Daily House Price Indexes: Construction, Modeling, and Longer-Run Predictions*

Tim Bollerslev[†], Andrew Patton[‡], and Wenjing Wang[§]

First Draft: August 2012

This Version: June 20, 2013

Abstract

We construct *daily* house price indexes for ten major U.S. metropolitan areas. Our calculations are based on a comprehensive database of several million residential property transactions and a standard repeat-sales method that closely mimics the procedure used in the construction of the popular monthly Case-Shiller house price indexes. Our new daily house price indexes exhibit similar characteristics to other daily asset prices, with mild autocorrelation and strong conditional heteroskedasticity, which are well described by a relatively simple multivariate GARCH type model. The sample and model-implied correlations across house price index returns are low at the daily frequency, but rise monotonically with the return horizon, and are all commensurate with existing empirical evidence for the existing monthly and quarterly house price series. A simple model of daily house price index returns produces forecasts of monthly house price changes that are superior to various alternative forecast procedures based on lower frequency data, underscoring the informational advantages of our new more finely sampled daily price series.

*We thank Pat Bayer, Jia Li, George Tauchen and seminar participants at the Duke Financial Econometrics Lunch Group for helpful comments and suggestions.

[†]Department of Economics, Duke University, Durham, NC 27708, and NBER and CREATES, boller@duke.edu, 919-660-1846.

[‡]Department of Economics, Duke University, Durham, NC 27708, andrew.patton@duke.edu.

[§]Department of Economics, Duke University, Durham, NC 27708, wenjing.wang@duke.edu.

"There are many ways to measure changes in house prices, but the Standard & Poor's/Case-Shiller index has become many economists' favored benchmark in recent years."

Wall Street Journal, September 25, 2012

1 Introduction

For many U.S. households their primary residence represents their single largest financial asset holding.[1] Consequently, changes in housing valuations importantly affect households' saving and spending decisions, and in turn the overall growth of the economy. Indeed, a number of studies (see, e.g., Campbell and Cocco, 2007; Case et al., 2005, 2011; Davis and Heathcote, 2005, among others) have argued that the wealth effect of the housing market for aggregate consumption is significantly larger than that of the stock market. The recent economic crisis, which arguably originated with the precipitous drop in housing prices beginning in 2006, directly underscores this point.

Meanwhile, compared to most other financial asset classes data on residential real estate valuations, especially at the aggregate level, are notoriously poor, and only available at relatively low monthly or quarterly frequencies. Set against this background, we provide a new set of *daily* house price indexes for ten major U.S. metropolitan areas. To the best of our knowledge, this represents the first set of house price indexes at the daily frequency. Our construction is based on a comprehensive database consisting of all publicly recorded residential property transactions. We show that the dynamic dependencies in the new daily housing price series closely mimic those of other aggregate asset price indexes, and that these dynamic dependencies along with the cross-city correlations are well described by a standard multivariate GARCH type model. This relatively simple daily model in turn allows for the construction of improved longer-run monthly and quarterly housing

[1] At the aggregate level, estimates based on the Federal Reserve Flow of Fund Accounts put the total value of the U.S. residential real estate market at $16 trillion at the end of 2011, compared with $18 trillion for the U.S. stock market as estimated by the Center for Research in Security Prices.

price forecasts compared to the forecasts based only on existing monthly and/or quarterly indexes.

Our house price indexes are based on the same "repeat-sales" methodology as the popular S&P/Case-Shiller indexes, currently published at a monthly frequency (see Shiller, 1991, for further documentation).[2] This low frequency of reporting ignores the potential information available in the daily records of housing transactions, and is likely to induce "aggregation biases" if the true index changes at a higher frequencies than the measurement period (see, e.g., the discussion in Calhoun et al., 1995). Further along these lines, aggregating the indexes to lower frequencies also reduces their volatility, thereby underestimating the true risk of the housing market.

Our construction of "high frequency" house price indexes is related to recent work on real-time estimates of the macro economy (Evans, 2005; Aruoba et al., 2009, among others). Like these authors, we are motivated by the potential usefulness of more timely information about the state of the economy and economic activity. More timely house prices, in particular, are also of direct interest to policy makers, central banks, developers and lenders, as well as, of course, potential buyers and sellers.

In parallel to the daily returns on most other broadly defined asset classes, we find that our new daily house price indexes exhibit only mild predictability in the mean, but strong evidence of volatility clustering. We show that the volatility clustering within and across the different house price indexes can be satisfactorily described by a relatively simple multivariate GARCH model. The correlation between the daily returns on the city indexes is much lower than the correlation observed for the existing monthly return indexes. However, as we temporally aggregate the daily returns to a monthly or quarterly frequency, we find that the correlations increase to levels consistent with the ones observed for existing lower frequency indexes. Going one step further, we also document that the new daily indexes result in improved forecasts over longer monthly horizons for both the

[2] A similar approach is also used in the construction of the quarterly Office of Federal Housing Enterprise Oversight house price indexes.

composite and city-specific housing returns, thus directly underscoring the informational advantages over the existing monthly published indexes.

The rest of the paper is organized as follows. The next section provides a review of house price index construction, and introduces the S&P/Case-Shiller methodology that we employ in our analysis. Section 3 describes the data and the practical construction of our new daily prices series. Section 4 briefly summarizes the dynamic and cross-sectional dependencies in the daily series, and presents our simple multivariate GARCH-type model designed to account for these dependencies. We also show how the low cross-city correlations estimated by the daily model are consistent with the much stronger longer-run monthly and quarterly correlations observed across cities in the existing indexes. Section 5 demonstrates how the new daily series and our modeling thereof may be used in more accurately forecasting the corresponding longer-run returns. Section 6 concludes.

2 Construction of house price indexes

The construction of house price indexes is plagued by two major difficulties. Firstly, houses are heterogeneous assets; each house is a unique asset, in terms of its location, characteristics, maintenance status, etc., which will affect its price. House price indexes aim to measure the price movements of a hypothetical house of average quality, with the assumption that average quality remains the same across time. In reality, average quality has been increasing over time, because newly-built houses tend to be of higher quality and more in line with current households' requirements than older houses. Detailed house qualities are not always available or not directly observable, so when measuring average house prices, it is difficult to take the changing average qualities of houses into consideration.

The second major difficulty is sale infrequency.[3] The price of a house is not observed until an actual transaction occurs. Related to that, the houses sold at each point in time

[3] For example, the average time interval between two successive transactions of the same property is about six years in Los Angeles, based on our data set described in detail in Section 3 below.

might not be a representative sample of the overall housing stock. These difficulties may in turn result in various biases.

2.1 Index methodologies

Three main methodologies have been used to overcome the above-mentioned difficulties in the constructing reliable house price indexes (see, e.g., the surveys by Rappaport, 2007; Ghysels et al., 2013). One approach is to rely on the median value of all transaction prices in a given period. The National Association of Realtors employ this methodology and publishes median prices of existing home sales monthly for both the national and four Census regions. The median price index has the obvious advantage of calculation simplicity, but it does not control at all for heterogeneity of the houses actually sold.

A second, more complicated, approach uses a hedonic technique, which prices the "average quality" house by explicitly pricing its specific attributes. This method controls the heterogeneity of houses. The U.S. Census Bureau constructs its Constant Quality (Laspeyres) Price Index of New One-Family Houses Sold using the hedonic method and publishes its index quarterly. However this method requires much richer data than are typically available.

A third approach is based on a repeat sales model. This is the method used by Standard & Poor's and the Office of Federal Housing Enterprise Oversight (OFHEC).[4] The repeat sales model was originally introduced by Bailey, Muth, and Nourse (1963), and subsequently modified by Case and Shiller (1989). The specific model currently used to construct the S&P/Case-Shiller indexes was proposed by Shiller (1991).[5]

As their name suggests, repeat sales models estimate house price changes by looking

[4]As Meese and Wallace (1997) point out, repeat-sales models can be viewed as a special case of hedonic models, assuming that the attributes, and the shadow prices of the attributes, of the houses do not change between sales. Thus, if the additional house characteristic data were widely available, it would clearly be preferable to use a hedonic pricing model.

[5]Chauvet et al. (2013) propose an interesting, but quite different, type of housing index based on internet search data for queries related to investor distress.

at repeated transactions of the same house.[6] This provides some control for the heterogeneity in the characteristics of houses, while only requiring data on transaction prices and dates. The basic models, however, are subject to some strong assumptions (see, e.g., the discussion in Cho, 1996; Rappaport, 2007). Firstly, it is assumed that the quality of a given house remains unchanged over time. In practice, of course, the quality of most houses changes through aging, maintenance or reconstruction. This in turn causes a so-called "renovation bias." Secondly, repeat sales indexes exploit information only from houses that have been sold at least twice during the sampling period. This subset of all houses is assumed to be representative of the entire housing stock, possibly resulting in a "sample-selection bias." In addition, as previously noted, these indexes are also subject to a possible "aggregation bias" if the true average house price fluctuates within the estimation window.[7]

2.2 The S&P/Case-Shiller methodology

Our new daily home price indexes are designed to mimic the popular S&P/Case-Shiller house price indexes for the "typical" prices of single-family residential real estate. They are based on a repeat sales method and the transaction dates and prices for all houses that sold at least twice during the sample period.[8] Specifically, for a house j that sold at times

[6] Cho (1996) provides a good survey on various repeat sales models, and Ghysels et al. (2013) discuss house price index construction with a view to forecasting. Extensions to the basic repeat-sales model include Shiller (1993), Case and Quigley (1991), Quigley (1995), Goetzmann and Spiegel (1997), Nagaraja et al. (2011) and others. Although these extensions have advantages in one aspect or another, the basic repeat sales model is still the most computationally efficient and least data demanding method, and is widely used by many companies and institutions.

[7] Calhoun et al. (1995) compare repeat sales indexes over annual, semiannual, quarterly as well as monthly intervals, and conclude that aggregation bias arises for all intervals greater than one month. By analogy, if the true housing values fluctuates within months, the standard monthly indexes are likely to be biased.

[8] If a given house sold more than twice, then only the non-overlapping sale pairs are used. For example, a house that sold three times would generate two sale pairs from the first and second transaction, and the second and third transaction; the pair formed by the first and third transaction is not included.

s and t at prices $H_{j,s}$ and $H_{j,t}$, the repeat sales model postulates that,

$$\beta_t H_{j,t} = \beta_s H_{j,s} + \sqrt{2}\sigma_w w_{j,t} + \sqrt{(t-s)}\sigma_v v_{j,t}, \qquad 0 \leq s < t \leq T, \tag{1}$$

with the value of the house price index at time τ is defined by the inverse of β_τ. The last two terms on the right-hand side account for "errors" in the sale pairs, with $\sqrt{2}\sigma_w w_{j,t}$ representing the "mispricing error," and $\sqrt{(t-s)}\sigma_v v_{j,t}$ representing the "interval error." Mispricing errors are included to allow for imperfect information between buyers and sellers, potentially causing the actual sale price of a house to differ from its "true" value. The interval error represents a possible drift over time in the value of a given house away from the overall market trend, and is therefore scaled by the (square root of the) length of the time-interval between the two transactions. The error terms $w_{j,t}$ and $v_{j,t}$ are assumed independent and identically standard normal distributed.

The S&P/Case-Shiller model and the corresponding error structure naturally lend itself to estimation by a three-stage generalized least square type procedure (for additional details, see Case and Shiller, 1987). The base period of the S&P/ Case-Shiller indexes is January 2000. All index values prior to the base period are estimated simultaneously. After the base period, the index values are estimated using a chain-weighting procedure that conditions on all previous values. This chain-weighting procedure is used to prevent revisions of previously published index values. Finally, and importantly, the indexes are smoothed by repeating a given transaction in three successive months, so that, for example, the December index is based on a triplication of the sales that occurred in October, November and December.[9]

[9]Quoting from the Index Construction Section of S&P/Case-Shiller Home Price Index Methodology: "The indices are calculated monthly, using a three-month moving average algorithm ... The index point for each reporting month is based on sales pairs found for that month and the preceding two months."

3 Daily house price indexes

Following most studies in the academic literature, as well as reports in the popular press, we will focus our analysis on the ten largest Metropolitan Statistical Areas (MSAs), as measured in the year 2000. Table 1 shows the counties included in the calculation for each of the ten MSAs. For conciseness, we will refer to the largest county in a given MSA as the label for that area.

[INSERT TABLE 1 ABOUT HERE]

3.1 Data and data cleaning

The transaction data used in our daily index estimation is obtained from DataQuick, a property information company. This database contains detailed transactions of more than one hundred million properties in the United States. For most of the areas, the historical transaction records extends from the late 1990s to 2012, with some large metropolitan areas, such as Boston and New York, having transactions recorded as far back as 1987. Properties are uniquely identified by property IDs, which enable us to identify sale pairs. We rely U.S. Standard Use Codes contained in the DataQuick database to identify transactions of single-family residential homes.

Our data ”cleaning” rules are based on the same filters used by S&P/Case-Shiller and Caplin et al. (2008) in the construction of their monthly indexes. In brief, we remove any transaction that are not “arms-length,” using a flag for such transactions available in the database. We also remove transactions with “unreasonably” low or high sale prices (below $5000 or above $100 million, and those generating an average annual return of below -50% or above 100%), as well as any sales pair with an interval of less than six months. Sale pairs are also excluded if there are indications that major improvements have been made between the two transactions, although such indications are not always present in the database. For the Los Angeles MSA, for example, this yields a total of 877,885 ”clean” sale pairs, representing an average of 180 *daily* sale pairs over the estimation period. Table 2 further summarizes the data for all of the ten MSAs.

[INSERT TABLE 2 ABOUT HERE]

3.2 Practical estimation

The monthly S&P/Case-Shiller repeat-sales indexes are estimated using equation (1). This approach is not computationally feasible at the daily frequency, as it involves the simultaneous estimation of several thousand parameters.[10] To overcome this difficulty, we use an expanding-window estimation procedure: conditional on a start-up period, we begin by estimating daily index values for the final month in an initial sample, imposing the constraint that all of the earlier months have only a single *monthly* index value. Restricting the daily values to be the same within each month for all but the last month drastically reduces the dimensionality of the estimation problem. We then expand the estimation period by one month, thereby obtaining daily index values for the new "last" month. We continue this expanding estimation procedure through to the end of our sample period. Finally, following the S&P/Case-Shiller methodology, we normalize all of the individual indexes to 100 based on their average values in the year 2000.

One benefit of the estimation procedure we adopt is that it is possible to formally test whether the "raw" daily price series actually exhibit significant within month variation. In particular, following the approach used by Calhoun et al. (1995) to test for "aggregation biases" in excess of one month, we test the null hypothesis that the estimates of $\beta_{i,\tau}$ for MSA i are the same for all days τ within a given calendar month against the alternative that these estimates differ within the month. All of these tests strongly reject the null for all of the ten metropolitan areas; further details concerning the actual test results are available upon request. Importantly, as we describe in detail below, this statistically

[10] The daily time spans for the ten MSAs range from a low of 2837 for Washington, D.C. to a high of 4470 days for New York.

significant intra-monthly variation also translates into economically meaningful variation and corresponding gains in forecast accuracy compared to the forecasts based on coarser monthly index values only.

3.3 Noise filtering

The raw daily house price indexes discussed above are invariably subject to measurement errors. To help alleviate this problem, it is useful to further clean the data in an effort to extract more accurate estimates of the the true latent daily price series. We rely on a standard Kalman filter-based approach for doing so.[11]

Specifically, let $P_{i,t}$ denote the true latent index for MSA i at time t. We assume that the "raw" price indexes constructed in the previous section, $P^*_{i,t} = 1/\beta_{i,t}$, are related to the true latent price indexes by,

$$\log P^*_{i,t} = \log P_{i,t} + \eta_{i,t}, \tag{2}$$

where the $\eta_{i,t}$ measurement errors are assumed to be serially uncorrelated. For simplicity of the filter, we will further assume that the true daily price index follows a random walk with drift,

$$r_{i,t} \equiv \Delta \log P_{i,t} = \mu_i + u_{i,t}, \tag{3}$$

where $\eta_{i,t}$ and $u_{i,t}$ are mutually uncorrelated. It follows readily by substitution that,

$$r^*_{i,t} \equiv \Delta \log P^*_{i,t} = r_{i,t} + \eta_{i,t} - \eta_{i,t-1}. \tag{4}$$

Combining (3) and (4), this in turn implies an MA(1) error structure for the "raw" returns, with the value of the MA(1) coefficient determined by the variances of $\eta_{i,t}$ and $u_{i,t}$, σ^2_η and σ^2_u, respectively. This simple MA(1) structure is consistent with the sample auto-

[11]This mirrors the use of filtering techniques for extracting the true latent price process from high-frequency intraday data contaminated by market microstructure noise explored in the financial econometrics literature; see, e.g., Owens and Steigerwald (2006).

correlations for the raw return series for all of the ten cities reported in Figure 1. Except for the first-order autocorrelations, which are close to -0.5 for all of the cities, none of the higher order autocorrelations are significantly different from zero on a systematic basis.

[INSERT FIGURE 1 ABOUT HERE]

Interpreting equations (3) and (4) as a simple state-space system, μ, σ_η^2 and σ_u^2 may easily be estimated by standard (quasi-)maximum likelihood methods. This also allows for the easy filtration of of the "true" daily returns $r_{i,t}$ from $r^*_{i,t}$ by a standard Kalman filter; see, e.g., Hamilton (1994).[12] The resulting estimates reported in Table 3 imply that the σ_η/σ_u noise-to-signal ratios for the daily index returns range from a low of 6.48 (Los Angeles) to a high of 15.18 (Boston), underscoring the importance of filtering out the noise.

[INSERT TABLE 3 ABOUT HERE]

The corresponding filtered estimates of the latent "true" daily price series for Los Angeles are depicted in Figure 2. For comparison, we also include the raw daily prices and the monthly S&P/Case-Shiller index. Looking first at the top panel for the year 2000, the figure clearly illustrates how the filtered daily index mitigates the noise in the raw price series. At the same time, the filtered prices also point to discernable within month variation compared to the step-wise constant monthly S&P/Case-Shiller index.

The bottom panel for the full 1995-2012 sample period tells a similar story. The visual differences between the filtered daily series and the monthly S&P/Case-Shiller index are obviously less glaring on this scale. Nonetheless, the considerable (excessive) variation in the raw daily prices coming from the noise is still evident.

The full-sample plots for the same three price series for all of the ten MSAs shown in Figure 3 further corroborate these same ideas. The relatively simple Kalman filter-based

[12]The Kalman filter implicitly assumes that $\eta_{i,t}$ and $u_{i,t}$ are *iid* normal. If the assumption of normality is violated, the filtered estimates are interpretable as best linear approximations.

approach effectively cleans out the noise in the raw daily prices. We will consequently refer to and treat the filtered series as *the* daily house price indexes in the sequel.[13]

Before turning to our empirical analysis and modeling of the dynamic dependencies in the daily series, it is instructive to more formally contrast the information inherent in the daily indexes with the traditional monthly S&P/Case-Shiller index.

[INSERT FIGURES 2 AND 3 ABOUT HERE]

3.4 Comparisons with the monthly S&P/Case-Shiller index

Like the monthly S&P/Case-Shiller indexes, our daily house price indexes are based on all publicly available property transactions. However, the complicated non-linear transformations of the data used in the construction of the indexes prevent us from expressing the monthly indexes as explicit functions of the corresponding daily indexes. Instead, as a simple way to help gauge the relationship between the indexes, and the potential loss of information in going from the daily to the monthly frequency, we consider the linear projection of the monthly S&P/Case-Shiller returns for MSA i, denoted $r_{i,t}^{S\&P}$, on 60 lagged values of the corresponding daily index returns,[14]

$$r_{i,t}^{S\&P} = \delta(L) r_{i,t} + \varepsilon_{i,t} = \sum_{j=0}^{59} \delta_j L^j r_{i,t} + \varepsilon_{i,t}, \tag{5}$$

where $L^j r_{i,t}$ refers to the daily return on the j^{th} day before the last day of month t. The inclusion of 60 daily lags match the three-month smoothing window used in the construction of the monthly S&P/Case-Shiller indexes, as discussed in Section 2. The true population coefficients in the linear $\delta(L)$ filter are, of course, unknown. However, they are readily estimated by ordinary least squares (OLS).

[13]The "smoothed" daily prices constructed from the full sample look almost indistinguishable from the filtered series shown in the figures. We purposely rely on filtered as opposed to smoothed estimates to facilitate the construction of meaningful forecasts.

[14]As discussed further below, all of the price series appear to be non-stationary. We consequently formulate the projection in terms of returns as opposed to the price levels.

The resulting estimates for $\delta_{j=0,\ldots,59}$ obtained from the single regression that pools the returns for all ten MSAs are reported in the top panel of Figure 4. Each of the individual coefficients are obviously subject to a fair amount of estimation error. At the same time, there is a clear pattern in the estimates for δ_j across lags. This naturally suggests the use of a polynomial approximation in j to help smooth out the estimation error. The solid line in the figure shows the resulting nonlinear least squares (NLS) estimates obtained from a simple quadratic approximation. The corresponding R^2s for the unrestricted OLS and the NLS fit $\hat{\delta}_j = 0.1807 + 0.0101j - 0.0002j^2$ are 0.860 and 0.851, respectively, indicating only a slight deterioration in te accuracy of the fit by imposing the quadratic approximation to the lag coefficients. Moreover, even though the monthly S&P/Case-Shiller returns are not an exact linear function of the daily returns, the simple relationship dictated by $\delta(L)$ accounts for the majority of the monthly variation.

[INSERT FIGURE 4 ABOUT HERE]

To further illuminate the feature of the approximate linear filter linking the monthly returns to the daily returns, consider the gain and the phase of $\delta(L)$,

$$G(\omega) = \left[\sum_{j=0}^{59} \sum_{k=0}^{59} \delta_j \delta_k cos(|j-k|\omega) \right]^{1/2}, \quad \omega \in (0, \pi), \tag{6a}$$

$$\theta(\omega) = tan^{-1} \left(\frac{\sum_{j=0}^{59} \delta_j sin(j\omega)}{\sum_{j=0}^{59} \delta_j cos(j\omega)} \right), \quad \omega \in (0, \pi). \tag{6b}$$

Looking first at the gains in Figure 4b and 4c, the unrestricted OLS estimates and the polynomial NLS estimates give rise to similar conclusions. The filter effectively down-weights all of the high-frequency variation (corresponding to periods less than around 70 days), while keeping all of the low-frequency information (corresponding to periods in excess of 100 days). As such, potentially valuable information for forecasting changes in house prices is obviously lost in the monthly aggregate. Further along these lines, Figure 4d and 4e show the estimates of $\frac{\theta(\omega)}{\omega}$, or the number of days that the filter shifts the daily returns back in time across frequencies. Although the OLS and NLS estimates differ somewhat for the very highest frequencies, for the lower frequencies (periods in excess of 60 days) the filter systematically shifts the daily returns back in time by about 30 days. This corresponds roughly to one-half of the three month (60 business days) smoothing window used in the construction of the monthly S&P/Case-Shiller index.

In sum, the monthly S&P/Case-Shiller indexes essentially “kill” all of the within quarter variation inherent in the new daily indexes, while delaying all of the longer-run information by more than a month. We turn next to a more detailed analysis of the actual time series properties of the new daily indexes, along with a simple model designed to conveniently describe the dependencies.

4 Time series modeling of daily housing returns

To facilitate the formulation of a multivariate model for all of the ten city indexes, we restrict our attention to the common sample period from June 2001 to September 2012. Excluding weekends and federal holidays, this leaves us with a total of 2,843 daily observations.

4.1 Summary statistics

Summary statistics for each of the ten daily series are reported in Table 4. The first Panel A gives the sample means and standard deviations for each of the index levels. Standard unit root tests clearly suggest that the price series are non-stationary, and as such the sample moments in Panel A need to be interpreted with care; further details concerning the unit root tests are available upon request. In the sequel, we focus on the easier-to-interpret daily return series.

[INSERT TABLE 4 ABOUT HERE]

The daily sample mean returns reported in Panel B are generally positive, ranging a low of -0.006 (Las Vegas) to a high of 0.015 (Los Angeles and Washington D.C.). The standard deviation of the most volatile daily returns 0.599 (Chicago) is double that of the least volatile returns 0.291 (New York). The first-order autocorrelations are fairly close to zero for all of the cities, but the Ljung-Box tests for up to tenth order serial correlation indicate significant longer-run dynamic dependencies in many of the series.

The corresponding results for the squared daily returns reported in Panel C indicate very strong dynamic dependencies. This is also immediately evident from the plot of the ten daily return series in Figure 5, which show a clear tendency for large returns in an absolute sense to be follow by other large absolute returns. This directly mirrors the ubiquitous volatility clustering widely documented in the literature for other daily speculative returns. Again, consistent with the extant empirical finance literature and the evidence reported for other financial asset classes, there is also a clear commonality in the volatility patterns across the ten series.

4.2 Modeling conditional mean dependencies

The summary statistics discussed above point to existence of some, albeit relatively mild, dynamic dependencies in the daily conditional means for most of the cities. Some of these dependencies may naturally arise from a common underlying dynamic factor that influences housing valuations nationally. In order to accommodate both city specific and national effects within a relatively simple linear structure, we postulate the following model for the conditional means of the daily returns,[15]

$$E_{t-1}(r_{i,t}) = c_i + \rho_{i1} r_{i,t-1} + \rho_{i5} r_{i,t-5} + \rho_{im} r^m_{i,t-1} + b_{ic} r^m_{c,t-1}, \tag{7}$$

[15] We do not seek to identify the absolute *best* time series model for each of the ten individual daily MSA indexes. Instead, we attempt to provide a relatively simple and easy-to-implement common parametric specification that fits all of the ten cities reasonably well.

where $r_{i,t}^m$ refers to the (overlapping) "monthly" returns defined by the summation of the corresponding daily returns,

$$r_{i,t}^m = \sum_{j=0}^{19} r_{i,t-j}, \tag{8}$$

and the composite (national) return $r_{c,t}$ is defined as a weighted average of the individual city returns,

$$r_{c,t} = \sum_{i=1}^{10} w_i r_{i,t}, \tag{9}$$

with the weights identical to the ones used in the construction of the composite ten city monthly S&P/Case Shiller index.[16] The own fifth lag of the returns is included to account for any weekly calendar effects. The inclusion of the own monthly returns and the composite monthly returns provides a parsimonious way of accounting for longer-run city-specific and common national dynamic dependencies. This particular formulation is partly motivated by the Heterogeneous Autoregressive (HAR) model originally proposed by Corsi (2009) for modeling so-called realized volatilities, and we will refer to it as an HAR-X model for short. We estimate this model for the conditional mean simultaneously with the model for the conditional variance described in the next section via quasi-maximum likelihood.

The estimation results in Table 5 reveal that the ρ_1 and ρ_5 coefficients associated with the own lagged returns are mostly, though not uniformly, insignificant when judged by the robust standard errors reported in parentheses. Meanwhile, the b_c coefficients associated with the composite monthly return are significant for nine out of the ten cities. Still, the one-day return predictability implied by the model is fairly modest, with the average daily R^2 across the ten cities equal to 0.024, ranging from a low of 0.007 (Denver) to a high of 0.049 (San Francisco). This is consistent with the low R^2s generally obtained from time series modeling of other daily financial returns.

The adequacy of the common specification for the conditional mean in equation (7)

[16] The specific values for each of the ten cities are 0.212, 0.074, 0.089, 0.037, 0.050, 0.015, 0.055, 0.118, 0.272, and 0.078, respectively, representing the total aggregate value of the housing stock in the ten MSAs in the year 2000.

is broadly supported by the tests for up to tenth order serial correlation in the residuals $\varepsilon_{i,t} \equiv r_{i,t} - E_{t-1}(r_{i,t})$ from the model reported in Panel C of Table 5. Only two of the tests are significant at the 5% level (San Francisco and Washington, D.C.). At the same time, the tests for serial correlation in the squared residuals $\varepsilon_{i,t}^2$ from the model, given in the bottom two rows, clearly indicate strong non-linear dependencies in the form of volatility clustering.

[INSERT TABLE 5 ABOUT HERE]

4.3 Modeling conditional variance and covariance dependencies

Numerous parametric specifications have been proposed in the literature for best describing volatility clustering in asset returns. Again, in an effort to keep our modeling procedures simple and easy-to-implement, we will rely on the popular GARCH(1,1) model (Bollerslev, 1986) for describing the dynamic dependencies in the conditional variances for all of the ten cities,

$$Var_{t-1}(r_{i,t}) \equiv h_{i,t} = \omega_i + \kappa_i \varepsilon_{i,t-1}^2 + \lambda_i h_{i,t-1}. \tag{10}$$

The results from estimating this model jointly with the the conditional mean model described in the previous section are reported in Panel B of Table 5 together with robust standard errors following Bollerslev and Wooldridge (1992) in parentheses.

The estimated GARCH parameters are all highly statistically significant and fairly similar across cities. Consistent with the results obtained for other daily financial return series, the estimates for the sum $\kappa+\lambda$ are all very close to unity (and just above for Chicago, at 1.002) indicative of a highly persistent, but eventually mean-reverting, time-varying volatility process.

The Wald tests for up to tenth order serial correlation in the resulting standardized residuals $\varepsilon_{i,t}/h_{i,t}^{1/2}$ reported in Panel C suggest that little predictability remains, with only one city (San Francisco) rejecting the null of no autocorrelation. The tests for serial correlation in the squared standardized residuals $\varepsilon^2_{i,t}/h_{i,t}$ reject the null for four cities, perhaps indicative of some remaining predictability in volatility not captured by this relatively simple model. However for the majority of cities the specification in equation (10) appears to provide a satisfactory fit. The dramatic reduction in the values of the test statistics for the squared residuals compared to the values reported in the second row of Panel C is particularly noteworthy.

The univariate HAR-X-GARCH models defined by equations (7) and (10) indirectly incorporate commonalities in the cross-city returns through the composite monthly returns $r_{c,t}$ included in the conditional means. The univariate models do not, however, explain the aforementioned commonalities in the volatilities observed across cities and the corresponding dynamic dependencies in the conditional covariances of the returns.

The Constant Conditional Correlation (CCC) model proposed by Bollerslev (1990) provides a particularly convenient framework for jointly modeling the ten daily return series by postulating that the temporal variation in the conditional covariances are proportional to the products of the conditional standard deviations. Specifically, let $\mathbf{r_t} \equiv [r_{1,t}, ..., r_{10,t}]'$ and $D_t \equiv diag\left\{h^{1/2}_{1t}, ..., h^{1/2}_{10,t}\right\}$ denote the 10×1 vector of daily returns and 10×10 diagonal matrix with the GARCH conditional standard deviations along the diagonal, respectively. The GARCH-CCC model for the conditional covariance matrix of the returns may then be succinctly expressed as,

$$Var_{t-1}(\mathbf{r_t}) = D_t R D_t, \tag{11}$$

where R is a 10×10 matrix with ones along the diagonal and the conditional correlations in the off-diagonal elements. Importantly, the R matrix may be efficiently estimated by the sample correlations for the 10×1 vector of standardized HAR-X-GARCH residuals; i.e., the estimates of $D_t^{-1}\left[\mathbf{r_t} - E_{t-1}(\mathbf{r_t})\right]$.

The resulting estimates for R are reported in Table 6. These daily correlations may seem surprisingly small, but as discussed further below, the model do imply much larger longer-run correlations.

We also experimented with the estimation of the Dynamic Conditional Correlation (DCC) model of Engle (2002), in which the time-varying conditional correlations are determined by a GARCH(1,1) structure for the vector of standardized residuals. The maximized value of the (quasi-) log-likelihood function for the DCC model of 85.229 is only slightly larger than the value of 85.176 obtained for the CCC model, and the GARCH parameters characterizing the temporal variation in R_t are also statistically insignificant. Hence, we conclude that the relatively simple multivariate HAR-X-GARCH-CCC model defined by equations (7), (10), and (11) provides a satisfactory fit to the joint dynamic dependencies in the conditional first and second order moments of the ten daily housing return series.

4.4 Temporal aggregation and housing return correlations

The estimated conditional correlations from the HAR-X-GARCH-CCC model for the daily index returns reported in Table 6 averages only 0.022. By contrast the unconditional correlations for the monthly S&P/Case Shiller indexes calculated over the same time period reported in Table 7 averages 0.708, ranging from 0.382 (Denver–Las Vegas) to 0.926 (Los Angeles–San Diego). This apparent discrepancy between the two sets of numbers, seemingly calls into question the integrity of our new daily indexes and/or the time-series models for describing the dynamic dependencies therein.

Of course, the conditional daily correlations and the unconditional monthly correlations are not directly comparable, as the conditioning and the temporal aggregation may affect covariances (the numerator) differently from variances (the denominator). Hence, in order to more directly compare the longer-run dependencies inherent in our new daily indexes to the more traditional monthly S&P/Case Shiller indexes, we aggregate our daily return indexes to a monthly level by summing the daily returns within a month (20 days).

The unconditional sample correlations for these new monthly returns are reported in the lower triangle of Panel B in Table 8. These numbers are obviously much closer to the unconditional correlations for the published S&P/Case Shiller indexes in Table 7. However, they are generally still below those values.

However, as previously noted, the monthly S&P/Case Shiller indexes are artificially "smoothed," by repeating each sale pair in the two months following the actual repeat sale. As such, a more meaningful comparison of the longer-run correlations inherent in our new daily indexes with the correlations in the S&P/Case Shiller indexes is afforded by the unconditional quarterly (60 days) correlations reported in the upper triangle of Panel B in Table 8. All of these sample correlations are very close to the ones in Table 7, thus indirectly corroborating our new daily index construction.

[INSERT TABLE 8 ABOUT HERE]

This, of course, says nothing about the validity of the HAR-X-GARCH-CCC model for the daily returns, and the previously noted very low daily conditional correlations for the model in Table 6. As a further indirect specification check for the model, we therefore also consider the model-implied longer-run correlations, and study how these compare with the sample correlations for the actual longer-run aggregate returns.

The HAR-X-GARCH-CCC model does not admit closed-form expressions for the lower-frequency implied correlations, but these are easy to approximate by numerical simulations. The top number in each element of Panels A and B of Table 8 gives the simulated median model-implied unconditional correlations for the daily, weekly, monthly, and quarterly return horizons, based on a total of 500 simulated sample paths. The bottom number in each element is the corresponding sample correlations for the actual longer-run aggregate returns.

Consistent with the conditional correlations from the model in Table 6, the daily unconditional correlations in Panel A are all close to zero. However, the unconditional correlations implied by the model gradually increase with the return horizon, and almost all of the quarterly correlations are in excess of one-half. Importantly, the longer-run model-implied correlations closely match their unconditional sample analogues reported in the upper triangular panels.

Further to this effect, Figure 6 presents the median model-implied and sample correlations for return horizons ranging from one-day to a quarter, along with the corresponding simulated 95% confidence intervals implied by the model for the Los Angeles–Boston and Los Angeles–New York city pairs. The model evidently provides a very good fit across all horizons, with the actual correlations well within the confidence bands. The corresponding plots for all of the 45 city pairs, presented in Figure 7, tell a similar story. The median simulated values and actual sample correlations reported in the upper triangular part of the figure are systematically close, and the sample values are well within the 95% confidence intervals implied by the model given in the lower triangular portion of the figure.

Taken as whole the results in Figures 6 and 7 clearly support the idea that the longer-run cross-city dependencies inherent in our new finer sample daily house price series are fully compatible with those in the published coarser monthly S&P/Case Shiller indexes. The results also confirm that the joint dynamic dependencies in the daily returns are well described by the relatively simple HAR-X-GARCH-CCC model developed above, in turn suggesting that this model could possibly be used in the construction of improved house price forecasts over longer monthly horizons.

[INSERT FIGURES 6 and 7 ABOUT HERE]

5 Forecasting housing returns

One of the major potential benefits from higher frequency data is the possibility of constructing more accurate forecasts, by using models that more quickly incorporate new information. In order to ascertain the potential improvements along this dimension afforded by the daily house price series and our modeling thereof, we consider a comparison of the forecasts from the daily HAR-X-GARCH-CCC model with different benchmark alternatives.

Specifically, consider the problem of forecasting the 20-day return (which we shall refer to as "monthly") on the house price index for MSA i,

$$r_{i,t}^{(m)} \equiv \sum_{j=0}^{19} r_{i,t-j} \tag{12}$$

for forecast horizons ranging from $h = 20$ days ahead to $h = 1$ day ahead.[17] When $h = 20$ this corresponds to a simple one-step ahead forecast for one-month returns, but for $h < 20$ an optimal forecast will contain a mixture of observed data and a forecast for the return over the remaining part of the month. We will use the period June 2001 to June 2009 as our in-sample period, and the period July 2009 to September 2012 as our out-of-sample period.[18] All of the model parameters are estimated once over the fixed in-sample period.

Our simplest benchmark forecast is based purely on end-of-month data, and is therefore *not* updated as the horizon shrinks. We will consider a simple AR(1) for these monthly returns,

$$r_{i,t}^{(m)} = \phi_0 + \phi_1 r_{i,t-20}^{(m)} + e_{i,t}.$$

As the forecast is not updated through the month, the forecast made at time $t-h$ is simply the AR(1) forecast made at time $t - 20$,

$$\hat{r}_{i,t-h}^{Mthly} = \hat{\phi}_0 + \hat{\phi}_1 r_{i,t-20}^{(m)}. \tag{13}$$

Our second benchmark forecast is again purely based on monthly data, but now we

[17]In the forecast literature, this is referred to as a "fixed event" forecast design, see Nordhaus (1987) for an early analysis of such problems.

[18]A preliminary version of this paper used an earlier vintage of the DataQuick database that ended in June 2009, which is how we chose our sample-split point. This preliminary version of the paper did not consider any out-of-sample comparisons, and so the results presented here are close to "true," rather than "pseduo," out-of-sample.

allow the forecaster to update the forecast at time $t-h$, which may be in the middle of a month. We model the incorporation of observed data by allowing the forecaster to take a linear combination of the monthly return observed on day $t-h$ and the one-month-ahead forecast made on that day,

$$\hat{r}_{i,t-h}^{Interp} = \left(1 - \frac{h}{20}\right) r_{i,t-h}^{(m)} + \frac{h}{20}\left(\hat{\phi}_0 + \hat{\phi}_1 r_{i,t-h}^{(m)}\right). \tag{14}$$

Our third forecast fully exploits the daily return information, by using the actual returns from time $t-19$ to $t-h$ as the first component of the forecast, as these are part of the information set at time $t-h$, and then using a "direct projection" method to obtain a forecast for the remaining h-day return based on the one-month return available at time $t-h$. Specifically,

$$\hat{r}_{i,t-h}^{Direct} = \sum_{j=h}^{19} r_{i,t-j} + \hat{\beta}_0^{(h)} + \hat{\beta}_1^{(h)} r_{i,t-h}^{(m)}, \tag{15}$$

where,

$$\sum_{j=0}^{h-1} r_{i,t-j} = \beta_0^{(h)} + \beta_1^{(h)} r_{i,t-h}^{(m)} + u_{i,t},$$

and the $\beta_0^{(h)}$ and $\beta_1^{(h)}$ coefficients are estimated from the relevant projection.

Finally, we consider a forecast based on the HAR-X-GARCH-CCC model presented in the previous section. Like the third forecast, this forecast uses the actual returns from time $t-19$ to $t-h$ as the first component, and then iterates the expression for the conditional daily mean returns in equation (7) forward to get forecasts for the remaining h days,

$$\hat{r}_{i,t-h}^{HAR} = \sum_{j=h}^{19} r_{i,t-j} + \sum_{j=0}^{h-1} \hat{E}_{t-h}\left[r_{i,t-j}\right]. \tag{16}$$

Given the construction of the target variable, we expect the latter three forecasts ("Interp", "Direct", "HAR") to all beat the "Mthly" forecast for all horizons less than 20 days. If intra-monthly returns have dynamics that differ from those of monthly returns, then we expect the latter two forecasts to beat the "Interp" forecast, as they both explicitly allow for this characteristic. Finally, if the HAR-X-GARCH-CCC model presented in the previous section provides a closer description of the true dynamics than a simple direct projection, then we would also expect the fourth forecast to beat the third.

Figure 8 shows the resulting Root Mean Squared Errors (RMSEs) for the four forecasts as a function of the forecast horizon, when evaluated over the July 2009 to September 2012 out-of-sample period. The first striking, though not surprising, feature is that exploiting higher frequency (intra-monthly) data leads to smaller forecast errors than a forecast based purely on monthly data. All three of the forecasts that use intra-monthly information out-perform the model based solely on end-of-month data. The only exception to this is for Las Vegas at the $h = 20$ horizon, where the HAR model slightly under-performs the monthly model.

Another striking feature of Figure 8 is that the more accurate modeling of the daily dynamic dependencies afforded by the HAR-X-GARCH-CCC model results in lower RMSEs across *all* forecasts horizons for eight of the ten cities. For San Francisco and Las Vegas the direct projection forecasts perform essentially as well as the HAR forecasts, and for Denver and Los Angeles the improvement of the HAR forecast is small (but positive for all horizons). For some of the cities (Boston, Miami and Washington D.C., in particular) the improvements are especially dramatic over longer horizons.

[INSERT FIGURE 8 ABOUT HERE]

The visual impression from Figure 8 is formally underscored by Diebold-Mariano tests, reported in Table 9. Not surprisingly, the HAR forecasts significantly outperform the monthly forecasts for horizons of 1, 5 and 10 days, for all ten cities and the composite index. At the one-month horizon, a tougher comparison for the model, the HAR forecasts are significantly better than the monthly model forecasts for four out of ten cities, as well as the composite index, and are never significantly beaten by the monthly model forecasts. Almost identical conclusions are drawn when comparing the HAR forecasts to the "interpolation" forecasts, supporting the conclusion that the availability of daily data clearly holds the promise of more accurate forecasts, particularly over shorter horizons, but also even at the monthly level.

The bottom row of each panel in Table 9 compares the HAR forecasts with those from a simple direct projection model. Such forecasts have often been found to perform well in comparison with "iterated" forecasts from more complicated models (see, e.g., Marcellino et al., 2006, for a recent comparison along these lines in the context of macroeconomic forecasting). By contrast, the Diebold-Mariano tests reported here suggest that the more complicated HAR forecasts generally perform better than the direct projection forecasts. For no city-horizon pair does the direct projection forecast lead to significantly lower out-of-sample forecast RMSE than the HAR forecasts, while for many city-horizon pairs the reverse is true. In particular, for Boston, Miami and Washington D.C., the HAR forecasts significantly beat the direct projection forecasts across all four horizons, and for the composite index this is true for all but the shortest horizon.

All-in-all, the results from our out-of-sample forecast analysis clearly suggest that the access to higher frequency daily data holds the promise of more accurate longer-run house price forecasts. Of course, more complicated models incorporating additional high-frequency information may give rise to even better forecasts than the relatively simple HAR-X-GARCH-CCC time series model developed here.

[INSERT TABLE 9 ABOUT HERE]

6 Conclusion

We present a set of new *daily* house price indexes for ten major U.S. Metropolitan Statistical Areas spanning the period from June 2001 to September 2012. The index construction is based on the same repeat sales method underlying the popular monthly and quarterly S&P/Case-Shiller indexes (Shiller, 1991), along with a comprehensive database of several million publicly recorded residential property transactions. We demonstrate that the dynamic dependencies in the new daily housing price series closely mimic those of other financial asset prices, and that these dynamic dependencies along with the cross-city correlations are well described by a standard multivariate GARCH-type model. We also show that this relatively simple daily model allows for the construction of improved longer-run weekly and monthly housing price forecasts compared to the forecasts based solely on existing coarser monthly price indexes.

The new "high frequency" house price indexes developed here open the possibility for many other applications. Most directly, by providing more timely estimates of movements in the housing market, the daily series should be of immediate interest to policy makers and central banks, alike. In a related context, the series may also prove useful in further studying the microstructure of the housing market, as well as in the estimation of more structural equilibrium-based models for the housing market (see, e.g. Bajari et al., 2013). At a broader level, combining the daily house price series with other daily estimates of economic activity should afford better and more up-to-date insights into changes in the macro economy. Along these lines, the series also hold the promise for the construction of more accurate forecasts for many *other* macro economic and financial time series. We leave all of these issues for future research.

References

Aruoba, S. Borağan, Francis X. Diebold, and Chiara Scotti (2009), "Real-time measurement of business conditions." *Journal of Business & Economic Statistics*, 27, 417–427.

Bailey, Martin J., Richard F. Muth, and Hugh O. Nourse (1963), "A regression method for real estate price index construction." *Journal of the American Statistical Association*, 58, 933–942.

Bajari, Patrick, Phoebe Chan, Dirk Krueger, and Daniel Miller (2013), "A dynamic model of housing demand: Estimation and policy implications." *International Economic Review*, 54, 409–442.

Bollerslev, Tim (1986), "Generalized autoregressive conditional heteroskedasticity." *Journal of Econometrics*, 31, 307–327.

Bollerslev, Tim (1990), "Modelling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model." *Review of Economics and Statistics*, 72, 498–505.

Bollerslev, Tim and Jeffrey M. Wooldridge (1992), "Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances." *Econometric Reviews*, 11, 143–172.

Calhoun, Charles A., Peter Chinloy, and Isaac F. Megbolugbe (1995), "Temporal aggregation and house price index construction." *Journal of Housing Research*, 6, 419–438.

Campbell, John Y. and João F. Cocco (2007), "How do house prices affect consumption? evidence from micro data." *Journal of Monetary Economics*, 54, 591 – 621.

Caplin, Andrew, Sumit Chopra, John Leahy, Yann Lecun, and Trivikrmaman Thampy (2008), "Machine learning and the spatial structure of house prices and housing returns." Working paper, New York University.

Case, Bradford and John M. Quigley (1991), "The dynamics of real estate prices." *Review of Economics and Statistics*, 73, 50–58.

Case, Karl E., John M. Quigley, and Robert J. Shiller (2005), "Comparing wealth effects: The stock market versus the housing market." *Advances in Macroeconomics*, 5, 1–32.

Case, Karl E., John M. Quigley, and Robert J. Shiller (2011), "Wealth effects revisited 1978-2009." Working paper, National Bureau of Economic Research.

Case, Karl E. and Robert J. Shiller (1987), "Prices of single-family homes since 1970: New indexes for four cities." *New England Economic Review*, 45–56.

Case, Karl E. and Robert J. Shiller (1989), "The efficiency of the market for single-family homes." *American Economic Review*, 79, 125–137.

Chauvet, Marcelle, Stuart Gabriel, and Chandler Lutz (2013), "Fear and loathing in the housing market: Evidence from search query data." Working paper, University of California, Los Angeles.

Cho, Man (1996), "House price dynamics: A survey of theoretical and empirical issues." *Journal of Housing Research*, 7, 145–172.

Corsi, Fulvio (2009), "A simple approximate long-memory model of realized volatility." *Journal of Financial Econometrics*, 7, 174–196.

Davis, Morris A. and Jonathan Heathcote (2005), "Housing and the business cycle." *International Economic Review*, 46, 751–784.

Engle, Robert F. (2002), "Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models." *Journal of Business & Economic Statistics*, 20, 339–350.

Evans, Martin D.D. (2005), "Where are we now? Real-time estimates of the macro economy." *International Journal of Central Banking*, 1, 127–175.

Ghysels, Eric, Alberto Plazzi, Walter Torous, and Rossen Valkanov (2013), "Forecasting real estate prices, in Graham Elliott and Allan Timmermann (eds.)." *Handbook of Economic Forecasting*, forthcoming.

Goetzmann, William N. and Matthew Spiegel (1997), "A spatial model of housing returns and neighborhood substitutability." *Journal of Real Estate Finance and Economics*, 14, 11–31.

Marcellino, Massimiliano, James H Stock, and Mark W Watson (2006), "A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series." *Journal of Econometrics*, 135, 499–526.

Meese, Richard A and Nancy E Wallace (1997), "The construction of residential housing price indices: A comparison of repeat-sales, hedonic-regression, and hybrid approaches." *Journal of Real Estate Finance and Economics*, 14, 51–73.

Nagaraja, Chaitra H., Lawrence D. Brown, and Linda H. Zhao (2011), "An autoregressive approach to house price modeling." *Annals of Applied Statistics*, 5, 124–149.

Nordhaus, William D (1987), "Forecasting efficiency: concepts and applications." *Review of Economics and Statistics*, 69, 667–674.

Owens, John P and Douglas G Steigerwald (2006), "Noise reduced realized volatility: A kalman filter approach." *Advances in Econometrics*, 20, 211–227.

Quigley, John M. (1995), "A simple hybrid model for estimating real estate price indexes." *Journal of Housing Economics*, 4, 1–12.

Rappaport, Jordan (2007), "A guide to aggregate house price measures." *Economic Review, Federal Reserve Bank of Kansas City*, 41–71.

Shiller, Robert J. (1991), "Arithmetic repeat sales price estimators." *Journal of Housing Economics*, 1, 110–126.

Shiller, Robert J. (1993), "Measuring asset values for cash settlement in derivative markets: Hedonic repeated measures indices and perpetual futures." *Journal of Finance*, 48, 911–931.

Standard & Poor's (2009), "S&P/Case-Shiller Home Price Indices Methodology."

Table 1: Metropolitan Statistical Areas (MSAs)

MSA	Represented counties	Counties in our indexes
Los Angeles-Long Beach-Santa Ana, CA Metropolitan Statistical Area (Los Angeles)	Los Angeles CA, Orange CA	Los Angeles CA, Orange CA
Boston-Cambridge-Quincy, MA-NH Metropolitan Statistical Area (Boston)	Essex MA, Middlesex MA, Norfolk MA, Plymouth MA, Suffolk MA, Rockingham NH, Strafford NH	Essex MA, Middlesex MA, Norfolk MA, Plymouth MA, Suffolk MA, Rockingham NH, Strafford NH
Chicago-Naperville-Joliet, IL Metropolitan Division (Chicago)	Cook IL, DeKalb IL, Du Page Il, Kane IL, Kendal IL, McHenry IL, Will IL, Grundy IL	Cook IL, DeKalb IL, Du Page Il, Kane IL, Kendal IL, McHenry IL, Will IL, Grundy IL
Denver-Aurora, CO Metropolitan Statistical Area (Denver)	Adams CO, Arapahoe CO, Broomfield CO, Clear Creek CO, Denver CO, Douglas CO, Elbert CO, Gilpin CO, Jefferson CO, Park CO	Adams CO, Arapahoe CO, Broomfield CO, Clear Creek CO, Denver CO, Douglas CO, Elbert CO, Gilpin CO, Jefferson CO, Park CO
Miami-Fort Lauderdale-Pompanp Beach, FL Metropolitan Statistical Area (Miami)	Broward FL, Miami-Dade FL, Palm Beach FL	Broward FL, Miami-Dade FL, Palm Beach FL
Las Vegas-Paradise, NV Metropolitan Statistical Area (Las Vegas)	Clark NV	Clark NV
San Diego-Carlsbad-San Marcos, CA Metropolitan Statistical Area (San Diego)	San Diego CA	San Diego CA
San Francisco-Oakland-Fremont, CA Metropolitan Statistical Area (San Francisco)	Alameda CA, Contra Costa CA, Marin CA, San Francisco CA, San Mateo CA	Alameda CA, Contra Costa CA, Marin CA, San Francisco CA, San Mateo CA

Table 1: Continued

MSA	Represented counties	Counties in our indices
New York City Area (New York)	Fairfield CT, New Haven CT, Bergen NJ, Essex NJ, Hudson NJ, Hunterdon NJ, Mercer NJ, Middlesex NJ, Monmouth NJ, Morris NJ, Ocean NJ, Passaic NJ, Somerset NJ, Sussex NJ, Union NJ, Warren NJ, Bronx NY, Dutchess NY, Kings NY, Nassau NY, New York NY, Orange NY, Putnam NY, Queens NY, Richmond NY, Rockland NY, Suffolk NY, Westchester NY, Pike PA	Fairfield CT, New Haven CT, Bergen NJ, Essex NJ, Hudson NJ, Hunterdon NJ, Mercer NJ, Middlesex NJ, Monmouth NJ, Morris NJ, Ocean NJ, Passaic NJ, Somerset NJ, Sussex NJ, Union NJ, Warren NJ, Bronx NY, Dutchess NY, Kings NY, Nassau NY, New York NY, Orange NY, Putnam NY, Queens NY, Richmond NY, Rockland NY, Suffolk NY, Westchester NY
Washington-Arlington-Alexandria, DC-VA-MD-WV Metropolitan Statistical Area (Washington, D.C.)	District of Columbia DC, Calvert MD, Charles MD, Frederick MD, Montgomery MD, Prince Georges MD, Alexandria City VA, Arlington VA, Clarke VA, Fairfax VA, Fairfax City VA, Falls Church City VA, Fauquier VA, Fredericksburg City VA, Loudoun VA, Manassas City VA, Manassas Park City VA, Prince William VA, Spotsylvania VA, Stafford VA, Warren VA, Jefferson WV	District of Columbia DC, Calvert MD, Charles MD, Frederick MD, Montgomery MD, Prince Georges MD, Alexandria City VA, Arlington VA, Fairfax VA, Falls Church City VA, Fauquier VA, Fredericksburg City VA, Loudoun VA, Manassas City VA, Manassas Park City VA, Prince William VA, Spotsylvania VA, Stafford VA

Note: The table reports the counties included in the ten Metropolitan Statistical Areas (MSAs) underlying the S&P/Case-Shiller indexes. The name of each MSA is abbreviated by that of its major city or county, as indicated in parenthesis.

Table 2: Data summary

Los Angeles	Boston	Chicago	Denver	Miami	Las Vegas	San Diego	San Francisco	New York	Washington, D.C.
Panel A: Data availability									
Full sample start date									
04/01/88	02/01/87	01/08/96	02/01/98	02/01/97	04/01/88	04/01/88	04/01/88	02/01/87	01/10/96
Daily index start date									
03/01/95	03/01/95	09/01/99	05/03/99	04/01/98	01/03/95	01/02/96	01/03/95	01/03/95	06/01/01
Full sample end date									
10/23/12	10/11/12	10/12/12	10/17/12	10/15/12	10/17/12	10/23/12	10/18/12	10/23/12	10/23/12
Panel B: Transactions									
Total transactions									
10,285,770	2,121,471	3,948,706	1,672,669	3,689,159	2,236,138	2,845,804	3,778,446	5,943,114	2,168,018
Single family residential housing transactions									
5,970,536	1,141,930	1,886,433	1,000,785	1,366,745	1,479,872	1,584,732	2,331,860	2,951,031	1,055,537
Arms-length transaction									
2,562,884	975,964	1,157,215	672,512	935,985	915,408	755,440	1,031,261	2,307,079	759,752
After excluding transaction value <= 5000 or >= 100,000,000									
2,555,165	917,039	1,156,042	671,605	935,178	913,682	754,106	1,030,384	2,271,467	757,675
After excluding houses sold only once									
1,980,740	638,577	659,732	475,481	668,552	729,365	579,152	757,379	1,234,074	459,842
After excluding transactions happen within 6 months									

1,627,149	532,761	561,945	374,045	576,810	645,869	510,450	688,869	1,026,836	325,251
After excluding >= 2× standard deviations and >= 6× median transaction values									
1,578,869	514,356	543,038	360,944	561,805	628,790	494,894	665,537	999,284	313,777
				Panel C: Sale pairs					
Total pairs									
939,476	294,101	292,737	198,608	321,358	378,093	296,985	397,229	544,326	162,693
After excluding renovation/reconstruction between two sales									
899,573	286,760	292,737	187,977	287,790	226,701	244,059	350,500	540,235	151,203
After excluding abnormal annual returns (less than -50% or more than 100%)									
878,017	272,858	277,160	181,633	281,393	221,877	239,232	341,878	512,251	143,481
After excluding sale pairs with second transaction on weekends									
878,002	272,727	277,095	180,504	277,442	221,876	239,215	341,858	508,860	143,433
After excluding sale pairs with second transaction on federal holidays									
877,885	272,414	277,079	180,003	276,676	221,554	239,041	341,469	508,548	143,431
Average *daily* sale pairs for the daily index estimation period									
180	55	84	53	77	49	51	70	109	49

Table 3: Noise filter estimates

	Los Angeles	Boston	Chicago	Denver	Miami	Las Vegas	San Diego	San Francisco	New York	Washington, D.C.
μ	0.018	0.019	0.002	0.009	0.013	0.001	0.020	0.018	0.016	0.017
	(0.006)	(0.007)	(0.001)	(0.008)	(0.010)	(0.000)	(0.006)	(0.007)	(0.005)	(0.009)
σ_η	2.457	5.888	4.668	3.779	4.113	5.362	3.746	4.925	4.349	4.612
	(0.039)	(0.140)	(0.155)	(0.139)	(0.076)	(0.212)	(0.058)	(0.105)	(0.132)	(0.108)
σ_u	0.379	0.388	0.593	0.327	0.497	0.568	0.407	0.525	0.376	0.501
	(0.022)	(0.034)	(0.057)	(0.040)	(0.035)	(0.056)	(0.022)	(0.029)	(0.034)	(0.037)
σ_η/σ_u	6.478	15.180	7.866	11.544	8.273	9.448	9.204	9.376	11.576	9.200

Note: Quasi Maximum Likelihood Estimates (QMLE) with robust standard errors in parentheses.

Table 4: Daily summary statistics

	Los Angeles	Boston	Chicago	Denver	Miami	Las Vegas	San Diego	San Francisco	New York	Washington, D.C.
					Panel A: Level					
Mean	177.764	145.561	128.901	118.049	162.896	136.511	164.473	137.702	159.450	170.039
Std. dev.	41.121	13.381	21.631	4.605	48.351	48.568	34.058	27.169	25.877	34.830
					Panel B: Returns					
Mean	0.015	0.008	-0.002	0.003	0.006	-0.006	0.010	0.005	0.011	0.015
Std.dev.	0.347	0.351	0.599	0.303	0.428	0.370	0.387	0.509	0.291	0.502
AR(1)	-0.059	0.047	0.008	-0.018	-0.034	0.061	-0.005	-0.113	0.049	-0.018
LB(10)	67.877	21.935	24.362	16.838	17.742	59.549	15.065	269.509	13.335	24.977
					Panel C: Squared returns					
Mean	0.121	0.123	0.358	0.092	0.183	0.137	0.150	0.259	0.085	0.252
Std. dev.	0.200	0.260	1.269	0.242	0.336	0.369	0.270	0.616	0.170	0.607
AR(1)	0.113	0.102	0.075	0.021	0.107	0.071	0.037	0.042	0.042	0.132
LB(10)	182.307	109.914	102.316	33.414	445.189	85.348	50.715	179.632	53.109	106.434

Note: The table reports summary statistics for each of the ten MSAs for the common June 2001 to September 2012 sample period, for a total of 2,843 daily observations. AR(1) denotes the first order autocorrelation coefficient. LB(10) refers to the Ljung-Box portmanteau test for up to tenth order serial correlation.

Table 5: Daily HAR-X-GARCH models

$$r_{i,t} = c_i + \rho_{i,1} r_{i,t-1} + \rho_{i,5} r_{i,t-5} + \rho_{i,m} r^m_{i,t-1} + b_{i,c} r^m_{c,t-1} + \varepsilon_{i,t}$$
$$\varepsilon_{i,t} | \Omega_{t-1} \sim N(0, h_{i,t})$$
$$h_{i,t} = \omega_i + \kappa_i \varepsilon^2_{i,t-1} + \lambda_i h_{i,t-1}$$

	Los Angeles	Boston	Chicago	Denver	Miami	Las Vegas	San Diego	San Francisco	New York	Washington, D.C.
						Panel A: Mean				
c ($\times 10^{-2}$)	1.710	-0.302	0.094	-0.074	1.152	-0.111	0.240	-0.222	0.908	1.245
	(0.678)	(0.769)	(0.163)	(5.338)	(0.942)	(0.368)	(3.221)	(0.223)	(0.538)	(0.884)
ρ_1	-0.080	0.030	0.005	-0.015	-0.034	0.004	-0.037	-0.094	0.040	0.012
	(0.020)	(0.022)	(0.011)	(0.052)	(0.020)	(0.016)	(0.020)	(0.018)	(0.020)	(0.024)
ρ_5	0.054	0.009	-0.006	0.010	-0.006	0.006	-0.036	0.160	0.004	0.032
	(0.020)	(0.017)	(0.010)	(0.101)	(0.032)	(0.039)	(0.022)	(0.022)	(0.017)	(0.020)
ρ_m	-0.014	-0.014	-0.023	-0.011	-0.008	0.017	-0.013	-0.014	-0.029	-0.035
	(0.007)	(0.005)	(0.007)	(0.008)	(0.006)	(0.004)	(0.006)	(0.006)	(0.006)	(0.007)
b_c	0.059	0.039	0.049	0.020	0.060	0.035	0.060	0.056	0.054	0.084
	(0.009)	(0.007)	(0.008)	(0.018)	(0.008)	(0.007)	(0.010)	(0.009)	(0.006)	(0.010)
R^2	0.039	0.018	0.009	0.007	0.027	0.044	0.030	0.049	0.033	0.027

Table 5: Continued

	Los Angeles	Boston	Chicago	Denver	Miami	Las Vegas	San Diego	San Francisco	New York	Washington, D.C.
					Panel B: Variance					
ω ($\times 10^{-2}$)	0.013	0.230	0.075	0.215	0.016	0.014	0.024	0.023	0.041	0.067
	(0.015)	(0.074)	(0.058)	(0.103)	(0.014)	(0.013)	(0.028)	(0.026)	(0.023)	(0.043)
κ	0.020	0.056	0.056	0.034	0.013	0.017	0.014	0.016	0.026	0.032
	(0.008)	(0.010)	(0.009)	(0.012)	(0.003)	(0.006)	(0.007)	(0.006)	(0.005)	(0.006)
λ	0.979	0.926	0.946	0.943	0.986	0.982	0.985	0.983	0.969	0.965
	(0.009)	(0.012)	(0.009)	(0.017)	(0.002)	(0.006)	(0.008)	(0.007)	(0.006)	(0.007)
$\kappa + \lambda$	0.999	0.982	1.002	0.977	0.999	0.999	0.999	0.999	0.995	0.998
					Panel C: Serial correlation tests					
$\varepsilon_{i,t}$	16.325	10.934	15.178	11.144	8.952	18.086	8.953	25.641	7.133	18.906
	(0.091)	(0.363)	(0.126)	(0.346)	(0.537)	(0.054)	(0.537)	(0.004)	(0.713)	(0.042)
$\varepsilon_{i,t}^2$	92.430	62.011	56.910	22.875	150.471	46.849	41.513	72.156	36.577	36.247
	(0.000)	(0.000)	(0.000)	(0.011)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\varepsilon_{i,t} h_{i,t}^{-1/2}$	11.003	11.878	15.071	14.344	6.576	20.148	7.677	18.762	6.386	12.855
	(0.357)	(0.293)	(0.130)	(0.158)	(0.765)	(0.028)	(0.660)	(0.043)	(0.782)	(0.232)
$\varepsilon_{i,t}^2 h_{i,t}^{-1}$	12.511	24.289	24.616	25.424	9.426	4.946	16.156	40.312	8.650	11.998
	(0.252)	(0.007)	(0.006)	(0.005)	(0.492)	(0.895)	(0.095)	(0.000)	(0.566)	(0.285)

Note: Panel A and B report Quasi Maximum Likelihood Estimates (QMLE) of HAR-X-GARCH models with robust standard errors in parentheses. Panel C reports Wald test statistics for up to tenth order serial correlation in the (squared) residuals and standardized residuals, with corresponding p-values in parentheses.

Table 6: Daily HAR-X-GARCH-CCC correlations

	Los Angeles	Boston	Chicago	Denver	Miami	Las Vegas	San Diego	San Francisco	New York	Washington, D.C.
Los Angeles	1.000	-0.001	-0.004	-0.009	0.049	0.033	0.056	0.177	0.011	0.031
Boston		1.000	-0.013	0.014	0.023	0.025	0.041	0.023	-0.010	0.022
Chicago			1.000	-0.004	0.018	0.006	-0.014	0.036	0.048	0.015
Denver				1.000	0.030	0.024	0.031	-0.004	-0.022	0.023
Miami					1.000	0.017	0.025	0.038	0.038	0.018
Las Vegas						1.000	0.028	0.030	-0.023	0.026
San Diego							1.000	0.054	0.002	0.011
San Francisco								1.000	0.005	0.023
New York									1.000	0.025
Washington, D.C.										1.000

Note: Conditional daily correlations estimated from the HAR-X-GARCH-CCC model.

Table 7: Unconditional correlations of monthly S&P/Case-Shiller index returns

	Los Angeles	Boston	Chicago	Denver	Miami	Las Vegas	San Diego	San Francisco	New York	Washington, D.C.
Los Angeles	1.000	0.651	0.658	0.543	0.870	0.875	0.926	0.835	0.778	0.881
Boston		1.000	0.767	0.749	0.527	0.495	0.672	0.693	0.725	0.773
Chicago			1.000	0.679	0.637	0.544	0.567	0.688	0.818	0.762
Denver				1.000	0.398	0.382	0.545	0.693	0.496	0.666
Miami					1.000	0.799	0.782	0.743	0.795	0.802
Las Vegas						1.000	0.819	0.663	0.684	0.748
San Diego							1.000	0.833	0.712	0.839
San Francisco								1.000	0.659	0.855
New York									1.000	0.816
Washington, D.C.										1.000

Note: The correlations are based on the same June 2001 to September 2012 sample period used in the estimation of the daily HAR-X-GARCH-CCC model.

Table 8: Unconditional return correlations for different return horizons

	Los Angeles	Boston	Chicago	Denver	Miami	Las Vegas	San Diego	San Francisco	New York	Washington, D.C.
					Panel A: Daily (lower triangle) and Weekly (upper triangle)					
Los Angeles	–	0.117 0.065	0.061 0.124	0.066 0.073	0.197 0.219	0.172 0.250	0.198 0.240	0.280 0.309	0.164 0.145	0.156 0.204
Boston	0.017 0.026	–	0.033 0.068	0.068 0.128	0.139 0.130	0.133 0.121	0.143 0.063	0.118 0.054	0.105 0.128	0.120 0.129
Chicago	0.002 0.019	−0.007 −0.001	–	0.025 0.108	0.077 0.149	0.058 0.064	0.049 0.042	0.084 0.148	0.102 0.115	0.068 0.089
Denver	0.001 −0.003	0.023 0.031	−0.002 −0.003	–	0.105 0.100	0.092 0.110	0.100 0.090	0.060 0.106	0.053 0.006	0.084 0.090
Miami	0.072 0.069	0.047 0.043	0.024 0.046	0.044 0.047	–	0.173 0.239	0.178 0.214	0.165 0.176	0.187 0.169	0.150 0.183
Las Vegas	0.060 0.077	0.051 0.049	0.015 0.032	0.038 0.027	0.053 0.054	–	0.165 0.209	0.147 0.162	0.123 0.060	0.142 0.173
San Diego	0.077 0.072	0.059 0.053	−0.006 0.022	0.045 0.042	0.056 0.060	0.058 0.065	–	0.171 0.263	0.148 0.169	0.137 0.127
San Francisco	0.183 0.235	0.037 0.038	0.037 0.065	0.006 −0.003	0.057 0.060	0.052 0.068	0.069 0.066	–	0.138 0.137	0.136 0.151
New York	0.032 0.041	0.011 0.000	0.047 0.061	−0.009 −0.002	0.065 0.063	0.010 −0.002	0.027 0.029	0.024 0.031	–	0.149 0.088
Washington, D.C.	0.047 0.045	0.038 0.034	0.017 0.024	0.032 0.041	0.041 0.038	0.049 0.034	0.033 0.027	0.038 0.038	0.044 0.043	–

Table 8: Continued

	Los Angeles	Boston	Chicago	Denver	Miami	Las Vegas	San Diego	San Francisco	New York	Washington, D.C.
Panel B: Monthly (lower triangle) and Quarterly (upper triangle)										
Los Angeles	–	0.634 0.621	0.530 0.602	0.463 0.506	0.730 0.852	0.600 0.837	0.731 0.906	0.724 0.834	0.759 0.747	0.733 0.856
Boston	0.382 0.348	–	0.451 0.655	0.400 0.559	0.616 0.507	0.533 0.522	0.624 0.673	0.594 0.623	0.643 0.735	0.627 0.688
Chicago	0.266 0.344	0.207 0.320	–	0.323 0.502	0.519 0.612	0.417 0.510	0.513 0.567	0.500 0.667	0.572 0.767	0.532 0.675
Denver	0.251 0.355	0.210 0.254	0.138 0.293	–	0.457 0.370	0.391 0.317	0.454 0.557	0.416 0.625	0.458 0.411	0.456 0.513
Miami	0.493 0.619	0.384 0.277	0.274 0.355	0.271 0.239	–	0.591 0.797	0.696 0.769	0.669 0.754	0.734 0.761	0.697 0.801
Las Vegas	0.395 0.633	0.328 0.322	0.210 0.233	0.229 0.201	0.404 0.547	–	0.589 0.782	0.558 0.657	0.599 0.659	0.582 0.708
San Diego	0.497 0.626	0.388 0.307	0.260 0.276	0.266 0.351	0.468 0.570	0.400 0.497	–	0.678 0.822	0.731 0.711	0.694 0.824
San Francisco	0.511 0.623	0.334 0.288	0.253 0.404	0.216 0.427	0.424 0.527	0.343 0.417	0.435 0.600	–	0.700 0.663	0.677 0.791
New York	0.505 0.478	0.384 0.415	0.318 0.427	0.247 0.149	0.499 0.496	0.383 0.354	0.480 0.430	0.431 0.394	–	0.738 0.761
Washington, D.C.	0.469 0.603	0.366 0.375	0.277 0.385	0.253 0.309	0.444 0.515	0.368 0.444	0.433 0.551	0.414 0.486	0.478 0.437	–

Note: Model-implied correlations are upper numbers and data-based correlations are in smaller font just below. Daily, weekly, monthly and quarterly horizons correspond to 1, 5, 20, 60 days respectively.

Table 9: Diebold-Mariano forecast comparison tests

	Composite	Los Angeles	Boston	Chicago	Denver	Miami	Las Vegas	San Diego	San Francisco	New York	Washington, D.C.
					Panel A: One-day-ahead ($h = 1$)						
Mthly v.s. HAR	9.240	8.337	7.378	10.060	9.845	8.680	9.981	9.929	8.067	8.981	9.142
Interp v.s. HAR	8.707	10.171	7.623	6.242	9.249	11.415	8.569	10.786	7.865	8.609	10.293
Direct v.s. HAR	1.599	1.381	2.943	-0.176	1.224	2.785	0.126	-0.276	3.139	-0.012	2.173
					Panel B: One-week-ahead ($h = 5$)						
Mthly v.s. HAR	4.956	4.458	3.876	5.412	5.126	5.087	6.682	6.581	4.258	5.268	4.981
Interp v.s. HAR	4.071	2.964	4.856	5.466	6.724	5.882	4.501	4.761	5.349	4.588	5.304
Direct v.s. HAR	4.495	1.200	3.580	1.514	1.141	2.669	-0.298	0.768	-0.373	0.562	3.212
					Panel C: Two-weeks-ahead ($h = 10$)						
Mthly v.s. HAR	4.544	2.751	3.799	6.647	4.343	4.078	5.204	5.847	3.453	5.261	4.392
Interp v.s. HAR	4.372	1.478	3.617	4.586	4.042	3.333	2.489	3.598	2.954	2.973	3.798
Direct v.s. HAR	5.668	0.828	3.567	2.640	0.763	2.585	-0.214	1.342	-0.381	0.964	3.563
					Panel D: One-month-ahead ($h = 20$)						
Mthly v.s. HAR	–	–	–	–	–	–	–	–	–	–	–
Interp v.s. HAR	6.762	0.623	3.553	4.117	0.830	2.211	-0.511	1.777	0.941	1.909	4.268
Direct v.s. HAR	–	–	–	–	–	–	–	–	–	–	–

Note: The table reports the Diebold-Mariano test statistics for equal forecast accuracy against the alternative that the HAR forecast outperforms the other three forecasts, Mthly, Interp and Direct. The tests are based on the out-of-sample period from July 2009 to September 2012. The Mthly, Interp and Direct models are all identical when $h = 20$, so only one set of test statistics are reported in Panel D.

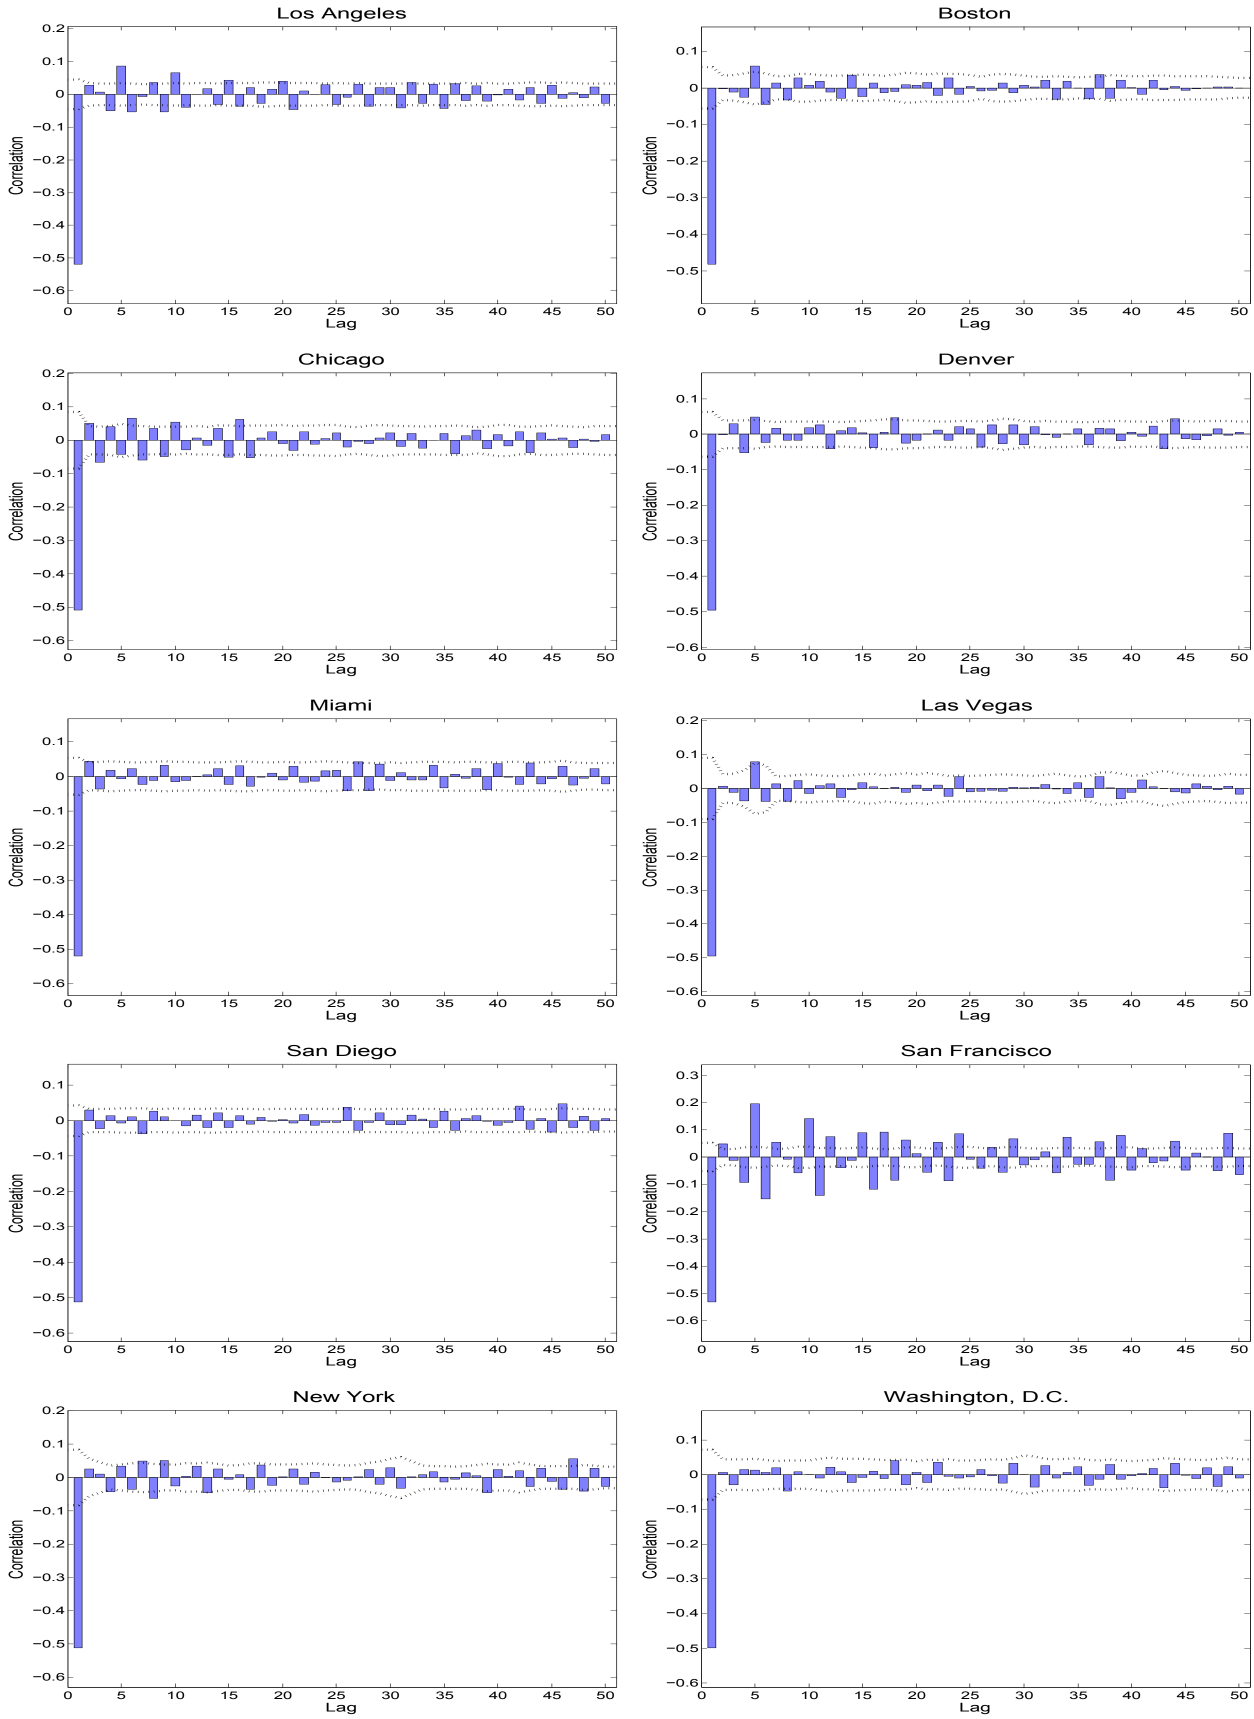

Figure 1: Sample autocorrelations for raw daily index returns

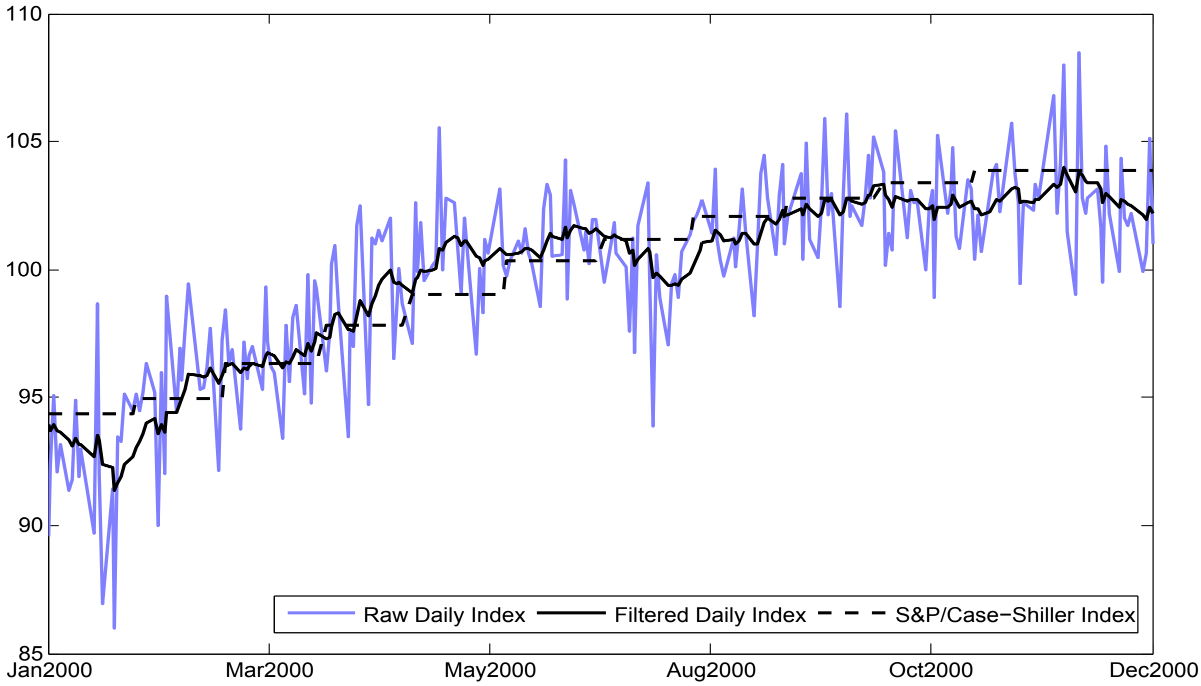

(a) January 3, 2000 to December 29, 2000

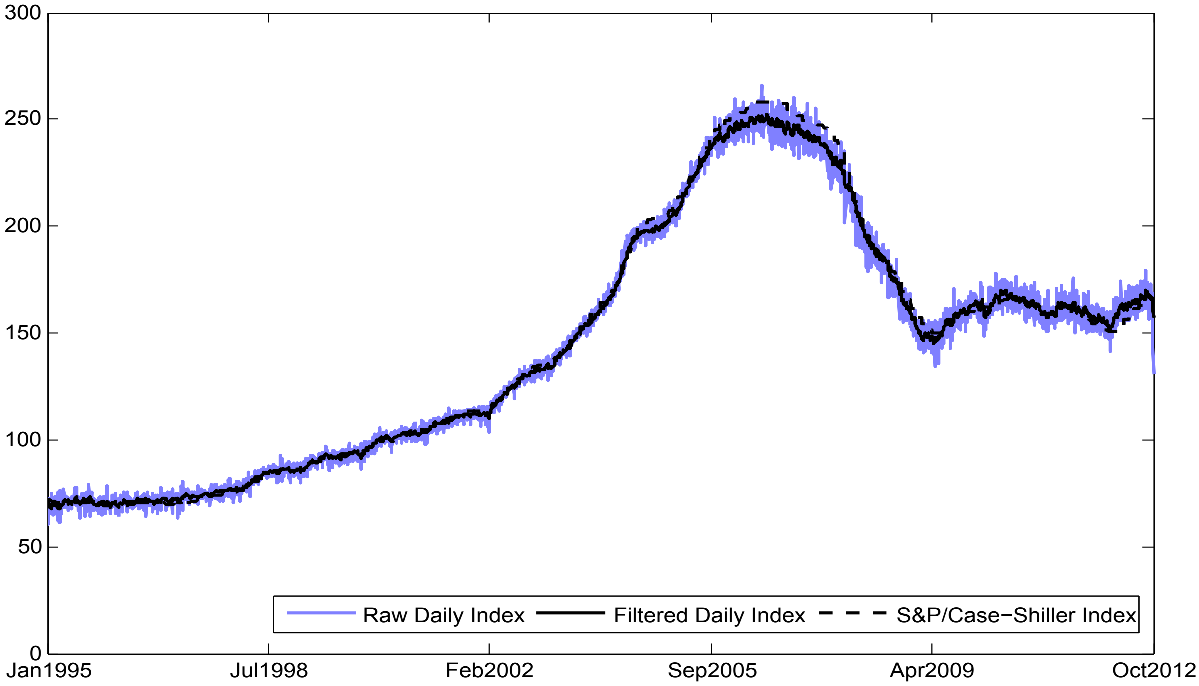

(b) January 3, 1995 to October 23, 2012

Figure 2: Raw and filtered daily house price indexes for Los Angeles

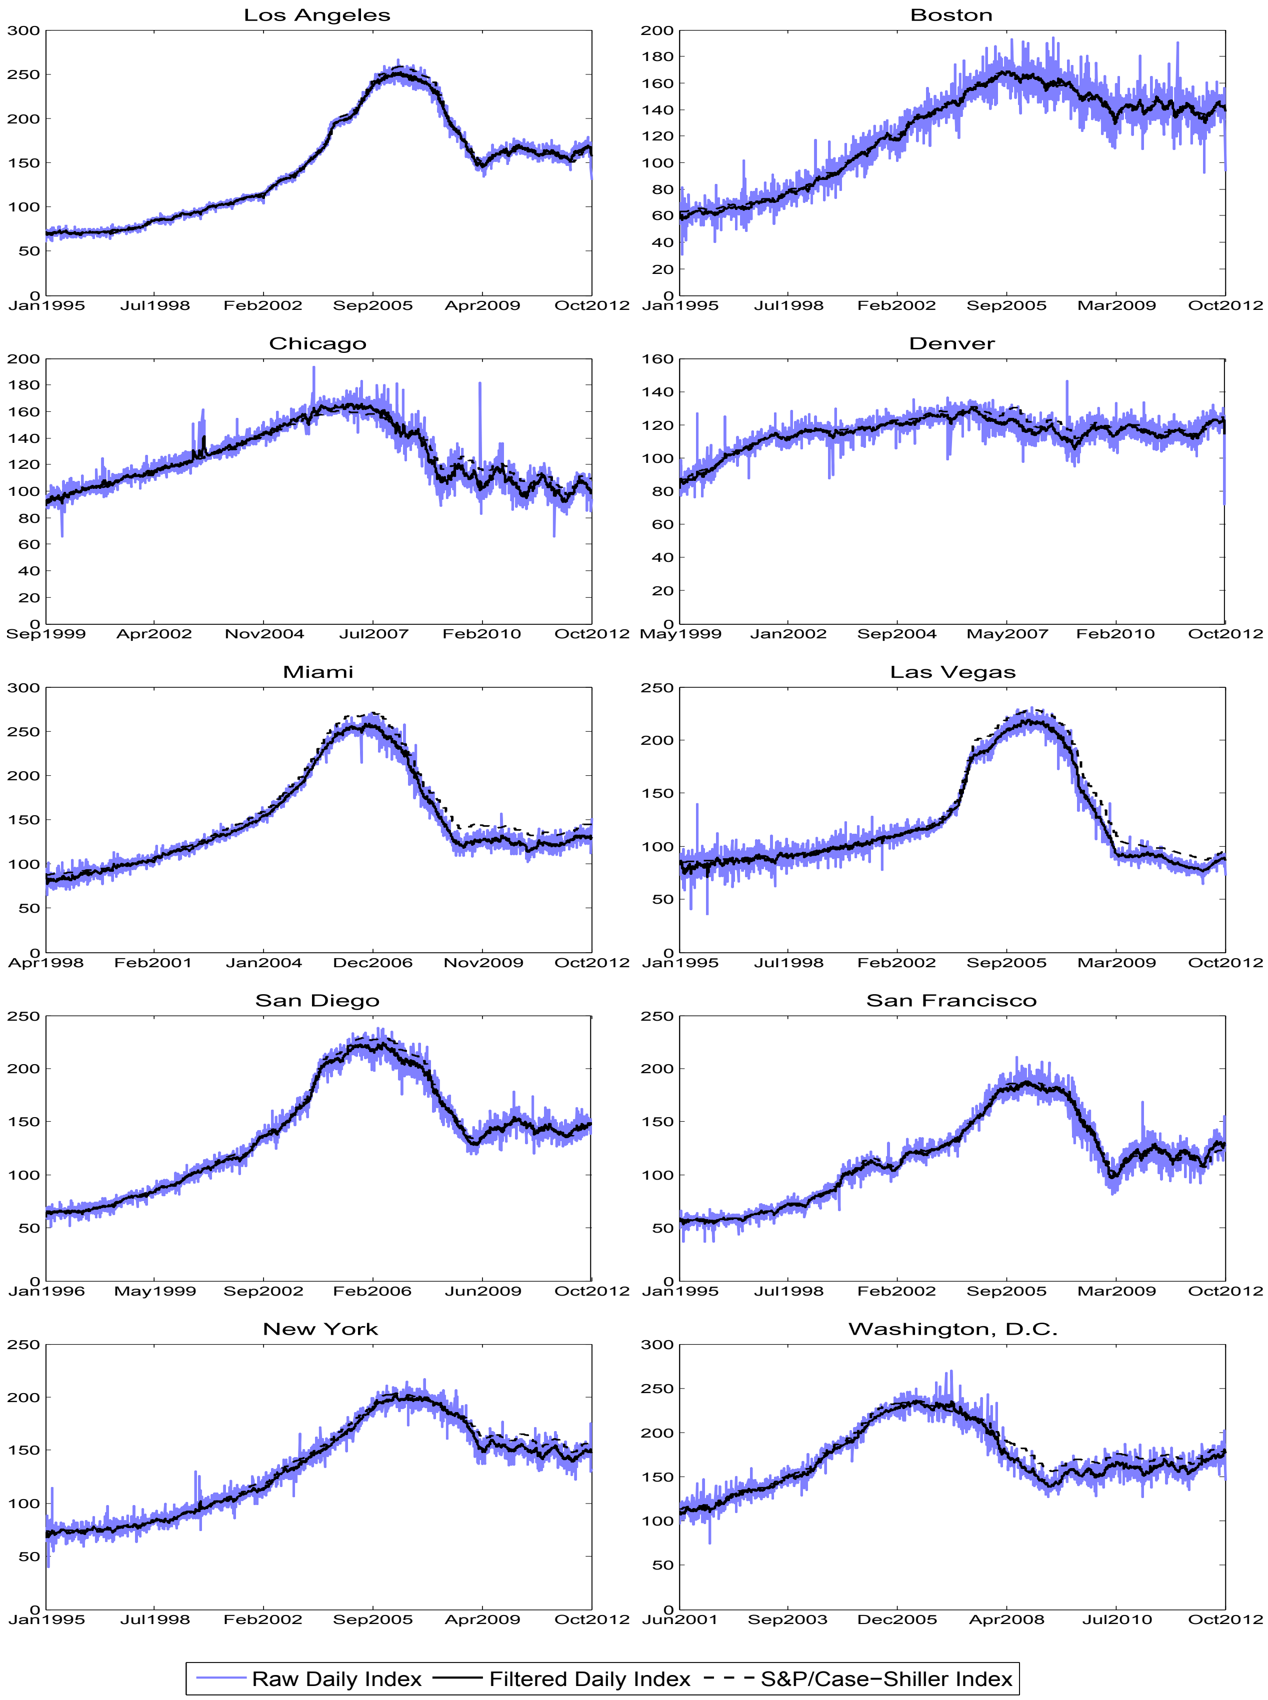

Figure 3: Raw and filtered daily house price indexes for ten MSAs

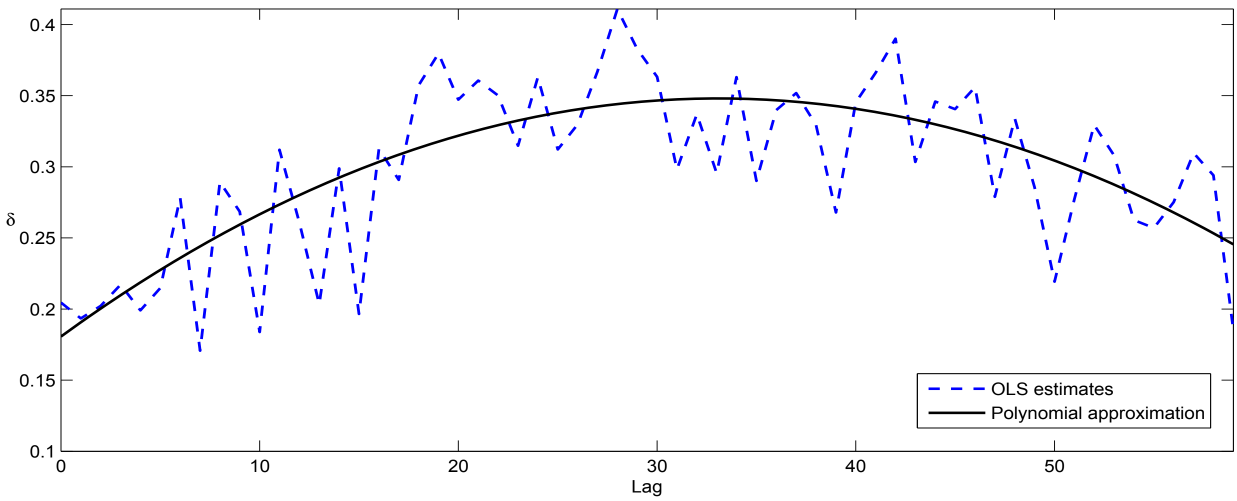

(a) Estimated $\delta(L)$ filter coefficients

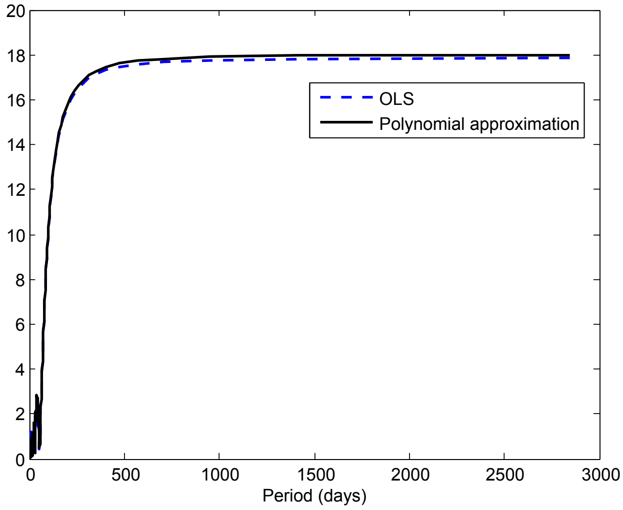

(b) Gain (all periods)

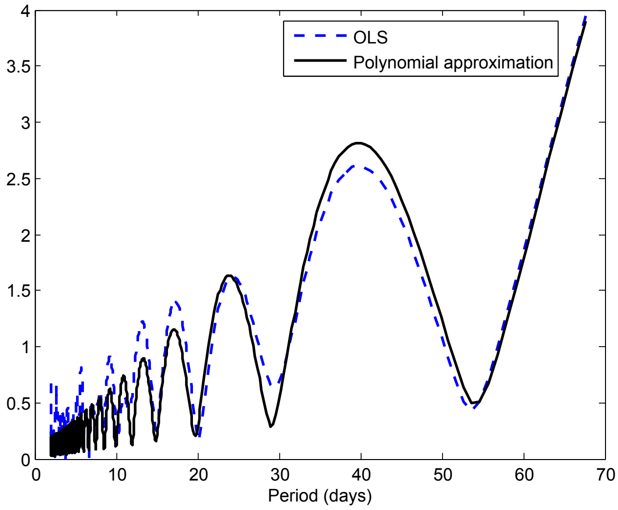

(c) Gain (shorter-run periodicities)

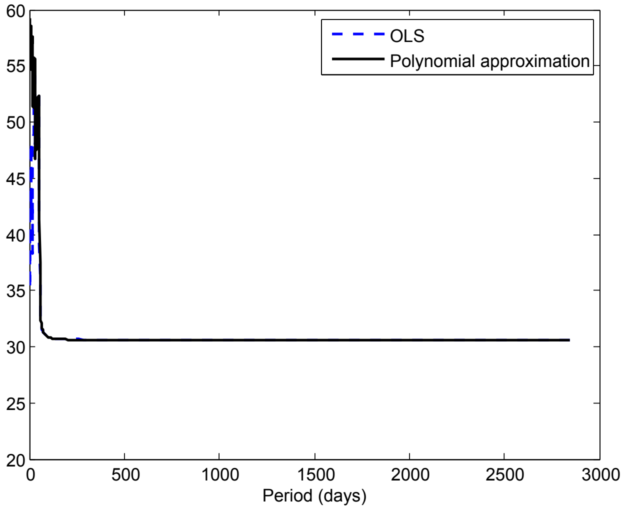

(d) Shift (all periods)

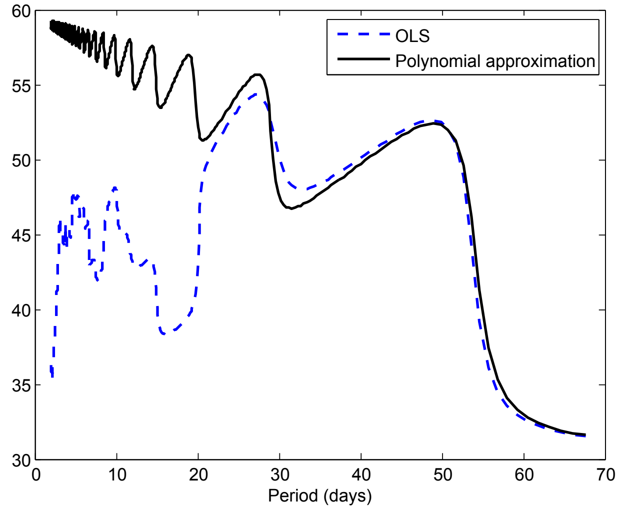

(e) Shift (shorter-run periodicities)

Figure 4: Characteristics of the $\delta(L)$ filter

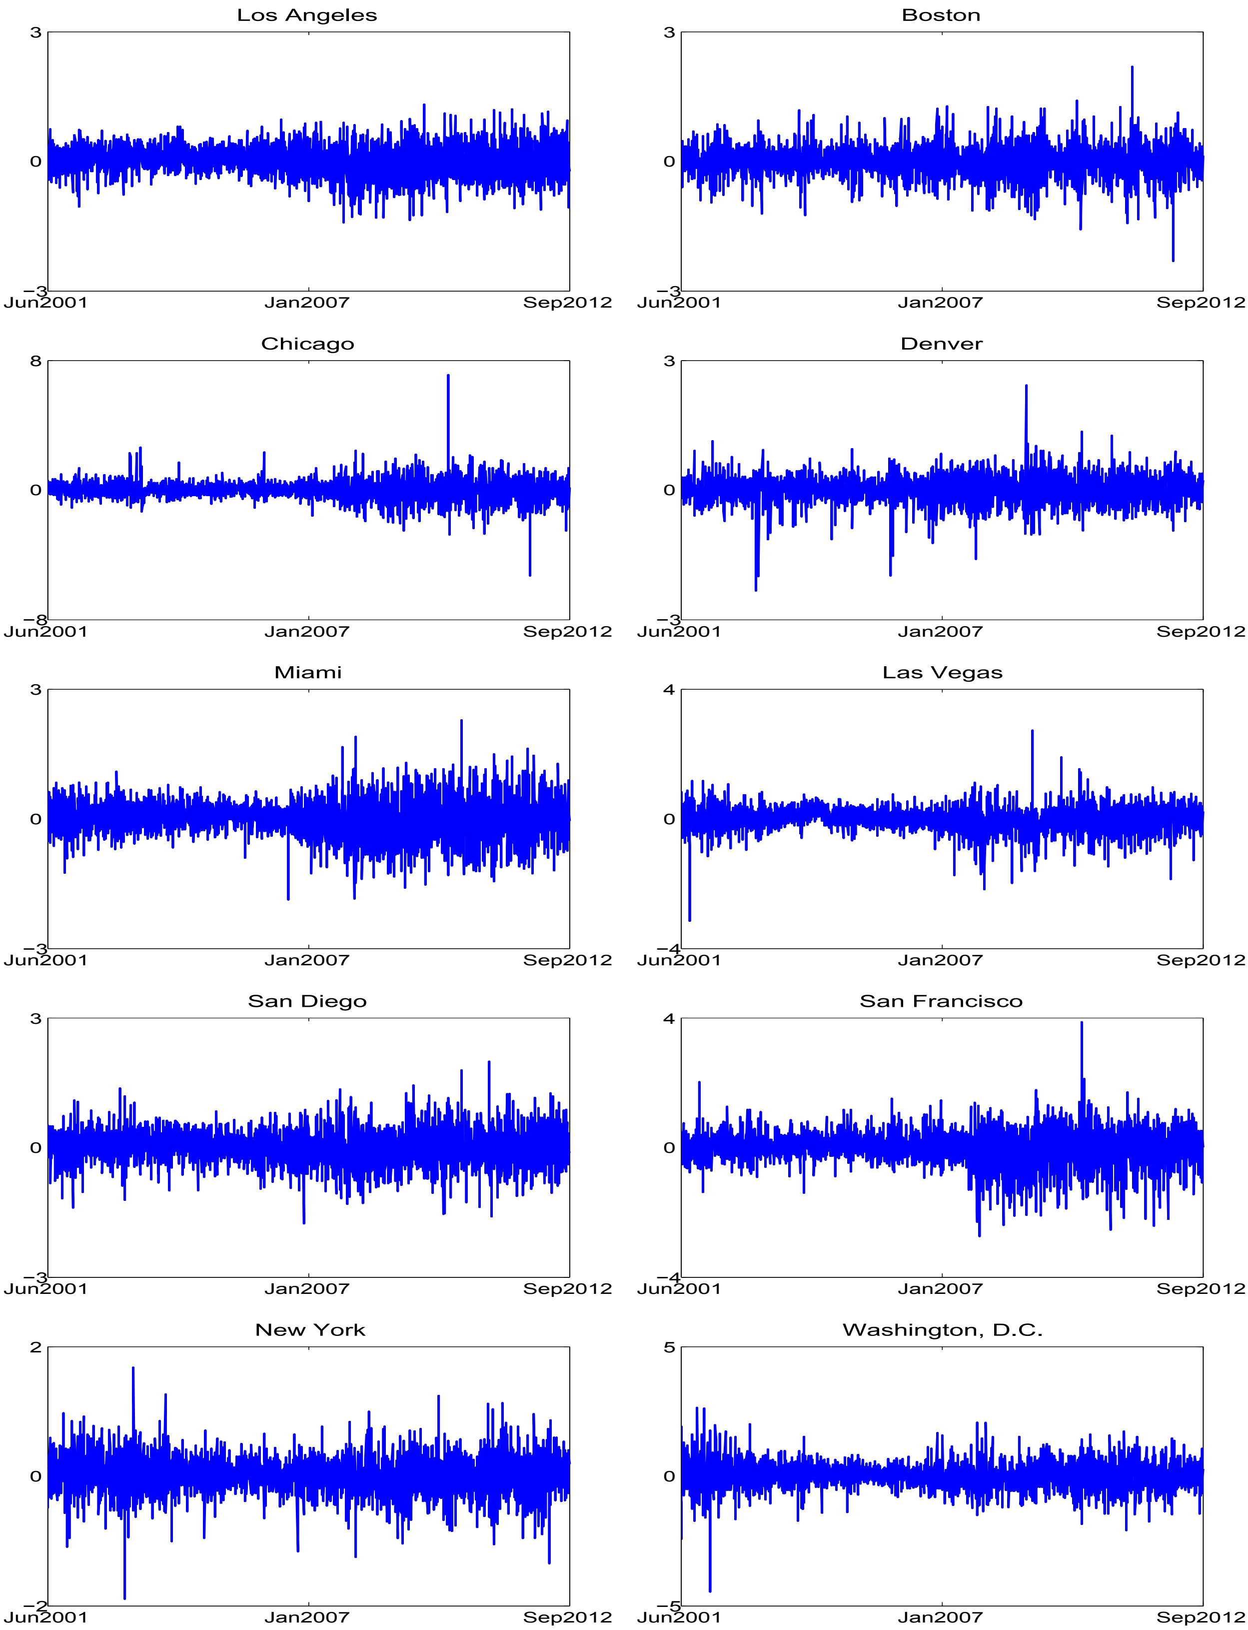

Figure 5: Daily housing returns

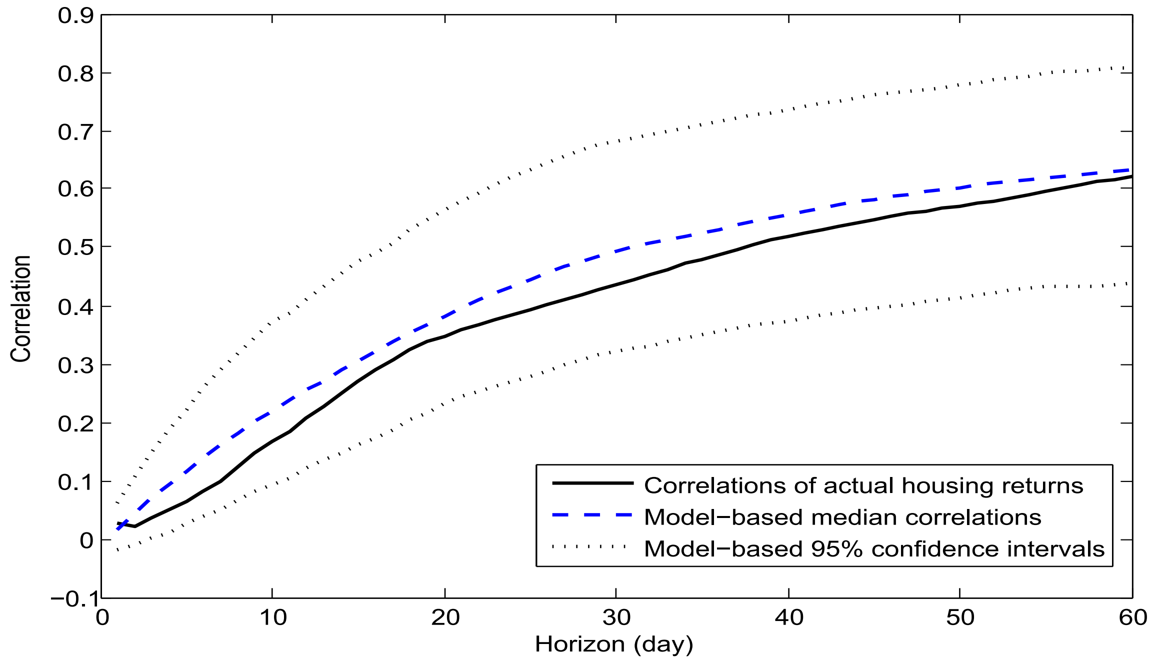

(a) Los Angeles and Boston

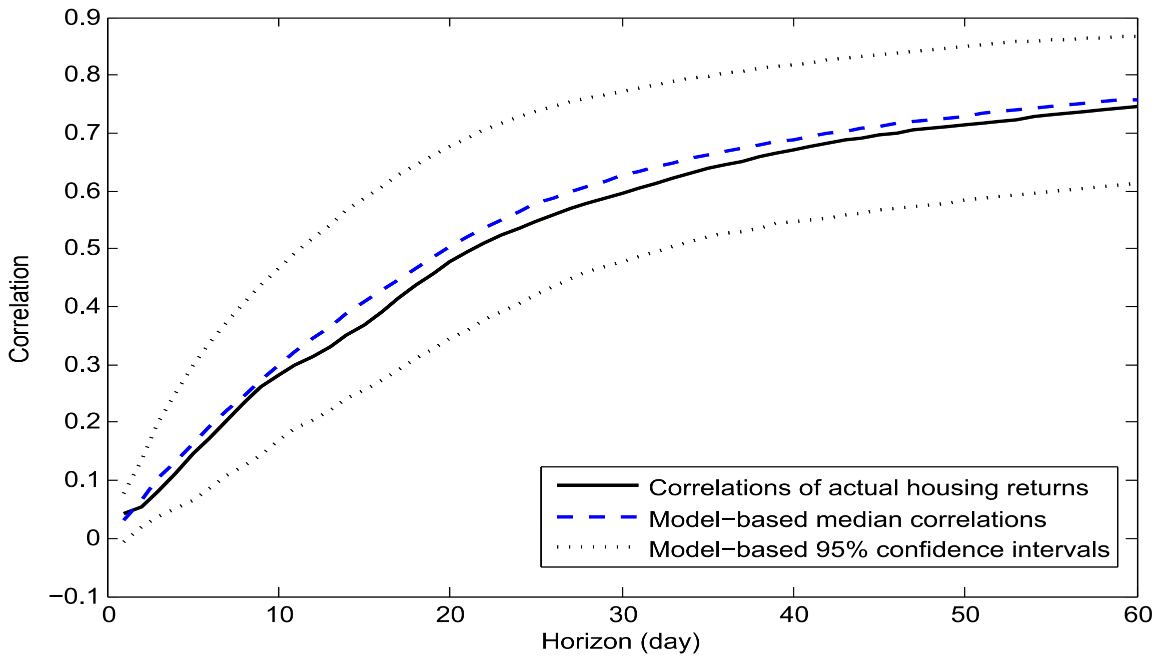

(b) Los Angeles and New York

Figure 6: Unconditional return correlations as a function of return horizon

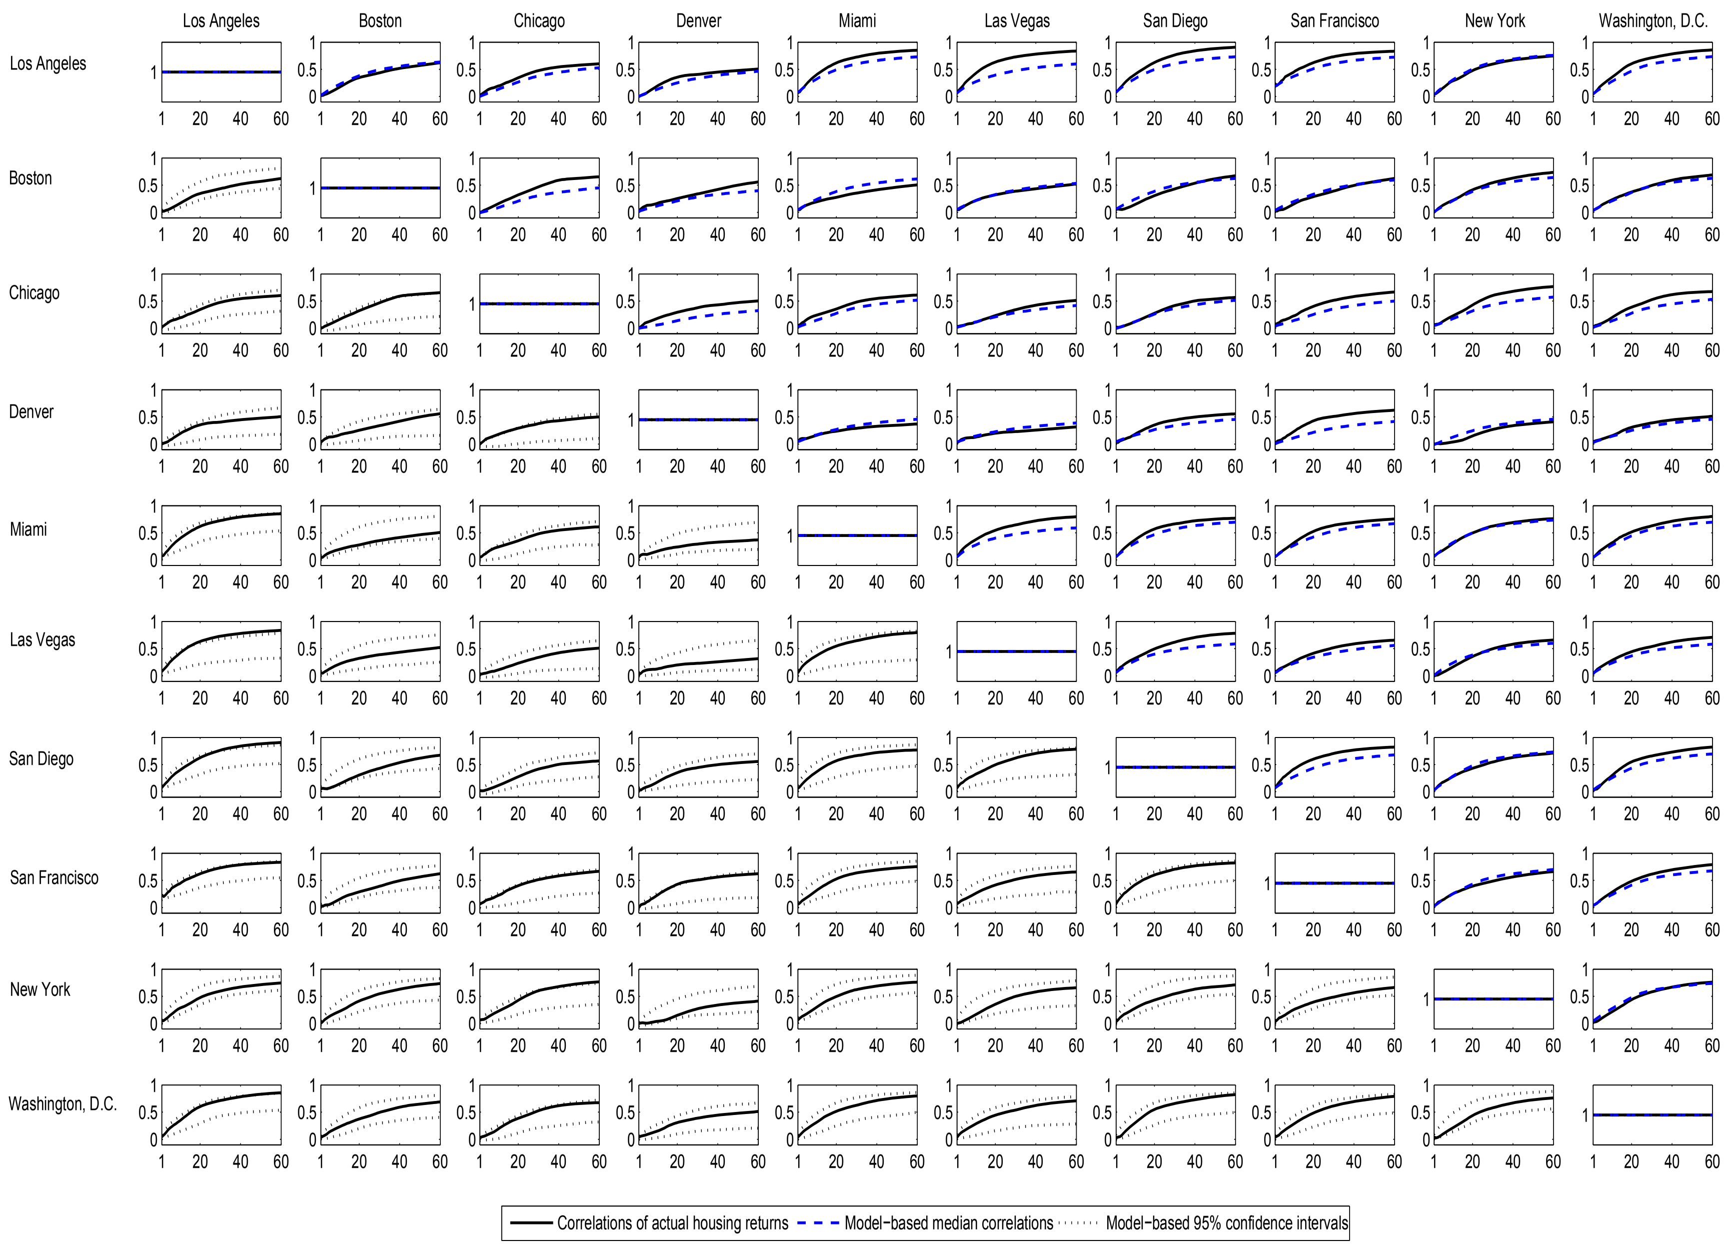

Figure 7: Unconditional return correlations as a function of return horizon

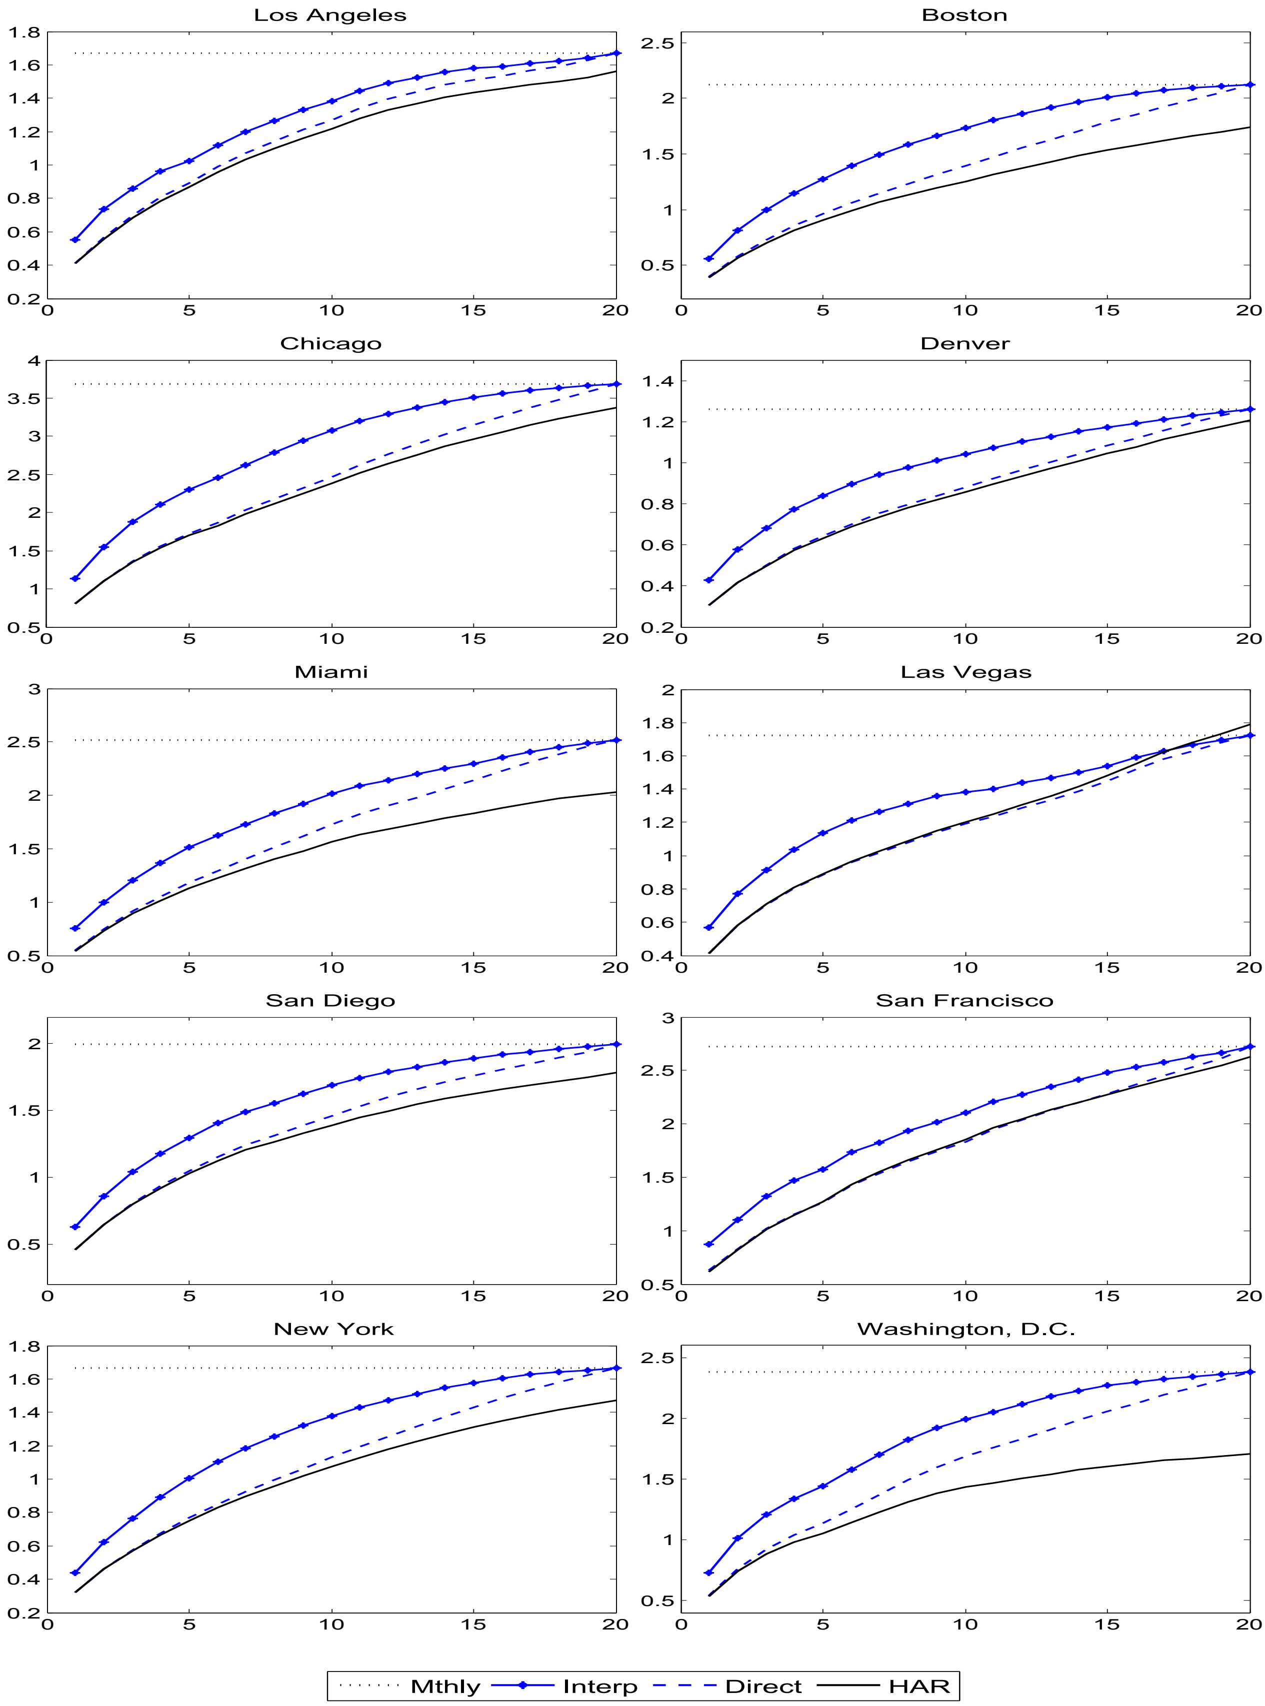

Figure 8: Forecast RMSEs as a function of forecast horizon