹⁵ In real forecasting applications, the time dimension is indeed of particular importance since the identification of temporal dependencies is the main purpose. These dependencies are

¹⁴A review of the literature on spatiotemporal modeling is provided by Liu (2013).

¹⁵The reason for the minor interest in real forecasting might be that traditional applications have been ex-post analyses such as housing index construction.

measured by including spatial and temporal lags, while other informational inefficiencies are deliberately not exploited, thereby keeping methods rather simple. Hence, only two out of the five models incorporate temporal dependencies in our model comparison on in-sample estimation and forecast evaluation.

3 Methodology

3.1 Prediction Approaches

We compare the predictive power of the models using three different prediction approaches: *in-sample*, *cross-validation*, and *one-day-ahead forecasting*. As in other studies, the in-sample prediction is an overall model fit, while the cross-validation technique serves as an out-of-sample robustness test. In the in-sample prediction, the estimation is performed over the whole sample period. Cross-validation technique, in contrast, is based on splitting the data set into a training and test sample using a random re-sampling procedure. Finally, unlike in most hedonic prediction analyses, we also test the forecasting accuracy of the models.¹⁶ Many studies on housing prediction do not account for the temporal dimension and therefore make out-of-sample analyses in the form of cross-validation.¹⁷ In contrast, our data comes with a date of day, which allows us to model the temporal dependence. Therefore, the forecasting application is a real one-day-ahead forecast. For that purpose, we define a rolling estimation window of 500 days. In Section 4.5 we demonstrate that varying the window size does not substantially change the results. Based on the estimation window, we fit the models and forecast all rental prices observed on the subsequent day. Note that the rolling estimation window does, to some degree, account for temporal parameter heterogeneity.

[INSERT TABLE 1 HERE]

¹⁶Note that Liu (2013) makes one-step-ahead forecasts based on annual frequency.

¹⁷For instance, Bourassa, Cantoni, and Hoesli (2010) do not account for the temporal dimension.

In all three prediction approaches, we use several accuracy measures based on the residual \hat{u} from a regression on the logarithm of the rental price. An important and widely applied error variation measure is the mean squared error, defined as $MSE = n^{-1} \sum_{i=1}^n \hat{u}_i^2$. In order to get an impression of the size of the prediction, we report the square root of the MSE , i.e., the root mean squared error $RMSE = \sqrt{MSE}$. As a similar measure, but with less weight on the tails, we also present the mean absolute error $MAE = n^{-1} \sum_{i=1}^n \text{abs}(\hat{u}_i)$.

3.2 Model Specifications

The subsequent sections introduce the models that we consider in our comparative study on in-sample estimation and forecast evaluation. The specification of these models follows the incremental enhancement of a simple baseline model up to a spatiotemporal autoregressive (*STAR*) model. All models can be embedded in the general framework formulated in Equation (1). The *baseline* model is a simple benchmark model in terms of a hedonic pricing function, where the rental price is a linear combination of physical apartment characteristics stacked in the vector X . In addition, yearly time dummy variables absorb the temporal fixed-effects in the price variations, i.e.,

$$\textbf{Baseline:} \quad p_i = \alpha(s, t) + \beta X_i + \varepsilon_i, \quad \text{where } \alpha(s, t) = \alpha_t, \quad (4)$$

with $\alpha_t = \text{const}$ for $t = 1, \dots, T$.

A more advanced specification is the *ZIP* model, which accounts for spatial heteroscedasticity in a simple way. In particular, predefined sub-markets constituted by ZIP codes are included in the form of dummy variables:

$$\textbf{ZIP:} \quad p_i = \alpha(s, t) + \beta X_i + \varepsilon_i, \quad \text{where } \alpha(s, t) = \alpha_t + \alpha_s, \quad (5)$$

with $\alpha_t = \text{const}$ for $t = 1, \dots, T$ and $\alpha_s = I(s \in R_k)$ for $k = 1, \dots, K$ (ZIP codes).

The *TREE* model has the same functional form as the *ZIP* model:

$$\mathbf{TREE:} \quad p_i = \alpha(s, t) + \beta X_i + \varepsilon_i, \quad \text{where } \alpha(s, t) = \alpha_t + \alpha_s, \quad (6)$$

with $\alpha_t = \text{const}$ for $t = 1, \dots, T$ and $\alpha_s = f(\text{lat}, \text{lon}) = \sum_{m=1}^M c_m I(\{\text{lat}_i, \text{lon}_i\} \in R_m)$ for $m = 1, \dots, M$ (regions), where the prediction in each region is a constant c_m .

Hence, the only difference from the *ZIP* model constitutes the definition of the sub-markets. As outlined in Sub-section 2.2, the use of predefined boundaries may be problematic. Since the information about the precise location of apartments is given, we can optimize the sub-market homogeneity through statistical methods. For this purpose, we use the classification and regression tree (CART) approach.¹⁸ The basic principle of a tree is the partitioning of the predictor space such that predictions are kept constant at each partition. For our sub-market selection, we first estimate the *baseline* model, from which we derive the residuals \hat{u}_b . In the second step, the idea is to predict this residual surface by partitioning the coordinate space into M partitions. On each partition, the prediction is a (local) constant. This second task is solved by a regression tree (see Sub-section 3.3 for a detailed description of the *TREE* regression approach). As an additional restriction, price predictions, based on a model with regional indicator variables, require a sufficient number of data points in each region. Hence, sub-markets must be specified so that enough data is available in each partition for prediction purposes. In order to avoid over-fitting and to make results comparable to the baseline model, we restrict the number of sub-markets to the number of *ZIP* code regions, i.e., the *ZIP* and *TREE* model have the same number of parameters. This allows us to evaluate the marginal improvement derived from the two different sub-market definitions.

¹⁸Since Breiman, Friedman, Olshen, and Stone (1984) regression trees have gained popularity and have been extensively applied in computational statistics. The CART approach is also used by Clapp and Wang (2006) for sub-market identification.

3.3 Classification and Regression Tree (CART)

In the following sub-section we demonstrate the use of regression trees for sub-market construction. As mentioned above, for this purpose the residuals are estimated from a regression of rental price on the physical apartment attributes. Then, a regression tree is used to predict these residuals using the coordinates as the only predictor variables. Concepts and notations of this second step are largely based on Hastie, Tibshirani, and Friedman (2009).

We first make the assumption that the price of an apartment depends on physical characteristics X and the price of the local amenities A . Thus, the rental price of dwelling i can be written as

$$p_i = \alpha + \beta X_i + A_i + \varepsilon_i, \quad (7)$$

where ε_i is an error term. Assume that the local amenities are unobserved and depend on the coordinates (lat and lon), $A_i = A(lat_i, lon_i)$. We define

$$u_i \equiv A(lat_i, lon_i). \quad (8)$$

The goal of the sub-market construction is to identify M homogeneous partitions of the whole market, such that

$$A(lat_i, lon_i) = A_m, \quad \text{for all } \{lat_i, lon_i\} \in R_m, m = 1..M, \quad (9)$$

where $R_m, m = 1..M$ define partitions of the Euclidian space. In order to identify these sub-markets, we run a regression of the rental price p on the physical characteristics and denote the residual from this regression by u . Next, we use the regression tree to make a prediction for u_i by using the two predictors (lat_i, lon_i) , where $i = 1, \dots, N$.

The idea of the regression tree is to partition the Euclidian space into M regions,

where the prediction in each region is a constant c_m :

$$f(lat, lon) = \sum_{m=1}^M c_m I(\{lat_i, lon_i\} \in R_m). \quad (10)$$

The goal is to identify the splitting variable (lat or lon) and the splitting points such that the prediction error is minimized. This problem is solved by a recursive algorithm.¹⁹ The illustration in Figure 2 shows 60 tree-based sub-markets resulting from this procedure for the canton Zurich. The left panel shows the tree structure, which is pruned to only 7 splits for illustrative purposes. The right panel shows the resulting sub-markets, where each leaf of the tree corresponds to each partition on the map, which represents one sub-market.

[INSERT FIGURE 2 HERE]

3.4 A Spatiotemporal Autoregressive (*STAR*) Model

To derive the spatiotemporal autoregressive (*STAR*) model presented in Pace, Barry, Clapp, and Rodriguez (1998), we start with a simple hedonic model without spatial effects:

$$p_t = \alpha_t + x_t \beta_t + u_t, \quad (11)$$

where p_t is the price, x_t is a set of exogenous factors, and β_t contains a set of parameters. As the time subscripts indicate, we assume heterogeneity over time. Note that this simple hedonic model does not account for potential dependence in the error term u_t . To resolve this shortcoming, we assume structural dependence in the error term of the following form:

¹⁹For details we refer the reader to Hastie, Tibshirani, and Friedman (2009).

$$u_t = W_t u_t + \varepsilon_t, \quad (12)$$

where u is a white noise process and W is the weighting matrix. Following Liu (2013), the *STAR* model can be written in compact form as:

$$(I - W_t)r_{n,t} = (I - W_t)x_{n,t}\beta_t + \varepsilon_t. \quad (13)$$

In a standard spatial model, W_t would be the spatial weighting matrix. The *STAR* model, however, generalizes the weighting matrix to a set of spatial and temporal lags as well as combinations of these lags. In their most general specification, Pace, Barry, Clapp, and Rodriguez (1998) suggest that W_t be of the form:

$$W_t = \phi_{S,t}S + \phi_{T,t}T + \phi_{ST,t}ST + \phi_{TS,t}TS, \quad (14)$$

with S and T being the spatial and temporal lag operator, respectively. By sequentially applying these operators additional lags can be constructed. For instance, ST is the spatial lag of the temporal lag and TS the temporal lag of the spatial lag. Substituting (14) in (13) leads to the general regression equation:²⁰

$$\begin{aligned} p = & \alpha_t + x\beta_t + Sx\beta_{S,t} + Tx\beta_{T,t} + STx\beta_{ST,t} \\ & + TSx\beta_{TS,t} + \phi_{S,t}Sp + \phi_{T,t}Tp + \phi_{ST,t}STp + \phi_{TS,t}TS p + \varepsilon. \end{aligned} \quad (15)$$

This is the unrestricted form of a STAR model, where all spatial and temporal lags as well as the corresponding interactions are included. Note that x is a vector of k explanatory variables. The term β_{ST} is, for instance, a vector of k coefficients corresponding to the spatial lag of the temporal lag of the explanatory variables. In summary, the spatiotemporal autoregressive (*STAR*) model includes all combinations of

²⁰For illustrative purposes, we have simplified the notations by leaving out subscripts: $x_{n,t} \equiv x$ and $p_{n,t} \equiv p$.

linear spatial and temporal lags:

$$\begin{aligned}
\textbf{STAR:} \quad p_{t,s} = & \alpha(t, s) + \beta X_{t,s} \\
& + \phi^S X_{t,s-\sigma} + \phi^T X_{t-\tau,s} + \phi^{ST} X_{t-\tau,s-\sigma} \\
& + \psi^S p_{t,s-\sigma} + \psi^T p_{t-\tau,s} + \psi^{ST} p_{t-\tau,s-\sigma} + \varepsilon_{t,s}
\end{aligned} \tag{16}$$

A spatial weighting matrix is in general an $N \times N$ matrix with each element representing a bilateral weight, which is defined as a proximity measure. However, the observations with a non-zero (positive) weight are restricted for two reasons: First, observations in the future are not expected to have an effect and are thus excluded. Second, the spatial dependence is diminishing along the space. As a simplification and for computational reasons, only a predefined number of observations enters the weighting matrix, namely the closest previous m_S neighbors. Using this restriction, the spatial weighting matrix reduces to a sparse matrix with only $N \times m_S$ non-zero elements. Similarly, the temporal weighting matrix defines the weights given to past observations in order to compute the temporal lag. The number of prior observations considered is restricted to m_T .

The final model we introduce is the temporal autoregressive (*TAR*) model, which is a truncated version of the *STAR* model. Accordingly, the components correspond to those of the *STAR* model. It is an enhancement compared to the *TREE* model, since it includes a temporal moving average of rental prices.²¹ This model accounts for global time trends in the rental market in terms of autocorrelation:

$$\begin{aligned}
\textbf{TAR:} \quad p_{t,s} = & \alpha(t, s) + \beta X_{t,s} \\
& + \phi^T X_{t-\tau,s} + \psi^T p_{t-\tau,s} + \varepsilon_{t,s}
\end{aligned} \tag{17}$$

²¹The moving average in this context refers to *rolling average* and is not to be confused with a *moving-average (MA) process*.

4 Empirical Results

4.1 Data

We test our models using an exclusive data set consisting of 28,728 offered rental apartments between 2002 and 2014 in the canton of Zurich, Switzerland. The data pool stems from a Swiss rental market online platform. The available information about the rental objects and the descriptive statistics are shown in Table 2. Note that the rental price and surface of living area (*surface*) are in log terms. Panel D in Table 2 shows that the number of objects containing all relevant information increased from 2002 to 2009 and has been relatively stable since then. Because the sample ends in the first half-year of 2014, far fewer observations are available for the last year.

[INSERT TABLE 2 HERE]

The offered rent in housing and apartment markets is in general not the same as the contract price. In particular, if the offered price for a dwelling is too high, the contract may have not been finalized. We do not directly observe whether a contract associated with an advertisement is concluded or not. However, by filtering and eliminating the objects that have been re-advertised within a short time period, we are left with a good proxy for the market price.²² After dropping the re-advertisements, we have more than 28,000 observations at our disposal, which can be identified as market contract prices.²³ Figure 3 shows a map of the canton of Zurich with observations indicated by grey dots. The illustrated administrative borders are defined by ZIP codes.

[INSERT FIGURE 3 HERE]

²²In particular, if the same apartment occurs on the platform within 30 days, it is identified as a re-advertisement. In that case, only the last price of the (series of) re-advertisements is considered as a contract price. The previous advertisements are ignored in the forecasting application.

²³For an extensive discussion of listing prices, contract prices, and market prices, we refer to Knight (2002), Anglin, Rutherford, and Springer (2003), and Allen, Rutherford, and Thomson (2009).

With regard to our comparative study of prediction models, there are several advantageous features of the rental market in the canton of Zurich. First, the density of observations is relatively high, because the area is small (91 square kilometers) and the number of observations is large. This is because in Switzerland the percentage of renters is high. In the canton of Zurich, particularly, the home-ownership rate was as low as 7% in 2005 (see Statistik Stadt Zürich (2005)). The resulting high density of observations allows us to better identify the temporal and spatial structure.

4.2 Preliminary Analysis: Spatial Dependence

The use of methods including spatial components is only appropriate if rental prices obtain spatial dependence. The magnitude of the spatial dependence of the rental price is measured by the lags in the *TAR* and *STAR* model. The presence of spatial dependence can also be captured by a variogram. Particularly, we would expect spatial dependence in the residuals of the *baseline* model, which contains no spatial elements. Therefore, we estimate the empirical variogram for the residuals resulting from the regression of this benchmark model and compare it with the variograms for the residuals of the alternative models. For this purpose, we briefly introduce the variogram function, which represents the variance between observations conditional on a specific distance. Consider the set of all location pairs (s_i, s_j) with a distance h , denoted by $N(h)$. The semivariogram is then defined by the function $\gamma(h)$, with

$$2\gamma(h) = |N(h)|^{-1} \text{Var}(u(s_i) - u(s_j)), \quad (18)$$

where $|N(h)|$ is the number of distinct pairs in $N(h)$ and $u(s)$ the residuals from the hedonic regression. If the correlation between the residuals is non-constant for different distances, that is evidence of spatial dependence. In terms of the variogram function $\gamma(h)$, this means that the functional value is low (high) for small (large) distances. The

semivariogram for the residuals of the *baseline* model is shown in Figure 4. Because the semivariogram increases with distance, the residual indeed exhibits spatial dependence. Thus, the consideration of spatial elements might be appropriate. Nevertheless, the question of whether location dummies or spatial lags are more suitable to account for spatial heterogeneity is an empirical one. Yet, by including spatial components in the regression, the semivariogram should become flatter. The next sub-section delineates whether and how the spatial dependence will decrease with increasing complexity in the spatial specification.

[INSERT FIGURE 4 HERE]

4.3 In-Sample Estimation

In this sub-section, the coefficients derived from in-sample estimations are discussed with regard to their economic and statistical significance; the results on prediction accuracy are presented in the next sub-section. For the in-sample prediction, we run an OLS regression of the rental prices on a set of apartment attributes as well as spatial and temporal lags. Table 3 shows the regression results of the five models. The model complexity and the number of parameters increase from left to right, i.e., from the *baseline* to the *STAR* model.

Panel A of Table 3 shows that the most important physical attributes of the apartments have highly significant effects and expected signs. The effect with the highest significance is that for the apartment’s surface of living area. Since both rental price and surface of living area are in logarithms, the coefficients can be interpreted as elasticities. Thus, an increase of one percent in the apartment size is associated with a rental price increase of 0.54% to 0.63%, depending on the model. The size of this effect is comparable to findings from other studies on dwelling prices. The number of rooms also has an increasing effect on the rental price, i.e., dividing the same apartment area into more rooms increases the rent. However, the effect is small and ranges between 3.2% and 6.9% per

room, depending on the model specification.

An interesting finding that reflects the problem of spatial heterogeneity involves the coefficients for *Parking* and *Garage*. The former variable indicates that outside parking is available, while the latter indicates that a parking garage belongs to the apartment. The coefficients for these variables are significantly positive in all models, except in the baseline model. The explanation for this result is the following: in central regions where the availability of parking spaces is rare, prices for locations tend to be higher. This is a good example of how spatial heterogeneity leads to biased estimates. Closely related to this finding are the coefficients for apartments' age. Indeed, the *baseline* model would suggest that old apartments (70 years and more) are more expensive than new buildings. Again, this can be attributed to the fact that buildings in central locations are older on average, particularly those in historic city centers. Finally, the type of apartment does have a large effect on the rental price, which is presumably the result of the specific taste of households.

[INSERT TABLE 3 HERE]

Concerning the spatial and temporal dependence, the regression coefficients of the corresponding lags are shown in Panel B of Table 3. First and most notably, the coefficient for the spatial lag of the explained variables, denoted by $[lag(S) \ln(Rent)]$, is highly significant. The magnitude of the effect is 0.816, indicating an almost one-to-one spillover effect of spatially lagged objects. This finding is largely in line with Liu (2013), who estimates a coefficient between 0.87 and 0.93 (depending on the year) for the spatial dependence in house prices.

Second, the spillover effect of spatially lagged objects is also present for the apartments' attributes. The corresponding coefficients are those starting with $lag(S)$. All effects (except the *Parking* variable) are significantly negative. This finding is in line with the local competition concept, which states that the availability of objects with favorable

characteristics (i.e., large area, many rooms, and an elevator) has negative effects on rental prices. Therefore, we can conclude that the spatial effects estimated in the *STAR* model reflect reasonable market dynamics. Particularly, competition is taking place both through lower prices and better attributes of objects close to the apartment at hand.

In contrast to the highly significant spatial lags, the temporal lags $[lag(T)]$ are almost negligible. Similarly, the cross-effects $[lag(TS)]$ and $[lag(ST)]$ are of minor importance with the exception of the temporal lag of the spatial lag. The highly significant negative sign indicates that price spillovers are only temporary, i.e., the effect on the price is mean-reverting. In particular, the temporal lag is calculated on the basis of the last 180 observations ($m_T = 180$), which corresponds to an average time of approximately 31.3 days.

To gain more insights about the remaining spatial dependence of the *ZIP*, *TREE*, *TAR*, and *STAR* model, we estimate the semivariogram of the regression residuals. Figure 5 shows the corresponding semivariograms. A comparison to the *baseline* model (shown in Figure 4) reveals that the spatial dependence decreases incrementally with increasing spatial elements, with the exception of the *TREE* and *TAR* model, which include by construction only temporal components. These findings are largely in line with our expectations. In addition, it demonstrates that the *TREE* model has already captured a large portion of the spatial dependence. This result shows that a clear distinction between spatial dependence and spatial heterogeneity can empirically be diluted.

[INSERT FIGURE 5 HERE]

4.4 Forecasting Performance

The in-sample analysis suggests the presence of spatial heterogeneity. On the one hand, this is an important assumption for the consistency of the coefficients as the effect of the variable *Parking availability* has illustrated. On the other hand, accounting for spatial

heterogeneity can improve the prediction accuracy, which is the subject of this section. In order to prevent model overfitting, we additionally test the predictive power out-of-sample. Here, we use two different approaches: cross-validation and one-day-ahead forecasting. In the cross-validation approach, we create a random test sample, which consists of 15% of the observations. Since we still include region and year dummy variables, the random sampling is based on homogeneous groups made up of all year-region combinations. The rest (85%) of the data is used for the estimation. This specification allows a quasi-out-of-sample prediction, while spatial and temporal heterogeneity can still be incorporated in the form of time and region dummy variables. Panels A and B of Table 4 show the in-sample and cross-validation prediction results.

For the one-day-ahead forecasts, we use a rolling estimation window of 500 days. The reason for employing this approach is to ensure that strictly historical information is used for the estimation of the models, which is particularly important in practical prediction applications. The results of the out-of-sample predictions are shown in Panel C of Table 4.

[INSERT TABLE 4 HERE]

The overall fit of the models (Panel A) in terms of R^2 is similar in all three prediction approaches. The *baseline* model with only physical attributes fits the data quite well, with an R^2 of 0.72. Adding location dummy variables in the *ZIP model* raises this number to 0.81, i.e., by approximately 9 percentage points. As expected, the ZIP code classification already captures a substantial part of the spatial discrepancies in rental prices. This finding is in line with Bourassa, Hamelink, Hoesli, and MacGregor (1999), where sub-market choice using sophisticated statistical methods brought low improvement of the prediction accuracy. The use of tree-based dummy variables, however, increases the R^2 to about 0.85, while by comparison the *TAR* model is not better in terms of R^2 than the *TREE* model. This suggests that the past general rental price level, captured by a moving

average of the rents, does not contain much information for current rental prices.²⁴ The inclusion of spatial lags in the *STAR* model, in contrast, raises the R^2 to 0.88 (an increase of 3 percentage points). These results of the overall fit of the models are comparable to other similar studies.²⁵

A similar pattern to that of the R^2 can be identified for prediction measures *RMSE* and *MAE*. We focus on these measures to compare different prediction approaches. Most notably, the cross-validation results (Panel B) are almost the same as for the in-sample fit, which confirms the high robustness of the in-sample results. While the prediction errors in the *baseline* model are large (21.4% *RMSE* and 15.7% *MAE*), these figures substantially decrease to 13.8% and 9.7% for the *STAR* model.

In the one-day-ahead forecasting application (Panel C), the results of all models improve slightly. The reason for this improvement is the rolling estimation window, which accounts to some extent for temporal heterogeneity of all coefficients in all models. The magnitude of the improvement, however, is small. For instance, the *baseline* model improves in terms of *RMSE* by only 1 percentage point. Similar improvements are obtained for the other models. The highest relative improvement, even though small, is achieved for the *STAR* model, which corresponds to our expectation. In particular, note that the spatial lags are based on past but local observations. Therefore, capturing the local market trend seems to be most important when forecasting out-of-sample, i.e., when time-fixed effects are not feasible.

4.5 Robustness Tests

The forecast evaluations in the previous sub-section are derived on the basis of a specific choice of model parameters. These parameters mainly refer to the size of the rolling

²⁴The moving average in this context means *rolling average* and is not to be confused with a *moving-average (MA) process*.

²⁵For instance, the R^2 in the findings of Liu (2013) is 0.78 for the baseline model and 0.88 for the *STAR* model (e.g., for the year 2007).

window, the number of neighbors in the determination of spatial and temporal lags, and the regression tree parameter for sub-market definitions. In this section, we vary these parameters in order to determine the sensitivity of the prediction accuracy in the one-day-ahead forecast. The *RMSE* of the robustness tests, i.e., for different parameter values, are summarized in Table 5.

[INSERT TABLE 5 HERE]

First, we vary the length of the estimation window for the one-day-ahead forecasts in Panel A of Table 5. In the case of high temporal heterogeneity, we would expect higher accuracy for a shorter estimation window. Surprisingly, the *RMSE* for all models is slightly lower for a time window of 600 instead of 500 days; however, the differences are small. In addition, increasing the estimation window further means that we need more observations at the short end of the data. Hence, we refrain from increasing the estimation window further. In contrast, when the rolling estimation window is *decreased* to 400 days, the *RMSE* is marginally lower for all models.

Second, we address the arbitrary choice of parameters in the *STAR* model. As mentioned in sub-section 3.2, the number of observations to calculate the temporal and spatial lags in the *STAR* model is chosen according to Pace, Barry, Clapp, and Rodriquez (1998). Panel B of Table 5 presents the *RMSE* resulting from a variety of different parameter combinations of m_T and m_S .²⁶ Deviations in the parameters from their initial values ($m_T=180$ and $m_S=30$) do not change the outcomes notably. While the change in the spatial lag parameter m_S is negligible for the prediction results, the parameter m_T slightly changes the *RMSE*. Nevertheless, the prediction results are highly robust with respect to the choice of these parameters. In particular, the *RMSE* varies within less than 0.06 percentage points.

²⁶Note again that the parameter m_T denotes the number of observations to calculate the time lag, and m_S the number of observations to calculate the spatial lag.

A third source of arbitrary parameter choice is the number of defined sub-markets in the *TREE* model.²⁷ Panel C of Table 5 shows the *RMSE* for different choices of the regression tree complexity, which results in different numbers of sub-markets. As the resulting values indicate, the *RMSE* does not change substantially even with very few sub-markets. This indicates that the macro location is most important for differences in rental prices.

5 Conclusion

We test the forecast performance of five different hedonic model specifications with varying complexity of temporal and spatial structure for the residential rental market. The baseline model with no spatial and temporal elements primarily served as a benchmark. By fitting the models to rental apartment data from the canton of Zurich, Switzerland, and evaluating the forecast accuracy on daily rental price data, we reveal several important findings: With regard to the hedonic prices, we demonstrate the importance of accounting for spatial heterogeneity. Particularly, if only physical attributes are included in the model, the high dependence of certain variables on the location (e.g., available parking) can cause large biases in the estimated coefficients. Similarly, the prediction results are substantially worse in the absence of location dummy variables. Indeed, we show that including sub-market dummy variables constituted by ZIP codes highly improves the prediction results. Since sub-markets defined by ZIP codes do not systematically account for the specific spatial structure, we test a more sophisticated sub-market definition based on empirical spatial heterogeneity. Using this we construct sub-markets based on the residual values of the *baseline* model using a regression tree. We find that the construction of more sophisticated sub-markets substantially improves the prediction accuracy.

Including temporal lags (in the *STAR* model) reveals two further interesting find-

²⁷Note that in our previous estimations we have set the number of sub-markets equal to the number of regions in the *ZIP* model.

ings: First, we find empirical evidence for a functioning local market competition. In particular, the dependence on the rental price of neighboring objects is positive, while the dependence on favorable attributes of neighboring objects is negative. This result suggests that the presence of apartment offerings in the neighborhood with low rental prices and favorable (high) attributes affects the rental price of an object negatively. Second and more importantly for the purpose of this paper, including spatial dependence increases the prediction accuracy considerably. Particularly, for the *STAR* compared to the *TREE* model the *RMSE* decreased from 0.155 to 0.138 in the in-sample-prediction and cross-validation. The highest improvement, however, is found in the one-day-ahead forecast, where the *RMSE* decreases from 0.152 to 0.133. This finding is in line with the hypothesized expectation that local dynamics are important in the one-day-ahead forecasting. The *STAR* model is therefore the best candidate for empirical applications where no ex-ante information is available.

Another important outcome of the one-day-ahead forecasting application is the fact that the prediction accuracy is better than in the in-sample and cross-validation predictions. Hence, it follows that rental prices exhibit, to some degree, temporal heterogeneity, which is accounted for by the rolling estimation window. For real world applications, e.g., for the purpose of price setting by real estate developers or valuation purposes by appraisers, all models except the baseline model are suitable in terms of prediction accuracy. In particular, the prediction errors in terms of the root mean squared error (*RMSE*) are less than 18%, and the mean absolute error (*MAE*) is less than 13%. The best prediction is achieved by the *STAR* model in the one-day-ahead forecasting with an *RMSE* of 13.3% and a *MAE* of 9.7%. These results are highly robust to variations in the model parameters. For instance, changing the rolling estimation window length has almost no implications for the prediction results. Most notably, however, the reduction in the number of sub-markets in the *TREE* model from 60 to 10 increases the prediction error by only 0.1 percentage points.

The temporal structure, particularly the temporal autocorrelation, has only minor effects. Therefore, the *TAR* model does not bring (marginal) improvements. For out-of-sample forecasting applications, the temporal dimension merits greater attention in future research by, for instance, testing autoregressive dynamics in hedonic prices (under temporal heterogeneity). In addition, it may be promising to use more sophisticated forecasting models than simple temporal lags in terms of moving averages. However, the empirical evidence in this study shows that dynamics in the residential market are *local*. It is therefore questionable whether including *global* trends can indeed improve prediction results.

References

- ADAMS, Z., AND R. FÜSS (2010): “Macroeconomic determinants of international housing markets,” *Journal of Housing Economics*, 19(1), 38–50.
- ALLEN, M., R. RUTHERFORD, AND T. THOMSON (2009): “Residential Asking Rents and Time on the Market,” *Journal of Real Estate Finance and Economics*, 38(4), 351–365.
- ANGLIN, P. M., AND R. GENCAY (1996): “Semiparametric Estimation of a Hedonic Price Function,” *Journal of Applied Econometrics*, 11(6), 633–648.
- ANGLIN, P. M., R. RUTHERFORD, AND T. M. SPRINGER (2003): “The Trade-Off Between the Selling Price of Residential Properties and Time-on-the-Market: The Impact of Price Setting,” *Journal of Real Estate Finance and Economics*, 26(1), 95–111.
- ANSELIN, L. (2010): “Thirty Years of Spatial Econometrics,” *Papers in Regional Science*, 89(1), 3–25.
- BASU, S., AND T. G. THIBODEAU (1998): “Analysis of Spatial Autocorrelation in House Prices,” *Journal of Real Estate Finance and Economics*, 17(1), 61–85.
- BOURASSA, S., E. CANTONI, AND M. HOESLI (2007): “Spatial Dependence, Housing Submarkets, and House Price Prediction,” *Journal of Real Estate Finance and Economics*, 35(2), 143–160.
- BOURASSA, S. C., E. CANTONI, AND M. HOESLI (2010): “Predicting House Prices with Spatial Dependence: A Comparison of Alternative Methods,” *Journal of Real Estate Research*, 32(2), 139–160.
- BOURASSA, S. C., F. HAMELINK, M. HOESLI, AND B. D. MACGREGOR (1999): “Defining Housing Submarkets,” *Journal of Housing Economics*, 8(2), 160–183.
- BOURASSA, S. C., M. HOESLI, AND V. S. PENG (2003): “Do Housing Submarkets Really Matter?,” *Journal of Housing Economics*, 12(1), 12–28.
- BREIMAN, L., J. FRIEDMAN, R. OLSHEN, AND C. STONE (1984): *Classification and Regression Trees*. Chapman & Hall, New York.
- CAN, A. (1990): “The Measurement of Neighborhood Dynamics in Urban House Prices,” *Economic Geography*, 66(3), 254–272.

- (1992): “Specification and Estimation of Hedonic Housing Price Models,” *Regional Science and Urban Economics*, 22(3), 453–474.
- CLAPP, J. M., AND Y. WANG (2006): “Defining Neighborhood Boundaries: Are Census Tracts Obsolete?,” *Journal of Urban Economics*, 59(2), 259–284.
- DES ROSIERS, F., AND M. THERIAULT (1996): “Rental Amenities and the Stability of Hedonic Prices: A Comparative Analysis of Five Market Segments,” *Journal of Real Estate Research*, 12(1), 17–36.
- DIEWERT, E., H. SAEED, AND M. SILVER (2009): “Hedonic Imputation versus Time Dummy Hedonic Indexes,” in *Price Index Concepts and Measurement*, ed. by W. E. Diewert, J. S. Greenlees, and C. R. Hulten, pp. 161–196. University of Chicago Press.
- DIEWERT, W. E., J. DE HAAN, AND R. HENDRIKS (2015): “Hedonic Regressions and the Decomposition of a House Price Index into Land and Structure Components,” *Econometric Reviews*, 34(1-2), 106–126.
- DUBE, J., D. LEGROS, AND S. THANOS (2014): “Putting Time into Space: Expectations and Causality Issues in the Price Determination Process,” Conference Paper, Conference of the International Association for Applied Econometrics.
- DUBIN, R. A. (1998): “Predicting House Prices Using Multiple Listings Data,” *Journal of Real Estate Finance and Economics*, 17(1), 35–59.
- EUROSTAT (2015): “Data Explorer - Distribution of Population by Tenure Status, Type of Household and Income Group,” Discussion paper, Eurostat.
- FAHRLÄNDER, S. S. (2008): “Indirect Construction of Hedonic Price Indexes for Private Properties,” *Swiss Journal of Economics and Statistics (SJES)*, 144(IV), 607–630.
- GOODMAN, A. C. (1981): “Housing Submarkets within Urban Areas: Definitions and Evidence,” *Journal of Regional Science*, 21(2), 175–185.
- GOODMAN, A. C., AND T. G. THIBODEAU (1998): “Housing Market Segmentation,” *Journal of Housing Economics*, 7(2), 121–143.
- (2003): “Housing Market Segmentation and Hedonic Prediction Accuracy,” *Journal of Housing Economics*, 12(3), 181–201.
- HARTER-DREIMAN, M. (2004): “Drawing inferences about housing supply elasticity from house price responses to income shocks,” *Journal of Urban Economics*, 55(2), 316–337.

- HASTIE, T., R. TIBSHIRANI, AND J. FRIEDMAN (2009): *The Elements of Statistical Learning*, vol. 2. Springer.
- HORT, K. (1998): “The determinants of urban house price fluctuations in Sweden 1968-1994,” *Journal of Housing Economics*, 7(2), 93–120.
- KNIGHT, J. R. (2002): “Listing Price, Time on Market, and Ultimate Selling Price: Causes and Effects of Listing Price Changes,” *Real Estate Economics*, 30(2), 213–237.
- LIU, X. (2013): “Spatial and Temporal Dependence in House Price Prediction,” *Journal of Real Estate Finance and Economics*, 47(2), 341–369.
- MACLENNAN, D. (1977): “Some thoughts on the nature and purpose of house price studies,” *Urban Studies*.
- MALPEZZI, S. (2003): “Hedonic Pricing Models: A Selective and Applied Review,” *Section in Housing Economics and Public Policy: Essays in Honor of Duncan MacLennan*.
- MARTINS-FILHO, C., AND O. BIN (2005): “Estimation of Hedonic Price Functions via Additive Nonparametric Regression,” *Empirical Economics*, 30(1), 93–114.
- MCCCLUSKEY, W., M. MCCORD, P. DAVIS, M. HARAN, AND D. MCILHATTON (2013): “Prediction Accuracy in Mass Appraisal: A Comparison of Modern Approaches,” *Journal of Property Research*, 30(4), 239–265.
- MICHAELS, R. G., AND V. K. SMITH (1990): “Market segmentation and valuing amenities with hedonic models: The case of hazardous waste sites,” *Journal of Urban Economics*, 28(2), 223–242.
- MUNNEKE, H. J., AND B. A. SLADE (2001): “A Metropolitan Transaction-Based Commercial Price Index: A Time-Varying Parameter Approach,” *Real Estate Economics*, 29(1), 55–84.
- PACE, R. K., R. BARRY, J. M. CLAPP, AND M. RODRIQUEZ (1998): “Spatiotemporal Autoregressive Models of Neighborhood Effects,” *Journal of Real Estate Finance and Economics*, 17(1), 15–33.
- PACE, R. K., R. BARRY, O. W. GILLEY, AND C. F. SIRMANS (2000): “A method for spatial-temporal forecasting with an application to real estate prices,” *International Journal of Forecasting*, 16(2), 229–246.

- PACE, R. K., AND O. W. GILLEY (1997): “Using the Spatial Configuration of the Data to Improve Estimation,” *Journal of Real Estate Finance and Economics*, 14(3), 333–340.
- ROSEN, S. (1974): “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition,” *Journal of Political Economy*, 82(1), 34–55.
- SIRMANS, G. S., AND J. D. BENJAMIN (1991): “Determinants of Market Rent,” *Journal of Real Estate Research*, 6(3), 357–380.
- SIRMANS, S. G., D. A. MACPHERSON, AND E. N. ZIETZ (2005): “The Composition of Hedonic Pricing Models,” *Journal of Real Estate Literature*, 13(1), 1–44.
- STATISTIK STADT ZÜRICH (2005): “Veränderung der Wohnverhältnisse in der Stadt Zürich zwischen den Eidg. Volkszählungen von 1990 und 2000,” *Analysen*.
- VALENTE, J., S. WU, A. GELFAND, AND C. SIRMANS (2005): “Apartment Rent Prediction Using Spatial Modeling,” *Journal of Real Estate Research*, 27(1), 105–136.

Table 1: Overview of Model Specifications

This table lists the different hedonic model specifications with varying complexity of temporal and spatial structure. We compare the predictive power of the models using three different prediction approaches: *in-sample*, *cross-validation*, and *one-day-ahead forecasting*. The table shows the in-sample predictions as an overall model fit, including all observations over the whole sample period. For all model specifications, we further conduct an in-sample cross-validation, which is based on splitting the data set into a training and a test sample using a random re-sampling procedure. We also test the forecasting accuracy of the models based on a real one-day-ahead forecast. For that purpose, we define a rolling estimation window of 500 days.

Model	Temporal Structure		Spatial Structure	
	<i>Heterogeneity</i>	<i>Dependence</i>	<i>Heterogeneity</i>	<i>Dependence</i>
Baseline	time dummies	no	no	no
ZIP	time dummies	no	ZIP dummies	no
TREE	time dummies	no	TREE dummies	no
TAR	time dummies	temporal lags	TREE dummies	no
STAR	time dummies	temporal lags	TREE dummies	spatial lags

Table 2: Descriptive Statistics

This table reports the descriptive statistics of rental prices for the canton Zurich. The data pool stems from a Swiss rental market online platform. Note that the rental price and living area (*surface*) in Panel A are in log values. Panels B and C show the proportion of binary housing attributes and apartment types covered in the sample. Panel D reports the number of observations over the entire sample period 2002 to 2014. Because only the first months of 2014 are included, fewer observations are available for that year.

PANEL A: Descriptive Statistics of Continuous Variables					
Variable	Mean	Std.Dev.	Q 5%	Median	Q 95%
log(rental price)	7.218	0.755	5.908	7.222	8.503
log(surface)	3.505	1.781	1.000	3.500	5.500
number of rooms	4.403	0.432	3.5	4.5	5.0
PANEL B: Descriptive Statistics of Binary Variables					
Percentage of Observations					
Variable	yes	no	Total		
Elevator	43.80	56.20	100.00		
Parking	30.92	69.08	100.00		
Garage	53.58	46.42	100.00		
PANEL C: Descriptive Statistics of Apartment Type					
Apartment Type	Percentage of Obs.	Age	Percentage of Obs.		
Standard	82.22	less than 4 years	18.29		
Duplex	6.58	4 to 14 years	18.46		
Attic flat	3.15	15 to 29 years	18.45		
Roof flat	4.31	30 to 43 years	17.29		
Studio	0.35	44 to 70 years	17.10		
Single Room	0.33	more than 70 years	10.41		
Furnished flat	1.67	Total	100.00		
Terrace flat	0.41				
Single flat	0.23				
Loft	0.54				
other	0.21				
Total	100.00				
PANEL D: Observations by Year					
Year	Observations	Proportion of the Sample			
2002	637	2.2%			
2003	1102	3.8%			
2004	1557	5.4%			
2005	2085	7.3%			
2006	2281	7.9%			
2007	2445	8.5%			
2008	2691	9.4%			
2009	3074	10.7%			
2010	2975	10.4%			
2011	3092	10.8%			
2012	3053	10.6%			
2013	2878	10.0%			
2014	858	3.0%			
Total	28728	100.0%			

Table 3: In-Sample Estimation

This table reports the coefficients of the *baseline*, *ZIP*, *TREE*, *TAR*, and *STAR* regressions in Panel A. The coefficients of the spatial and temporal lags are shown in Panel B. The spatial and temporal lags are indicated by $lag(S)$ and $lag(T)$, respectively. Cross-effects are indicated by two lag indicators: the spatial lag of the temporal lag is denoted by $lag(ST)$ and the temporal lag of the spatial lag by $lag(TS)$. Variable $log(Rent)$ refers to lags of the rental price as the endogenous variable.

	Baseline		ZIP		TREE		TAR		STAR	
Variables	coeff.	t-stat.	coeff.	t-stat.	coeff.	t-stat.	coeff.	t-stat.	coeff.	t-stat.
PANEL A: Coefficients of Physical Attributes										
log(Surface)	0.628	95.90	0.566	103.50	0.537	111.30	0.537	111.30	0.536	112.60
Rooms	0.032	14.80	0.059	32.20	0.065	40.50	0.065	40.50	0.069	43.80
Elevator	0.068	22.80	0.033	13.30	0.020	9.20	0.020	9.20	0.034	13.70
Parking	-0.031	-11.30	0.002	0.60	0.005	2.30	0.005	2.20	0.007	3.40
Garage	-0.001	-0.20	0.013	5.40	0.019	8.90	0.019	8.90	0.027	12.00
Apartment	0.178	6.30	0.170	7.30	0.162	7.90	0.162	7.90	0.138	7.50
Duplex	0.204	7.10	0.205	8.60	0.199	9.50	0.199	9.50	0.164	8.80
Attic flat	0.372	12.80	0.352	14.60	0.334	15.70	0.334	15.70	0.286	15.10
Roof flat	0.205	7.10	0.223	9.30	0.213	10.10	0.214	10.10	0.178	9.50
Studio	0.192	5.50	0.146	5.00	0.103	4.00	0.103	4.00	0.073	3.20
Single Room	0.111	3.00	0.041	1.30	-0.022	-0.80	-0.022	-0.80	-0.049	-2.10
Furnished flat	0.484	15.90	0.416	16.50	0.355	16.00	0.356	16.00	0.282	14.20
Terrace flat	0.327	9.40	0.316	11.00	0.301	11.90	0.302	11.90	0.253	11.20
Single Flat	0.229	5.60	0.201	6.00	0.195	6.60	0.195	6.60	0.178	6.70
Loft	0.234	6.90	0.270	9.70	0.272	11.10	0.273	11.10	0.228	10.40
Age 5-14	-0.078	-18.30	-0.076	-21.30	-0.074	-23.70	-0.074	-23.60	-0.060	-21.30
Age 15-29	-0.120	-26.80	-0.126	-33.80	-0.139	-42.10	-0.138	-42.00	-0.104	-34.40
Age 30-43	-0.154	-32.50	-0.177	-44.60	-0.191	-55.10	-0.191	-55.00	-0.133	-40.50
Age 44-70	-0.099	-19.70	-0.168	-39.30	-0.189	-50.10	-0.189	-50.10	-0.142	-40.30
Age 70+	0.096	17.30	-0.014	-3.00	-0.102	-23.30	-0.102	-23.30	-0.095	-23.80
Q2	0.000	-0.10	0.007	2.20	0.007	2.40	0.007	2.30	0.004	1.40
Q3	-0.002	-0.40	0.005	1.50	0.004	1.40	0.003	1.10	0.002	0.80
Q4	0.000	-0.10	0.010	3.20	0.010	3.60	0.010	3.10	0.007	2.40
PANEL B: Coefficients of Spatial and Temporal Lags										
$lag(T)$ ln(Surface)							-0.066	-1.00	-0.067	-0.90
$lag(T)$ Rooms							0.016	0.80	0.011	0.50
$lag(T)$ Elevator							-0.017	-0.70	-0.040	-1.20
$lag(T)$ Parking							0.003	0.10	0.001	0.10
$lag(T)$ Garage							0.007	0.30	-0.037	-1.20
$lag(S)$ ln(Surface)									-0.504	-45.10
$lag(S)$ Rooms									-0.048	-14.70
$lag(S)$ Elevator									-0.060	-15.90
$lag(S)$ Parking									0.005	1.10
$lag(S)$ Garage									-0.033	-7.90
$lag(ST)$ ln(Surface)									0.415	3.00
$lag(ST)$ Rooms									-0.022	-0.50
$lag(ST)$ Elevator									-0.074	-1.20
$lag(ST)$ Parking									-0.093	-1.50
$lag(ST)$ Garage									0.097	1.80
$lag(TS)$ ln(Surface)									-0.037	-0.30
$lag(TS)$ Rooms									0.015	0.40
$lag(TS)$ Elevator									0.069	1.30
$lag(TS)$ Parking									-0.008	-0.20
$lag(TS)$ Garage									0.095	1.80
$lag(T)$ log(Rent)							0.080	1.70	0.099	1.40
$lag(S)$ log(Rent)									0.816	85.40
$lag(ST)$ log(Rent)									-0.702	-11.60
$lag(TS)$ log(Rent)									0.026	0.30
Time Dummies	yes		yes		yes		yes		yes	
Location Dummies	no		ZIP		TREE		TREE		TREE	
R-squared	0.722		0.810		0.853		0.853		0.883	
# parameters	35		94		94		102		120	

Table 4: Forecast Accuracy

This table reports the results on prediction accuracy for *baseline*, *ZIP*, *TREE*, *TAR*, and *STAR* models. Panel A shows the in-sample prediction results, Panel B the cross-validation prediction results, and Panel C the one-day-ahead forecasting results. For the one-day-ahead forecasting, the reported R -squared is the pseudo R -squared (calculated as the squared correlation between forecasted and actual value).

PANEL A: In-Sample Prediction (Global Estimation)					
	Baseline	ZIP	TREE	TAR	STAR
Root Mean Squared Error (RMSE)	0.214	0.176	0.155	0.155	0.138
Mean Absolute Error (MAE)	0.157	0.127	0.113	0.113	0.097
R -squared	0.722	0.810	0.853	0.853	0.883
Prediction error less than..					
20%	0.726	0.807	0.850	0.850	0.888
15%	0.604	0.697	0.743	0.744	0.802
10%	0.438	0.523	0.568	0.568	0.644
5%	0.230	0.284	0.310	0.310	0.375
PANEL B: Out-of-Sample Prediction (Random Sampling)					
	Baseline	ZIP	TREE	TAR	STAR
Root Mean Squared Error (RMSE)	0.214	0.177	0.155	0.155	0.138
Mean Absolute Error (MAE)	0.157	0.128	0.113	0.113	0.097
R -squared	0.720	0.809	0.852	0.852	0.883
Prediction error less than..					
20%	0.724	0.807	0.850	0.850	0.888
15%	0.602	0.696	0.743	0.744	0.802
10%	0.437	0.523	0.567	0.567	0.645
5%	0.230	0.284	0.309	0.309	0.375
PANEL C: One-Day-Ahead Forecast (Rolling Window)					
	Baseline	ZIP	TREE	TAR	STAR
Root Mean Squared Error (RMSE)	0.204	0.169	0.152	0.152	0.133
Mean Absolute Error (MAE)	0.149	0.123	0.111	0.112	0.095
R -squared	0.714	0.804	0.842	0.841	0.879
Prediction error less than..					
20%	0.746	0.817	0.852	0.850	0.893
15%	0.630	0.706	0.745	0.744	0.809
10%	0.464	0.538	0.570	0.567	0.648
5%	0.242	0.296	0.314	0.311	0.374

Table 5: Robustness Tests

This table shows the results of the robustness tests for varying model parameters. All tests are performed in the one-day-ahead forecasting application. In Panel A, we report the *root mean squared error* (*RMSE*), with varying length (in days) of the rolling estimation window. Panel B shows the *RMSE* when the number of neighbors (m_S and m_T) for the determination of spatial and temporal lags are chosen differently. m_S and m_T denote the number of observations to calculate the spatial and time lag, respectively. Panel C shows the prediction accuracy with varying regression tree parameters and corresponding number of tree-defined sub-markets. The values in bold are those used in the forecasting application of Section 4 (standard values).

PANEL A: RMSE with Varying Length of Rolling Estimation Window						
Estimation Window (Days)	Baseline	ZIP	TREE	TAR	STAR	
400	0.204	0.170	0.152	0.153	0.135	
500	0.204	0.169	0.152	0.152	0.133	
600	0.203	0.169	0.151	0.151	0.132	
PANEL B: RMSE of <i>STAR</i> Model with Varying m_S and m_T						
		m_T				
		100	140	180	220	260
m_S	10	0.1320	0.1326	0.1327	0.1325	0.1326
	20	0.1322	0.1328	0.1329	0.1326	0.1327
	30	0.1322	0.1328	0.1329	0.1326	0.1327
	40	0.1322	0.1328	0.1329	0.1326	0.1327
	50	0.1322	0.1328	0.1329	0.1326	0.1327
PANEL C: RMSE of <i>TREE</i> Model with Varying Tree Complexity Parameter						
Complexity Parameter		Number of Sub-Markets		RMSE		
20.00		5		0.1705		
5.00		10		0.1522		
1.20		30		0.1517		
0.55		60		0.1512		

Figure 1: Illustration of Time and Space Data Structure

This figure shows the aggregation of discrete time and space in the data structure. The left panel illustrates the aggregation of time and space into discrete partitions by including simple time and location dummies to account for potential *heterogeneity*. The right panel illustrates the modeling of spatial and temporal *dependency*. The dark shaded circle indicates the current observation which depends on temporally lagged ($t1$, $t2$, and $t3$) or spatially lagged ($s1$, $s2$, $s3$, and $s4$) observations. T , S and s , t stand for time and space and spatially and temporally lagged observations, respectively.

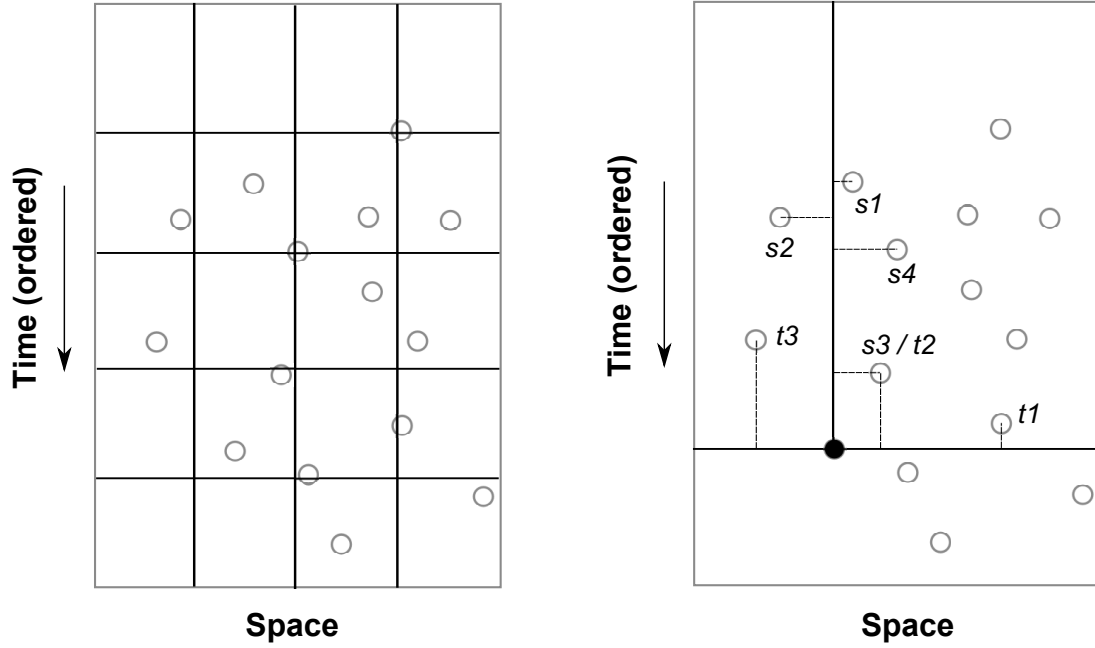


Figure 2: Map of Regression-Tree-Based Sub-Markets of Canton Zurich

The figure illustrates the sub-market identification according to the regression tree approach based on residual location values. In this approach, the splitting variable (*lat* or *lon*) and the splitting points are determined in such a way that the prediction error is minimized. The illustration shows 60 tree-based sub-markets resulting from this procedure for the canton Zurich. The left panel shows the reduced tree structure for 7 splits. The right panel shows the resulting sub-markets with each partition on the map corresponding to each leaf of the tree representing a sub-market.

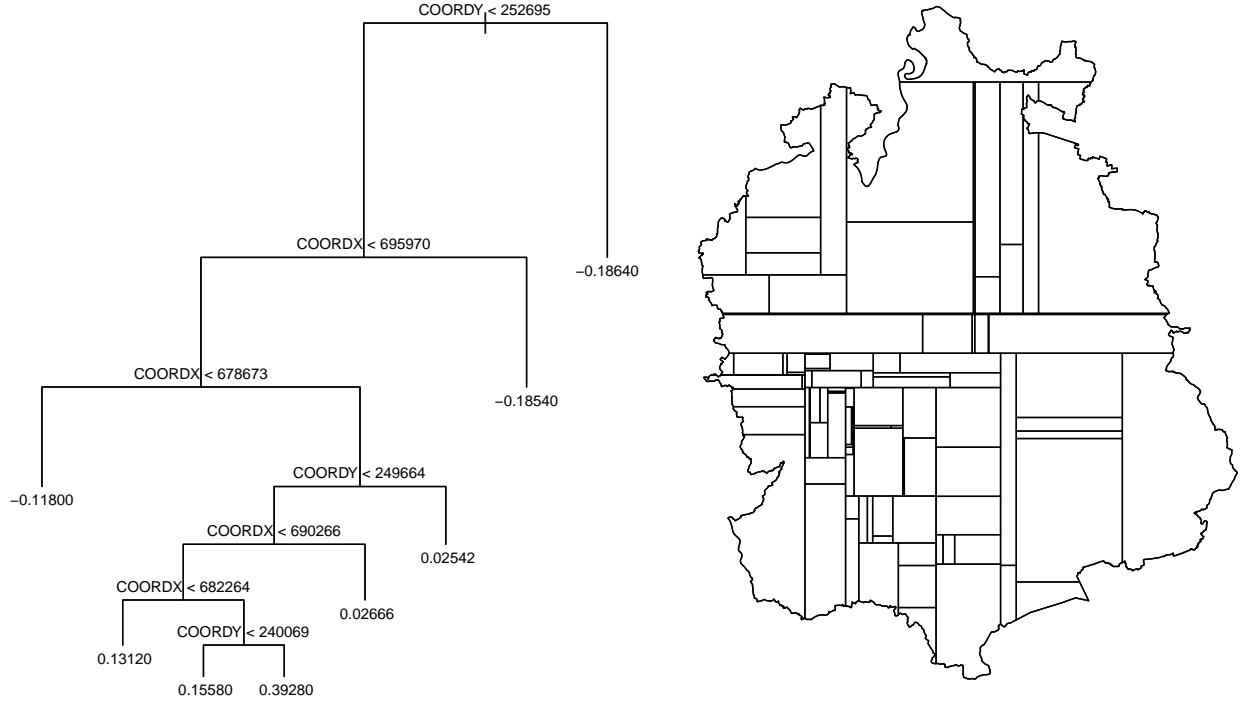


Figure 3: Map of Study Area

The figure shows the map of the canton Zurich with the distribution of observations indicated by the grey dots (approx. 28,000 observations). It also lists the six towns with the highest population. The canton Zurich accommodates 171 municipalities and covers an area of 91 square kilometers. The administrative borders of these municipalities, shown in the graph, are defined by ZIP codes.

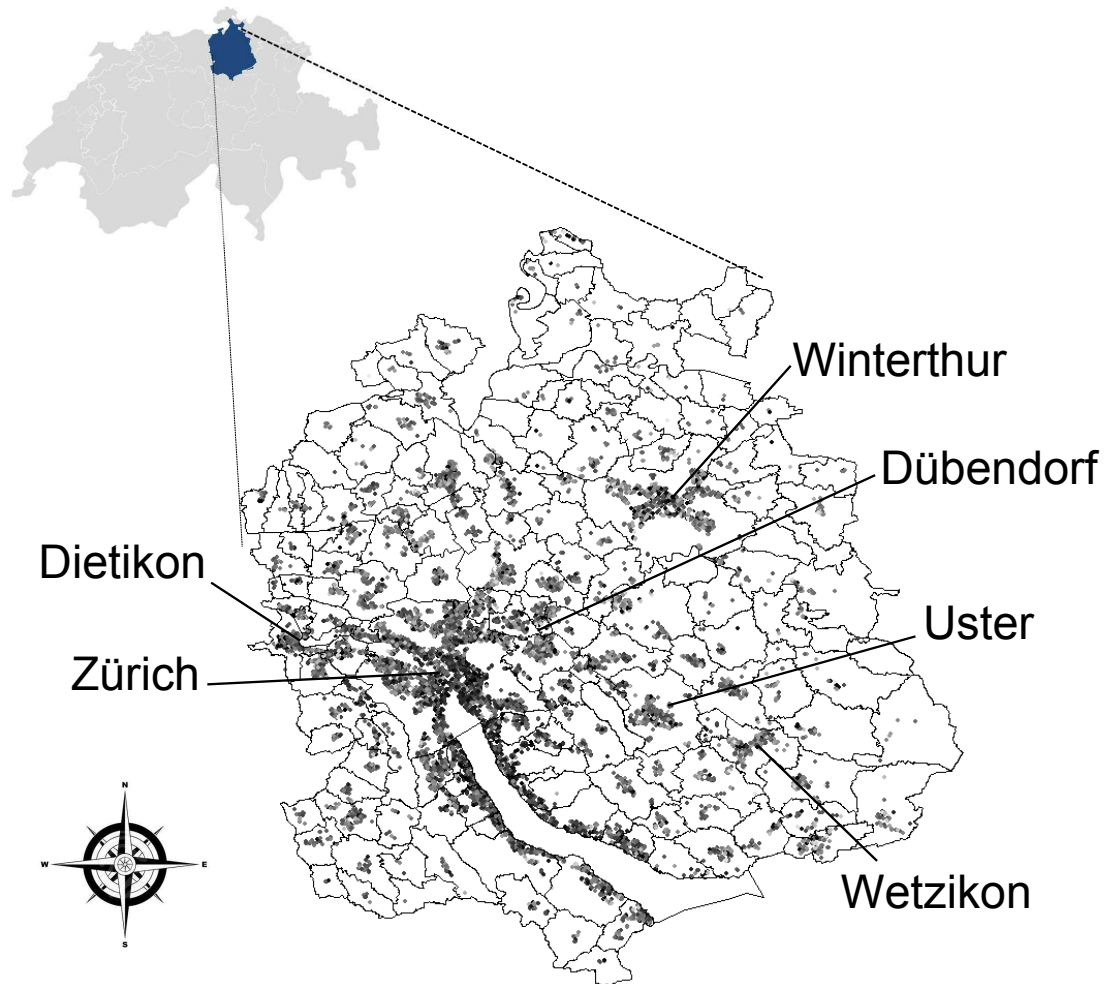


Figure 4: Empirical Semivariogram of the Baseline Model

This figure shows the semivariogram estimated from the residuals of the *baseline* model. The increasing pattern, which reflects spatial dependence as a function of distance (in km) in the semivariogram, indicates that spatial autocorrelation is left in the residuals.




Figure 5: Empirical Semivariogram of Alternative Models

The figure shows the semivariograms estimated from the regression residuals of the alternative models *ZIP*, *TREE*, *TAR*, and *STAR*. For the *ZIP* and *TREE* specifications the spatial dependence decreases with increasing number of spatial elements. Due to the consideration of only temporal components, the *TREE* and *TAR* models fail to adequately capture spatial dependence.

