See how the land lies: Land valuation using spatial models*

Jacqueline Seufert, Geert Goeyvaerts & Sven Damen

January 2023

Abstract

Economists have been advocating for a land tax rather than a regular property tax. There are, however, several challenges to value land for tax purposes. Indeed, data on vacant land transactions are scarce, land and structure are conventionally traded in a bundle and it is hard to capture all factors that determine the value of land. We propose to use a new Bayesian spatial model and apply the model to the universe of vacant and improved land sales from Belgium in 2018. Our results indicate that vacant land prices are substantially more difficult to predict than house prices. However, the predictive performance of the spatial model improves considerably in comparison to a regular linear hedonic approach. Models that combine data from vacant and improved land are unable to improve the predictive accuracy.

*Seufert: Department of Economics, KU Leuven. E-mail: jacqueline.seufert@kuleuven.be. Goeyvaerts: Department of Economics, KU Leuven. E-mail: geert.goeyvaerts@kuleuven.be. Damen: Department of Accountancy and Finance, University of Antwerp. E-mail: sven.damen@uantwerp.be. The authors would like to thank seminar participants at KU Leuven and University of Antwerp. Funding from Research Foundation Flanders (FWO), project number S005718N, is gratefully acknowledged.

1 Introduction

Economists have been advocating for a land tax rather than a regular property tax on the value of the house as a regular property tax may distort investments in structure (George 1879, Anas 2015, Yang 2014). However, most countries tax both land and structure as only three OECD countries levy a pure land tax in 2014 (Blöchliger & Kim 2016). A major challenge for the implementation of a land tax is that the land value of each plot is not directly observed. Therefore, methods are required to be able to accurately assess the value of each plot of land.

There are, however, several challenges to value land. First, data on vacant land transactions are scarce. Indeed, vacant land is rarely transacted and does not follow a uniform pattern across space (Larson & Shui 2022). In addition, it is often difficult for researchers to obtain data on land sales. Second, land and improvement are conventionally traded in a bundle, rendering the disentanglement of both components arduous. These methods typically need additional data on the cost of structures to be able to decompose the total price into the structure and land components.[1] Third, it is hard to capture all factors that determine the value of land. A lot of potential covariates such as distance to amenities, zoning plans or local infrastructure are not easy to measure or cannot be obtained at all. These unobserved variables complicate the accurate assessment of land values.

To tackle these challenges, we use a Bayesian spatial model which enjoys great popularity in spatial settings (Besag et al. 1991, Cressie 1993, Diggle & Ribeiro 2008) and use different types of spatial approaches to model the land values. We apply the model to the universe of vacant and improved land sales from Belgium in 2018, which includes both urban and rural transactions.

[1] A large literature has proposed different methodologies to disentangle the value of land and structure, see Diewert et al. (2015), Kuminoff & Pope (2013), Davis & Palumbo (2008), Davis et al. (2017, 2021) among others.

Our methodology and data allow us to answer several research questions. First, we study whether it is more difficult to predict vacant land prices in comparison to improved land. The results indicate that the predictive accuracy is worse for models that predict vacant land. Therefore, our findings confirm the above-described challenges to predict vacant land prices. Second, we study whether the spatial model that we propose is able to improve the accuracy of a standard linear model. The results indicate that the spatial component of the novel statistical approach is able to capture a large portion of the spatial variation. All classic metrics of predictive accuracy improve in the spatial model in comparison to a regular linear specification. In our models to predict vacant land, the R-squared of the spatial model is twice as high in comparison to a linear specification. Third, because we have access to both vacant and improved land sales, we are able to study whether models that are trained on improved land are able to accurately predict vacant land. The results, however, indicate that our models that are trained on improved land or a combination of vacant and improved land are unable to improve the predictive accuracy of the models that are trained on vacant land.

We contribute to the existing literature in several ways. First, we are the first to propose to use a Bayesian approach using the Integrated Nested Laplace Approach. The approach is able to accommodate a large class of models and is often faster than alternative algorithms. Moreover, we include a spatial component in our model which we find to improve predictive performance considerably in comparison to classic hedonic approaches that are common in the existing literature that studies (near-)vacant land (Haughwout et al. 2008, Nichols et al. 2013, Albouy et al. 2018, Barr et al. 2018, Gedal & Ellen 2018).

A second contribution is due to the new data that we bring to the table. A considerable portion of the existing literature either uses the Maricopa data set for the metropolitan area of Phoenix, AZ (Clapp et al. 2021, Clapp & Lindenthal 2020, Larson & Shui 2022, Bourassa & Hoesli 2022) or the

CoStar COMPS data set for metropolitan areas across the U.S. (Albouy et al. 2018, Haughwout et al. 2008, Nichols et al. 2013). Our data includes both urban and rural land in order to give full insight into land values across different locations. Indeed, a number of papers only concentrate on urban vacant land (Albouy et al. 2018, Clapp et al. 2021, Gedal & Ellen 2018). Studies that only focus on land in urban locations may suffer from a lack of generalizability. Another advantage of our data is that it includes both vacant and improved land. Therefore, we are able to compare valuation of land from vacant land models with models estimated on both structure and land.

Third, we cross-validate our results thoroughly. Through this, we can make more confident statements about the quality of our models. Despite a substantial literature on land prices, very few papers evaluate the out-of-sample predictive accuracy of their models. The literature that reports an out-of-sample root mean squared error (RMSE) shows a wide range for the predictive accuracy (Wentland et al. 2021, Davis et al. 2021, Larson & Shui 2022, Clapp et al. 2021, Albouy et al. 2018). The substantial heterogeneity in predictive accuracy further illustrates the difficulty in predicting land prices. In contrast to this literature, we also apply a spatial-leave-one-out-cross-validation that avoids underestimation of the prediction error due to spatial autocorrelation. In addition we also quantify the predictive uncertainty thereby giving further insight about the quality of our prediction.

The paper is structured as follows. In Section 2 we provide a deeper discussion of existing land valuation models. We describe the data that we will use in Section 3 and the methodology in Section 4. The results are presented in Section 5.

2 Background on land valuation

The existing literature has taken different approaches to appraise land. For land underneath improvements, there are three different methods to assess the value of the land. First, the teardown approach which infers the value from real estate purchased with the objective to demolish the existing improvement on the plot (Gedal & Ellen 2018, Dye & McMillen 2007). However, teardowns are usually sparse. They also might not be representative of a more exhaustive set of parcels within an area. Besides the problem of infrequent and heavily localized transactions, this approach also ignores the potentially high costs of demolition (Clapp & Lindenthal 2020).

The second approach is the residual approach which derives land value from the difference between the total market value of the property and the depreciated cost of replacing the improvement on the parcel (Larson & Shui 2022, Davis et al. 2021, Davis & Palumbo 2008). Withal, demolition costs are not always available. Moreover, this approach makes the assumption that each component of the total residential real estate evolves independently. Clapp et al. (2021) argues that this does not hold for urban land as its value depends heavily on surrounding infrastructure. They also point out that the real-world application of this approach is flawed as it is only applicable to relatively new improvements.

Third, there is the hedonic approach which assesses the land value through running a regression of the sales price of the property on land and improvement characteristics (Kuminoff & Pope 2013, Diewert et al. 2015). Yet, hedonic approaches for land underneath improvements tend to attribute the land value solely to the marginal contribution of the surface of the plot and fixed effects per municipality (Kuminoff & Pope 2013, Wentland et al. 2020). Often, distance to amenities and access is ignored. Approaches for estimating vacant land range from basic hedonic regressions (Barr et al. 2018, Combes et al. 2019) to more advanced statistical techniques (Albouy et al. 2018, Barr et al. 2018), often using a more exhaustive selection of locational covariates. Howbeit, there is no method that can address the discussed selection issues and unobserved covariates in a satisfactory manner. A portion of the unexplained variation might be accounted for through the inclusion of a spatial component to the linear predictor. Larson & Shui (2022) applies a novel approach by using a spatial interpolation method called Kriging, first used by Davis et al. (2021), to estimate vacant land prices on an aggregated county-level. They found that the modelling of a spatial relationships between counties improves their predictive power. Basu & Thibodeau (1998) argues that regular hedonic models for housing prices are only more accurate if the unexplained variation of prices is uncorrelated, otherwise spatial procedures are more suitable. Yet, the assumption of knowing the true data generating spatial process is unrealistic. Therefore, the inclusion of spatial random effects is crucial for increasing the predictive quality of the model. In addition, assuming spatially uncorrelated error terms is not plausible for these kind of data and therefore, the independence assumption is violated. This can lead to obtain confounding parameter uncertainty estimates and therefore, report misleading results (Fieberg et al. 2021). Despite this finding, most papers refrain from using spatial methods. We believe that capturing spatial variation adds value. Section 4 provides further reasoning for our approach. Papers such as Wentland et al. (2021), Davis et al. (2021) strive to provide estimates at a fine spatial level. Still, more granular fixed effects effects might lead to overfitting since there are few observation per group (Wentland et al. 2021).

One way to mitigate this issue is to turn to Bayesian statistics as it mitigates the effect of potential outliers, as further explained in Section 4. One example which applies Bayesian statistics to estimate land values is the paper by Albouy et al. (2018). They apply an empirical Bayesian approach rather than a full Bayesian approach for the sake of computational speed. The downside of this approach is that it underestimates the uncertainty of the posterior distribution for city- and time-specific effects through using a fixed point estimate rather than a distribution. This defies one of the major advantages of Bayesian statistics, the accurate quantification of uncertainty. With recent developments in Bayesian statistics, such methods such as R-INLA by Rue et al. (2009), offer a faster alternative than established full-Bayesian methods, such as Markov Chain Monte Carlo, without a loss in accuracy.

Moreover, empirical Bayes aims to choose the best value for the shrinkage parameter for predicting the full data set. However, for the sake of generalizability of the model, cross-validation is a better approach. It aims to choose the value for the shrinkage parameter which is most appropriate in predicting the validation set given a training set (Van Erp et al. 2019). Thus, a regular cross-validation usually is better suited for a generalization of the results, especially with regards to the evaluation of predictive performance.

The last concern related to the complexity of land prices is the evaluation of the predictive power. As the main goal of many papers is to give accurate prediction of land values, a rigorous cross-validation is necessary. This resampling method uses different partitions of the data to train a model on several iterations and then test the predictive performance. The aim of this technique is to raise flags to issues such as overfitting or selection bias. The key is to evaluate the model's performance on unknown data and its ability to generalize to this data set (Browne 2000). Despite a quite substantial literature on land prices, very few papers evaluate the predictive accuracy of their models. Metrics like the out-of-sample root mean squared error (RMSE) are usually the standard values reported. For land prices, the overall predictive performance of the models is moderate at best. Wentland et al. (2021) uses Machine Learning techniques to reduce the out-of-sample RMSE by 75% in comparison to regular hedonic approaches for land underneath improvements. Also, for land underneath improvements, but across counties, Davis et al. (2021) finds a median out-of-sample RMSE of 40%. Larson & Shui (2022) reports an out-of-sample value of 66% for the

Kriging model for vacant land. Also, Clapp et al. (2021) obtains an in-sample RMSE of 97% - 100% and an out-of-sample RMSE of 99% - 100%, both for land underneath improvements. Albouy et al. (2018)'s analysis yields an out-of-sample value between 97% and 128%, depending on the covariates included. While Davis et al. (2021) and Larson & Shui (2022) use an 80% - 20% split for training and test set, Albouy et al. (2018) applies a more methodologically solid form of a leave-one-out cross-validation. The other papers, however, do not reveal their strategy for predictive evaluation. Again, the literature shows a wide range of RMSE value further illustrating the difficulty in predicting land prices.

3 Data

For our analysis, we investigate both vacant and improved land in Belgium. Our data combines different sources. The main strand of data is the universe of sales transactions, obtained from the Belgian Federal Public Service Finance. This data set contains information about sales of vacant and improved plots. In our case, improved plots include only single-family homes. For our analysis, we restrict the data set to the year 2018, as this selection provides a lot of variation and the data quality drastically improved after administrative changes in 2015 (Moerkerke 2017). The variables of interest can be categorized as follows: Housing characteristics (number of bedrooms, number of bathrooms, etc.), transaction characteristics (sales price, date of transaction, etc.) and plot characteristics (surface, classification of land, etc.).

While improved land is easy to identify in the data set, we need define vacant land more carefully. We classify those parcels as vacant for which the construction type is undefined, but the parcel is characterized as building land. Moreover, we exclude vacant land parcels which are either split or assembled in the following years.

For additional regressors, we join supplementary data sets: Open Street Map (distances to amenities, water, etc.), Malaria Atlas Project (travel time to nearest cities) (Weiss et al. 2018), the Belgian Statistical Office (Statbel, additional covariates).

Our paper also takes the shape of the plot into account. In previous work, Asabere & Harvey (1985) and Glumac et al. (2019) find that for urban land, more regular-shaped plots are sold at a premium while irregular shapes lead to a discounted sales price. However, regularity is difficult to assess: Existing papers account for the effect of shape in a mono-dimensional approach (Glumac et al. 2019) or use opaque, somewhat arbitrary measures provided in the original data set (Asabere & Harvey 1985, Gedal & Ellen 2018). Demetriou et al. (2013) proposes an index which includes different shape and boundary indices as a single metric will lead to spurious classification of a plot. Following his idea loosely, we use the following covariates in order to correctly assess the shape:

1. The number of vertices of the polygon: While the obvious desirable number of vertices is four, a slightly larger number can still guarantee a rather regular shape. A larger deviation, however, indicates a rather complex polygon.

2. The standard deviation of angles at vertices: Both reflex and acute angles can lead to a polygon of unattractive shape.

3. The standard deviation of the distance of vertices to the centroid of the polygon: Assuming an ideal shape should be close to a square, the vertices should lie on a circle whose center is the centroid of the polygon.

4. The standard deviation of the edge length: Given the ideal shape of a square, the edge length should not differ a lot.

5. The ratio of the bounding box: We calculate the ratio of the longer

side to the shorter side of the bounding box. An ideal ratio amounts to 1. The further away the value from 1, the less regular the shape.

While these indices only capture only one aspect of regularity, all of them combined paint a comprehensive picture of whether a plot has a desirable, rectangular shape. For the calculation of these metrics, we use the official Cadaster shapefile of 2018. All of our data points are observed at the centroid of the polygons, thus we obtain latitude and longitude of each observation. All covariates which relate to the plot, such as measures of regularity or surface, are available for every parcel. However, some covariates are only available for improved land. For a full list of covariates, see Table 1.

Our data set comprises 8,720 observations of vacant land and 73,846 parcels of improved land. However, we still need to cull extreme values which might e.g. arise from misclassification of vacant agricultural land as vacant building land. In order to remove outliers, we restrict the range of vacant land prices to the one of improved land which reduces the number of observations to 6,077. Our analysis is run on the reduced sample, but the results for the full sample are available in the Appendix (see Tables 12, 14, 13).

The average improved land costs 246,050€ whereas vacant land averages to 176,720€. Tables 2 and 3 show the skewed distribution of vacant land prices in comparison to improved land prices. Furthermore, vacant land parcels are on average larger (1,386.93 m^2) than improved land parcels (553.61 m^2). In terms of regularity, the standard deviation of the edge length stands out. There is much more variance in the edge length of vacant plots as compared to improved plots. Also, vacant land appears to be further away from any sort of amenity. Figure 1 illustrates the difference in both data sets: Whereas vacant land is sparse, particularly in Wallonia, and does not display a strong regional variation in prices, we have a large amount of improved land with peaks in price at the coast, Flemish metropolitan areas and the region bordering Luxembourg.

4 Methodology

4.1 Modeling land values

In this section, we lay the methodological groundwork for our analysis. We use the most common method for land appraisal (Colwell et al. 1983), the market approach, in order compare similar plots. This approach is based on hedonic price models.

We apply this approach to explore the predictive performance for different models for vacant land, land underneath improvements and improved land. For all scenarios, we use a linear model as the baseline. Furthermore, we extend our models by a spatial component in the linear predictor. Observations which lie closely together in space are likely to display similar price values. This could lead to have a high degree of residual spatial autocorrelation if the variation of the prices is not mostly explained by the fixed effects. A spatial model will consider this spatial autocorrelation in order to disentangle the general trend, driven by covariates without a spatial structure, from the exclusively spatial random variation. Thus, this allows us to account for the missed covariates which are spatially correlated. Over the last decades, real estate economics have put more emphasis on the inclusion of a spatial component (Clapp et al. 2002, Moralı & Yılmaz 2020) as ignoring it violates model assumptions. The spatial autocorrelation in the error terms leads to a biased estimation of error variance. Whereas regression coefficients might remain unbiased (Anselin & Griffith 1988), recent findings suggest that this is not always the case (Dupont et al. 2022). Through the violation of independence (Cressie 1993), significance tests and assessments of model fit might be misleading. Also, the uncertainty estimates might become unreliable (Anselin & Griffith 1988).

There are two exploratory approaches that are prevalent in the economic literature: K-nearest-neighbors (KNN) and Geographically Weighted Regression (GWR). KNN is a straightforward, deterministic method as it finds the closest neighbors in space to a certain observation. One assumes that points far away have less influence than close points. This approach is e.g. used to assess the impact of surrounding amenities on rents in Berlin (Schäfer & Hirsch 2017). Bourassa et al. (2010) also uses a two-stage process incorporating nearest neighbors' residuals in the second stage to capture spatial dependency. Yet, they find that classic geostatistical approaches work better. The downsides of this approach include sensitivity to outliers, limited applicability for sparse data and the disregard for spatial autocorrelation (Longley et al. 2005). GWR (Brunsdon et al. 1998), however, takes into account spatial autocorrelation and is one of the most popular approaches to take into account for spatial dependencies. Nevertheless, this method has some shortcomings: GWR estimates a set of values of coefficients for every observation by using all data within a certain bandwidth and weighting by distance. This makes results highly dependent on the choice of the bandwidth. Land prices in Belgium - both vacant and improved - are influenced by municipal and regional juridiction. Therefore, it is more crucial to take into account the relationship between neighboring municipalities rather than taking an average within a circle ignoring borders. GWR is more of an explanatory method to detect non-stationarity, and its applicability as a prediction tool is controversial (Wheeler & Calder 2007). Also, since GWR computes location-specific parameter estimates, it is a computationally intense method in comparison to other available techniques. More appropriate techniques for including spatial components are spline-based approaches (Wood 2011) or Bayesian models (Gelfand et al. 2003).

4.2 Bayesian spatial models

In our case, we opt for the latter option as Bayesian models offer many advantages. Bayesian statistics is an approach that involves updating prior beliefs with new data. This type of models enjoys great popularity in spatial settings (Besag et al. 1991, Cressie 1993, Diggle & Ribeiro 2008). One benefit is the high flexibility, thereby allowing us to implement complex models in a straightforward manner. An example of such complex models involves hierarchical models, which are a commonly used to model spatial structure. Rather than correcting the uncertainty between levels as required in a frequentist setting, we can build the hierarchical structure directly into the prior specification. Moreover, we can obtain a full representation of parameter uncertainty through the posterior distribution. This is beneficial for the quantification of uncertainty of estimates (Kruschke & Vanpaemel 2015). Another key advantage is that missing data and outliers are considered random variables, for whom we obtain a posterior distribution: As we estimate missing data through the posterior distribution, we can obtain estimates for, e.g. a municipality where we do not observe any vacant land, through the sales in other municipalities. Also in the case of outliers, the advantage of this method is evident: Rather than point-estimating an outlier as in a frequentist method, the Bayesian approach updates its prior belief, e.g. based on other municipalities, to take this observation to a certain degree into account and thus, mitigates its effect.

Due to our Bayesian framework, all parameters are considered random. Therefore, the typical frequentist terminology of fixed and random effects no longer holds. In order to keep the terminology accurate, we will use the Bayesian definition of these effects. In the Bayesian setting, a fixed effect is for covariates which affects all observations in the same way. It usually has a vague prior, and we estimate each parameter independently. On the other hand, a random effect serves the purpose of introducing additional structure by modeling the parameters for each level as being drawn from a distribution. Random effects take into account variation and usually have a more informative prior distribution. In our case, we use spatial random effects to model spatial structure or differences between municipalities. In our spatial models, we have structured random effects and unstructured random effects. Structured random effects can capture spatial or temporal dependencies. The spatially structured effect can uncover remaining spatial structure that is unexplained by the model as long as the structure adheres to Tobler's first law of geography: "everything is related to everything else, but near things are more related than distant things." (Tobler 1970). Unstructured effects allow us to model unstructured variability via e.g. an *iid*-normal effect. Moreover, our models also contain an unstructured effect which can be considered as a classical error term. This error term captures e.g. measurement errors or the effect of covariates which are weakly correlated spatially.

4.3 Integrated Nested Laplace Approach

For the implementation of our Bayesian framework, we choose the Integrated Nested Laplace Approach (INLA) (Rue et al. 2009). This method applies numerical integration and Laplace approximations for approximate Bayesian inference and is implemented in the `R` package `R-INLA`. The advantages of `R-INLA` lie in the ability to accommodate a large class of models (e.g. Generalized Linear Mixed Models or Generalized Additive Mixed Models) (Rue et al. 2009) and also in computational speed. `R-INLA` is often faster than approaches relying on Markov Chain Monte Carlo (MCMC) algorithms and also provides diagnostic metrics which are straightforward to interpret and reproducible (Wang et al. 2018). Over the last decade, `R-INLA` has grown to be an established method in environmental statistics (Crewe & McCracken 2015, Bowler et al. 2015) and for disease mapping (Python et al. 2021, Bhatt et al. 2015), but it remains vastly unknown in real estate economics. For details on the method, see Rue et al. (2009). In this

section, we only focus on the pivotal building blocks of `R-INLA` for spatial models.
A baseline spatial model can be described as follows : The data y_i are conditionally independent given the predictor η_i

$$y_i|\eta_i, \boldsymbol{\theta_0} \sim \pi(y_i|\eta_i, \boldsymbol{\theta_0}), \tag{1}$$

where $\boldsymbol{\theta_0}$ is the first part of the set of hyperparameters, usually relating to parameters of the selected likelihood π.
We define the predictor η as

$$\eta_i = \beta_0 + \Sigma_{j=1}^{k} \beta_j x_{ij} + \Sigma_{l=1}^{L} u_{li} \tag{2}$$

, where β_0 is the intercept, $\boldsymbol{\beta}$ describes the fixed effects of covariates $\boldsymbol{x}$ and $\boldsymbol{u}$ denotes the random effects.

Also, the random vector u has a normally distributed prior:

$$\mathbf{u}|\boldsymbol{\theta_L} \sim N(\mathbf{0}, \mathbf{Q}\left(\boldsymbol{\theta_L}\right)^{-1}) \tag{3}$$

The precision matrix $\mathbf{Q}$ depends on the hyperparameters $\boldsymbol{\theta_L}$ (Bakka et al. 2018). We use two different types of spatial approaches in order to model different kinds of land values: The Intrinsic Conditionally Autoregressive (ICAR) model and the Stochastic Partial Differential Equations (SPDE) approach. While the former model deals with a discretely indexed spatial component, the latter captures a continuous one.

4.4 Intrinsic Conditionally Autoregressive Model

Our first model specification deals with lattice data which describes data usually observed within administrative boundaries. The spatial adjacency is mostly expressed through a neighborhood matrix containing non-zero values at the intersection of rows and columns of adjacent areas. This matrix serves the purpose of modeling spatially correlated random effects via a multivariate Gaussian distribution whose precision matrix depends on the neighborhood matrix. This normal distribution has the mean of the averaged neighbors' random effects, with variance proportional to one over the number of neighbors so that a larger number of neighbors reduces variability (Besag 1974, Besag et al. 1991). In particular, this approach models the random effect values u_a on a set of regions $a = 1 \dots m$ conditionally on its neighboring regions. We define regions which share a common border as neighbors. The conditional distribution for u_a is

$$u_a | \mathbf{u}_{-a}, \tau_u \sim N\left(\frac{1}{d_a} \sum_{b \sim a} u_b, \frac{1}{d_a} \frac{1}{\tau_u}\right), \tag{4}$$

where $b \sim a$ means that regions a and b are neighboring regions, d_a is the number of neighbors and τ_u denotes the precision parameter. The joint distribution is then described by

$$\mathbf{u} | \tau_u \sim N\left(0, \frac{1}{\tau_u} \mathbf{Q}^{-1}\right), \tag{5}$$

$$\text{where } Q_{ab} = \begin{cases} d_a, & \text{if } a = b \\ -1 & a \sim b \end{cases} \tag{6}$$

(Bakka et al. 2018).

We apply the ICAR structure to random effects in models which use vacant land as the training set. The reasons are two-fold: On the one hand, the results of the cross-validation suggests that a discretely indexed random effect is more appropriate for vacant land than one with a continuous index. On the other hand, vacant land is more dependent on municipal jurisdiction than improved land as the administration determines the future building potential of vacant land. Due to the skewed nature of the sales prices of vacant land (see Figure 2), we opt for a logistic distribution rather than a Gaussian as it has heavier tails and thus, can better account for extreme values. We set up the model as follows:

$$\begin{aligned} \log(\texttt{sales\_price}_i) &\sim Logistic(\eta_i, \tau_\epsilon^{-1}) \\ \mathbb{E}(\eta_i) &= \beta_0 + \Sigma_{j=1}^{k} \beta_j x_{ija} + u_a, \end{aligned} \tag{7}$$

where the sales price of each parcel i follows a logistic distribution with mean of the linear predictor η_i and precision τ_ϵ. Furthermore, u_a denotes the random effect per municipality a. In order to assess the performance of the spatial field, we compare the spatial model with the linear version:

$$\begin{aligned} \log(\texttt{sales\_price}_i) &\sim Logistic(\eta_i, \tau_\epsilon^{-1}) \\ \mathbb{E}(\eta_i) &= \beta_0 + \Sigma_{j=1}^{k} \beta_j x_{ija} \end{aligned} \tag{8}$$

4.5 Stochastic Partial Differential Equations

Our second model specification tackles continuously indexed spatial models. In this case, our precision matrix $\mathbf{Q}$ is in general no longer sparse which renders computations unviable, especially for large data sets (Bakka et al. 2018). Lindgren et al. (2011) proposes an approach where a Gaussian field with a Matérn correlation, expressed as a Gaussian Markov Random Field (GMRF), is the solution to a stochastic partial differential equation. A GMRF is a discretely indexed Gaussian field whose full conditional distributions depend only on a set of neighbours to each location. In our case, we assume a stationary GMRF for simplicity. The related Matérn covariance function is defined as:

$$Cov(\xi(s_i), \xi(s_j)) = \frac{\sigma_\xi^2}{\Gamma(\lambda)} 2^{1-\lambda} (\kappa ||s_i - s_j||)^\lambda K_\lambda(\kappa ||s_{i-} s_j||), \tag{9}$$

where $||s_i - s_j||$ denotes the Euclidean distances between the two locations $s_i, s_j \in R^2$, σ_ξ^2 is the marginal variance of the Gaussian random field,

K_λ is the modified Bessel function of second kind and of order $\lambda > 0$ which quantifies the smoothness of the field and $\kappa > 0$ is the scaling parameter related to the range ρ_ξ which descries the distance between two locations above we assume the spatial correlation to become negligible. Through the representation of a Gaussian Field as a GMRF, we obtain a sparse representation of the spatial effect, making computation feasible. Another benefit of the SPDE method is that allows to discretize the study area, even for for irregularly distributed point data. For further details, see Lindgren et al. (2011). We define our spatial random effect as $¸(\mathbf{s}) \sim N(0, \mathbf{Q}(ˇ, \mathbf{ø}_{\mathbf{ffl}}))^{-1}$.

In this case, we apply the SPDE model for improved land. The logged sales price follows a normal distribution (see 2). In order to take differences across municipalities into account, we add an *iid* - random effect per municipality a. Again, the results from cross-validation, discussed in the next section, suggest that this is the most appropriate model. This model is defined in the following manner:

$$\begin{aligned} \log(\texttt{sales\_price}_i) &\sim N(\eta_i, \tau_\epsilon^{-1}) \\ \mathbb{E}(\eta_i) &= \beta_0 + \Sigma_{j=1}^{k} \beta_j x_{ija} + \\ &\quad \xi_i(s_i, \rho_\xi, \sigma_\xi) + \zeta_a, \end{aligned} \tag{10}$$

where the sales price of parcel i follows a Gaussian with mean η_i and precision τ_ϵ. Moreover, s_i is location i, σ_ϵ^2 denotes the marginal variance of the GMRF, ρ_ϵ describes range of the GMRF and ζ_a stands for the *iid* random effect per municipality.

The linear version of this model is:

$$\begin{aligned} \log(\texttt{sales\_price}_i) &\sim N(\eta_i, \tau_\epsilon^{-1}) \\ \mathbb{E}(\eta_i) &= \beta_0 + \Sigma_{j=1}^{k} \beta_j x_{ija} \end{aligned} \tag{11}$$

4.6 Cross-validation

Cross-validation is one of the most commonly used resampling methods to estimate the true prediction error of models and compare predictive performance.

As alluded in the previous section, we rigorously cross-validate our models. We do not only apply regular Leave-one-out-cross-validation (LOOCV), but we also apply Spatial-leave-one-out-cross-validation (SLOOCV).

The reason for spatially cross-validating a model lies in the high risk of overfitting for complex models like ours (Lucas et al. 2020). Model selection processes may choose excessively complex models leading to underestimation of the prediction error (Mosteller & Tukey 1977). In particular, we are exposed to a high risk of underestimating the model error if training and testing sets are geographically close, but predictions take place far from the training locations (Lucas et al. 2020). When the test set is drawn to be close to the training set, the independence of both sets could be jeopardized in the presence of spatial autocorrelation (Hastie et al. 2009). This would lead to overly optimistic estimates of our prediction errors and might even produce inaccurate conclusions (Roberts et al. 2017, Hastie et al. 2009).

Ignoring spatial structure when splitting into training and testing set causes models to seem more reliable than they are. This allures us to have more faith in the model's predictive power than it deserves.

The procedure for SLOOCV is the following (Le Rest et al. 2014):

1. Remove one observation from the training set.

2. Remove all observation within a buffer. The remaining data points constitute the training set.

3. Predict at the location of the removed observation.

These steps are repeated k-times with k being lower or equal to the number of observations.

5 Results

In a first step, we tackle three different tasks:

1. The comparison of predictive power between models for vacant and improved land prices
2. The prediction of vacant land prices
3. The prediction of land prices underneath improvements and land shares

For each of our tasks, we specify our models in two versions: One including and another one excluding the spatial component in order to evaluate the impact of spatial structure on predictive performance.

5.1 Comparison of vacant and improved land price models

5.1.1 Comparison of baseline models

While it is relatively straightforward to predict improved land, the prediction of vacant land is more complicated. The reasons for this intricacy are manifold: Not only is the land component more volatile in comparison to the improvement component in residential land value, but the infrequent transactions of vacant further complicate the prediction (**?**). Moreover, improved land is usually transacted as a bundle of the land underneath the improvement and the actual improvement (Clapp et al. 2021). The improvement characteristics themselves might also provide some indication of the land value, such as larger houses with more bedrooms in affluent areas. In order to illustrate the difference in predictive performance, we first compare a model which estimates vacant land prices with a model which estimates improved land prices.

We predict vacant land prices with the ICAR-method (Model 1) and improved land prices with the SPDE-method (Model 2). For Model 1, we use all available covariates for vacant land and for Model 2, we extend this set of covariates by the improvement-specific variables. Table 4 presents the results of the two models. For both models, the superiority of spatial models in predictive performance in comparison to a regular linear specification is evident: All classic metrics of predictive accuracy (Mean average error (MAE), RMSE, Relative root mean squared error (rRMSE)) indicate lower values for the model incorporating a spatial component. Moreover, the goodness-of-fit is much lower for both models for the linear specifications. Model 1 appears to benefit even stronger from the spatial component than Model 2. It appears that the plot-specific covariates do not have a strong impact, but that the spatial component captures quite a large portion of the unexplained spatial variation.

When comparing the two models to each other, Model 2 exhibits a larger R-squared and also a lower relative RMSE than Model 1. This illustrates our assumption that it is easier to predict improved land than vacant land. The reasons for this difference in predictive performance can stem from multiple sources: On the one hand, the improved land data set is considerably larger than the vacant land set. On the other hand, through the improvement characteristics the model gains explanatory power (see also Model 6 in Table 11). Therefore, the improvement might be the key component which facilitates the prediction of improved land relative to vacant land. Furthermore, a proportion of the unexplained variation in vacant land prices might be due to unobserved factors which we can not capture through a spatial structure. This could be e.g. a plot-specific covariate which does not behave similarly to other plots close in space or a covariate that is subject to temporal autocorrelation.

In order to cross-validate our results, we apply LOOCV and SLOOCV with 400 iterations (for computational reasons). The radius for SLOOCV was chosen as the ten-fold of the average distance between two points for each data set, improved and vacant land. Through setting different radii and comparing the results, we ensured that not too many nor too few points were removed. We compare the findings with the in-sample predictions. For both data sets (vacant and improved land), we evaluate three model specification: Linear regression, SPDE and ICAR. Table 5 summarizes our validation through MAE, RMSE and rRMSE. For both improved and vacant land, the linear specification performs the worst in all three validation methods across every validation metric in comparison with any of the spatial models. Moreover, the SLOOCV yields more prudent metrics across all model specifications, validation methods and validation metrics. The difference between SLOOCV and LOOCV is more notable for the case of improved land. This implies that we might need to take the evaluation of the LOOCV with a grain of salt due to potential overfitting. For vacant land, the two sampling validation approaches do not differ as starkly. This might be due to a lower degree of residual autocorrelation. While for improved land, SPDE yields a higher prediction accuracy than ICAR, the opposite holds true for vacant land. The decrease in rRMSE for SPDE with SLOOCV, the most conservative validation method, amounts to 18.9% in comparison to the linear model. For vacant land, the increase for ICAR equals 20.7%. The difference in SLOOCV between the best models for each data sets indicates that the rRMSE for improved land is still 17.5% lower than for vacant land. This further illustrates the difficulty of predicting land prices. Improved land seems to require a fine spatial effect to take into account the spatial variation resulting from clustering and variation within municipalities. An ICAR would omit too much of the within-municipality effect. However, vacant land calls for an ICAR specification to better account for variation between municipalities. The spatial field from an SPDE would smooth too much of this variation due to the sparse spatial structure. For the case of SPDE for improved land and ICAR for vacant land, the in-sample RMSE and rRMSE are higher in the in-sample validation than in the cross-validation. Whereas the difference in the second case is of negligible size when taking the standard deviation of the observed data into account, it has considerable size in the first one. The large difference for the SPDE model in comparison to the linear model stems from the fact that the improved land faces a substantial amount of residual autocorrelation. This could cause the LOOCV estimates to be overly optimistic.

In order to assess the quality of our predictions, it is not only necessary to consider prediction accuracy, but also prediction uncertainty. Through Bayesian statistics, we can easily evaluate the uncertainty of our predictions. While most research is concerned with with predictive accuracy, predictive models can come short in quantifying predictive uncertainty. For policy-makers, uncertainty measures are of particular interest as new data might differ from the existing data due to sampling bias or non-stationarity (Ovadia et al. 2019). Taking predictive uncertainty into account can help to take more prudent decisions for zoning, approval of real estate loans or assessment of tax bases. It is imperative to tell a cautionary tale about land price assessment. Therefore, we compare the distribution of uncertainty between our models to get a better understanding whether uncertainty follows a similar pattern across models. We consider the density of the predictive standard deviation (SD) divided by the predictive mean to adjust for outliers, obtained from the posterior distribution. Figure 3 displays the results. At first sight, it might appear that the linear models have much lower uncertainty than the spatial counterparts. However, recall that the linear specification may violate the assumption of independence and leads to estimates of uncertainty which are on average too small. Therefore, our spatial models might give a more realistic estimation of uncertainty. Through the inclusion of a spatial effect, the uncertainty becomes more nuanced. While a part of the density concentration around a smaller uncertainty is due to the higher sample size for Model 2, another part can probably be accounted for through the strong spatial clustering of the dependent variable coinciding with autocorrelation of the model, and a more normal distribution of improved land prices. For Model 1, the variance of the uncertainty is larger. This is probably due to the skewed distribution of vacant land prices which cannot be easily captured through the model, the sparse spatial pattern which is difficult to account for, and the additional unexplained non-spatial variation driving up uncertainty.

In order to further evaluate the distribution of predictive standard deviation, we use the Hellinger distance. This metric is commonly used in information theory to measure similarity between two probability distributions, p and q, and is defined as $v_{pq} \in [0, 1]$: when the measure is 0 the two probability distributions are identical and when it is 1 they are completely dissimilar. Moreover, we use the demeaned version of this metric, $\widetilde{v_{pq}}$. While v_{pq} assesses the overlap between two distributions, $\widetilde{v_{pq}}$ evaluates the overall similarity in shape between the two distributions. The first line of Table 6 indicates that the distribution of the standard deviation of Model 2 and Model 1 do not overlap a lot. The similarity in shape is also not too large (0.46). Thus, it appears that uncertainty for vacant land price models and improved land price models might be driven by different factors.

5.1.2 Analysis of Market Thickness

To assess the difference between Model 2 and Model 1 more accurately with respect to sample size and spatial autocorrelation, we conduct an additional analysis. We explore how the spatial thickness of the market impacts the predict performance. We take two samples of size $n = 6077$ (same size as vacant land). We expect that predictive uncertainty is lower in a thick market. While the first sample is drawn randomly and exhibits strong spatial clustering, the second sample contains the improved parcels closest in space to the vacant ones, thus displaying spatial sparsity. We hereon refer to Model 2 applied to the former as *Model 2 clustered* and applied to the latter *Model 2 sparse*. Table 7 shows that there is barely any difference between both versions of Model 2 across all metrics. Table 8 illustrates the results of the validation. We also look at a standard of spatial autocorrelation, Moran's I: It measures how one observation is similar to the ones surrounding it. It is defined on $[-1, 1]$, where -1 is perfect dispersion, 0 is perfect randomness and 1 is perfect clustering of similar values. While sales prices exhibit some level of positive autocorrelation between parcels in all three data sets, it is much stronger for improved land and somewhat similar for both subsets. For vacant land, the spatial autocorrelation of prices is a bit higher on the municipality level (0.08 vs 0.06 on parcel level). This suggests that municipal difference could matter more than within-municipal differences.

There is no strong difference in in-sample metrics for both versions of Model 2, but they both predict better than Model 2 with the full set of improved land. For the two sampling methods, *Model 2 clustered* indicates a better performance, potentially due to the stronger autocorrelation in prices. This finding implies that the variation in sales prices are well captured through the spatial effect and the fixed effects in both versions of Model 2. Again, even with a smaller set for improved land, Model 1 performs worse than both versions of Model 2. Also, the metrics of LOOCV for both versions of Model 2 are lower than the in-sample implying an overly optimistic assessment of the predictive accuracy. The difference in metrics between LOOCV and SLOOCV in comparison to the linear version is of considerable size for both versions of Model 2, hinting at the presence of residual spatial autocorrelation. Keep in mind that even though the data for Model 1 and *Model 2 sparse* display a similar spatial pattern, this does not imply similar autocorrelation of residuals. Considering the linear specification, *Model 2 clustered* performs better than *Model 2 sparse*. *Model 2 sparse* appears to benefit more from the inclusion of a spatial field. This suggests that the fixed effects of *Model 2 sparse* capture less of the spatial variation. For Model 1, the predictive improvement also implies that the fixed effects cannot explain much of the spatial variation, even to a lesser extent than for *Model 2 sparse*. Sample size does not seem to drive the difference in performance between Model 1 and Model 2, rather it is due to a different level of spatial autocorrelation and explanatory power of fixed effects.

For this task, we again consider the density of predictive uncertainty across the models (see Figure 5). *Model 2 clustered* displays a high density around low uncertainty. The skewed nature of the density plots suggests that there are more outliers in comparison to *Model 2 sparse* which displays a more normal distribution of uncertainty, albeit centered around a higher level. The skewed shape of *Model 2 clustered* is probably due to some extent due to the stronger spatial autocorrelation. Figure 4 shows that SD increases with distance from clustering and in spatially sparse regions as expected. Furthermore, there is more variation in uncertainty for Model 1 further illustrating the difficulty of predicting land prices. The Hellinger distance suggests that both versions of Model 2 exhibit a high similarity in shape and some amount of overlap. Model 1 displays some similarity in shape with both versions of Model 2, but less overlap with *Model 2 clustered* than with *Model 2 sparse* (see Table 9. In addition, we explore the relationship between the predictive uncertainty between Model 1 and *Model 2 sparse*. We find a strong (0.82) correlation between the two uncertainty (see Figure 6). We interpret this finding as an indication that land is the main driver of uncertainty in prediction of real estate values. While the non-spatial residual variation amounts to 0.43 on average for Model 1, it equals 0.07 for *Model 2 sparse*. This suggests that vacant land prices display a higher degree of non-spatial variation. In this case, we find predictive accuracy and predictive uncertainty are related to each other.

Additionally, we investigate the impact of regularity of parcel shapes on predictive uncertainty. For this, we define a non-parametric index:

$$\psi_i = \frac{\Sigma_r^l R_{ri}}{max(\Sigma_r^l R_{ri}) - min(\Sigma_r^l R_{ri})}, \tag{12}$$

where R_{ri} is the rank of regularity measure $r = 1 \dots 5$ of plot i. ψ_i is normalized to a range of $[0, 1]$ such that 1 means the most regular shape (a perfect square) and 0 is a polygon of irregular shape.Figure 7 shows that there is no large difference in the regularity across the three data sets. Taking the predictive uncertainty into account, Figure 8 displays a similar pattern for the three data sets: The relationship between regularity and uncertainty is U-shaped implying that very regular or irregular plots exhibit lower standard deviation in comparison to the remaining parcels. This implies that we can be more certain about our predictive performance if the parcel has an extraordinary shape in terms of regularity.

5.2 Prediction of vacant land prices

Considering the difficulty of observing and predicting vacant land, the question arises whether we could use an improved land price model to predict vacant land.

We therefore introduce two new models: Model 3 estimates the full data set containing both vacant and improved land. Model 4 estimates only improved land. Subsequently, we make use of these estimates to predict vacant land. For these models, we set the improvement-specific covariates equal 0 if the land in the training set is vacant and to their true value otherwise. In the testing set, we also set the improvement-specific covariates equal to 0 to create as-if vacant land. We then compare our findings with the original model predicting vacant land, Model 1. Table 10 shows the findings. Here, the spatial specification again performs overall better than the linear one across all models, with the exception of ϕ in Model 4. For both Model 3 and Model 4, the R-squared increases with the inclusion of a spatial component. Overall, Model 1 appears to perform better across all metrics in comparison to the other two models. Among the linear specifications, there is a negligible difference in R-squared for all models. Also, the difference between the predictive metrics for Model 1 and the other models is rather small in a linear setting. The larger difference and improvement in performance emerges from the specification of the spatial component. This implies that the specification of the spatial structure is able to explain the unexplained variation better via the ICAR-specification rather than through the SPDE-model. While Model 3 performs slightly better through the additional vacant land in the training data, the difference between Model 3 and Model 4 are no longer pronounced when including the spatial component. Therefore, the results suggest that Model 1 is a more suitable approach when predicting vacant land. Overall, Model 1 performs 10% better than Model 3 and 14% than Model 4, considering the rRMSE. Thus, we should refrain from using an improved land price model for vacant land. Spatial structure, explanatory power of covariates and distribution of sales prices differ too much between vacant and improved land as that an improved land price model could be appropriate. We have more unexplained variation in the case of vacant land which we are simply unable to account for.

In terms of uncertainty, we observe again that the linear models seem to underestimate the uncertainty across all models (see Figure 9). The figure shows that Model 3 and Model 4 yield lower predictive uncertainty than Model 1. Model 3 and Model 4 are fairly certain about their predictions as they were trained improved land or a combination of improved and vacant land. In contrast, Model 1 displays a wider variation of uncertainty. Line 2-4 of Table 6 display the Hellinger distances. Model 3 and 4 are very similar in shape and exhibit a large overlap. Model 1, however is fairly different from both Model 5 and 6 and does not display a great deal of overlap. Nonetheless, these findings require particular prudence as they are misleading. The resulting spatial field assumes a much stronger spatial autocorrelation than the one that is present in vacant land. Moreover, fixed effects can explain more of the variation of improved land. Thus, using the same fixed effects and spatial field of Model 3 or Model 4 yields overly optimistic estimates with regards to uncertainty, unaware of the potential bias caused through model mis-specification. Thus, we advise the reader to always interpret uncertainty with caution as it is highly dependent on accurate model specification and the correctness of assumptions. In this case, it is safer to give more attention to the in-sample validation.

In order to take advantage of the spatial analysis, we further investigate the estimates obtained from the spatial field. Figure 10 displays the mean and the standard deviation of the spatial fields of Model 1, Model 2 and Model 3 as these models cover all different training sets. The plots of the mean are in line with our expectations, especially for Model 2 and Model 3. Metropolitan areas, the coastline and the border region to Luxembourg have a positive effect whereas Wallonia and in general, the east of the country displays a negative effect. For Model 1, the pattern is similar, albeit less pronounced. There is a clear divide in the sign of the spatial effect between Flanders and Wallonia. Moreover, in some municipalities, such as Quevy, the fixed effects are notably inept to explain prices such that the process is highly dominated by the spatial effect. Also, the spatial effect for Model 3 is slightly smoother due to the inclusion of vacant land. The spatial field smooths more in order to account for the lower autocorrelation in vacant land prices. With regards to the uncertainty, the standard deviation is a lot higher in Wallonia, partly due to the sparse data. Overall, we observe that uncertainty is low when the amount of observations is ample.

5.3 Prediction of land prices underneath improvements

Disentangling the value of residential property into an improvement element and a land element is of paramount relevance for policy-makers. On the one hand, it is essential to be able to break down the total value of a property into the two building blocks for an accurate municipal tax base. On the other hand, improvements and land constitute two different elements on national balance sheets. Besides, when valuating a property, the decomposition in two components is pivotal as improvements depreciate over time unlike land (Diewert et al. 2015).

However, despite the beliefs of the most adamant proponent of a land tax, **?**, the estimation of land underneath improvements is notoriously difficult. We cannot assume that land underneath improvements behaves in a similar manner as vacant land. **?** show evidence that vacant land values is not an appropriate proxy for land underneath improvement. According to them, vacant parcels are sold at a premium in comparison to land underneath improvements due to unobserved differences between both types of land.

In order to further investigate this finding, we assess two models estimating land underneath improvements. Model 5 is trained on vacant land via an ICAR-model and then predicts the land underneath improvement. Model 6 is trained on improved land via an SPDE-model and also predicts the land underneath improvement. For this endeavour, we use the full set of covariates. For vacant land in the training set in Model 5, we set the improvement-specific covariates equal to 0. We also do this for the testing set in both models. The training set in Model 6, however, keeps the observed properties of the improvement.

Table 11 displays the results. In this case, the spatial model trumps the linear model in Model 6, but not necessarily in Model 5. One explanation for this could be that the spatial effect in Model 5 is on a municipality level and thus, smooths the spatial within-municipality variation of improved land prices too strongly. Furthermore, we observe that the R-squared increases in Model 6. If we upgrade our model from containing only plot-specific covariates to including a spatial component, the model gains roughly as much explanatory power as if we were to add the improvement-specific covariates. Since we compare the predictions with the improvement-specific covariates equal to 0 to the true values which take into account the real covariate values, the metrics of predictive performance are a lot higher. Comparing the linear specification of Model 5 to the spatial one of Model 6, there is no overly stark difference in the metrics assessing predictive accuracy. The only exception is the R-squared. While it may seem counterintuitive at first sight, RMSE and R-squared describe two different things. R-squared is more concerned with the covariates actually predicting, in contrast to the outcome variable merely having low variance and being simple to predict, even without the covariate. The spatial model appears to make poorer predictions in comparison to the linear version. However, the higher R-squared implies a roughly constant bias, potentially caused through the spatial field. The predictors appear to still be somewhat able to explain the observed value. This is in line with our hypothesis that the municipal effect in Model 5 captures the some of the unexplained variation, but systematically wrong so in this case. Moreover, ϕ is also similar in both models. Since ϕ denotes the mean ratio of predicted values to true values, we can conclude that, on average, land underneath improvements makes up for 35-64% of the total residential value. This finding is highly dependent on the model specification. While Model 6 indicates a better fit of the model to the data, the linear version of Model 5 suggest a more realistic land share. These results deserve additional attention as we suspect that the land share varies with age of the improvement and probably fluctuates across regions of different levels of affluence. However, this investigation is beyond the scope of this paper. Due to to the similarity in predictive metrics, we conclude that it is possible to predict land underneath improvements with a regular improved land price model. It is not necessary to estimate via a vacant land price model. This alleviates the struggle of estimating land underneath improvements when there is no vacant land available. From our findings, opposing **?**, vacant land is a good enough proxy for land underneath improvement as regular improved land.

References

Albouy, D., Ehrlich, G. & Shin, M. (2018), 'Metropolitan land values', *Review of Economics and Statistics* **100**(3), 454–466.

Anas, A. (2015), *The Effects of Land Value Taxation in a Computable General Equilibrium Model*, Lincoln Institute of Land Policy.

Anselin, L. & Griffith, D. A. (1988), 'Do spatial effecfs really matter in regression analysis?', *Papers in Regional Science* **65**(1), 11–34.

Asabere, P. K. & Harvey, B. (1985), 'Factors influencing the value of urban land: Evidence from halifax-dartmouth, canada', *Real Estate Economics* **13**(4), 361–377.

Bakka, H., Rue, H., Fuglstad, G.-A., Riebler, A., Bolin, D., Illian, J., Krainski, E., Simpson, D. & Lindgren, F. (2018), 'Spatial modeling with r-inla: A review', *Wiley Interdisciplinary Reviews: Computational Statistics* **10**(6), e1443.

Barr, J., Smith, F. H. & Kulkarni, S. J. (2018), 'What's manhattan worth? a land values index from 1950 to 2014', *Regional science and urban economics* **70**, 1–19.

Basu, S. & Thibodeau, T. G. (1998), 'Analysis of spatial autocorrelation in house prices', *The Journal of Real Estate Finance and Economics* **17**(1), 61–85.

Besag, J. (1974), 'Spatial interaction and the statistical analysis of lattice systems', *Journal of the Royal Statistical Society. Series B (Methodological)* **36**(2), 192–236.
URL: *http://www.jstor.org/stable/2984812*

Besag, J., York, J. & Mollié, A. (1991), 'Bayesian image restoration, with two applications in spatial statistics', *Annals of the Institute of Statistical Mathematics* **43**(1), 1–20.
URL: *https://link.springer.com/article/10.1007/bf00116466*

Bhatt, S., Weiss, D., Cameron, E., Bisanzio, D., Mappin, B., Dalrymple, U., Battle, K., Moyes, C., Henry, A., Eckhoff, P. et al. (2015), 'The effect of malaria control on plasmodium falciparum in africa between 2000 and 2015', *Nature* **526**(7572), 207–211.

Blöchliger, H. & Kim, J., eds (2016), *Fiscal Federalism 2016*, OECD.
URL: *https://www.oecd-ilibrary.org/governance/fiscal-federalism-2016*$_9$789264254053 − *en*

Bourassa, S. C. & Hoesli, M. (2022), 'Hedonic, residual, and matching methods for residential land valuation', *Swiss Finance Institute Research Paper* (22-55).

Bourassa, S., Cantoni, E. & Hoesli, M. (2010), 'Predicting house prices with spatial dependence: a comparison of alternative methods', *Journal of Real Estate Research* **32**(2), 139–160.

Bowler, D. E., Haase, P., Kröncke, I., Tackenberg, O., Bauer, H.-G., Brendel, C., Brooker, R. W., Gerisch, M., Henle, K., Hickler, T. et al. (2015), 'A cross-taxon analysis of the impact of climate change on abundance trends in central europe', *Biological Conservation* **187**, 41–50.

Browne, M. W. (2000), 'Cross-validation methods', *Journal of mathematical psychology* **44**(1), 108–132.

Brunsdon, C., Fotheringham, S. & Charlton, M. (1998), 'Geographically weighted regression', *Journal of the Royal Statistical Society: Series D (The Statistician)* **47**(3), 431–443.

Clapp, J. M., Cohen, J. P. & Lindenthal, T. (2021), 'Are estimates of rapid growth in urban land values an artifact of the land residual model?', *The Journal of Real Estate Finance and Economics* pp. 1–49.

Clapp, J. M., Kim, H.-J. & Gelfand, A. E. (2002), 'Predicting spatial patterns of house prices using lpr and bayesian smoothing', *Real Estate Economics* **30**(4), 505–532.

Clapp, J. M. & Lindenthal, T. (2020), 'The valuation of urban land: Comparison and critique of three cama methods', *University of Connecticut School of Business Research Paper* (20-09).

Colwell, P. F., Cannaday, R. E. & Wu, C. (1983), 'The analytical foundations of adjustment grid methods', *Real Estate Economics* **11**(1), 11–29.

Combes, P.-P., Duranton, G. & Gobillon, L. (2019), 'The costs of agglomeration: House and land prices in french cities', *The Review of Economic Studies* **86**(4), 1556–1589.

Cressie, N. A. C. (1993), *Statistics for spatial data*, Wiley series in probability and mathematical statistics, rev. ed. edn, John Wiley, New York [etc].

Crewe, T. & McCracken, J. (2015), 'Long-term trends in the number of monarch butterflies (lepidoptera: Nymphalidae) counted on fall migration at long point, ontario, canada (1995–2014)', *Annals of the Entomological Society of America* **108**(5), 707–717.

Davis, M. A., Larson, W. D., Oliner, S. D. & Shui, J. (2021), 'The price of residential land for counties, zip codes, and census tracts in the united states', *Journal of Monetary Economics* **118**, 413–431.

Davis, M. A., Oliner, S. D., Pinto, E. J. & Bokka, S. (2017), 'Residential land values in the washington, dc metro area: New insights from big data', *Regional Science and Urban Economics* **66**, 224–246.

Davis, M. A. & Palumbo, M. G. (2008), 'The price of residential land in large us cities', *Journal of Urban Economics* **63**(1), 352–384.

Demetriou, D., See, L. & Stillwell, J. (2013), 'A parcel shape index for use in land consolidation planning', *Transactions in GIS* **17**(6), 861–882.

Diewert, W. E., Haan, J. d. & Hendriks, R. (2015), 'Hedonic regressions and the decomposition of a house price index into land and structure components', *Econometric Reviews* **34**(1-2), 106–126.

Diggle, P. & Ribeiro, P. J. (2008), *Model-Based Geostatistics*, Springer.

Dupont, E., Wood, S. N. & Augustin, N. H. (2022), 'Spatial+: a novel approach to spatial confounding', *Biometrics* .

Dye, R. F. & McMillen, D. P. (2007), 'Teardowns and land values in the chicago metropolitan area', *Journal of urban economics* **61**(1), 45–63.

Fieberg, J., Signer, J., Smith, B. & Avgar, T. (2021), 'A 'how to'guide for interpreting parameters in habitat-selection analyses', *Journal of Animal Ecology* **90**(5), 1027–1043.

Gedal, M. & Ellen, I. G. (2018), 'Valuing urban land: Comparing the use of teardown and vacant land sales', *Regional Science and Urban Economics* **70**(3), 190–203.

Gelfand, A. E., Kim, H.-J., Sirmans, C. & Banerjee, S. (2003), 'Spatial modeling with spatially varying coefficient processes', *Journal of the American Statistical Association* **98**(462), 387–396.

George, H. (1879), *Progress and poverty: An inquiry into the cause of industrial depressions, and of increase of want with increase of wealth. The remedy*, D. Appleton.

Glumac, B., Herrera-Gomez, M. & Licheron, J. (2019), 'A hedonic urban land price index', *Land Use Policy* **81**, 802–812.

Hastie, T., Tibshirani, R., Friedman, J. H. & Friedman, J. H. (2009), *The elements of statistical learning: data mining, inference, and prediction*, Vol. 2, Springer.

Haughwout, A., Orr, J. & Bedoll, D. (2008), 'The price of land in the new york metropolitan area', *Current Issues in Economics and Finance* **14**(3).

Kruschke, J. K. & Vanpaemel, W. (2015), 'Bayesian estimation in hierarchical models', *The Oxford handbook of computational and mathematical psychology* pp. 279–299.

Kuminoff, N. V. & Pope, J. C. (2013), 'The value of residential land and structures during the great housing boom and bust', *Land Economics* **89**(1), 1–29.

Larson, W. D. & Shui, J. (2022), 'Land valuation using public records and kriging: Implications for land versus property taxation in cities', *Journal of Housing Economics* **58**, 101871. - If land value taxation is almost universally held to be so much betterfor society compared to a general property tax, then why has it beenso rarely implemented?2One major reason is the difficulty measuringland's value.
- vacant land sells for a premium on a per-acre basis relative to landunderneath structures
- while tract-level estimates of land value are different whenmeasured using vacant land versus land under structures, they arehighly correlated in levels. However, due to noisiness in the series,growth rate associations are weak
- aggregation rectifies the tract-level noise issue, and the county-level growth rate is similar, facilitating the use of one measure as an instrument or proxy for the other.

Le Rest, K., Pinaud, D., Monestiez, P., Chadoeuf, J. & Bretagnolle, V. (2014), 'Spatial leave-one-out cross-validation for variable selection in the presence of spatial autocorrelation', *Global ecology and biogeography* **23**(7), 811–820.

Lindgren, F., Rue, H. & Lindström, J. (2011), 'An explicit link between gaussian fields and gaussian markov random fields: the stochastic partial differential equation approach', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(4), 423–498.

Longley, P. A., Goodchild, M. F., Maguire, D. J. & Rhind, D. W. (2005), *Geographic information systems and science*, John Wiley & Sons.

Lucas, T., Python, A. & Redding, D. (2020), 'Graphical outputs and spatial cross-validation for the r-inla package using inlautils', *arXiv preprint arXiv:2004.02324* .

Moerkerke, J. (2017), 'The use of the imola template in deed systems'.

Moralı, O. & Yılmaz, N. (2020), 'An analysis of spatial dependence in real estate prices', *The Journal of Real Estate Finance and Economics* pp. 1–23.

Mosteller, F. & Tukey, J. W. (1977), *Data analysis and regression: A second course in statistics*, Addison-wesley series in behavioral science-quantitative methods, Addison-Wesley, Reading, Mass.

Nichols, J. B., Oliner, S. D. & Mulhall, M. R. (2013), 'Swings in commercial and residential land prices in the united states', *Journal of Urban Economics* **73**(1), 57–76.

Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., Dillon, J., Lakshminarayanan, B. & Snoek, J. (2019), 'Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift', *Advances in neural information processing systems* **32**.

Python, A., Bender, A., Blangiardo, M., Illian, J. B., Lin, Y., Liu, B., Lucas, T. C., Tan, S., Wen, Y., Svanidze, D. et al. (2021), 'A downscaling approach to compare covid-19 count data from databases aggregated at different spatial scales', *Journal of the Royal Statistical Society: Series A (Statistics in Society)* .

Roberts, D. R., Bahn, V., Ciuti, S., Boyce, M. S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J. J., Schröder, B., Thuiller, W., Warton, D. I., Wintle, B. A., Hartig, F. & Dormann, C. F. (2017), 'Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure', *Ecography* **40**(8), 913–929.

Rue, H., Martino, S. & Chopin, N. (2009), 'Approximate bayesian inference for latent gaussian models by using integrated nested laplace approximations', *Journal of the royal statistical society: Series b (statistical methodology)* **71**(2), 319–392.

Schäfer, P. & Hirsch, J. (2017), 'Do urban tourism hotspots affect berlin housing rents?', *International Journal of Housing Markets and Analysis* .

Tobler, W. R. (1970), 'A computer movie simulating urban growth in the detroit region', *Economic geography* **46**(sup1), 234–240.

Van Erp, S., Oberski, D. L. & Mulder, J. (2019), 'Shrinkage priors for bayesian penalized regression', *Journal of Mathematical Psychology* **89**, 31–50.

Wang, X., Yue, Y. & Faraway, J. J. (2018), *Bayesian regression modeling with INLA*, Chapman and Hall/CRC.

Weiss, D. J., Nelson, A., Gibson, H. S., Temperley, W., Peedell, S., Lieber, A., Hancher, M., Poyart, E., Belchior, S., Fullman, N., Mappin, B., Dalrymple, U., Rozier, J., Lucas, T. C. D., Howes, R. E., Tusting, L. S., Kang, S. Y., Cameron, E., Bisanzio, D., Battle, K. E., Bhatt, S. & Gething, P. W. (2018), 'A global map of travel time to cities to assess inequalities in accessibility in 2015', *Nature* **553**(7688), 333–336.

Wentland, S. A., Ancona, Z. H., Bagstad, K. J., Boyd, J., Hass, J. L., Gindelsky, M. & Moulton, J. G. (2020), 'Accounting for land in the united states: Integrating physical land cover, land use, and monetary valuation', *Ecosystem Services* **46**, 101178.

Wentland, S., Moulton, J. & Cornwall, G. (2021), 'For what it's worth: Measuring land value in the era of big data and machine learning', *Working paper* .

Wheeler, D. C. & Calder, C. A. (2007), 'An assessment of coefficient accuracy in linear regression models with spatially varying coefficients', *Journal of Geographical Systems* **9**(2), 145–166.

Wood, S. N. (2011), 'Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(1), 3–36.

Yang, Z. (2014), *The Effects of the Two-Rate Property Tax: What Can We Learn from the Pennsylvania experience?*, Lincoln Institute of Land Policy.

A Tables

Variable	Description	Source	Improved Land only
regularity1	Standard deviation of the edge length	Cadaster shapefile 2018	
regularity2	Distance of vertices to centroid of plot	Cadaster shapefile 2018	
regularity3	Number of points	Cadaster shapefile 2018	
regularity4	Ratio of bounding box	Cadaster shapefile 2018	
regularity5	Standard deviation of angles at vertices	Cadaster shapefile 2018	
builtsurface	Proportion of built surface per municipality 2020	Statbel	
log(population)	Population per km^2 in 2020 (logarithm)	Statbel	
log(travel)	Travel time to nearest city in meters (logarithm)	Malaria Atlas Project	
log(distance1)	Distance to water in meters (logarithm)	Open Street Map	
log(distance2)	Distance to amenities in meters (logarithm)	Open Street Map	
log(distance3)	Distance to leisure location in meters (logarithm)	Open Street Map	
log(distance4)	Distance to railway station in meters (logarithm)	Open Street Map	
upperfloors	Number of stories	Belgian Federal Public Service Finance	x
log(salesprice)	Sales price of plot in 1000€ (logarithm)	Belgian Federal Public Service Finance	x
log(surface)	Surface of the plot (logarithm)	Belgian Federal Public Service Finance	x
garages	Number of garages	Belgian Federal Public Service Finance	x
bathrooms	Number of bathrooms	Belgian Federal Public Service Finance	x
habitablerooms	Number of habitable rooms	Belgian Federal Public Service Finance	x
housesurface	Surface of the building	Belgian Federal Public Service Finance	x
age	Age of the building	Belgian Federal Public Service Finance	x
facades	Number of the facades of the building	Belgian Federal Public Service Finance	x
renovation	1 if building was renovated, 0 ow	Belgian Federal Public Service Finance	x
centralheating	1 if building has central heating, 0 ow	Belgian Federal Public Service Finance	x
habitableattic	1 if building has habitable attic, 0 ow	Belgian Federal Public Service Finance	x
improvedland	1 if the data set contains improved land, 0 ow	Belgian Federal Public Service Finance	x

Table 1: Overview of the variables

Statistic	N	Mean	St. Dev.	Min	Median	Max
builtsurface	6,077	212.87	98.31	16.82	197.38	578.16
distance1	6,077	459.90	339.48	0.00	382.43	2,898.89
distance2	6,077	353.65	324.00	2.56	254.15	3,733.05
distance3	6,077	1,159.09	899.12	26.38	906.46	6,594.09
distance4	6,077	4,030.10	3,640.87	51.26	2,989.10	29,056.02
population	6,077	1,189.43	1,456.55	0.00	772.00	26,613.00
regularity1	6,077	15.67	8.99	0.06	13.66	98.05
regularity2	6,077	3.72	4.89	0.0000	2.20	78.22
regularity3	6,077	8.65	6.50	4.00	7.00	318.00
regularity4	6,077	1.89	0.93	1.00	1.62	11.77
regularity5	6,077	36.46	12.93	0.23	40.48	79.79
salesprice	6,077	176.72	145.43	50.00	138.00	2,000.00
surface	6,077	1,386.93	4,948.07	14.00	720.00	255,931.00
travel	6,077	7.73	7.04	0.00	6.00	48.00

Table 2: Summary statistics for vacant land

Statistic	N	Mean	St. Dev.	Min	Median	Max
age	73,843	74.98	41.82	0.00	66.00	168.00
bathrooms	73,828	0.85	0.45	0.00	1.00	3.00
builtsurface	73,846	237.43	111.39	16.82	225.92	676.59
centralheating	73,846	0.63	0.48	0.00	1.00	1.00
distance1	73,846	449.16	334.56	0.00	368.88	3,104.41
distance2	73,846	280.47	295.04	1.04	187.44	4,336.79
distance3	73,846	922.88	850.43	4.21	654.07	8,639.20
distance4	73,846	3,059.84	3,207.96	19.38	1,947.87	32,668.09
facades	73,835	2.88	0.84	2.00	3.00	4.00
garages	73,785	0.68	0.60	0.00	1.00	5.00
habitableattic	73,846	0.35	0.48	0.00	0.00	1.00
habitablerooms	73,832	5.67	1.41	1.00	5.00	10.00
population	73,846	2,277.43	2,956.83	0.00	1,334.00	29,219.00
regularity1	73,846	9.85	7.28	0.00	7.79	91.52
regularity2	73,846	4.09	3.37	0.00	3.45	68.61
regularity3	73,846	10.58	4.85	4.00	10.00	561.00
regularity4	73,846	2.10	1.35	1.00	1.69	34.23
regularity5	73,846	41.93	11.18	0.11	45.52	124.06
renovation	73,846	0.29	0.45	0.00	0.00	1.00
salesprice	73,846	246.05	127.14	50.00	227.00	2,000.00
surface	73,846	553.61	550.37	7.00	360.00	4,000.00
surfaceuseful	73,846	163.26	58.45	16.00	155.00	1,745.00
travel	73,615	5.99	7.16	0.00	4.00	53.00
upperfloors	73,816	1.70	0.58	0.00	2.00	5.00

Note: A lower N indicates missing values.

Table 3: Summary statistics for improved land

Specification	Metric	Model 1	Model 2
spatial	MAE	59.57	48.97
linear	MAE	74.54	64.22
spatial	R^2	0.43	0.67
linear	R^2	0.17	0.46
spatial	ϕ	1.08	1.04
linear	ϕ	1.12	1.07
spatial	RMSE	120.55	86.85
linear	RMSE	137.04	104.38
spatial	rRMSE	0.83	0.68
linear	rRMSE	0.94	0.82

Notes: All metrics reported relate to the testing set and are in-sample. MAE: Mean average error, R^2: R-squared, ϕ: Average ratio of observed values to predicted values, RMSE: Root mean squared error, rRMSE: Relative root mean squared error

Table 4: Prediction of vacant vs. improved land prices

Data set	Specification	Validation	MAE	RMSE	rRMSE
Improved land	Linear	In-sample	64.22	104.38	0.82
		LOOCV	75.34	105.04	0.85
		SLOOCV	78.18	108.45	0.88
	SPDE	In-sample	48.97	86.85	0.68
		LOOCV	44.87	68.37	0.55
		SLOOCV	54.71	91.99	0.74
	ICAR				
		LOOCV	49.06	75.30	0.61
		SLOOCV	59.26	95.53	0.77
Vacant land	Linear	In-sample	75.02	136.85	0.94
		LOOCV	78.3	112.47	1.03
		SLOOCV	82.44	115.18	1.05
	SPDE				
		LOOCV	64.81	96.82	0.88
		SLOOCV	69.09	100.01	0.91
	ICAR	In-sample	59.57	120.55	0.83
		LOOCV	61.46	89.47	0.82
		SLOOCV	67.61	95.41	0.87

Notes: analysis run on 400 randomly drawn observations for both Leave-One-Out validations. In-sample relates to the metrics calculated on the testing set in Table **??**. LOOCV is regular Leave-One-Out-Cross-Validation whereas SLOOCV is Spatial-Leave-One-Out-Cross-Validation. MAE: Mean average error, RMSE: Root mean squared error, rRMSE: Relative root mean squared error

Table 5: Validation of Model 1 and Model 2

Model	v_{pq}	$\widetilde{v_{pq}}$
$p = Model\ 1, \quad q = Model\ 2$	0.89	0.46
$p = Model\ 1, \quad q = Model\ 3$	0.72	0.40
$p = Model\ 1, \quad q = Model\ 4$	0.71	0.41
$p = Model\ 3, \quad q = Model\ 4$	0.04	0.01

Notes: v_{pq} refers to the Hellinger distance between the SD of the predictions for Model p and q. $\widetilde{v_{pq}}$ refers to the Hellinger distance between the demeaned SD of the predictions Model p and q.

Table 6: Hellinger Distance of the predictive SD for Model 1, 3, 4

Metric	*Model 2 clustered*	*Model 2 sparse*
MAE	47.24	48.88
R^2	0.66	0.66
ϕ	1.04	1.04
RMSE	76.24	77.36
rRMSE	0.62	0.62

Notes: Two subsets were taken: A random draw of $n = 6077$ from the improved land data set and a draw of the improved land data set such that it selects improved parcels closest in space to vacant parcels, also $n = 6077$. Model 2 is then run on the first subset *(Model 2 clustered)* and the second subset *(Model 2 sparse)*.

All metrics reported relate to the testing set and are in-sample. MAE: Mean average error, R^2: R-squared, ϕ: Average ratio of observed values to predicted values, RMSE: Root mean squared error, rRMSE: Relative root mean squared error

Table 7: Prediction of improved land prices on a subset

Model	Specification	Validation	MAE	RMSE	rRMSE	Moran's I
Model 1	Linear	In-sample	75.02	136.85	0.94	0.06
		LOOCV	78.3	112.47	1.03	
		SLOOCV	82.44	115.18	1.05	
	SPDE					
		LOOCV	64.81	96.82	0.88	
		SLOOCV	69.09	100.01	0.91	
	ICAR	In-sample	59.57	120.55	0.83	
		LOOCV	61.46	89.47	0.82	
		SLOOCV	67.61	95.41	0.87	
Model 2 clustered	Linear	In-sample	56.83	87.06	0.71	0.29
		LOOCV	60.02	94.88	0.66	
		SLOOCV	64.22	104.05	0.73	
	SPDE	In-sample	47.24	76.24	0.62	
		LOOCV	49.98	80.58	0.56	
		SLOOCV	56.84	86.58	0.70	
Model 2 sparse	Linear	In-sample	63.36	95.13	0.76	0.23
		LOOCV	62.22	90.52	0.70	
		SLOOCV	64.16	94.58	0.79	
	SPDE	In-sample	48.88	77.36	0.62	
		LOOCV	47.14	66.93	0.52	
		SLOOCV	56.47	86.86	0.73	

Notes: Two subsets were taken: A random draw of $n = 6077$ from the improved land data set and a draw of the improved land data set such that it selects improved parcels closest in space to vacant parcels, also $n = 6077$. Model 2 is then run on the first subset *(Model 2 clustered)* and the second subset *(Model 2 sparse)*.
analysis run on 400 randomly drawn observations for both Leave-One-Out validations. In-sample relates to the metrics calculated on the testing set in Table 7. LOOCV is regular Leave-One-Out-Cross-Validation whereas SLOOCV is Spatial-Leave-One-Out-Cross-Validation. MAE: Mean average error, RMSE: Root mean squared error, rRMSE: Relative root mean squared error

Table 8: Validation of Model 1, *Model 2 sparse*, *Model 2 clustered*

Model	v_{pq}	$\widetilde{v_{pq}}$
$p = Model\ 1, \quad q = Model\ 2\ (sparse)$	0.50	0.39
$p = Model\ 1, \quad q = Model\ 2\ (clustered)$	0.78	0.42
$p = Model\ 2\ (sparse), \quad q = Model\ 2\ (clustered)$	0.6	0.09

Notes: v_{pq} refers to the Hellinger distance between the SD of the predictions for Model p and q. $\widetilde{v_{pq}}$ refers to the Hellinger distance between the demeaned SD of the predictions Model p and q.

Table 9: Hellinger Distance of the predictive SD for Model 1, *Model 2 sparse*, *Model 2 clustered*

Specification	Metric	Model 1	Model 3	Model 4
spatial	MAE	59.57	68.97	71.10
linear	MAE	75.02	76.99	78.41
spatial	R^2	0.43	0.35	0.30
linear	R^2	0.17	0.16	0.16
spatial	ϕ	1.08	1.11	1.08
linear	ϕ	1.14	1.14	0.99
spatial	RMSE	120.55	130.43	134.18
linear	RMSE	136.85	138.35	145.19
spatial	rRMSE	0.83	0.92	0.95
linear	rRMSE	0.94	0.98	1.02

Notes: All metrics reported relate to the testing set and are in-sample. MAE: Mean average error, R^2: R-squared, ϕ: Average ratio of observed values to predicted values, RMSE: Root mean squared error, rRMSE: Relative root mean squared error

Table 10: Prediction of vacant land prices

Specification	Metric	Model 5	Model 6
spatial	MAE	125.43	115.57
linear	MAE	120.05	127.26
spatial	R^2	0.31	0.45
linear	R^2	0.18	0.16
spatial	ϕ	0.58	0.64
linear	ϕ	0.35	0.64
spatial	RMSE	164.14	155.19
linear	RMSE	161.57	171.54
spatial	rRMSE	1.29	1.22
linear	rRMSE	1.27	1.35

Notes: All metrics reported relate to the testing set and are in-sample. MAE: Mean average error, R^2: R-squared, ϕ: Average ratio of observed values to predicted values, RMSE: Root mean squared error, rRMSE: Relative root mean squared error

Table 11: Prediction of land prices underneath improvement

B Pictures

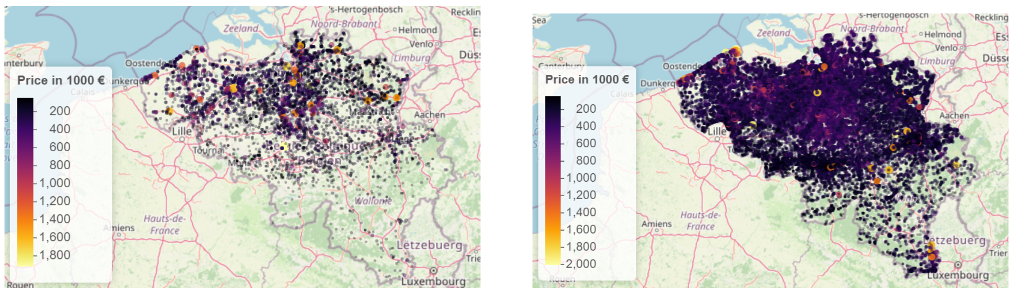

(a) Vacant land (b) Improved land

Figure 1: Land prices across Belgium

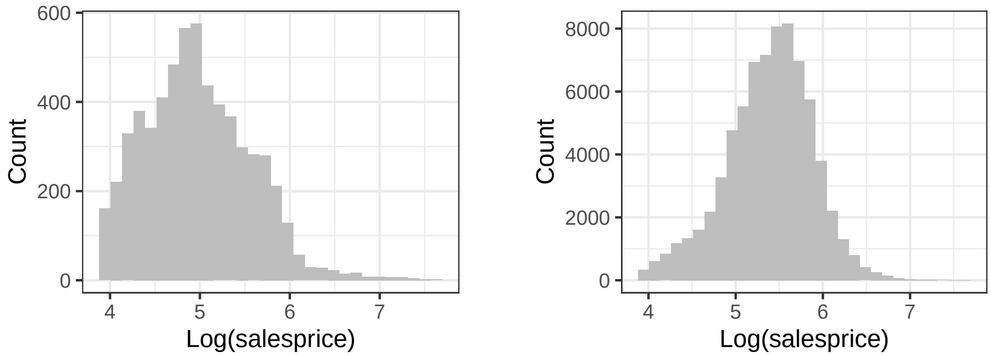

(a) Vacant land (b) Improved land

Figure 2: Histograms of land prices on the log-scale

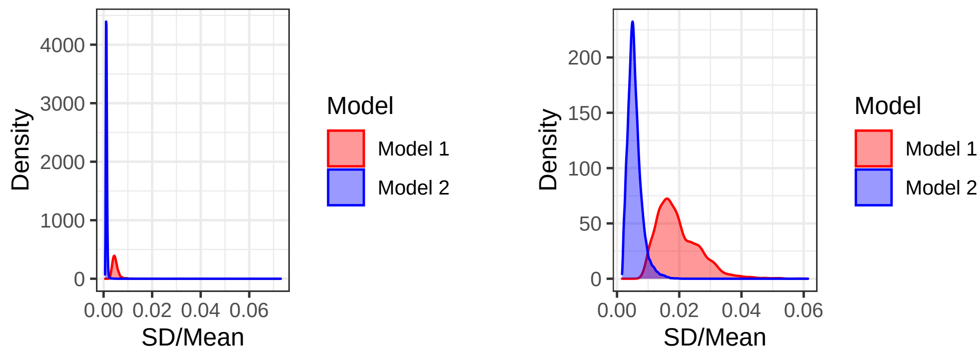

(a) Linear (b) Spatial

Figure 3: Density of predictive SD/ predictive Mean for Model 1 and Model 2 (linear and spatial)

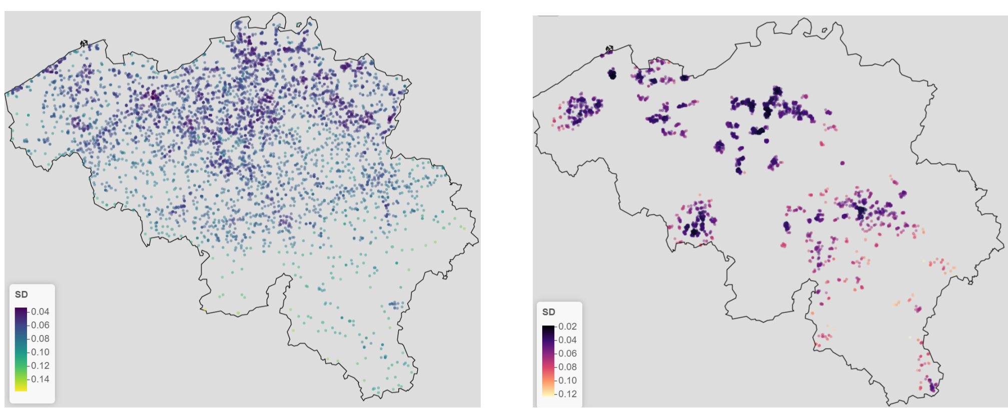

(a) Sparse (b) Clustered

Figure 4: Predictive SD for sparse and clustered subset of improved land

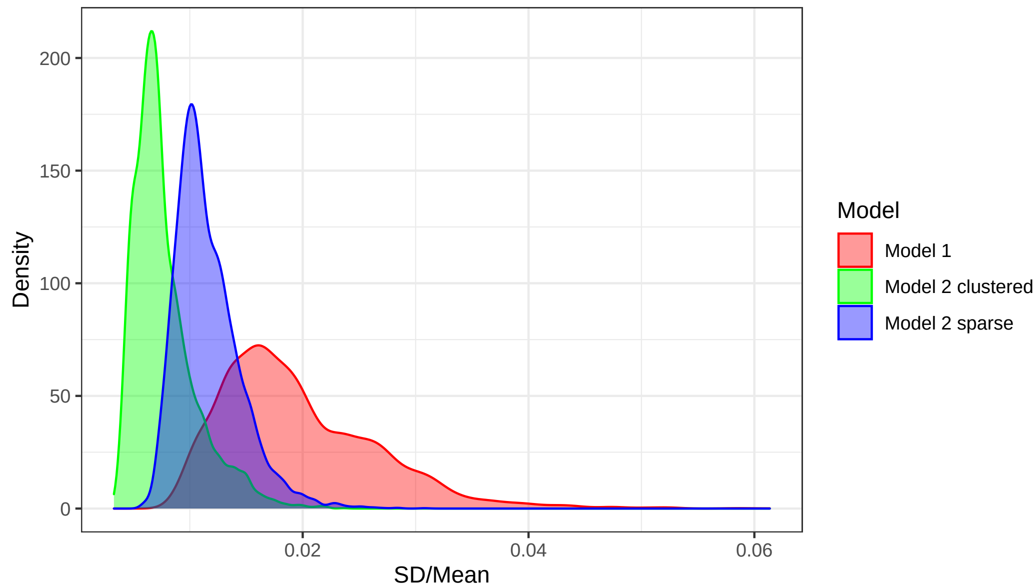

Figure 5: Density of predictive SD/ predictive Mean for Model 1, *Model 2 sparse*, *Model 2 clustered*

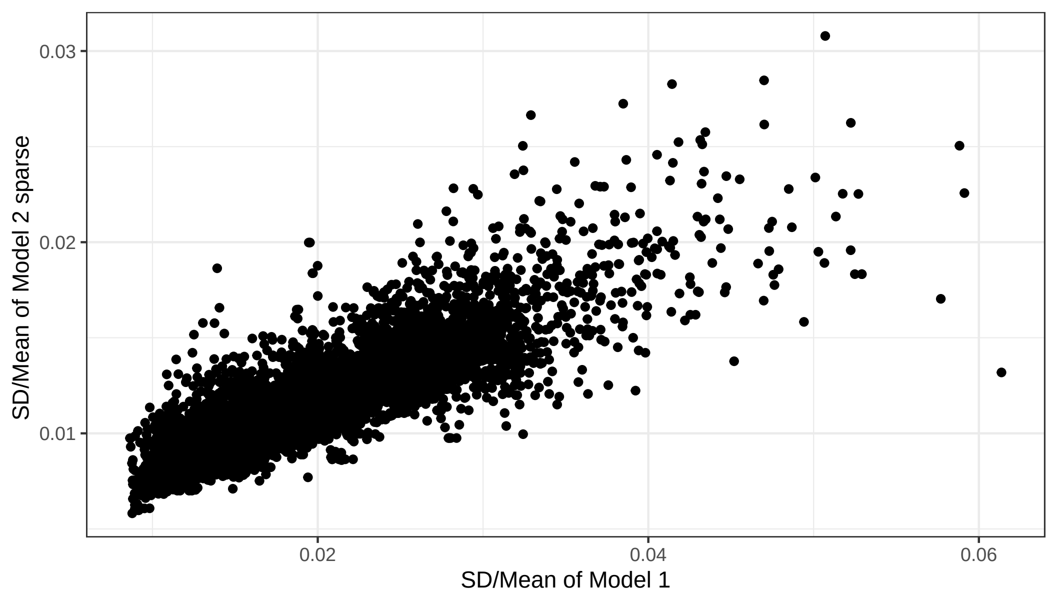

Figure 6: Scatterplot of predictive SD/predictive Mean of Model 1 vs *Model 2 sparse*

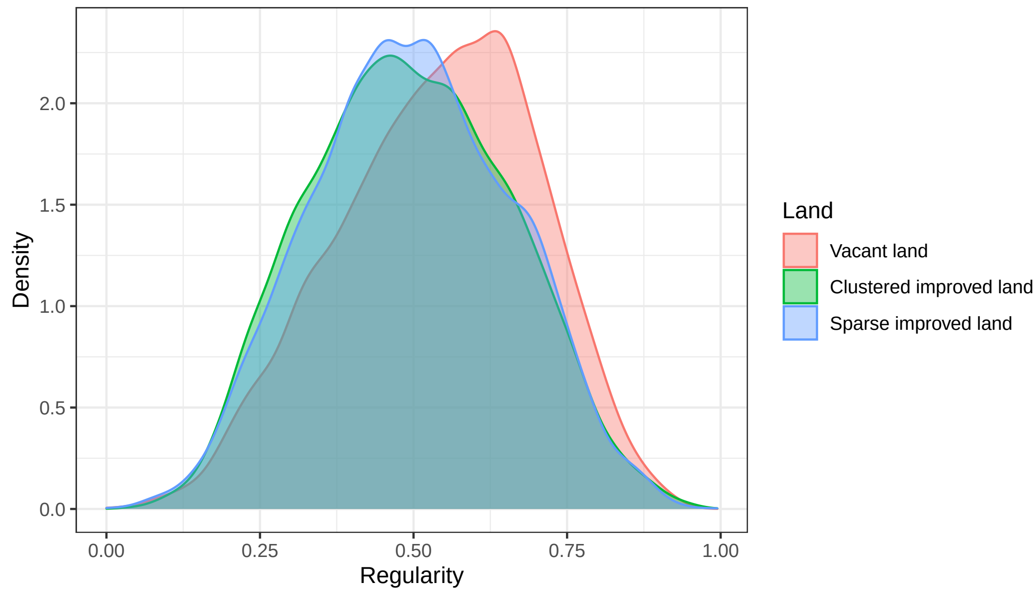

Figure 7: Density of Regularity of vacant land, sparse improved land and clustered improved land

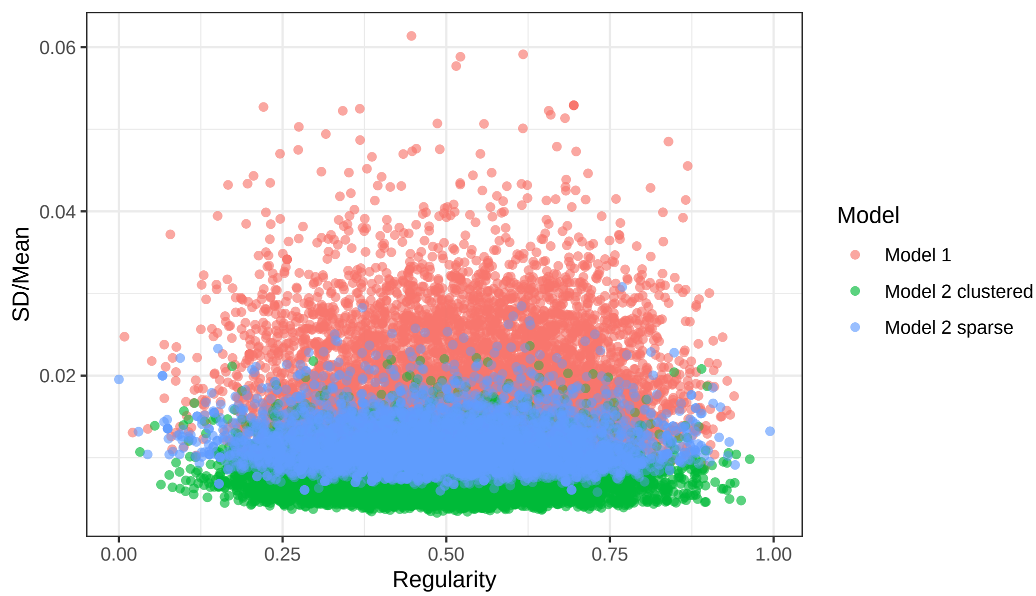

Figure 8: Distribution of predictive SD/ predictive Mean by regularity for Model 1, *Model 2 sparse*, *Model 2 clustered*

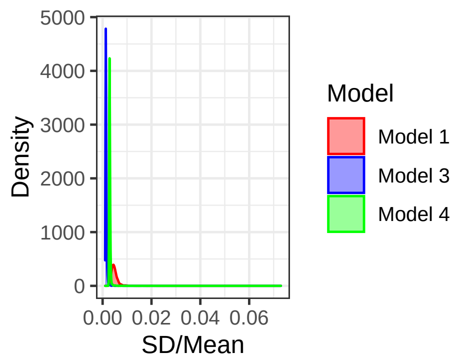

(a) Sparse

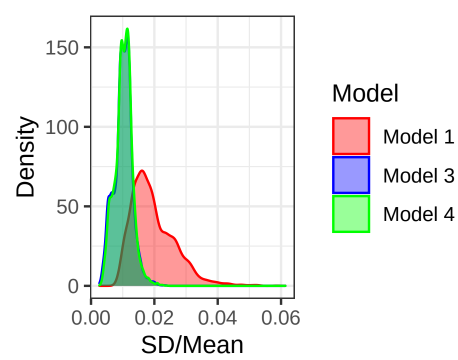

(b) Clustered

Figure 9: Density of predictive SD/ predictive Mean for Model 1, Model 3, Model 4

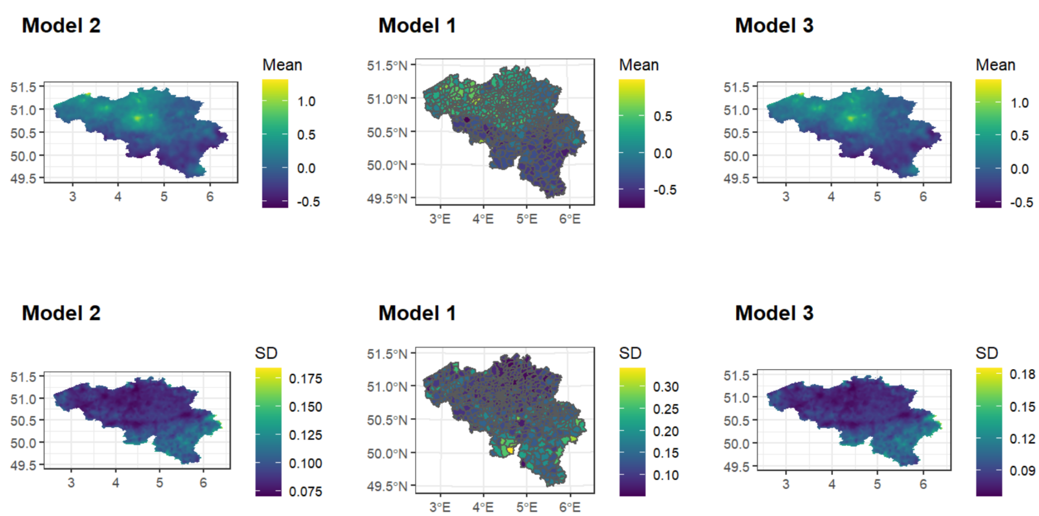

Figure 10: Predictive Mean and SD for Model 2, Model 1, Model 3

C Additional results

Specification	Metric	Model 1	Model 2
spatial	MAE	75.40	48.97
linear	MAE	85.56	64.22
spatial	R^2	0.52	0.67
linear	R^2	0.43	0.46
spatial	ϕ	3.24	1.04
linear	ϕ	3.19	1.07
spatial	RMSE	315.33	86.85
linear	RMSE	319.52	104.38
spatial	rRMSE	0.96	0.68
linear	rRMSE	0.98	0.82

Notes: Analysis run on the full data set with outliers included. All metrics reported relate to the testing set and are in-sample. MAE: Mean average error, R^2: R-squared, ϕ: Average ratio of observed values to predicted values, RMSE: Root mean squared error, rRMSE: Relative root mean squared error

Table 12: Prediction of vacant land vs. improved land prices - full data se

Specification	Metric	Model 1	Model 3	Model 4
spatial	MAE	75.40	86.50	86.59
linear	MAE	85.56	91.00	90.50
spatial	R^2	0.52	0.51	0.30
linear	R^2	0.43	0.43	0.33
spatial	ϕ	3.24	3.77	15.84
linear	ϕ	3.19	4.20	15.12
spatial	RMSE	315.33	324.25	319.78
linear	RMSE	319.52	327.86	323.47
spatial	rRMSE	0.96	0.99	0.98
linear	rRMSE	0.98	1.00	0.99

Notes: Analysis run on the full data set with outliers included. All metrics reported relate to the testing set and are in-sample. MAE: Mean average error, R^2: R-squared, ϕ: Average ratio of observed values to predicted values, RMSE: Root mean squared error, rRMSE: Relative root mean squared error

Table 13: Prediction of vacant land prices - full data se

Specification	Metric	Model 5	Model 6
spatial	MAE	168.12	115.57
linear	MAE	176.36	127.25
spatial	R^2	0.34	0.45
linear	R^2	0.25	0.16
spatial	ϕ	0.33	0.64
linear	ϕ	0.31	0.64
spatial	RMSE	199.07	155.19
linear	RMSE	208.67	171.54
spatial	rRMSE	1.57	1.22
linear	rRMSE	1.64	1.35

Notes: Analysis run on the full data set with outliers included. All metrics reported relate to the testing set and are in-sample. MAE: Mean average error, R^2: R-squared, ϕ: Average ratio of observed values to predicted values, RMSE: Root mean squared error, rRMSE: Relative root mean squared error

Table 14: Prediction of land prices underneath improvement - full data set