



2011

An Autoregressive Approach to House Price Modeling

Chaitra H. Nagaraja

Lawrence D. Brown
University of Pennsylvania

Linda Zhao
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/statistics_papers



Part of the [Applied Statistics Commons](#)

Recommended Citation

Nagaraja, C. H., Brown, L. D., & Zhao, L. (2011). An Autoregressive Approach to House Price Modeling. *Annals of Applied Statistics*, 5 (1), 124-149. <http://dx.doi.org/10.1214/10-AOAS380>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/statistics_papers/364
For more information, please contact repository@pobox.upenn.edu.

An Autoregressive Approach to House Price Modeling

Abstract

A statistical model for predicting individual house prices and constructing a house price index is proposed utilizing information regarding sale price, time of sale and location (ZIP code). This model is composed of a fixed time effect and a random ZIP (postal) code effect combined with an autoregressive component. The former two components are applied to all home sales, while the latter is applied only to homes sold repeatedly. The time effect can be converted into a house price index. To evaluate the proposed model and the resulting index, single-family home sales for twenty US metropolitan areas from July 1985 through September 2004 are analyzed. The model is shown to have better predictive abilities than the benchmark S&P/Case–Shiller model, which is a repeat sales model, and a conventional mixed effects model. Finally, Los Angeles, CA, is used to illustrate a historical housing market downturn.

Keywords

housing index, time series, repeat sales

Disciplines

Applied Statistics

AN AUTOREGRESSIVE APPROACH TO HOUSE PRICE MODELING¹

BY CHAITRA H. NAGARAJA, LAWRENCE D. BROWN² AND LINDA H. ZHAO

US Census Bureau, University of Pennsylvania and University of Pennsylvania

A statistical model for predicting individual house prices and constructing a house price index is proposed utilizing information regarding sale price, time of sale and location (ZIP code). This model is composed of a fixed time effect and a random ZIP (postal) code effect combined with an autoregressive component. The former two components are applied to all home sales, while the latter is applied only to homes sold repeatedly. The time effect can be converted into a house price index. To evaluate the proposed model and the resulting index, single-family home sales for twenty US metropolitan areas from July 1985 through September 2004 are analyzed. The model is shown to have better predictive abilities than the benchmark S&P/Case–Shiller model, which is a repeat sales model, and a conventional mixed effects model. Finally, Los Angeles, CA, is used to illustrate a historical housing market downturn.

1. Introduction. Modeling house prices presents a unique set of challenges. Houses are distinctive, each has its own set of hedonic characteristics: number of bedrooms, square footage, location, amenities and so forth. Moreover, the price of a house, or the value of the bundle of characteristics, is observed only when sold. Sales, however, occur infrequently. As a result, during any period of time, out of the entire population of homes, only a small percentage are actually sold. From this information, our objective is to develop a practical model to predict prices from which we can construct a price index. Such an index would summarize the housing market and would be used to monitor changes over time. Including both objectives allows one to look at both micro and macro features of a market, from individual houses to entire markets. In the following discussion, we propose an autoregressive model which is a simple, but effective and interpretable, way to model house prices and construct an index. We show that our model outperforms, in a predictive sense, the benchmark S&P/Case–Shiller Home Price Index method when applied to housing data for twenty US metropolitan areas. We use these results to evaluate the proposed autoregressive model as well as the resulting house price index.

Received April 2009; revised June 2010.

¹Disclaimer: This report is released to inform interested parties of research and to encourage discussion. The views expressed on statistical issues are those of the authors and not necessarily those of the US Census Bureau.

²Supported in part by NSF Grant DMS-07-07033.

Key words and phrases. Housing index, time series, repeat sales.

A common approach for modeling house prices, called repeat sales, utilizes homes that sell multiple times to track market trends. Bailey, Muth and Nourse (1963) first proposed this method and Case and Shiller (1987, 1989) extended it to incorporate heteroscedastic errors. In both models, the log price difference between two successive sales of a home is used to construct an index using linear regression. The previous sale price acts as a surrogate for hedonic information, provided the home does not change substantially between sales. There is a large body of work focused on improving the index estimates produced by the Bailey, et al. approach. For instance, a modified form of the repeat sales model is used for the Home Price Index produced by the Office of Federal Housing Enterprise Oversight (OFHEO). Gatzlaff and Haurin (1997) suggest a repeat sales model that corrects for the correlation between economic conditions and the chance of a sale occurring. Alternatively, Shiller (1991) and Goetzmann and Peng (2002) propose arithmetic average versions of the repeat sales estimator as an alternative to the original geometric average estimator. The former work is used commercially by Standard and Poors to produce the S&P/Case–Shiller Home Price Index. We will be using this index in our analysis as it is the most well known.

Several criticisms have been made about repeat sales methods. Theoretically, for a house to be included in a repeat sales analysis, no changes must have been made to it; however, in practice, that is almost never the case. Furthermore, Englund, Quigley and Redfearn (1999) and Goetzmann and Speigel (1995) have commented on the difficulty of detecting such changes without the availability of additional information about the home. Goetzmann and Speigel, however, do propose an alternate model which corrects for the effect of changes to homes around the time the house is sold.

Even if homes which have changed are removed from the data set, an index constructed out of the remaining homes may still not reflect the true index value. Case and Quigley (1991) argue that houses age which has a depreciating effect on their price. Therefore, as Case, Pollakowski and Wachter (1991) write, repeat sales indices produce estimates of time effects confounded with age effects. Palmquist (1982) has suggested applying an independently computed depreciation factor to account for the impact of age.

In a sample period, out of the entire population of homes, only a small fraction are actually sold. A fraction of these sales are repeat sales homes with no significant changes. Recall that the remaining sales, those of the single sales homes, are omitted from the analysis. If repeat sales indices are used to describe the housing market as a whole, one would like the sample of repeat sales homes to have similar characteristics to all homes. If not, Case, Pollakowski and Wachter remark that the indices would be affected by sample selection bias. Englund, Quigley and Redfearn in a study of Swedish home sales, and Meese and Wallace (1997), in a study of Oakland and Fremont home sales, both found that repeat sales homes are indeed different from single sale homes. Both studies also observed that in addition to being older, repeat sales homes were smaller and more “modest” [Englund,

Quigley and Redfean (1999)]. Therefore, repeat sales indices seem to provide information only about a very specific type of home and may not apply to the entire housing market. However, published indices do not seem to be interpreted in that manner. Case and Quigley (1991) propose an alternative hybrid model that combines repeat sales methodology with hedonic information which makes use of all sales. While the index constructed with this method represents all home sales, it requires housing characteristics which may be difficult to collect on a broad scale.

We feel the repeat sales concept is valuable although the current models of this type have the issues described above. The proposed model applies the repeat sales idea in a new way to address some of the criticisms while still maintaining the simplicity and reduced data requirements that the original Bailey et al. method had. While our primary goal is prediction, we believe the resulting index could be a better general description of housing sales than traditional repeat sales methodology.

In our method, log prices are modeled as the sum of a time effect (index), a location effect modeled as a random effect for ZIP (postal) code, and an underlying first-order autoregressive time series [AR(1)]. This structure offers four advantages. First, the price index is estimated with all sales: single and repeat. Essentially, the index can be thought of as a weighted sum of price information from single and repeat sales. The latter component receives a much higher weight because more useful information is available for those homes. Second, the previous sale price becomes less useful the longer it has been since the last sale. The AR(1) series includes this feature into the model more directly than the Case–Shiller method. Third, metropolitan areas are diverse and neighborhoods may have disparate trends. We include ZIP code effects to model these differences in location.³ Finally, the proposed model is straightforward to interpret even while including the features described above. We believe the model captures trends in the overall housing market better than existing repeat sales methods and is a practical alternative.

We apply this model to data on single family home sales from July 1985 through September 2004 for twenty US metropolitan areas. These data are described in Section 2. The autoregressive model is outlined and estimation using maximum likelihood is described in Section 3; results are discussed in Section 4. For comparison, two alternative models are fit: a conventional mixed effects model and the method used in the S&P/Case–Shiller Home Price Index. As a quantitative way to compare the indices, the predictive capacity of the three methods are assessed in Section 5. In Section 6 we examine the case of Los Angeles, CA, where the proposed model does not perform as well. We end with a general discussion in Section 7.

³ZIP code was readily available in our data; other geographic variables at roughly this scale might have been even more useful had they been available.

TABLE 1
Metropolitan areas in the data

Ann Arbor, MI	Kansas City, MO	Minneapolis, MN	Raleigh, NC
Atlanta, GA	Lexington, KY	Orlando, FL	San Francisco, CA
Chicago, IL	Los Angeles, CA	Philadelphia, PA	Seattle, WA
Columbia, SC	Madison, WI	Phoenix, AZ	Sioux Falls, SD
Columbus, OH	Memphis, TN	Pittsburgh, PA	Stamford, CT

2. House price data. The data are comprised of single family home sales qualifying for conventional mortgages from the twenty US metropolitan areas listed in Table 1. These sales occurred between July 1985 and September 2004. Not included in these data are homes with prices too high to be considered for a conventional mortgage or those sold at subprime rates. Note, however, that subprime loans were not prevalent during the time period covered by our data. Similar data are used by Fannie Mae, Freddie Mac, and to construct the OFHEO Home Price Index.

For each sale, the following information is available: address with ZIP code, month and year of sale, and price. To ensure adequate data per time period, we divide the sample period into three month intervals for a total of 77 periods, or quarters. We make an attempt to remove sales which are not arm's length by omitting homes sold more than once in a single quarter. Given the lack of hedonic information, we have no way of determining whether a house has changed substantially between sales. Therefore, we do not filter our data to remove such houses.

Table 2 displays the number of sales and unique houses sold in the sample period for a selection of cities. Complete tables for all summaries in this section are provided in Appendix A. Observe that the total number of sales is always greater than the number of houses because houses can sell multiple times (repeat sales). Perhaps more illuminating is Table 3, where we count the number of times each house is sold. We see that as the number of sales per house increases, the number of houses reduces rapidly. Nevertheless, a significant number of houses sell more than twice. With a sample period of nearly twenty years, this is not unusual; how-

TABLE 2
Summary counts for a selection of cities

Metropolitan area	Sales	Houses
Stamford, CT	14,602	11,128
Ann Arbor, MI	68,684	48,522
Pittsburgh, PA	104,544	73,871
Los Angeles, CA	543,071	395,061
Chicago, IL	688,468	483,581

TABLE 3
Sale frequencies for a selection of cities

Metropolitan area	1 sale	2 sales	3 sales	4+ sales
Stamford, CT	8,200	2,502	357	62
Ann Arbor, MI	32,458	12,662	2,781	621
Pittsburgh, PA	48,618	20,768	3,749	718
Los Angeles, CA	272,258	100,918	18,965	2,903
Chicago, IL	319,340	130,234	28,369	5,603

ever, single sales are the most common despite the long sample period. The first column of Table 3 shows this clearly. Moreover, this pattern holds for all cities in our data. Finally, in Figure 1, we plot the median price across time for the subset of cities. This graph illustrates that both the cost of homes and the trends over time vary considerably across cities.

For all metropolitan areas in our data, the time of a sale is fuzzy, as there is often a lag between the day when the price is agreed upon and the day the sale is recorded (around 20–60 days). Theoretically, the true value of the house would have changed between these two points. Therefore, in the strictest sense, the sale price of the house does not reflect the price at the time when the sale is recorded. Dividing the year into quarters reduces the importance of this lag effect.

3. Model. The log house price series is modeled as the sum of an index component, an effect for ZIP code (as an indicator for location), and an AR(1) time series. The sale prices of a particular house are treated as a series of sales: $y_{i,1,z}, y_{i,2,z}, \dots, y_{i,j,z}, \dots$, where $y_{i,j,z}$ is the log sale price of the j th sale of the i th house in ZIP code z . Note that $y_{i,1,z}$ is defined as the first sale price in the *sample period*; as a result, both new homes and old homes sold for the first time in the sample period are indicated with the same notation.

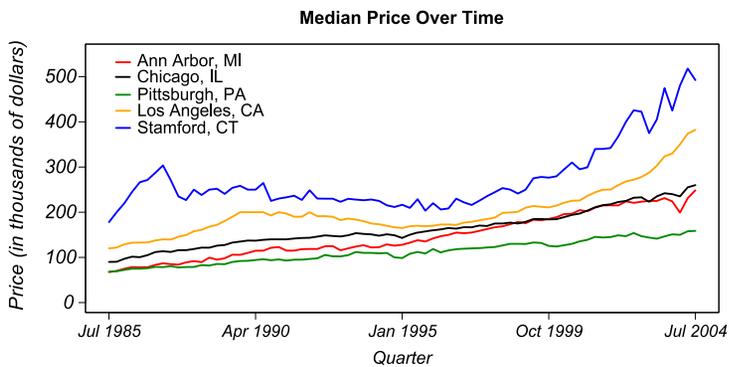


FIG. 1. *Median prices for a selection of cities.*

Let there be $1, \dots, T$ discrete time periods where house sales occur. Allow $t(i, j, z)$ to denote the time period when the j th sale of the i th house in ZIP code z occurs and let $\gamma(i, j, z) = t(i, j, z) - t(i, j-1, z)$, or the gap time between sales. Finally, there are a total of $N = \sum_{z=1}^Z \sum_{i=1}^{I_z} J_i$ observations in the data where there are Z ZIP codes, I_z houses in each ZIP code and J_i sales for a given house.

The log sale price $y_{i,j,z}$ can now be described as follows:

$$(1) \quad \begin{aligned} y_{i,1,z} &= \mu + \beta_{t(i,1,z)} + \tau_z + \varepsilon_{i,1,z}, & j &= 1, \\ y_{i,j,z} &= \mu + \beta_{t(i,j,z)} + \tau_z + \phi^{\gamma(i,j,z)}(y_{i,j-1,z} - \mu - \beta_{t(i,j-1,z)} - \tau_z) \\ &\quad + \varepsilon_{i,j,z}, & j &> 1, \end{aligned}$$

where:

1. The parameter $\beta_{t(i,j,z)}$ is the log price index at time $t(i, j, z)$. Let β_1, \dots, β_T denote the log price indices, assumed to be fixed effects.
2. ϕ is the autoregressive coefficient and $|\phi| < 1$.
3. τ_z is the random effect for ZIP code z . $\tau_z \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_\tau^2)$ where τ_1, \dots, τ_Z are the ZIP code random effects which are distributed normally with mean 0 and variance σ_τ^2 and where i.i.d. denotes independent and identically distributed.
4. We impose the restriction that $\sum_{t=1}^T n_t \beta_t = 0$ where n_t is the number of sales at time t . This allows us to interpret μ as an overall mean.
5. Finally, let

$$\varepsilon_{i,1,z} \sim \mathcal{N}\left(0, \frac{\sigma_\varepsilon^2}{1 - \phi^2}\right), \quad \varepsilon_{i,j,z} \sim \mathcal{N}\left(0, \frac{\sigma_\varepsilon^2(1 - \phi^{2\gamma(i,j,z)})}{1 - \phi^2}\right),$$

and assume that all $\varepsilon_{i,j,z}$ are independent.

Note that there is only one process for the series $y_{i,1,z}, y_{i,2,z}, \dots$. The error variance for the first sale, $\sigma_\varepsilon^2/(1 - \phi^2)$, is a marginal variance. For subsequent sales, because we have information about previous sales, it is appropriate to use the conditional variance (conditional on the previous sale), $\sigma_\varepsilon^2(1 - \phi^{2\gamma(i,j,z)})/(1 - \phi^2)$, instead. For more details refer to the supplemental article [Nagaraja, Brown and Zhao (2010)].

The underlying series for each house is given by $u_{i,j,z} = y_{i,j,z} - \mu - \beta_{t(i,j,z)} - \tau_z$. We can rewrite this series as $u_{i,j,z} = \phi^{\gamma(i,j,z)}u_{i,j-1,z} + \varepsilon_{i,j,z}$ where $\varepsilon_{i,j,z}$ is as given above. This autoregressive series is stationary, given a starting observation $u_{i,1,z}$, because $E[u_{i,j,z}] = 0$, a constant, where $E[\cdot]$ is the expectation function, and the covariance between two points depends only on the gap time and not on the actual sale times. Specifically, $\text{Cov}(u_{i,j,z}, u_{i,j',z}) = \sigma_\varepsilon^2 \phi^{(t(i,j',z) - t(i,j,z))}/(1 - \phi^2)$ if $j < j'$. Therefore, the covariance between a pair of sales depends *only* on the gap time between sales. Consequently, the time of sale is uninformative for the underlying series, only the gap time is required. As a result, the autoregressive series $u_{i,j,z}$ where i and z are fixed and $j \geq 1$ is a Markov process.

The autoregressive component adds two important features to the model. Intuitively, the longer the gap time between sales, the less useful the previous price should become when predicting the next sale price. For the model described in (1), as the gap time increases, the autoregressive coefficient decreases by construction ($\phi^{\gamma(i,j,z)}$), meaning that sales prices of a home with long gap times are less correlated with each other. (See Remark 3.1 at the end of this section for additional discussion on the form of ϕ .) Moreover, as the gap time increases, the variance of the error term increases. This indicates that the information contained in the previous sale price is less useful as the time between sales grows.

To fit the model, we formulate the autoregressive model in (1) in matrix form:

$$(2) \quad \mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\boldsymbol{\tau} + \boldsymbol{\varepsilon}^*,$$

where \mathbf{y} is the vector of log prices and \mathbf{X} and \mathbf{Z} are the design matrices for the fixed effects $\boldsymbol{\beta} = [\mu\beta_1 \cdots \beta_{T-1}]'$ and random effects $\boldsymbol{\tau}$, respectively. Then, the log price can be modeled as a mixed effects model with autocorrelated errors, $\boldsymbol{\varepsilon}^*$, and with covariance matrix \mathbf{V} .

We apply a transformation matrix \mathbf{T} to the model in (2) to simplify the computations; essentially, this matrix applies the autoregressive component of the model to both sides of (2). It is an $N \times N$ matrix and is defined as follows. Let $t_{(i,j,z),(i',j',z')}$ be the cell corresponding to the (i, j, z) th row and (i', j', z') th column. Then,

$$(3) \quad t_{(i,j,z),(i',j',z')} = \begin{cases} 1, & \text{if } i = i', j = j', z = z', \\ -\phi^{\gamma(i,j)}, & \text{if } i = i', j = j' + 1, z = z', \\ 0, & \text{otherwise.} \end{cases}$$

As a result, $\mathbf{T}\boldsymbol{\varepsilon}^* \sim \mathcal{N}(\mathbf{0}, \frac{\sigma_\varepsilon^2}{1-\phi^2} \text{diag}(\mathbf{r}))$ where $\text{diag}(\mathbf{r})$ is a diagonal matrix of dimension N with the diagonal elements \mathbf{r} being given by

$$(4) \quad r_{i,j,z} = \begin{cases} 1, & \text{when } j = 1, \\ 1 - \phi^{2\gamma(i,j)}, & \text{when } j > 1. \end{cases}$$

Using the notation from (1), let $\boldsymbol{\varepsilon} = \mathbf{T}\boldsymbol{\varepsilon}^*$. Finally, we restrict $\sum_{t=1}^T n_t \beta_t = 0$ where n_t is the number of sales at time t . Therefore, $\beta_T = -\frac{1}{n_T} \sum_{t=1}^{T-1} n_t \beta_t$.

The likelihood function for the transformed model is

$$(5) \quad L(\boldsymbol{\theta}; \mathbf{y}) = (2\pi)^{-N/2} |\mathbf{V}|^{-1/2} \times \exp\left\{-\frac{1}{2}(\mathbf{T}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}))' \mathbf{V}^{-1} (\mathbf{T}(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}))\right\},$$

where $\boldsymbol{\theta} = \{\boldsymbol{\beta}, \sigma_\varepsilon^2, \sigma_\tau^2, \phi\}$ is the vector of parameters, N is the total number of observations, \mathbf{V} is the covariance matrix, and \mathbf{T} is the transformation matrix. We can split \mathbf{V} into a sum of the variance contributions from the time series and the random effects. Specifically,

$$(6) \quad \mathbf{V} = \frac{\sigma_\varepsilon^2}{1-\phi^2} \text{diag}(\mathbf{r}) + (\mathbf{TZ})\mathbf{D}(\mathbf{TZ})',$$

where $\mathbf{D} = \sigma_\tau^2 \mathbf{I}_Z$ and \mathbf{I}_Z is an identity matrix with dimension $Z \times Z$.

We use the coordinate ascent algorithm to compute the maximum likelihood estimates (MLE) of θ for the model in (1). This iterative procedure maximizes the likelihood function with respect to each group of parameters while holding all other parameters constant. The algorithm terminates when the parameter estimates have converged according to the specified stopping rule. Bickel and Doksum (2001) include a proof showing that, for models in the exponential family, the estimates computed using the coordinate ascent algorithm converge to the MLE. The proposed model, however, is a member of the differentiable exponential family; therefore, as Brown (1986) states, the proof does not directly apply. Nonetheless, we find empirically that the likelihood function is well behaved, so the MLE appears to be reached for this case as well. Empirical evidence of convergence can be found in the supplemental article [Nagaraja, Brown and Zhao (2010)].

We outline Algorithm 1 below. The equations for updating the parameters and random effects estimates are given in Appendix B.

To predict a log price, we substitute the estimated parameters and random effects into (1):

$$(7) \quad \hat{y}_{i,j,z} = \hat{\mu} + \hat{\beta}_{t(i,j,z)} + \hat{\tau}_z + \hat{\phi}^{\gamma(i,j,z)} (y_{i,j-1,z} - \hat{\mu} - \hat{\beta}_{t(i,j-1,z)} - \hat{\tau}_z).$$

We then convert $\hat{y}_{i,j,z}$ to the price scale (denoted as $\hat{Y}_{i,j,z}$) using

$$(8) \quad \hat{Y}_{i,j,z}(\sigma^2) = \exp \left\{ \hat{y}_{i,j,z} + \frac{\sigma^2}{2} \right\},$$

where σ^2 denotes the variance of $y_{i,j,z}$. The additional term $\sigma^2/2$ approximates the difference between $E[\exp\{X\}]$ and $\exp\{E[X]\}$ where $E[\cdot]$ is the expectation function. We must adjust the latter expression to approximate the conditional mean of the response, \mathbf{y} . We improve the efficiency of our estimates by using the adjustment stated in Shen, Brown and Zhi (2006). In (8), σ^2 can be estimated from the

Algorithm 1 Autoregressive (AR) model fitting algorithm.

1. Set a tolerance level ϵ (possibly different for each parameter).
 2. Initialize the parameters: $\theta^0 = \{\beta^0, \sigma_\epsilon^{2,0}, \sigma_\tau^{2,0}, \phi^0\}$.
 3. For iteration k ($k = 0$ when the parameters are initialized),
 - (a) Calculate β^k using (19) in Appendix B with $\{\sigma_\epsilon^{2,k-1}, \sigma_\tau^{2,k-1}, \phi^{k-1}\}$.
 - (b) Compute $\sigma_\epsilon^{2,k}$ by computing the zero of (20) using $\{\beta^k, \sigma_\tau^{2,k-1}, \phi^{k-1}\}$.
 - (c) Compute $\sigma_\tau^{2,k}$ by calculating the zero of (21) using $\{\beta^k, \sigma_\epsilon^{2,k}, \phi^{k-1}\}$.
 - (d) Find the zero of (22) to compute ϕ^k using $\{\beta^k, \sigma_\epsilon^{2,k}, \sigma_\tau^{2,k}\}$.
 - (e) If $|\theta_i^{k-1} - \theta_i^k| > \epsilon$ for any $\theta_i \in \theta$, repeat step 3 after replacing θ^{k-1} with θ^k .
Otherwise, stop (call this iteration K).
 4. Solve for β_T by computing: $\hat{\beta}_T = -\frac{1}{n_T} \sum_{t=1}^{T-1} n_t \hat{\beta}_t^K$.
 5. Plug in $\{\beta^K, \sigma_\epsilon^{2,K}, \sigma_\tau^{2,K}, \phi^K\}$ to compute the estimated values for τ using (23).
-

mean squared residuals (MSR), where $\text{MSR} = \frac{1}{N} \sum_{i=1}^N (y_{i,j,z} - \hat{y}_{i,j,z})^2$ and N is the total number of observations used to fit the model. Therefore, the log price estimates, $\hat{y}_{i,j,z}$, are converted to the price scale by

$$(9) \quad \hat{Y}_{i,j,z} = \exp \left\{ \hat{y}_{i,j,z} + \frac{\text{MSR}}{2} \right\}.$$

Goetzmann (1992) proposes a similar transformation for the index values computed using a traditional repeat sales method. Calhoun (1996) suggests applying Goetzmann's adjustment when using an index value to predict a particular house price. For the autoregressive model, the standard error of the index is sufficiently small that the efficiency adjustment has a negligible impact on the estimated index. Therefore, we simply use $\exp\{\hat{\beta}_t\}$ to convert the index to the price scale. Finally, we rescale the vector of indices so that the first quarter has an index value of 1.

REMARK 3.1. The autoregressive coefficient form, $\phi^{\gamma(i,j,z)}$, deserves further explanation. For each house indexed by (i, z) , let $t_1(i, z) = t(i, 1, z)$ denote the time of the initial sale. Conditioning on the (unobserved) values of the parameters $\{\mu, \beta_t, \sigma_\varepsilon^2, \sigma_\tau^2\}$ and on the values of the random ZIP code effects, $\{\tau_z\}$, let $\{u_{i,z;t} : t = t_1(i, z), t_1(i, z) + 1, \dots\}$ be an underlying AR(1) process. To be more precise, $u_{i,z;t}$ is a conventional, stationary AR(1) process defined by

$$(10) \quad u_{i,z;t} = \begin{cases} \varepsilon_{i,1,z}, & \text{if } t = t_1(i, z), \\ \phi u_{i,z;t-1} + \varepsilon_{i,1,z}, & \text{if } t > t_1(i, z), \end{cases}$$

where if $t = t(i, j, z)$, then $\varepsilon_{i,z;t(i,j,z)} = \varepsilon_{i,j,z}$ and otherwise $\varepsilon_{i,z;t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \frac{\sigma_\varepsilon^2}{1-\phi^2})$. Then the observed log sale prices are given by $\{y_{i,j,z}\}$ where $u_{i,z;t(i,j,z)} = y_{i,j,z} - (\mu + \beta_{t(i,j,z)} + \tau_z)$. The values of $u_{i,z;t}$ are to be interpreted as the *potential sale price adjusted* by $\{\mu, \beta_t, \sigma_\varepsilon^2, \sigma_\tau^2\}$ of the house indexed by (i, z) if the house were to be sold at time t .

For housing data like ours, the value of the autoregressive parameter ϕ for this latent process will be near the largest possible value, $\phi = 1$. Consequently, if the underlying process were actually an observed process from which one wanted to estimate ϕ , then estimation of ϕ could be a delicate matter. However, sales generally occur with fairly large gap times and so the values of $\phi^{\gamma(i,j,z)}$ occurring in the data will generally not be close to 1. For that reason, conventional estimation procedures perform satisfactorily when estimating ϕ . We provide empirical evidence for this in Section 4 and in the supplemental article [Nagaraja, Brown and Zhao (2010)].

4. Estimation results. To fit and validate the autoregressive (AR) model, we divide the observations for each city into training and test sets. The test set contains all final sales for homes that sell three or more times. Among homes that sell twice, the second sale is added to the test set with probability 1/2. As a result, the test

TABLE 4
Parameter estimates for the AR model

Metropolitan area	$\hat{\mu}$	$\hat{\phi}$	$\hat{\sigma}_\varepsilon^2$	$\hat{\sigma}_\tau^2$
Ann Arbor, MI	11.6643	0.993247	0.001567	0.110454
Atlanta, GA	11.6882	0.992874	0.001651	0.070104
Chicago, IL	11.8226	0.992000	0.001502	0.110683
Columbia, SC	11.3843	0.997526	0.000883	0.028062
Columbus, OH	11.5159	0.994807	0.001264	0.090329
Kansas City, MO	11.4884	0.993734	0.001462	0.121954
Lexington, KY	11.6224	0.996236	0.000968	0.048227
Los Angeles, CA	12.1367	0.981888	0.002174	0.111708
Madison, WI	11.7001	0.994318	0.001120	0.023295
Memphis, TN	11.6572	0.994594	0.001120	0.101298
Minneapolis, MN	11.8327	0.992008	0.001515	0.050961
Orlando, FL	11.6055	0.993561	0.001676	0.046727
Philadelphia, PA	11.7106	0.991767	0.001679	0.183495
Phoenix, AZ	11.7022	0.992349	0.001543	0.106971
Pittsburgh, PA	11.3408	0.992059	0.002546	0.103488
Raleigh, NC	11.7447	0.993828	0.001413	0.047029
San Francisco, CA	12.4236	0.985644	0.001788	0.056201
Seattle, WA	11.9998	0.989923	0.001658	0.039459
Sioux Falls, SD	11.6025	0.995262	0.001120	0.032719
Stamford, CT	12.5345	0.987938	0.002294	0.093230

set for each city contains roughly 15% of the sales. The remaining sales (including single sales) comprise the training set. Table 8 in Appendix A lists the training and test set sizes for each city. We fit the model on the training set and examine the estimated parameters. The test set will be used in Section 5 to validate the AR model against two alternatives.

In Table 4, the estimates for the overall mean μ (on the log scale), the autoregressive parameter ϕ , the variance of the error term σ_ε^2 , and the variance of the random effects σ_τ^2 are provided for each metropolitan area. As expected, the most expensive cities have the highest values of μ : Los Angeles, CA, San Francisco, CA, and Stamford, CT. In Figure 2, the indices for a sample of the twenty cities are provided. There are clearly different trends across cities.

The estimates for the AR model parameter ϕ are close to one. This is not surprising as the adjusted log sale prices, $u_{i,j,z}$, for sale pairs with short gap times are expected to be closer in value than those with longer gap times. It may be tempting to assume that since ϕ is so close to 1, the prices form a random walk instead of an AR(1) time series (see Remark 3.1). However, this is clearly not the case. Recall that ϕ enters the model not by itself but as $\phi^{\gamma(i,j,z)}$ where $\gamma(i,j,z)$ is the gap time. These gap times are high enough that the correlation coefficient $\phi^{\gamma(i,j,z)}$ is considerably lower than 1. The mean gap time across cities is around 22 quarters.

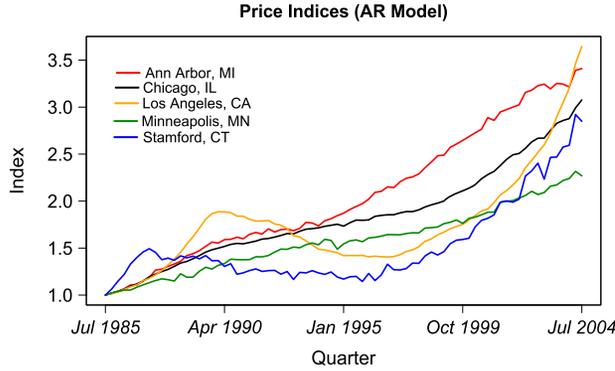


FIG. 2. The AR index for a selection of cities.

As an example, for Ann Arbor, MI, $\hat{\phi}^{22} = 0.993247^{22} \approx 0.8615$ which is clearly less than 1. Therefore, the types of sensitivity often produced as a consequence of near unit roots do not apply to our autoregressive model.

We have modeled the adjusted log prices, $u_{i,j,z} = y_{i,j,z} - \beta_{t(i,j,z)} - \tau_z$, as a latent AR(1) time series. Accordingly, for each gap time, $\gamma(i, j, z) = h$, there is an expected correlation between the sale pairs: ϕ^h . To check that the data support the theory, we compare the correlation between pairs of quarter-adjusted log prices at each gap length to the correlation predicted by the model.

First, we compute the estimated adjusted log prices $\hat{u}_{i,j,z} = y_{i,j,z} - \hat{\beta}_{t(i,j,z)} - \hat{\tau}_z$ for the training data. Next, for each gap time h , we find all the sale pairs $(\hat{u}_{i,j-1,z}, \hat{u}_{i,j,z})$ with that particular gap length. The sample correlation between those sale pairs produces an estimate of ϕ for gap length h . If we repeat this procedure for each possible gap length, we should obtain a steady decrease in the correlation as gap time increases. In particular, the points should follow the curve ϕ^h if the model is specified correctly.

In Figure 3, we plot the correlation of the adjusted log prices by gap time for Columbus, OH. Note that the computed correlations for each gap time were computed with varying quantities of sale pairs. Those computed with fewer than twenty sale pairs are plotted as blue triangles. We also overlay the predicted relationship between ϕ and gap time. The inverse relationship between gap time and correlation seems to hold well and we obtain similar results for most cities. One notable exception is Los Angeles, CA, which we discuss in Section 6.

5. Model validation. To show that the proposed AR model produces good predictions, we fit the model separately to each of the twenty cities and apply the fitted models to each test set. For comparison purposes, a mixed effects model along with the benchmark S&P/Case–Shiller model is applied to the data. The former model is a simple, but reasonable, alternative to the AR model. Both models are described below. In addition to the predictions, we compare the price indices and training set residuals.

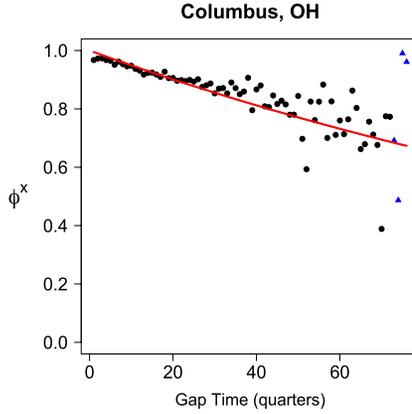


FIG. 3. Checking the AR(1) assumption for Columbus, OH.

The root mean squared error (RMSE)⁴ is used to evaluate predictive performance for each city in Section 5.3. We will see that the AR model provides the best predictions. In addition, we will show the results from Columbus, OH as a typical example.

5.1. *Mixed effects model.* A mixed effects model provides a very simple, but plausible, approach for modeling these data. This model treats the time effect (β_t) as a fixed effect, and the effects of house (α_i) and ZIP code (τ_z) are modeled as random effects. There is no time series component to this model. We describe the model as follows:

$$(11) \quad y_{i,j,z} = \mu + \alpha_i + \tau_z + \beta_{t(i,j,z)} + \varepsilon_{i,j,z},$$

where $\alpha_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_\alpha^2)$, $\tau_z \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_\tau^2)$, and $\varepsilon_{i,j,z} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_\varepsilon^2)$ for houses i from $1, \dots, I_z$, sales j from $1, \dots, J_i$, and ZIP codes z from $1, \dots, Z$. As before, μ is a fixed parameter and $\beta_{i,j,z}$ is the fixed effect for time. The estimates for the parameters $\theta = \{\mu, \beta, \sigma_\varepsilon^2, \sigma_\tau^2\}$ are computed using maximum likelihood estimation.

Finally, estimates for the random effects α and τ are calculated by iteratively calculating the following:

$$(12) \quad \hat{\alpha} = \left(\frac{\sigma_\varepsilon^2}{\sigma_\alpha^2} \mathbf{I}_I + \mathbf{W}'\mathbf{W} \right)^{-1} \mathbf{W}'(\mathbf{y} - \mathbf{X}\hat{\beta} - \mathbf{Z}\hat{\tau}),$$

$$(13) \quad \hat{\tau} = \left(\frac{\sigma_\varepsilon^2}{\sigma_\tau^2} \mathbf{I}_Z + \mathbf{Z}'\mathbf{Z} \right)^{-1} \mathbf{Z}'(\mathbf{y} - \mathbf{X}\hat{\beta} - \mathbf{W}\hat{\alpha}),$$

⁴RMSE = $\sqrt{\frac{1}{n} \sum_{k=1}^n (Y_k - \hat{Y}_k)^2}$, where Y is the sale price and n is the test set size.

where \mathbf{X} and \mathbf{W} are the design matrices for the fixed and random effects respectively and \mathbf{y} is the response vector. These expressions are derived using the method of computing BLUP estimators outlined by Henderson (1975).

To predict the log price, $\hat{y}_{i,j,z}$, we substitute the estimated values:

$$(14) \quad \hat{y}_{i,j,z} = \hat{\mu} + \hat{\beta}_{t(i,j,z)} + \hat{\alpha}_i + \hat{\tau}_z.$$

We use transformation (9) to convert these predictions back to the price scale. Finally, we construct a price index similar to the autoregressive case. Therefore, as in Figure 2, the values of $\exp\{\hat{\beta}_t\}$ are rescaled so that the price index in the first quarter is 1.

5.2. S&P/Case–Shiller model. The original Case and Shiller (1987, 1989) model is a repeat–sales model which expands upon the Bailey, Muth and Nourse (1963) setting by accounting for heteroscedasticity in the data due to the gap time between sales. Borrowing some of their notation, the framework for their model is

$$(15) \quad y_{i,t} = \beta_t + H_{i,t} + u_{i,t},$$

where $y_{i,t}$ is the log price of the sale of the i th house at time t , β_t is the log index at time t , and $u_{i,t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_u^2)$. The middle term, $H_{i,t}$, is a Gaussian random walk which incorporates the previous log sale price of the house. Location information, such as ZIP codes, are not included in this model. Like the Bailey, Muth and Nourse setup, the Case and Shiller setting is a model for differences in prices. Thus, the following model is fit:

$$(16) \quad y_{i,t'} - y_{i,t} = \beta_{t'} - \beta_t + \sum_{k=t+1}^{t'} v_{i,k} + u_{i,t'} - u_{i,t},$$

where $t' > t$. The random walk steps are normally distributed where $v_{i,k} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_v^2)$. Weighted least squares is used to fit the model to account for both sources of variation.

The S&P/Case–Shiller procedure follows in a similar vein but is fit on the price scale instead of the log price scale. The procedure is similar to the arithmetic index proposed by Shiller (1991) which we will describe next; however, full details are available in the S&P/Case–Shiller® Home Price Indices: Index Methodology (2009) report. Let there be S sale pairs, consisting of two consecutive sales of the same house, and T time periods. An $S \times (T - 1)$ design matrix \mathbf{X} , an $S \times (T - 1)$ instrumental variables (IV) matrix \mathbf{Z} , and an $S \times 1$ response vector \mathbf{w} are defined next. Let the subscripts s and t denote the row and column index respectively. Finally, let $Y_{s,t}$ be the sale price (not log price) of the house in sale pair s at time t . Therefore, in each sale pair, there will be two prices $Y_{s,t}$ and $Y_{s,t'}$

where $t \neq t'$. The matrices \mathbf{X} , \mathbf{Z} and vector \mathbf{w} where s indicates the row and t indicates the column are now defined as follows:

$$X_{s,t} = \begin{cases} -Y_{s,t}, & \text{if first sale of pair } s \text{ is at time } t, t > 1, \\ Y_{s,t}, & \text{if second sale of pair } s \text{ is at time } t, \\ 0, & \text{otherwise,} \end{cases}$$

$$Z_{s,t} = \begin{cases} -1, & \text{if first sale of pair } s \text{ is at time } t, t > 1, \\ 1, & \text{if second sale of pair } s \text{ is at time } t, \\ 0, & \text{otherwise,} \end{cases}$$

$$w_s = \begin{cases} Y_{s,t}, & \text{first sale of pair } s \text{ at time } 1, \\ 0, & \text{otherwise.} \end{cases}$$

The goal is to fit the model $\mathbf{w} = \mathbf{X}\mathbf{b} + \boldsymbol{\varepsilon}$ where $\mathbf{b} = (b_1 \dots b_T)'$ is the vector of the reciprocal price indices. That is, $B_t = 1/b_t$ is the price index at time t . A three-step process is implemented to fit this model. First, \mathbf{b} is estimated using regression with instrumental variables. Second, the residuals from this regression are used to compute weights for each observation. Finally, \mathbf{b} is estimated once more while applying the weights. This process, outlined in full in the S&P/Case-Shiller[®] Home Price Indices: Index Methodology report, is described below:

1. Estimate \mathbf{b} by running a regression using instrumental variables: $\hat{\mathbf{b}} = (\mathbf{Z}'\mathbf{X})^{-1} \times \mathbf{Z}'\mathbf{w}$.
2. Calculate the weights for each observation using the squared residuals from the first step. These weights are dependent on the gap time between sales. We denote the residual as $\hat{\varepsilon}_i$ which is an estimate of $u_{i,t'} - u_{i,t} + \sum_{k=1}^{t'-t} v_{i,k}$. The expectation of ε_i is $E[u_{i,t'} - u_{i,t} + \sum_{k=1}^{t'-t} v_{i,k}] = 0$ and the variance is $\text{Var}[u_{i,t'} - u_{i,t} + \sum_{k=1}^{t'-t} v_{i,k}] = 2\sigma_u^2 + (t' - t)\sigma_v^2$. To compute the weights for each observation, the squared residuals from the first step are regressed against the gap time. That is,

$$(17) \quad \hat{\varepsilon}_i^2 = \underbrace{\alpha_0}_{2\sigma_u^2} + \underbrace{\alpha_1}_{\sigma_v^2} (t' - t) + \eta_i,$$

where $E[\eta_i] = 0$. The reciprocal of the square root of the fitted values from the above regression are the weights. Using their notation, we denote this weight matrix by $\boldsymbol{\Omega}^{-1}$.

3. The final step is to estimate \mathbf{b} again while incorporating the weights, $\boldsymbol{\Omega}$: $\hat{\mathbf{b}} = (\mathbf{Z}'\boldsymbol{\Omega}^{-1}\mathbf{X})^{-1}\mathbf{Z}'\boldsymbol{\Omega}^{-1}\mathbf{w}$. The indices are simply the reciprocals of each element in \mathbf{b} for $t > 1$ and, by construction, $B_1 = 1$.

Finally, to estimate the prices in the test set, we simply calculate

$$(18) \quad \hat{Y}_{i,j} = \frac{\hat{B}_{t(i,j-1)}}{\hat{B}_{t(i,j)}} Y_{i,j-1},$$

where $Y_{i,j}$ is the price of the j th sale of the i th house and B_t is the price index at time t . We do not apply the correction proposed by Goetzmann when estimating prices because it is appropriate only for predictions on the log price scale. The S&P/Case–Shiller method is fit on the price scale so no transformation is required.

5.3. Comparing predictions. We fit all three models on the training sets for each city and predict prices for those homes in the corresponding test set. The RMSE for the test set observations is calculated in dollars for each model in order to compare performance across models. These results are listed in Table 5. The model with the lowest RMSE value for each city is shown in italicized font. Note that while the S&P/Case–Shiller method produces predictions directly on the price scale, the autoregressive and mixed effects models must be converted back to the price scale using (9). It is clear that the AR model performs better than the S&P/Case–Shiller model for all of the cities, reducing the RMSE by up to 21% in some cases; the AR model produces lower RMSE values when compared to the mixed effects model as well for nearly all cities, San Francisco, CA, being the only exception. Moreover, the AR model performs better under alternate loss functions as well, which we show in the supplemental article [Nagaraja, Brown and Zhao (2010)].

TABLE 5
Test set RMSE for three models (in dollars)

Metropolitan area	AR (local)	Mixed effects (local)	S&P/C–S
Ann Arbor, MI	<i>41,401</i>	46,519	52,718
Atlanta, GA	<i>30,914</i>	34,912	35,482
Chicago, IL	<i>36,004</i>	—	42,865
Columbia, SC	<i>35,881</i>	38,375	42,301
Columbus, OH	<i>27,353</i>	30,163	30,208
Kansas City, MO	<i>24,179</i>	25,851	—
Lexington, KY	<i>21,132</i>	21,555	21,731
Los Angeles, CA	<i>37,438</i>	—	41,951
Madison, WI	<i>28,035</i>	30,297	30,640
Memphis, TN	<i>24,588</i>	25,502	25,267
Minneapolis, MN	<i>31,900</i>	34,065	34,787
Orlando, FL	<i>28,449</i>	30,438	30,158
Philadelphia, PA	<i>33,246</i>	—	35,350
Phoenix, AZ	<i>28,247</i>	29,286	29,350
Pittsburgh, PA	<i>26,406</i>	28,630	30,135
Raleigh, NC	<i>25,839</i>	27,493	26,775
San Francisco, CA	<i>49,927</i>	<i>48,217</i>	50,249
Seattle, WA	<i>38,469</i>	41,950	43,486
Sioux Falls, SD	<i>20,160</i>	21,171	21,577
Stamford, CT	<i>57,722</i>	58,616	68,132

Note that the RMSE value is missing for Kansas City, MO for the S&P/Case–Shiller model. Some of the observation weights calculated in the second step of the procedure were negative, halting the estimation process. This is another drawback to some of the existing repeat sales procedures. Calhoun (1996) suggests replacing the sale specific error $u_{i,t}$ [as given in (16)] with a house specific error u_i ; however, this fundamentally changes the structure of the error term and, as a result, the fitting process. Furthermore, it is not implemented in the S&P/Case–Shiller methodology. Therefore, we do not apply it to our data.

Three values are also missing in Table 5 for the mixed effect model results. For these three cities, the iterative fitting procedure failed to converge. We can attribute this to the size of these data and, more importantly, that the data do not conform well to the mixed effects model structure.

Next, we will examine several diagnostic plots to assess whether the model assumptions are satisfied for each method. We begin by investigating the variance of the residuals. As the gap time increases, we expect a higher error variance indicating that the previous price becomes less useful over time. The proposed autoregressive model and the S&P/Case–Shiller model each incorporate this feature differently, using an underlying AR(1) time series and a random walk respectively. The mixed effects model, however, assumes a constant variance regardless of gap time. In Figure 4, for each model, we plot the variance of the predictions by gap time for the training set residuals.⁵ The expected variance by gap time values using the estimated parameters is then overlaid. The autoregressive and mixed effects models are fit on the log price scale, whereas the S&P/Case–Shiller model is fit on the price scale. Therefore, the residual plots are graphed on very different scales.

There are two features to note here. The first is that heteroscedasticity is clearly present: the variance of the residuals does in fact increase with gap time. The second feature is that while none of the methods perfectly model the heteroscedastic error, the mixed effects model is undoubtedly the worst. This pattern holds across all of the cities in the data set. Both the autoregressive and S&P/Case–Shiller models seem to have lower than expected variances in Figure 4.

For both the AR and mixed effects models, the random effects for ZIP codes are assumed to be normally distributed. As a diagnostic procedure, we construct the normal quantile plots of the ZIP code effects. The results are shown in Figure 5. Columbus, OH has a total of 103 ZIP codes, or random effects. We find the normality assumption appears to be reasonably satisfied for the mixed effects model but less so for the autoregressive model. Note, however, that each random effect is

⁵Note that for these three plots, the term “residual” indicates the usual statistical residual values produced by applying the model and comparing the predictions with the response vector. For the AR and mixed effects models, these residuals are identical to the predictions on the log price scale discussed in previous sections; however, for the S&P/C–S model, this is not the case.

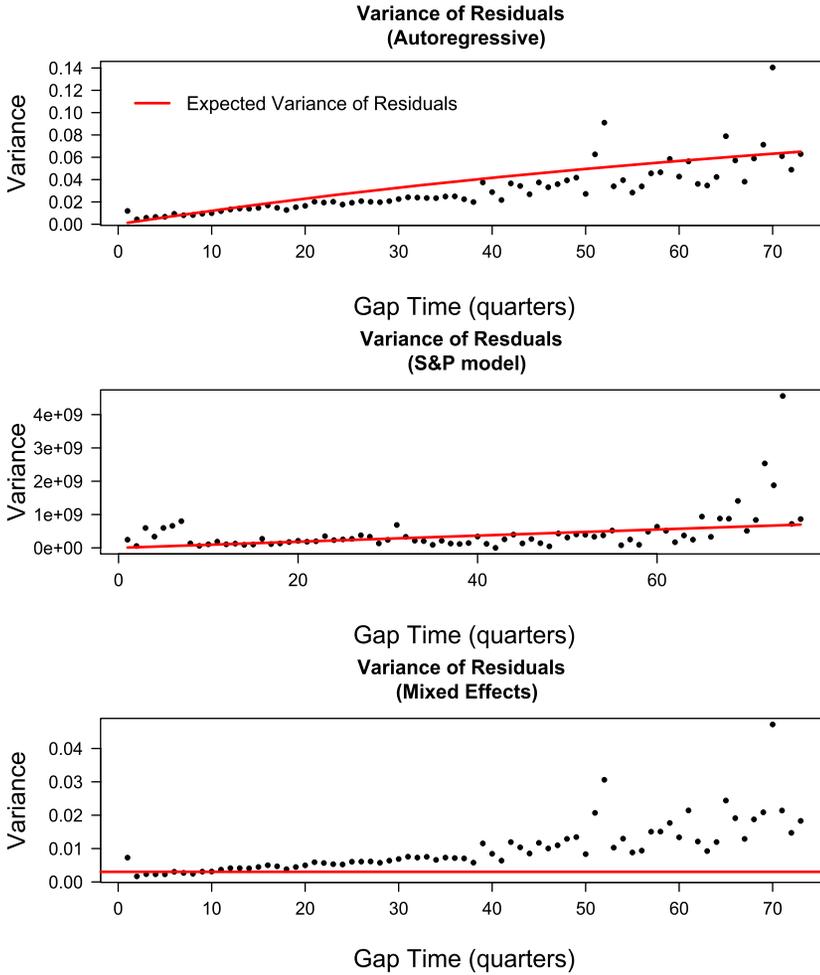


FIG. 4. Comparing the variance of the residuals for Columbus, OH.

estimated using a different number of sales. This interferes with the routine interpretation of these plots. In particular, the outliers in both plots correspond to ZIP codes containing ten or fewer sales. Across all metropolitan areas, the normality assumption seems to be well satisfied in some cases and not so well in others, but with no clear pattern we could discern as to the type of analysis, size of the data or geographic region. The supplemental article contains results of the Shapiro–Wilk test for normality [Nagaraja, Brown and Zhao (2010)].

In Figure 6, we plot four indices for Columbus, OH: the AR index, the mixed effects index, the S&P/Case–Shiller index, and the mean price index. The mean index is simply the average sale price at each quarter rescaled so that the first index value is 1. From the plot, we see that the autoregressive index is generally between the S&P/Case–Shiller index and the mean index at each point in time. The

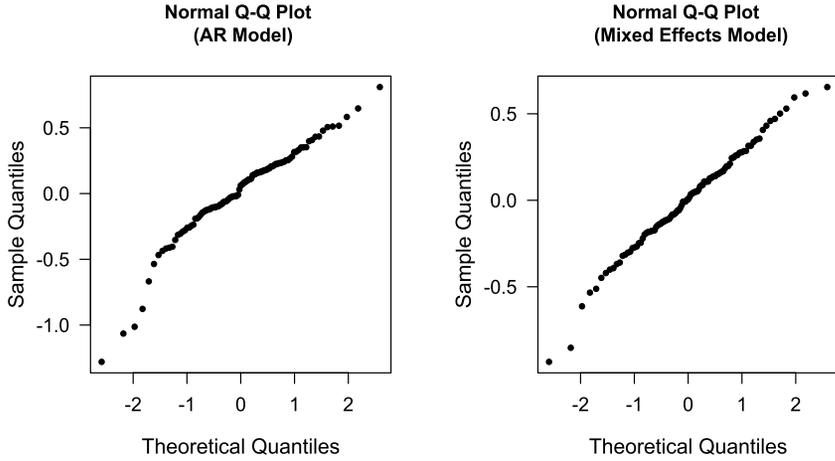


FIG. 5. Normality of ZIP code effects for Columbus, OH.

mean index treats all sales as single sales. That is, information about repeat sales is not included; in fact, no information about house prices is shared across quarters. The S&P/Case–Shiller index, on the other hand, only includes repeat sales houses. The autoregressive model, because it includes both single sales and repeat sales, is a mixture of the two perspectives. Essentially, the index constructed from the proposed model is a measure of the average house price placing more weight to those homes which have sold more than once.

6. The case of Los Angeles, CA. Even though the autoregressive model has a lower RMSE than the S&P/Case–Shiller model for Los Angeles, CA, it does not

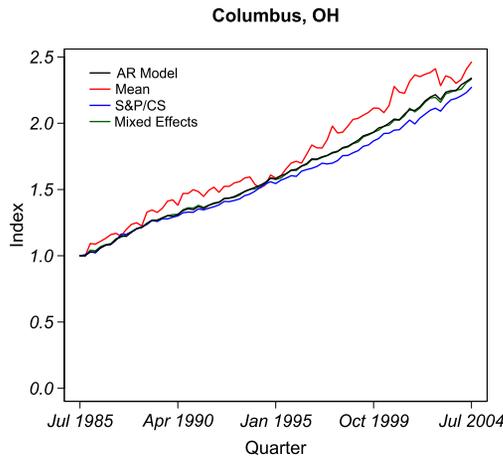


FIG. 6. House price indices for Columbus, OH.

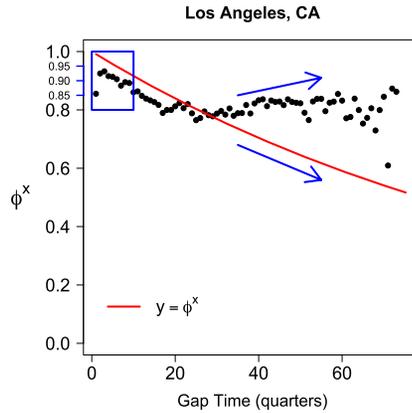


FIG. 7. *Problems with the assumptions.*

seem to fit the data well. If we examine Figure 7, a plot of the correlation against gap time, we immediately see two significant issues when what is expected (line) is compared with what the data indicate (dots). First, the value of ϕ is not as close to 1 as expected. Second, the pattern of decay, $\phi^{\gamma(i,j,z)}$, also does not follow the presumed pattern. We will focus on Los Angeles, CA, and discuss these two issues for the remainder of this section.

We expect ϕ to be close to 1; however, for Los Angeles, CA, this does not seem to be the case. In fact, according to the data, for short gap times, the correlation between sale pairs seems to be far lower than one. To investigate this feature, we examine sale pairs with gap times between 1 and 5 quarters more closely. In Figure 8, we construct a histogram of the quarters where the second sale occurred for this subset of sale pairs. We pair this histogram with a plot of the price index for Los Angeles, CA. Most of these sales occurred during the late 1980s and early 1990s. This corresponds to the same period when Sing and Furlong (1989) found that lenders were offering people mortgages where the monthly payment was *greater* than 33% of their monthly income. The threshold of 33% is set to help ensure that people will be able to afford their mortgage. Those persons with mortgages that exceed this percentage tend to have a higher probability of defaulting on their payments.

Bates (1989) found that a number of banks including the Bank of California and Wells Fargo were highly exposed to these risky investments, especially in the wake of the housing downturn during the early 1990s. If a short gap time is an indication that a foreclosure took place, this would explain why these sale pair prices are not highly correlated. We did observe, however, that other cities also experienced periods of decline, such as Stamford, CT (see Figure 2), but did not have anomalous autoregressive patterns like those in Figure 7 for Los Angeles, CA.

Even if this were not the case, the autoregressive model may not be performing well simply because there was a downturn in the housing market. Most of the

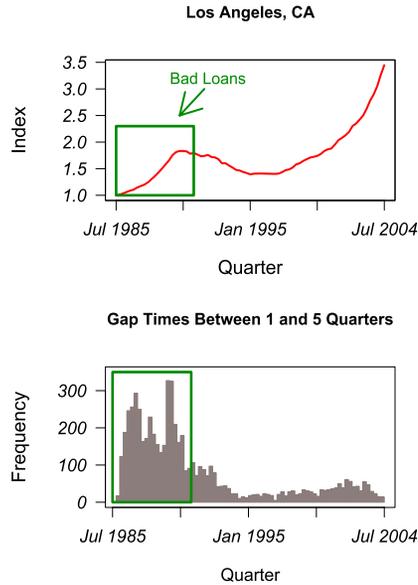


FIG. 8. Examining the housing downturn.

cities in our data cover periods where the indices are increasing—the model may be performing well only because of this feature. In the case of Los Angeles, CA, if we examine the period between January 1990 and December 1996 on Figure 8, the housing index was decreasing. However, if we calculate the RMSE of test set sales for this period only, we find that the autoregressive model still performs better than the S&P/Case–Shiller method. The RMSE values are \$32,039 and \$41,841, respectively. Therefore, the autoregressive model seems to perform better in a period of decline as well as in times of increase.

The second irregularity evident in Figure 7 is that the AR(1) process does not decay at the same rate as the model predicts. In 1978 California voters, as a protest against rising property taxes, passed Proposition 13 which limited how fast property tax assessments could increase per year. Galles and Sexton (1998) argue that Proposition 13 encouraged people to retain homes especially if they have owned their home for a long time. It is possible that this feature of Figure 7 is a long term effect of Proposition 13. On the other hand, it could be that California home owners tend to renovate their homes more frequently than others, reducing the decay in prices over time. However, we have no way of verifying either of these explanations given our data.

7. Discussion. Two key tasks when analyzing house prices are predicting sale prices of individual homes and constructing price indices which measure general housing trends. Using extensive data from twenty metropolitan areas, we have

compared our predictive method to two other methods, including the S&P/Case-Shiller Home Price Index. We find that on average the predictions using our method are more accurate in all but one of the twenty metropolitan areas examined.

Data such as ours often do not contain reliable hedonic information on individual homes, if at all. Therefore, harnessing the information contained in a previous sale is critical. Repeat sales indices attempt to do exactly that. Some methods have also incorporated ad hoc adjustments to take account of the gap time between the repeat sales of a home. In contrast, our model involves an underlying AR(1) time series which automatically adjusts for the time gap between sales. It also uses the home's ZIP code as an additional indicator of its hedonic value. This indicator has some predictive value, although its value is quite weak by comparison with the price in a previous sale if one has been recorded.

The index constructed from our statistical model can be viewed as a weighted average of estimates from single and repeat sales homes, with the repeat sales prices having a substantially higher weight. As noted, the time series feature of the model guarantees that this weight for repeat sales prices slowly decreases in a natural fashion as the gap time between sales increases.

Our results do not provide definitive evidence as to the value of our index when comparing with other currently available indices as a general economic indicator. Indeed, such a determination should involve a study of the economic uses of such indicators as well as an examination of their formulaic construction and their use for prediction of individual sale prices. We have not undertaken such a study, and so can offer only a few comments about the possible comparative values of our index.

As we have discussed, we feel it may be an advantage that our index involves all home sales in the data (subject to the naturally occurring weighting described above), rather than only repeat sales. Repeat sales homes are only a small, selected fraction of all home sales. Studies have shown that repeat sales homes may have different characteristics than single sale homes. In particular, they are evidently older on average, and this could be expected to have an effect on their sale price. Since our measure brings all home sales into consideration, albeit in a gently weighted manner, and since it provides improved prediction on average, it may produce a preferable index.

Another advantage of our model is that it remains easy to interpret at both the micro and macro levels, in spite of including several features inherent in the data. Future work seems desirable to understand anomalous features such as those we have discussed in the Los Angeles, CA, area. Such research may allow us to construct a more flexible model to accommodate such cases. For example, it could involve the inclusion of economic indicators which may affect house prices such as interest rates and tax rates and measures of general economic status such as the unemployment rate.

APPENDIX A: DATA SUMMARY

TABLE 6
Summary counts

City	No. sales	No. houses	No. houses per sale count			
			1	2	3	4+
Ann Arbor, MI	68,684	48,522	32,458	12,662	2,781	621
Atlanta, GA	376,082	260,703	166,646	76,046	15,163	2836
Chicago, IL	688,468	483,581	319,340	130,234	28,369	5,603
Columbia, SC	7,034	4,321	2,303	1,470	431	117
Columbus, OH	162,716	109,388	67,926	31,739	7,892	1,831
Kansas City, MO	123,441	90,504	62,489	23,706	3,773	534
Lexington, KY	38,534	26,630	16,891	7,901	1,555	282
Los Angeles, CA	543,071	395,061	272,258	100,918	18,965	2,903
Madison, WI	50,589	35,635	23,685	9,439	2,086	425
Memphis, TN	55,370	37,352	23,033	11,319	2,412	587
Minneapolis, MN	330,162	240,270	166,811	59,468	11,856	2,127
Orlando, FL	104,853	72,976	45,966	22,759	3,706	543
Philadelphia, PA	402,935	280,272	179,107	82,681	15,878	2,606
Phoenix, AZ	180,745	129,993	87,249	35,910	5,855	968
Pittsburgh, PA	104,544	73,871	48,618	20,768	3,749	718
Raleigh, NC	100,180	68,306	42,545	20,632	4,306	818
San Francisco, CA	73,598	59,416	46,959	10,895	1,413	149
Seattle, WA	253,227	182,770	124,672	47,406	9,198	1,494
Sioux Falls, SD	12,439	8,974	6,117	2,353	419	85
Stamford, CT	14,602	11,128	8,200	2,502	357	62

TABLE 7
Number of ZIP codes by city

City	No. ZIP codes
Ann Arbor, MI	57
Atlanta, GA	184
Chicago, IL	317
Columbia, SC	12
Columbus, OH	103
Kansas City, MO	179
Lexington, KY	31
Los Angeles, CA	280
Madison, WI	40
Memphis, TN	64
Minneapolis, MN	214
Orlando, FL	96
Philadelphia, PA	330
Phoenix, AZ	130

TABLE 7
(Continued.)

City	No. ZIP codes
Pittsburgh, PA	257
Raleigh, NC	82
San Francisco, CA	70
Seattle, WA	110
Sioux Falls, SD	30
Stamford, CT	23

TABLE 8
Training and test set sizes

City	Autoregressive model			S&P/Case–Shiller model	
	Training	Test	No. houses	Training pairs	No. houses
Ann Arbor, MI	58,953	9,731	48,522	10,431	9,735
Atlanta, GA	319,925	56,127	260,703	59,222	55,911
Chicago, IL	589,289	99,179	483,581	105,708	99,069
Columbia, SC	5,747	1,287	4,321	1,426	1,279
Columbus, OH	136,989	25,727	109,388	27,601	25,458
Kansas City, MO	107,209	16,232	90,504	16,705	16,092
Lexington, KY	32,705	5,829	26,630	6,075	5,748
Los Angeles, CA	470,721	72,350	395,061	75,660	72,338
Madison, WI	43,349	7,240	35,635	7,714	7,221
Memphis, TN	46,724	8,646	37,352	9,372	8,673
Minneapolis, MN	286,476	43,686	240,270	46,206	43,764
Orlando, FL	89,123	15,730	72,976	16,147	15,531
Philadelphia, PA	343,354	59,581	280,272	63,082	60,068
Phoenix, AZ	155,823	24,922	129,993	25,830	24,656
Pittsburgh, PA	89,762	14,782	73,871	15,891	14,956
Raleigh, NC	84,678	15,502	68,306	16,372	15,388
San Francisco, CA	66,527	7,071	59,416	7,111	6,948
Seattle, WA	218,741	34,486	182,770	35,971	34,304
Sioux Falls, SD	10,755	1,684	8,974	1,781	1,677
Stamford, CT	12,902	1,700	11,128	1,774	1,654

APPENDIX B: UPDATING EQUATIONS

In this section we provide the updating equations for estimating the parameters $\theta = \{\beta, \sigma_\varepsilon^2, \sigma_\tau^2, \phi\}$ in the autoregressive model (see Section 3). Observe that the covariance matrix \mathbf{V} is an $N \times N$ matrix where N is the sample size. Given the size of our data, it is simpler computationally to exploit the block diagonal structure of \mathbf{V} . Each block, denoted by $\mathbf{V}_{z,z}$, corresponds to observations in ZIP code z .

Computations are carried out on the ZIP code level and the updating equations provided below reflect this. For instance, \mathbf{y}_z and \mathbf{T}_z are the elements of the log price vector and transformation matrix respectively for observations in ZIP code z .

To start, an explicit expression for $\boldsymbol{\beta}$ can be formulated:

$$(19) \quad \hat{\boldsymbol{\beta}} = \left(\sum_{z=1}^Z (\mathbf{T}_z \mathbf{X}_z)' \mathbf{V}_{z,z}^{-1} \mathbf{T}_z \mathbf{X}_z \right)^{-1} \sum_{z=1}^Z (\mathbf{T}_z \mathbf{X}_z)' \mathbf{V}_{z,z}^{-1} \mathbf{T}_z \mathbf{y}_z.$$

Estimates must be computed numerically for the remaining parameters. As all of these are one-dimensional parameters, methods such as the Newton–Raphson algorithm are highly suitable. We first define $\mathbf{w}_z = \mathbf{y}_z - \mathbf{X}_z \boldsymbol{\beta}$ for clarity. To update σ_ε^2 , compute the zero of

$$(20) \quad 0 = - \sum_{z=1}^Z \text{tr}(\mathbf{V}_{z,z}^{-1} \text{diag}(\mathbf{r}_z)) + \sum_{z=1}^Z (\mathbf{T}_z \mathbf{w}_z)' \mathbf{V}_{z,z}^{-1} \text{diag}(\mathbf{r}_z) \mathbf{V}_{z,z}^{-1} (\mathbf{T}_z \mathbf{w}_z),$$

where $\text{tr}(\cdot)$ is the trace of a matrix and $\text{diag}(\mathbf{r})$ is as defined in (4). Similarly, to update σ_τ^2 , find the zero of

$$(21) \quad 0 = \sum_{z=1}^Z \text{tr}(\mathbf{V}_{z,z}^{-1} (\mathbf{T}_z \mathbf{1}_{n_z}) (\mathbf{T}_z \mathbf{1}_{n_z})') + \sum_{z=1}^Z (\mathbf{T}_z \mathbf{w}_z)' \mathbf{V}_{z,z}^{-1} (\mathbf{T}_z \mathbf{1}_{n_z}) (\mathbf{T}_z \mathbf{1}_{n_z})' \mathbf{V}_{z,z}^{-1} (\mathbf{T}_z \mathbf{w}_z),$$

where n_z denotes the number of observations in ZIP code z and $\mathbf{1}_k$ is a $(k \times 1)$ vector of ones.

Finally, to update the autoregressive parameter ϕ , we must calculate the zero of the function below:

$$(22) \quad 0 = - \sum_{z=1}^Z \text{tr} \left\{ \mathbf{V}_{z,z}^{-1} \left(\sigma_\tau^2 \left(\frac{\partial (\mathbf{T}_z \mathbf{1}_{n_z})}{\partial \phi} \right) (\mathbf{T}_z \mathbf{1}_{n_z})' + \sigma_\tau^2 (\mathbf{T}_z \mathbf{1}_{n_z}) \left(\frac{\partial (\mathbf{T}_z \mathbf{1}_{n_z})}{\partial \phi} \right)' + \frac{2\phi\sigma_\varepsilon^2}{(1-\phi^2)^2} \text{diag}(\mathbf{r}_z) + \frac{\sigma_\varepsilon^2}{1-\phi^2} \frac{\partial \text{diag}(\mathbf{r}_z)}{\partial \phi} \right) \right\} - \sum_{z=1}^Z \left(\frac{\partial \mathbf{T}_z}{\partial \phi} \mathbf{w}_z \right)' \mathbf{V}_{z,z}^{-1} (\mathbf{T}_z \mathbf{w}_z) - \sum_{z=1}^Z (\mathbf{T}_z \mathbf{w}_z)' \mathbf{V}_{z,z}^{-1} \left(\frac{\partial \mathbf{T}_z}{\partial \phi} \mathbf{w}_z \right) + \sum_{z=1}^Z \left[(\mathbf{T}_z \mathbf{w}_z)' \mathbf{V}_{z,z}^{-1} \left[\sigma_\tau^2 \left(\frac{\partial (\mathbf{T}_z \mathbf{1}_{n_z})}{\partial \phi} \right) (\mathbf{T}_z \mathbf{1}_{n_z})' \right. \right.$$

$$\begin{aligned}
& + \sigma_{\tau}^2 (\mathbf{T}_z \mathbf{1}_{n_z}) \left(\frac{\partial (\mathbf{T}_z \mathbf{1}_{n_z})}{\partial \phi} \right)' + \frac{2\phi\sigma_{\varepsilon}^2}{(1-\phi^2)^2} \text{diag}(\mathbf{r}_z) \\
& \quad + \frac{\sigma_{\varepsilon}^2}{1-\phi^2} \frac{\partial \text{diag}(\mathbf{r}_z)}{\partial \phi} \Big] \mathbf{V}_{z,z}^{-1} (\mathbf{T}_z \mathbf{w}_z) \Big].
\end{aligned}$$

After the estimates converge, we must estimate the random effects. We use Henderson's procedure to derive the Best Linear Unbiased Predictors (BLUP) for each ZIP code. His method assumes that the parameters in the covariance matrix, \mathbf{V} , are known; however, we use the estimated values. The formula is

$$\begin{aligned}
(23) \quad \hat{\tau}_z &= \left[\frac{2\hat{\sigma}_{\varepsilon}^2}{\hat{\sigma}_{\tau}^2} + (1 - \hat{\phi}^2) (\hat{\mathbf{T}}_z \mathbf{1}_z)' \text{diag}^{-1}(\mathbf{r}_z) (\hat{\mathbf{T}}_z \mathbf{1}_z) \right]^{-1} \\
& \quad \times ((1 - \hat{\phi}^2) (\hat{\mathbf{T}}_z \mathbf{1}_z)' \text{diag}^{-1}(\mathbf{r}_z) (\hat{\mathbf{T}}_z \hat{\mathbf{w}}_z)),
\end{aligned}$$

where $\text{diag}^{-1}(\hat{\mathbf{r}})$ is the inverse of the estimated diagonal matrix $\text{diag}(\mathbf{r})$.

Acknowledgments. The authors would like to thank the referees for their thorough and helpful comments.

SUPPLEMENTARY MATERIAL

Supplement to “An autoregressive approach to house price modeling” (DOI: [10.1214/10-AOAS380SUPP](https://doi.org/10.1214/10-AOAS380SUPP); .pdf). This supplement contains extra analysis on a variety of topics related to the paper from examining the convergence of the coordinate ascent algorithm, or applying alternate loss functions, to studying the impact of each feature included in the autoregressive (AR) model.

REFERENCES

- BAILEY, M. J., MUTH, R. F. and NOURSE, H. O. (1963). A regression method for real estate price index construction. *J. Amer. Statist. Assoc.* **58** 933–942.
- BATES, J. (1989). Survey cites four California banks with possibly risky realty loans. *Los Angeles Times* **1**, December 30.
- BICKEL, P. J. and DOKSUM, K. A. (2001). *Mathematical Statistics—Basic Ideas and Selected Topics, Vol. I*, 2nd ed. Prentice Hall, Englewood Cliffs, NJ.
- BROWN, L. D. (1986). *Fundamentals of Statistical Exponential Families with Applications in Statistical Decision Theory. IMS Lecture Notes Monogr. Ser. 9*. IMS, Hayward, CA. MR0882001
- CALHOUN, C. (1996). OFHEO house price indices: HPI technical description. Available at: <http://www.ofheo.gov>.
- CASE, B., POLLAKOWSKI, H. O. and WACHTER, S. M. (1991). On choosing among house price index methodologies. *AREUEA J.* **19** 286–307.
- CASE, B. and QUIGLEY, J. M. (1991). The dynamics of real estate prices. *Rev. Econ. Statist.* **73** 50–58.
- CASE, K. E. and SHILLER, R. J. (1987). Prices of single-family homes since 1970: New indexes for four cities. *N. Engl. Econ. Rev.* **Sept./Oct.** 45–56.
- CASE, K. E. and SHILLER, R. J. (1989). The efficiency of the market for single family homes. *Amer. Econ. Rev.* **79** 125–137.

- ENGLUND, P., QUIGLEY, J. M. and REDFEARN, C. L. (1999). The choice of methodology for computing housing price indexes: Comparisons of temporal aggregation and sample definition. *J. Real Estate Fin. Econ.* **19** 91–112.
- GALLES, G. M. and SEXTON, R. L. (1998). A tale of two tax jurisdictions: The surprising effects of California's Proposition 13 and Massachusetts' Proposition 2 1/2. *Amer. J. Econ. Soc.* **57** 123–133.
- GATZLAFF, D. H. and HAURIN, D. R. (1997). Sample selection bias and repeat-sales index estimates. *J. Real Estate Fin. Econ.* **14** 33–50.
- GOETZMANN, W. N. (1992). The accuracy of real estate indices: Repeat sales estimators. *J. Real Estate Fin. Econ.* **5** 5–53.
- GOETZMANN, W. N. and PENG, L. (2002). The bias of the RSR estimator and the accuracy of some alternatives. *Real Estate Econ.* **30** 13–39.
- GOETZMANN, W. N. and SPIEGEL, M. (1995). Non-temporal components of residential real estate appreciation. *Rev. Econ. Statist.* **77** 199–206.
- HENDERSON, C. R. (1975). Best linear unbiased estimation and prediction under a selection model. *Biometrics* **31** 423–447.
- MEESE, R. A. and WALLACE, N. E. (1997). The construction of residential housing price indices: A comparison of repeat-sales, hedonic-regression, and hybrid approaches. *J. Real Estate Fin. Econ.* **14** 51–73.
- NAGARAJA, C. H., BROWN, L. B. and ZHAO, L. H. (2010). Supplement to “An autoregressive approach to house price modeling.” DOI: [10.1214/10-AOAS380SUPP](https://doi.org/10.1214/10-AOAS380SUPP).
- PALMQUIST, R. B. (1982). Measuring environmental effects on property values without hedonic regression. *J. Urban Econ.* **11** 333–347.
- SHEN, H., BROWN, L. D. and ZHI, H. (2006). Efficient estimation of log-normal means with application to pharmacokinetic data. *Statist. Med.* **25** 3023–3038. MR2247080
- SHILLER, R. (1991). Arithmetic repeat sales price estimators. *J. Housing Econ.* **1** 110–126.
- SING, B. and FURLONG, T. (1989). Defaults feared if payments keep ballooning fast-rising interest rates causing concern for adjustable mortgages. *Los Angeles Times* **5**, March 29.
- S&P/CASE–SHILLER[®] HOME PRICE INDICES: INDEX METHODOLOGY (November 2009). Available at <http://www.standardandpoors.com>.

C. H. NAGARAJA
US CENSUS BUREAU
CENTER FOR STATISTICAL RESEARCH
AND METHODOLOGY
4600 SILVER HILL RD.
WASHINGTON, DC 20233
USA
E-MAIL: chaitra.nagaraja@census.gov

L. D. BROWN
L. H. ZHAO
STATISTICS DEPARTMENT
THE WHARTON SCHOOL
UNIVERSITY OF PENNSYLVANIA
400 JON M. HUNTSMAN HALL
3730 WALNUT ST.
PHILADELPHIA, PENNSYLVANIA 19104-6302
USA
E-MAIL: lbrown@wharton.upenn.edu
lzhao@wharton.upenn.edu