Australian National University
Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Nowcasting Transaction-Based House Price Indices Using Web-Scraped Listings and MIDAS Regression

CAMA Working Paper 45/2025
August 2025

Radoslaw Trojanek
Poznan University of Economics and Business

Luke Hartigan
University of Sydney
Centre for Applied Macroeconomic Analysis, ANU

Norbert Pfeifer
University of Graz

Miriam Steurer
University of Graz

Abstract

Timely transaction-based residential property price indices are crucial for effective monetary and macroprudential policy, yet transaction-based data often suffer from significant reporting delays. Online property platforms, by contrast, provide list prices of properties in real-time. This paper examines whether immediately available online list prices can improve timely nowcasts of transaction price movements. Using 16 years of micro-level data from Warsaw and Poznan, we construct quality-adjusted monthly list-price and quarterly transaction-price indices using the hedonic rolling-time-dummy method. We find that list-price indices consistently lead transaction-price indices by one to two months, with the strongest relationship in Warsaw's larger, more liquid market. Building on this lead-lag relationship, we develop a Mixed Data Sampling (MIDAS) regression framework to nowcast quarterly transaction-price growth using monthly list-price data. Our preferred MIDAS specifications reduce one-quarter-ahead root mean square error by approximately 16–23 percent for Warsaw and 5–15 percent for Poznan relative to standard autoregressive benchmarks. The predictive advantage is greatest when incorporating list-price data from the first or second month of the quarter, as third-month data introduce forward-looking noise. Our results show that properly constructed list-price indices can play an important role to provide early housing market signals, potentially enhancing the timeliness of policy responses.

Keywords

MIDAS regression, nowcasting, house price index, hedonic price index, macro-prudential supervision, online price data, rolling time dummy

JEL Classification

C43, E01, E31, R31

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

Nowcasting Transaction-Based House Price Indices Using Web-Scraped Listings and MIDAS Regression

Radoslaw Trojanek[a], Luke Hartigan[b], Norbert Pfeifer[c], and Miriam Steurer*[c]

[a]Poznan University of Economics and Business, Poland
[b]The University of Sydney and Centre for Applied Macroeconomic Analysis, Australia
[c]University of Graz, Austria

July 30, 2025

Abstract

Timely transaction-based residential property price indices are crucial for effective monetary and macroprudential policy, yet transaction-based data often suffer from significant reporting delays. Online property platforms, by contrast, provide list prices of properties in real-time. This paper examines whether immediately available online list prices can improve timely nowcasts of transaction price movements. Using 16 years of micro-level data from Warsaw and Poznan, we construct quality-adjusted monthly list-price and quarterly transaction-price indices using the hedonic rolling-time-dummy method. We find that list-price indices consistently lead transaction-price indices by one to two months, with the strongest relationship in Warsaw's larger, more liquid market. Building on this lead-lag relationship, we develop a Mixed Data Sampling (MIDAS) regression framework to nowcast quarterly transaction-price growth using monthly list-price data. Our preferred MIDAS specifications reduce one-quarter-ahead root mean square error by approximately 16–23 percent for Warsaw and 5–15 percent for Poznan relative to standard autoregressive benchmarks. The predictive advantage is greatest when incorporating list-price data from the first or second month of the quarter, as third-month data introduces forward-looking noise. Our results show that properly constructed list-price indices can play an important role to provide early housing market signals, potentially enhancing the timeliness of policy responses. JEL Codes: C43; E01; E31; R31
Keywords: MIDAS regression, Nowcasting, House price index, Hedonic price index, Macroprudential supervision, Online price data, Rolling Time Dummy

Acknowledgements:
The research was initially funded by the Polish National Agency for Academic Exchange under the Bekker programme (PPN/BEK/2019/1/00292/U/00001). Trojanek acknowledges partial financing by the National Science Centre of Poland under Grant number 2024/55/B/HS4/01541.
Hartigan acknowledges partial funding support provided by the Australian Research Council (DP230100959)

Declarations of interest: none

*Corresponding author: miriam.steurer@uni-graz.at

1 Introduction

Real-time monitoring of housing markets remains a major challenge for central banks and statistical agencies. Although transaction-based property price indices are considered the most reliable indicators of housing market trends, they are often released with considerable delays, sometimes taking several months to over a year due to slow administrative processing of transaction data (European Commission, 2018). As a result, policymakers are left without timely information on house price dynamics, which undermines their ability to respond effectively to risks of market overheating risks or sudden downturns (Leamer, 2007; Banti and Phylaktis, 2019; Choi and Hansz, 2021).

Our paper relates to a growing interest in the use of online data in price statistics. Several studies have shown that consumer price indices (CPIs) constructed with online data can closely approximate official CPIs (Cavallo, 2013; Cavallo and Rigabon, 2016; Cavallo et al., 2018; Bajari et al., 2023).[1] Similarly, the automation of web scraping of price data is now a high priority for the US CPI (National Academies of Sciences, Engineering, and Medicine, 2022). Inspired by these advances, we examine whether online real-estate listings (available in near real time) can similarly improve house price measurement.

However, housing markets are unlike consumer goods markets: each property is unique, prices are subject to negotiation, and there are asymmetric price adjustments. Sellers typically revise asking prices more slowly in downturns due to loss aversion and anchoring (Genesove and Mayer, 2001; Haurin et al., 2013; Shimizu et al., 2016). These behavioral dynamics are one reason why indices derived from list price data do not necessarily reflect actual market price movements (Pollakowski, 1995; Hoekstra et al., 2012; Shimizu et al., 2016; Anenberg and Laufer, 2017; Trojanek, 2018; Lyons, 2019; Wang et al., 2020; Kolbe et al., 2021; Pfeifer and Steurer, 2022). These factors, together with the significant administrative delays in the availability of transaction data, motivate our study of micro-level list and transaction data for two Polish cities, Warsaw and Poznań.

Our study makes three main contributions. First, we compile hedonic price indices using the longest available matched micro-level dataset of list and transaction price data, covering 16 years for Warsaw and Poznań. We construct quality-adjusted monthly list price indices and quarterly transaction price indices using the hedonic rolling-time-dummy method (Shimizu et al., 2010). The equivalence in method as well as market coverage ensures that any differences we find between

[1]During the COVID-19 period, the use of web-scraped price data in national CPIs increased, leading Eurostat to formulate guidelines for countries using web-scraped prices in the Harmonized Index of Consumer Prices (HICP) (Eurostat, 2020).

list and transaction price behavior are not due to differences in either market definition or index methodology.

Second, we provide new evidence on the timing relationship between list and transaction price indices by applying a modified Diewert distance metric (Diewert, 2002, 2009) to quantify the lead–lag structure between the indices. Specifically, we find that list-price indices consistently lead transaction-price indices by about one to two months, with this lead most pronounced in Warsaw, where the housing market is larger and more liquid.

Third, building on these insights, we introduce a Mixed Data Sampling (MIDAS) regression framework, originally proposed by Ghysels et al. (2004, 2007), for nowcasting residential property prices at the city level. We demonstrate that timely list-price indices can provide early and reliable signals of transaction-price dynamics. Our preferred MIDAS specifications reduce one-quarter-ahead RMSE by approximately 16-23 percent for Warsaw and 5-15 percent for Poznań relative to standard time series benchmarks. Our results suggest that incorporating list-price-based nowcasts into central bank and macroprudential monitoring frameworks could enhance policy response timeliness, especially during periods of rapid housing market change.

To our knowledge, only one prior MIDAS application has focused on residential real estate. A recent working paper by Garzoli et al. (2021) applies MIDAS to nowcast Italian national quarterly residential property price indices, combining a broad range of high-frequency market returns and macroeconomic variables, and reports significant RMSE gains over AR(1) and small-scale factor models. In contrast to their national approach and broad indicator set, we rely exclusively on micro-level list and transaction price data at the city level. Our list-price indices control for property quality and capture high-frequency house price signals specific to each urban housing market. By leveraging signals from both list and transaction prices in the same market, our approach provides a consistent and coherent basis for nowcasting final sale prices.

The remainder of the paper is structured as follows: Section 2 details our dataset and cleaning procedures. Section 3 describes the hedonic index methodology and results and presents evidence on the lead–lag relationship between list and transaction indices, using both a modified Diewert distance metric (Diewert, 2002, 2009) and cross-correlations to quantify timing differences. Section 4 evaluates the nowcasting performance of MIDAS models, and Section 5 concludes the paper.

2 Data

2.1 The temporal relationship between list and transaction price data

List price data can be downloaded from online listing portals and thus are available significantly earlier than transaction price data from registry offices. Figure 1 illustrates the availability of list-price data - downloaded in real time from the biggest Polish online portals (www.gratka.pl and www.otodom.pl) — with official transaction records from Polish land registry offices. Figure 1 shows the average weekly duration at each stage of the transaction process for Warsaw between 2008 and 2023.

Figure 1: Timeline of the transaction process for Warsaw

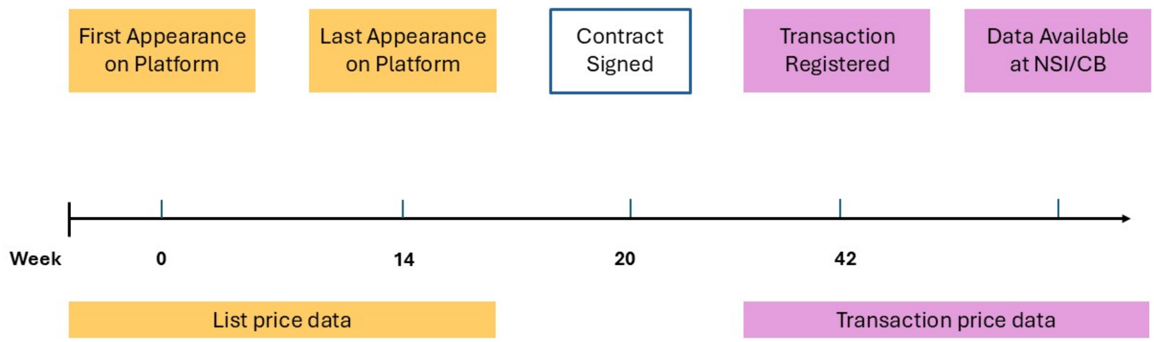

Figure 1 illustrates the typical steps involved in the selling process. The estimated times between the individual steps are based on our calculations based on data for Warsaw between 2008 and 2023. Note, we use the *Last Appearance on Platform* dates and prices as the basis for our list-price indices.

To ensure comparability, we ensured that the geo-location of properties in the list and transaction datasets coincided, thereby eliminating any differences in the results that might be due to regional market variations.

During the 16 year period (2008-2023), 162 015 properties were officially transacted in Warsaw, comparing with over 1.67 million list price entries during the same period (Table A1). Online portals make list prices accessible immediately, whereas land-registry data arrive only after the notarial deed is filed and processed.

Figure 1 shows the median lags at each stage for the Warsaw property market between 2008 and 2023: 6 weeks from negotiation close to notarial signing, and 22 weeks from signing to entry in the property price registry, summing up to a total median delay of 28 weeks. While this median masks year-to-year variation (e.g. higher backlogs during boom periods), the overall lag remains substantial.

2.2 Micro-level List- and Transaction Data for two Polish Cities

Our study is based on a unique long-term property database with micro-level transaction and list price data for two major Polish cities. We have 16 years of transaction and list price data (2008-2023:Q4) for the Polish capital Warsaw, as well as for Poznan (2008 - 2024:Q3), Poland's fifth largest city. This long observation period allows for a detailed analysis of the relationship between list and transaction price data.

The transaction data for both cities is sourced from the official Polish Property Price Register. We include only full-ownership transfers, excluding partial sales (e.g., cooperative rights). Each record provides sale price, unit size (in square meters), address, geospatial location, floor level, room count, building height, construction year, construction technology, garage/parking availability, and basement/storage presence. To enhance the transaction dataset, we integrated supplementary information from the official Polish Cadastre, which adds variables such as building height and construction year. Additionally, we calculate the distance to the city center based on each property's longitude and latitude. A detailed description of how this additional information was collected and incorporated into the transaction records can be found in Trojanek and Huderek-Glapska (2018).

Asking prices and property characteristics were collected through web scraping from Poland's major real estate portals (www.gratka.pl and www.otodom.pl) beginning in 2008 (Trojanek, 2021, 2025). The online descriptions of each property contain a variety of property characteristics, including property location (district, housing estate, and street), floor level, the type of ownership, size in square meters, the building construction technology, the parking facilities and the Quality of apartment.[2] For a detailed description of how the list price dataset was created, please refer to Trojanek (2025).

Following Shimizu et al. (2016) and Pfeifer and Steurer (2022), we retain only the final observed price per property, which aligns more closely with transaction prices and reduces upward bias from multiple downward revisions.[3] Table A1 illustrates that, although this step reduces raw

[2]The quality assessment of the apartments is based on the information provided in the listings and is categorized with values ranging from 1 to 4, where 1 indicates a property in need of renovation, 2 represents low quality, 3 signifies medium quality, and 4 denotes high quality.

[3]Sellers' pricing strategies influence list prices (Beracha and Seiler, 2014), and there is no guarantee that the property will sell, particularly if the stated list price is out of line with current market conditions. Consequently, properties that take a long time to sell are included in the list price dataset every period, while properties that sell very quickly appear only a few times. Without adjustments, this would result in an over-representation of expensive and atypical properties in the list-price dataset. We limit this effect by restricting the list price data to include only the very last appearance of each property on the platform. Since final listings represent sellers' last recorded asking prices, they typically align more closely with actual transaction prices than initial price statements (Shimizu et al., 2016; Anenberg and Laufer, 2017).

listings by roughly 50 percent, the list samples still remain substantially larger than the number of official transactions.

To exclude outliers and ensure comparability between the two data types, we exclude units smaller than 20 m^2 or larger than $300m^2$. We then drop hedonic outliers using Cook's distance regression, applying the usual threshold of $4/(n–k)$ within each rolling window, where n is the window's sample size and k is the number of regressors. Table A1 summarizes sample sizes and mean characteristics before and after cleaning. Notably, mean *price per* m^2 shifts by less than 5 percent across both datasets and cities, indicating that cleaning removes extreme observations without materially altering central tendencies.

Table 1: Summary of micro-dataset before and after cleaning

Warsaw				
	List data		Transaction data	
	Raw	After cleaning	Raw	After cleaning
Mean price/m^2	10,674.86	10,446.08	9,563.96	9,831.57
Mean area	62,38	59,67	53,71	53,42
Mean age	30,85	32,32	34,10	33,85
Observations	1,674,796	760,273	162,015	154,729
Poznan				
Mean price/m^2	6,403.82	6,934.77	6,239.02	6,264.14
Mean area	58,14	56,67	51,41	50,98
Mean age	33,71	35,90	40,83	40,51
Observations	338,164	133,026	50,891	44,384

3 Constructing Hedonic Price Indices from Transaction and List Price Data

3.1 Hedonic Regression

Originally going back to Court (1939) and later popularized by Rosen (1974), hedonic methods are the internationally preferred approach among National Statistical Institutes and international organizations (e.g., Eurostat or the IMF) for constructing quality-adjusted residential property price indices. By contrast, U.S. indices typically rely on repeat-sales methods (e.g. Case-Shiller).

In this section, we estimate hedonic indices for Warsaw and Poznan, following EU guidelines (Eurostat, 2013; European Systemic Risk Board, 2019). A variety of hedonic methods have been proposed (e.g. de Haan 2010; Hill 2013; Eurostat 2013), and—where data are plentiful—they yield nearly identical indices (Hill et al., 2018). In this analysis, we compile our residential property price indices using the hedonic rolling-time-dummy (RTD) method (Shimizu et al., 2010). This method is a straightforward extension of the hedonic time-dummy approach that avoids the need to revise the index as new data become available.

RTD hedonic models incorporate property characteristics and time dummies for consecutive periods with a fixed window length, for example, four quarters. Property characteristics – like location, size, quality, and distance to amenities – help to account for quality differences between properties. The bilateral price index for the most recent period is derived directly from the coefficients of the time dummies. When new data becomes available, the window is rolled forward, and the hedonic model is re-estimated.

The RTD method allows the index provider to select the window length. For smaller datasets, a longer window is preferable because more observations at each estimation improve the precision of the characteristic coefficients. Yet, longer windows may reduce the model's sensitivity to recent market trends, as older transactions remain in the estimation sample (Hill and Trojanek, 2022). Here, we re-estimate the hedonic model every quarter for our transaction price data (using a 4-quarter window) and monthly for our asking price data (using a 12-month window). In addition to the time dummies, we include the following variables in the regressions for the transaction dataset: *size, age, construction technology, garage (yes/no), district*, and *distance to city center*. As we do not have exact geospatial locations for the list price dataset, we cannot include the variables *garage* and *distance to city center* in the list price regressions. However, we can include an additional variable that describes the *quality* of the listed property as described on the platform.

Assuming that the first period in the window is period t, the first step is to estimate a semi-log hedonic model as follows:

$$\ln p_{\tau n} = \sum_{c=1}^{C} \beta_c z_{\tau cn} + \sum_{s=t+1}^{t+m} \delta_s d_{\tau sn} + \varepsilon_{\tau n}, \tag{1}$$

where n indexes the housing transactions that fall in the rolling window, $p_{\tau n}$ the transaction price of property n in period τ (where $t \leqslant \tau \leqslant t+m$), c indexes the set of available characteristics of the transacted properties, and ε is an identically, independently distributed error term with mean zero. The characteristics of the properties are given by $z_{\tau cn}$, while $d_{\tau sn}$ is a dummy variable that equals 1 when $\tau = s$, and zero otherwise.

Each time the RTD model is re-estimated, the updated coefficients allow for the calculation of the price change between the two most recent periods in the rolling window. For example, from (1) we obtain the price index:

$$\frac{P_{t+m}}{P_{t+m-1}} = \frac{\exp(\hat{\delta}^t_{t+m})}{\exp(\hat{\delta}^t_{t+m-1})}, \tag{2}$$

where P_{t+m-1} and P_{t+m} denote the level of the price index in periods $m-1$ and $t+m$, and $\hat{\delta}^t_{t+m-1}$ and $\hat{\delta}^t_{t+m}$ are the estimated values of the last two time dummies when t is the first period in the rolling window. Upon the arrival of new data for period $(t+m+1)$, the estimation window shifts forward by one period, and the model is re-estimated. The overall price index with period 1 as the base is calculated by chaining all the pairwise price indices (each obtained from a separate RTD model covering a different window) together as follows:[4]

$$\frac{P_{t+m}}{P_1} = \prod_{\tau=1}^{t-m} \left[\frac{\exp(\hat{\delta}^\tau_{\tau+m})}{\exp(\hat{\delta}^\tau_{\tau+m-1})} \right]. \tag{3}$$

We use a window length of 4 quarters for the quarterly transaction price indices and a window length of 12 months for the monthly list price indices. The same estimation procedure is applied to both Warsaw and Poznan. For the transaction-based indices, we include time dummies, the size of the property, the building age, the construction technology of the building, whether or not it includes a garage, and the distance to the city center. For list-price-based indices, we also account for the quality of the apartment but do not include the garage and distance to the city center. Table A1 and Table A2 in subsection Appendix B provide regression results for the last window of both indices and both cities. All included characteristics are highly significant, and the R^2 lies between 0.7 and 0.9. Figures A1–A3 depict the aggregate price indices for both cities and both types of data that we derive from these regression results. These figures show strong evidence of co-movement between aggregated list and transaction price indices in both cities (Pearson's $\rho > 0.98$) when they are both expressed at the same frequency (either quarterly or monthly). This high correlation underscores that list prices and transactions track each other closely once aligned in frequency, but it tells us nothing about their lead-lag relationship.

To better understand the timing of market information, we next examine the similarity in timing of changes between the two types of price indices, that is, their similarity in terms of growth rates. Figure 2 and Figure A3 illustrate the growth rates of the monthly list- and quarterly transaction price indices for each city. Due to their different frequencies, it is not easy to see if the list price index leads the transaction price index (and by how much).

[4] Chaining is necessary to maintain consistency over time despite rolling estimation windows.

Figure 2: Monthly List- and Quarterly Transaction Price Changes for Warsaw

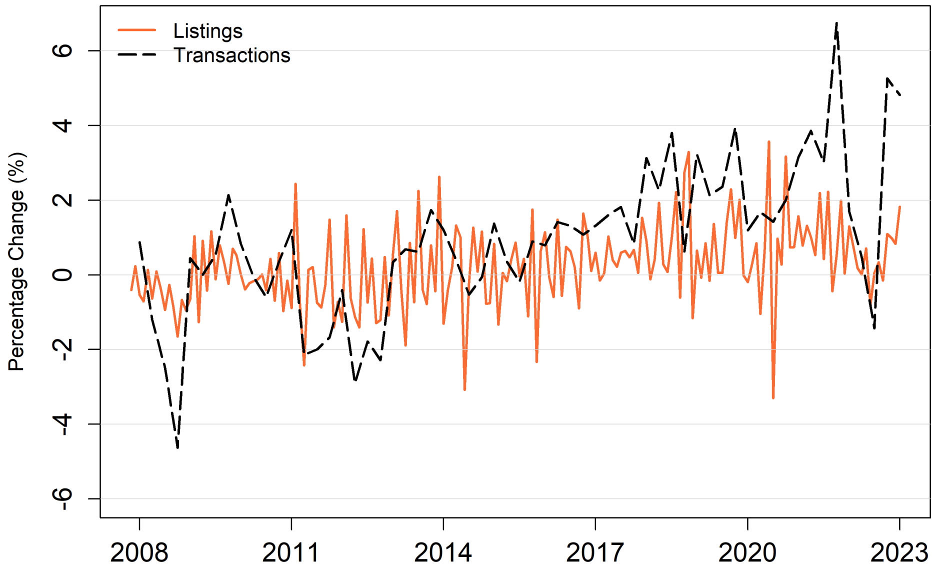

Figure 2 illustrates the monthly listing price and the quarterly transaction price growth rates for Warsaw between 2008 and 2023.

To make the relative movements of the transaction- and list price series better visible, we divide the monthly list-price series up into three separate (quarterly) series. We define series P_t^{L,m_1} as a quarterly version of the monthly list price index that takes its quarterly values from the aggregate price levels of the first month of each quarter. Similarly, P_t^{L,m_2} represents a quarterly list price index series consisting of the aggregate price levels at the second month in each quarter, and P_t^{L,m_3} of the values of the third month in each quarter.

To measure the potential lead or lag between list and transaction price indices, we apply a modified version of a distance metric proposed by Diewert (2002, 2009) for measuring the similarity of price vectors across countries. This Diewert-Metric (DWM) approach was first applied to identify price series lags in Hill et al. (2024). If defined in growth rates, as we do here, the DWM distance has a number of desirable axiomatic properties (Diewert, 2002; Hill et al., 2024). First, the metric treats the two indices symmetrically, and only takes the value 0 when the (lagged) period-to-period price changes are identical. Second, as we apply the DWM distance to growth rates, it is invariant to rescaling of either index.

We compare each of the shortened list price indices (P_t^{L,m_i}) with the quarterly transaction price index (P_t^{trans}). We then assessed each sub-series with the quarterly index using the following criteria:

Define growth factors $g_t = P_t/P_{t-1}$ for each series, then estimate

$$\mathrm{DWM_{gr}}(g^{\mathrm{L,m_i}}, g^{\mathrm{trans}}) = \frac{1}{n}\sum_t \Big(\frac{g_t^{\mathrm{L,m_i}}}{g_t^{trans}} + \frac{g_t^{trans}}{g_t^{\mathrm{L,m_i}}} - 2\Big),$$

and select the lag that minimizes the average proportional gap in growth rates.

In addition to the DWM, we also estimate the Pearson's ρ coefficient between each sub-series and the quarterly index (also shown in Table 2). Both DWM and Pearson's correlation coefficient reach the same conclusion: The best prediction of quarter-on-quarter transaction-price growth, occurs with the P_t^{L,m_1} list price series for Warsaw and the P_t^{L,m_2} list price series for Poznan. This strongly suggests that the list price series lead the quarterly transaction price series by one to two months. We attribute this to the six-week median delay from final listing to contract signature in Warsaw (Figure 1). Our result supports similar findings from the literature that also suggest that list prices lead transaction prices (Agarwal et al., 2014; Anenberg and Laufer, 2017; Lyons, 2019). Table 2 shows the similarity between the investigated price index series according to the DWM and Pearson's correlation, while Figures 3 and A4 illustrate the best alignments indicated by the DWM graphically.

Table 2: Optimal similarity between the growth rates of P_t^{trans} and P_t^{L,m_i}

City	Similarity Metric	month 1	month 2	month 3
Warsaw	DWM	**0.00028**	0.00032	0.00032
	Pearson	**0.677**	0.659	0.649
Poznan	DWM	0.00052	**0.00042**	0.00062
	Pearson	0.452	**0.549**	0.449

The results presented in Table 2 indicate two key findings. First, both the DWM and Pearson correlation coefficients show that the list and transaction price growth rates are more closely aligned in Warsaw compared to Poznan. Second, both metrics suggest that in Warsaw, the growth rates of the list price index for the first month align most closely with the quarterly growth rates of the transaction price index. In contrast, in Poznan, the second-month growth rates have the strongest alignment.

Our results reveal two key findings. First, both the DWM and Pearson correlation coefficients indicate that the growth rates of listing prices and transaction prices are more closely aligned in Warsaw than in Poznan. We believe this alignment is due to the more efficient property market in Warsaw, which has a higher number of transactions per period. Second, both metrics suggest that in Warsaw, the growth rates of the listing price index during the first month are most closely aligned with the quarterly growth rates of the transaction price index. In contrast, in Poznan, the strongest alignment occurs with the growth rates from the second month. This difference may again be attributed to the "bigger market effect" observed in Warsaw.

Figure 3: Warsaw: best list price alignment series according to DWM criterion

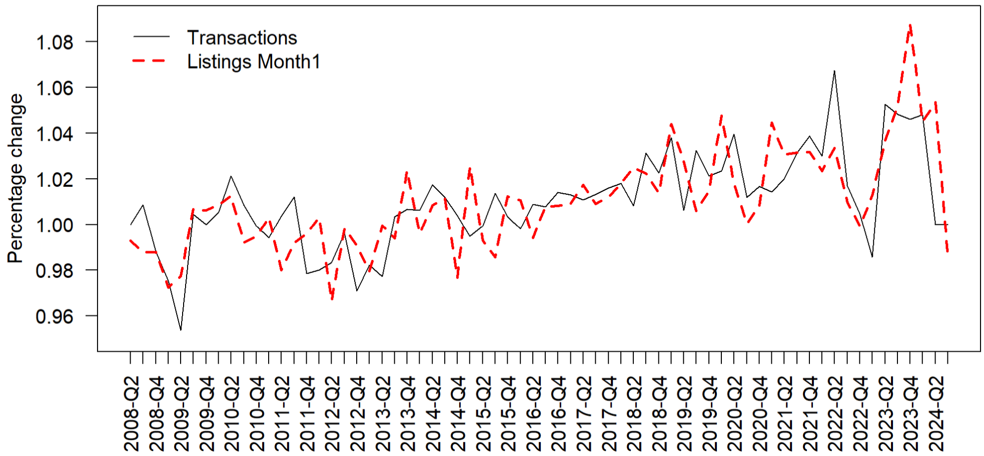

The results presented in Figure 3 illustrates the strong alignment between the growth rates of P_t^{trans} and P_t^{L,m_1} for Warsaw.

Our analysis indicates that the list-price and transaction-price indices not only move in the same direction but also closely align in terms of their growth rates over our sample period. This close alignment is partly due to our methodological choice to only consider each property's final listed price. By doing this, we minimize distortions from multiple downward revisions and avoid giving undue weight to 'hard-to-sell' properties that have unrealistic price expectations (see Shimizu et al. 2016 and Pfeifer and Steurer 2022). Equally important, Poland's market exhibited uninterrupted positive growth from 2008 to 2023. As sellers generally raise asking prices more readily than they lower them, we expect smaller gaps between asking and selling prices in such times of positive price movements. However, the literature indicates that during downturns, pronounced asymmetries can arise. Sellers' loss aversion and anchoring biases can delay downward asking price movement, leading to larger gaps between list and sale prices when markets weaken (e.g., Genesove and Mayer 2001, Haurin et al. 2013, Shimizu et al. 2016). Therefore, list-price indices cannot fully replace transaction-based measures, as they reflect both strategic pricing behavior and actual market clearing.

Instead, list prices serve as a complementary signal: they are available in real time and can anticipate transaction trends. Unlike transaction price indices, which are often released only quarterly and with an additional lag, we can estimate our list-price index monthly (and potentially even weekly as new offers appear). This higher frequency and faster availability should help detect turning points earlier. Naturally, higher frequency introduces more noise. Yet, by estimating hedonic indices on a rolling window (Shimizu et al. 2010), we effectively smooth small-sample volatility and preserve a robust signal.

In the next section, we leverage the greater timeliness of list price indices within a MIDAS nowcasting framework (Ghysels et al. 2004; Foroni et al. 2015) that directly links list-price indices to official transaction-price indices. This approach allows us to produce timely estimates of transaction-price growth rates ahead of official releases.

4 Predicting Quarterly Transaction Price Indices Using Monthly List Price Indices

We will now build on the finding that list-price indices lead transaction-price indices by one to two months (Section 3) and develop a MIDAS regression framework to nowcast quarterly transaction-price growth using monthly hedonic list-price indices.

4.1 Modelling mixed frequency data

Standard practice for modeling mixed-frequency data involves converting all series to the lowest frequency through temporal aggregation, typically by averaging the high-frequency variable within each low-frequency period (e.g., monthly values averaged to quarterly values). While simple to implement, this approach is wasteful because it throws out potentially valuable information about the timing of movements in the high-frequency series.

MIDAS regression offers a flexible way to model mixed-frequency data by directly incorporating high-frequency variables into low-frequency equations without resorting to simple temporal aggregation. The MIDAS approach applies distributed lag polynomials to summarize the effects of multiple high-frequency lags through a small number of parameters. This design helps avoid overfitting that would result from including many unrestricted lags while still capturing the timing structure of the high-frequency predictors. In our application, we consider normalized exponential Almon and beta lag functions, both of which allow for flexible shapes in the lag weights. Because these functions are nonlinear, the model is estimated using nonlinear least squares.[5]

A standard MIDAS model incorporating a single regressor and a single lag of the dependent variable is given by:

$$y_t = \alpha + \phi y_{t-1} + \beta W\left(L^{1/m};\theta\right) x_t^m + \epsilon_t \tag{4}$$

[5]We estimate our MIDAS models using the the BFGS optimization algorithm.

where $W\left(L^{1/m};\theta\right) = \sum_{k-0}^{K-1} W\left(k;\theta\right) L^{k/m}$; and $L^{1/m}$ is a high-frequency lag operator such that $L^{1/m}x_t^m = x_{t-1/m}^m$ with m indicating the higher sampling frequency of the explanatory variable. For example, $m = 3$ when x is monthly and y is quarterly.

The intercept is specified by α while ϕ reflects the degree of persistence in y and helps account for serial correlation in y. The coefficient β captures the overall effect of the high–frequency variable x on y and can be identified by normalizing the function $W\left(L^{1/m};\theta\right)$ to sum to one.

We define the variable x as the growth in the monthly list-price index, while the variable y represents the growth in the quarterly transaction-price index. We assume the residuals, denoted as ϵ_t form an independent and identically distributed sequence with zero mean and constant variance. Further, K represents the maximum lag length for the included high-frequency regressor.

The MIDAS model achieves parsimony by summarizing the high-frequency lag structure through $W(k;\theta)$, a set of weights defined by a small number of parameters $\theta = \{\theta_0, \theta_1, \ldots, \theta_j\}$ where $j \ll K$. This prevents overfitting when incorporating many lags. We consider two widely used functional forms for $W(k;\theta)$: the normalized exponential Almon lag function (Ghysels et al., 2004):

$$W\left(k;\theta\right) = \theta_0 \frac{\exp\left(\theta_1 k + \theta_2 k^2 + \cdots + \theta_j k^j\right)}{\sum_{k=1}^{K} \exp\left(\theta_1 k + \theta_2 k^2 + \cdots + \theta_j k^j\right)} \tag{5}$$

And the normalized beta function of Ghysels et al. (2007) given as:

$$W\left(k;\theta\right) = \theta_0 \frac{\left(\frac{k-1}{K-1}\right)^{\theta_1 - 1} \left(1 - \frac{k-1}{K-1}\right)^{\theta_2 - 1}}{\sum_{k=1}^{K} \left(\frac{k-1}{K-1}\right)^{\theta_1 - 1} \left(1 - \frac{k-1}{K-1}\right)^{\theta_2 - 1}} \tag{6}$$

These functions allow lag weights to capture flexible patterns, such as monotonically declining, hump-shaped, or plateau effects, depending on the data. Figure 4 illustrates the range of shapes that these polynomials can represent. Note, the x-axis specifies the lags of the high-frequency variable while the y-axis specifies the coefficient value for each high-frequency lag term. In our model selection, we compare these forms to determine which provides the best fit for our housing price data. We selected the normalized exponential Almon lag function because it provided a better fit to our data according to BIC.

Figure 4: MIDAS Polynomial Weighting Functions

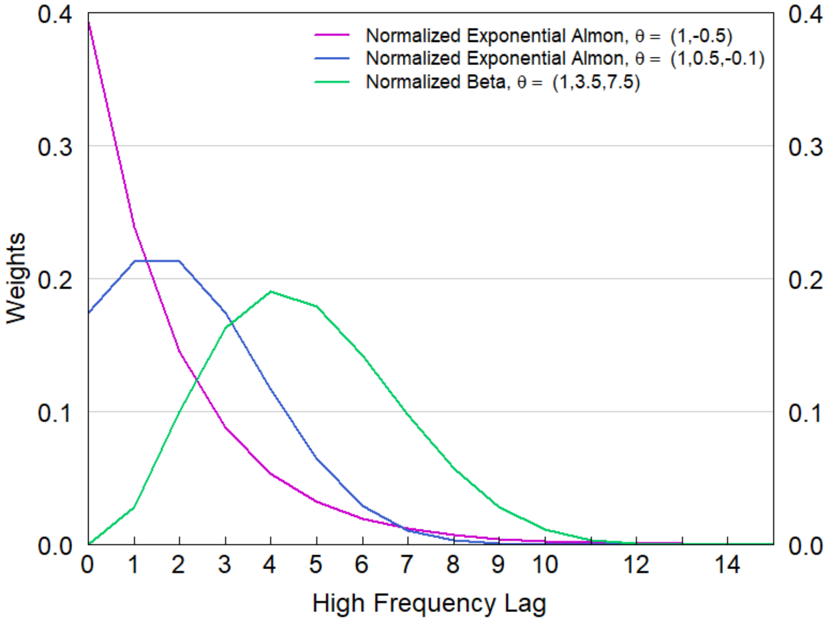

Figure 4 illustrates example polynomial weighting schemes generated by the normalized exponential Almon and normalized Beta functions for different parameterizations.

An alternative specification proposed by Foroni et al. (2015) is 'unrestricted MIDAS' (U-MIDAS). This method leaves the high-frequency lag coefficients unconstrained and can be estimated by OLS. This allows for greater flexibility in how the weights used in temporal aggregation are determined by the data. A U-MIDAS model with one explanatory variable and one lag of the dependent variable is given as:

$$y_t = \alpha + \phi y_{t-1} + B\left(L^{1/m}\right) x_t^m + \epsilon_t \tag{7}$$

where $B\left(L^{1/m}\right) = \sum_{k=0}^{K-1} \beta_k L^{k/m}$ and $L^{1/m}$ and the other parameters are defined as before.

Before we can specify a MIDAS model for nowcasting the quarterly growth in transaction price indices for both cities, we need to decide on two aspects about the model we intend to use. First, the functional constraints to implement. Second, the optimal maximum lag order K. We follow Foroni et al. (2015) and use the Bayesian Information Criterion (BIC) to evaluate a range of restricted and unrestricted MIDAS models. For the R-MIDAS models, we specify the normalized exponential Almon weighting function with $j = 2$ and $j = 3$ parameters. We also consider the normalized beta weighting function with $j = 3$ parameters. For all MIDAS model specifications, we consider four values for the maximum lag of the monthly explanatory variable x (i.e. $K \in$

$\{2,3,4,5\}$).[6] The results for Warsaw are presented in Table 3 while those for Poznan are in Table A3 in the Appendix.

For Warsaw, the BIC strongly prefers the R-MIDAS specification using the normalized exponential Almon weighting function with two parameters ($j = 2$) and maximum lag $K = 6$. For Poznan, the BIC also prefers the R-MIDAS model using the normalized exponential Almon weighting function with three parameters ($j = 3$) instead but the same maximum lag $K = 6$. Although, with Poznan, there were differences between which weighting function and number of parameters were selected across lags unlike with Warsaw, it always favored the normalized exponential Almon weighting function with three parameters.

The U-MIDAS model was found to be not optimal in either city. This is surprising since previous research indicates that U-MIDAS models perform comparably to R-MIDAS models when there is a modest frequency mismatch between the series, such as with monthly and quarterly data (Foroni et al., 2015). This result suggests that the functional constraints of the R-MIDAS models align well with our data.

Table 3: Warsaw: MIDAS Model Comparison

	Normalised Exponential Almon				Normalised Beta		Unrestricted	
	$j = 2$		$j = 3$		$j = 3$		MIDAS	
Lag	BIC	p-value	BIC	p-value	BIC	p-value	BIC	p-value
0:2	**236.21**	**0.54**	239.90	0.00	252.33	0.00	239.90	–
0:3	**228.76**	**0.71**	232.14	0.93	250.62	0.00	236.28	–
0:4	**231.93**	**0.28**	232.13	0.66	243.60	0.00	239.82	–
0:5	<u>**225.89**</u>	<u>**0.25**</u>	227.78	0.27	241.91	0.00	238.03	–

Table 3 shows the results of the in-sample model selection using the BIC. Bold values denote best model per lag. A bold and underline value denote best overall model. The best model is an normalized exponential Almon MIDAS with up to six lags. Even though we only have monthly updates, the unrestricted MIDAS is never optimal in our case. The p-values shows the results of the test of the null hypothesis that the restrictions on the MIDAS regression coefficients implied by the polynomial weighting function are valid and supported by the data. Failure to reject the null implies the functional restrictions are supported by the data. We find these restrictions are valid for our preferred model.

In addition to deciding on the most appropriate models based on the BIC, it is also possible to test the empirical adequacy of the polynomial weighting functions used with the R-MIDAS specifications under standard assumptions via a Wald-type test. The null hypothesis is that the functional restrictions are valid. Therefore, rejecting the null implies that the data do not support

[6] All estimation was done in R using the 'midasr' package of Ghysels et al. (2016).

the functional restrictions. By this test, both selected models are consistent with the data.[7]

Figure 5 shows the fitted values from the optimal R-MIDAS models for the quarterly growth in transaction prices in Warsaw (blue line) compared to the actual series (black line). The general fit for Warsaw seems reasonably good, with an R-squared statistic of 0.7. However, the model is unable to adequately capture the sharp increase in transaction prices prior to the COVID-19 crisis or the more recent moderation in growth at the end of the sample, where it overestimates by a significant amount. This result likely reflects the limits of reduced-form specifications in rapidly changing markets.

Figure 5: Warsaw: Restricted MIDAS Model Fit

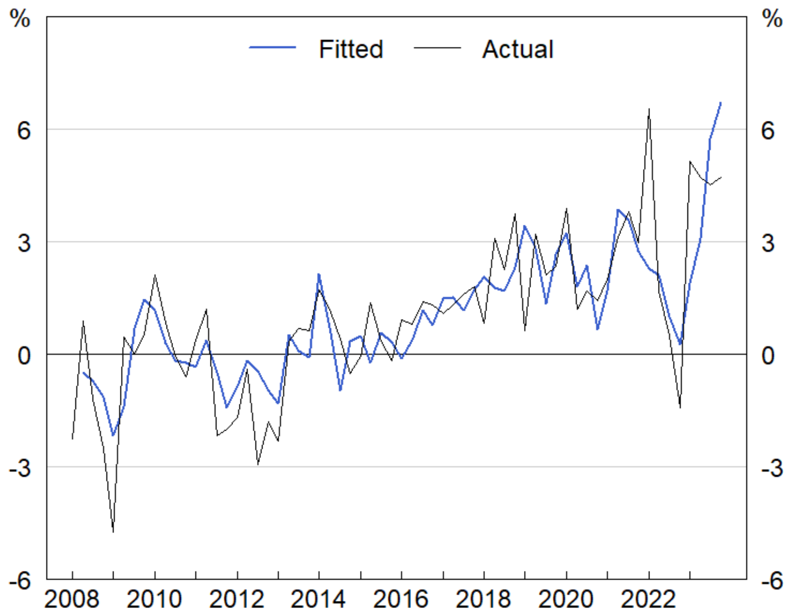

Figure 5 shows the fitted values from the optimal R-MIDAS models for Warsaw between 2008:Q1 to 2023:Q4 (blue line) compared to the actual transaction price series (black line).

Figure A5 in the Appendix shows the model fit for Poznan. In contrast to the Warsaw model, the model fit for Poznan is less favorable with an R-squared of 0.58, which confirms our previous comment about the two Poznan series showing less of a relationship compared to Warsaw.[8]

[7] When computing the Wald test for each R-MIDAS model specification, we use HAC standard errors, which are robust to unknown forms of heteroskedasticity and autocorrelation.

[8] We also investigated using a MIDAS-AR model with a common factor specification as in Clements and Galvão (2008) but the model fit as judged by the R-squared statistic for both cities was less favorable compared to the unrestricted MIDAS-AR model (Warsaw $R^2 = 0.52$; Poznan $R^2 = 0.53$).

4.2 Out-of-sample prediction evaluation

This section evaluates the nowcasting performance of MIDAS models using monthly list-price indices compared to standard time series forecasts. We illustrate performance through a rolling window and recursive pseudo-out-of-sample (OOS) exercises.

MIDAS regression offers a key advantage for mixed-frequency data: it enables predictions within periods as new high-frequency data points become available. In our setting, monthly list-price indices are released progressively over the quarter, allowing us to generate a sequence of predictions. Specifically, we produce (i) a forecast (FC) using all data from the previous quarter; (ii) nowcasts after the first (M1), second (M2), and third (M3) months of the quarter, as new list-price data are incorporated. The precise timing of forecast updates is illustrated in Figure 6.

Specifically, before the data on transaction price growth for t become available in quarter $t+1$, we have four releases of the monthly list price series. The first release incorporating monthly data up to $t-1$ in $t-2/3$ (i.e. first month of the current quarter), the second release of the list price series incorporating monthly data up to $t-2/3$ in $t-1/3$ (i.e. second month of the current quarter), the third release of the list price series incorporating monthly data up to $t-1/3$ in t (i.e. end of the current quarter). Finally, the fourth release of the list price series incorporating monthly data up to t in $t+1/3$ (i.e. first month of the next quarter).

Figure 6: Nowcasting Timeline

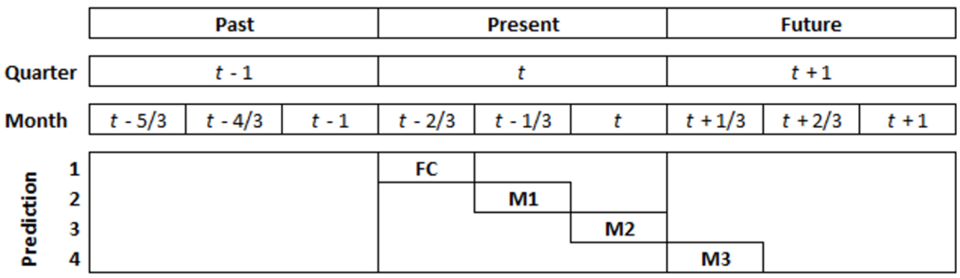

Figure 6 illustrates the timeline of forecast (FC) and nowcast updates (M1, M2, M3) within quarter t. Assuming prompt availability of transaction price data at the start of quarter t, the FC is generated using past quarter's transaction price index as well as list prices through $t-1$. The nowcasts (M1-–M3) progressively also incorporate updated monthly list price data from quarter t.

Across the OOS evaluation period we keep the R-MIDAS model specification for each city fixed. We compare the R-MIDAS models to standard benchmarks (AR(1), MA(1), AR(2), ARMA(1,1)).[9]

[9]Note that ARMA(1,1) is a flexible specification, as it can approximate both AR and MA processes of higher order. For example, an AR(1) process can be represented as an MA(∞), and an MA(1) process as an AR(∞); see Brockwell and Davis (1991).

Additionally, we include a quarter-average (QA) model. This model aggregates the monthly list-price series with equal weights across the three months of each quarter, making it a restricted form of R-MIDAS with uniform weighting. The QA model uses the aggregated list-price index for the current and previous quarter, making it similar in structure to nowcast M3.[10]

For both cities we train all models we consider on the first half of the available sample and use the remaining half to evaluate the performance of the various models. We split the data roughly in half to balance two goals: ensuring sufficient length for reliable parameter estimation and retaining enough periods for robust OOS evaluation.[11] This period also includes the COVID-19 crisis which serves as a test of the usefulness of timely data in prediction.

We evaluate the predictive accuracy using both recursive (expanding window) and rolling window OOS exercises. In the recursive setup, the estimation sample grows by one quarter each step; in the rolling window setup, the sample size is fixed (32 quarters for Warsaw, 34 for Poznan) and slides forward by one quarter at a time. For both cases, the model parameters are re-estimated each time.

For each quarter in the evaluation sample, we compute a forecast and three nowcasts depending on the available information set. For example, with Warsaw and for the initial evaluation quarter 2016:Q1, we compute a forecast using data up to 2015:Q4 (FC) and then a nowcast in 2016:M1 (M1), 2016:M2 (M2) and 2016:Q2 (M3). Additionally, we calculate forecasts for four standard time series models that rely solely on quarterly data.

We compare the Root Mean Square Errors (RMSEs) across different models as is common in the literature. To test statistical significance of the predictive accuracy of the competing models, we utilize the Model Confidence Set (MCS) procedure developed by Hansen et al. (2011). The MCS identifies models whose predictive accuracy is statistically indistinguishable at a given level of significance.

This is achieved by comparing loss differentials between pairs of models, through a sequence of tests. By this process the MCS enables the identification of the best performing models within a set, while also acknowledging the uncertainty associated with model selection. Note: In recursive estimation, we follow Hansen et al. (2011) and report pseudo-MCS results, as the stationarity assumption of loss differentials is unlikely to hold.

[10]Note, uniform weights are equivalent to the normalized exponential Almon function with the first parameter set to 1 and the rest set to 0 and the normalized beta function with all parameters set to 1.

[11]For Warsaw, the training sample is 2008:Q1–2015:Q4 while the evaluation sample is 2016:Q1–2023:Q4. For Poznan, the training sample is 2008:Q1–2016:Q2 while the evaluation sample is 2016:Q3–2024:Q3. The difference in sample split between Warsaw and Poznan comes form each city having a different number of observations.

The RMSEs for each model over the evaluation sample are presented in Table 4 for Warsaw and the rolling window OOS exercise and Table 5 for Warsaw and the recursive OOS exercise. In both tables, we compare RMSEs over three different horizons: the past three years, the past five years, and the full evaluation sample period. Both tables also highlight if a model is in the MCS with an asterisk. When computing the MCS, we follow the empirical example from Hansen et al. (2011) and set the significance level to 10 percent. The MCS is then referred to as $(90\%)\mathcal{M}$. We also include the MCS p-value used to decide on membership to the $(90\%)\mathcal{M}$ computed via a quadratic loss function with the $T_{R,\mathcal{M}}$ test statistic and using a stationary bootstrap with bloc length of 8 quarters and 5,000 replications.[12]

Table 4: Warsaw: Model Prediction Accuracy Comparison – Rolling Estimation

	AR(1)	MA(1)	AR(2)	ARMA(1,1)	FC	M1	M2	M3	QA
				Past 3-years					
RMSE	2.51	2.47	2.72	2.68	2.33	**2.19**	2.75	2.76	2.68
$(90\%)\mathcal{M}$						$*$			
p-value	0.00	0.04	0.00	0.00	0.04	1.00	0.00	0.00	0.00
				Past 5-years					
RMSE	2.18	2.18	2.27	2.25	2.13	**1.90**	2.26	2.27	2.19
$(90\%)\mathcal{M}$						$*$			
p-value	0.05	0.05	0.05	0.05	0.05	1.00	0.05	0.05	0.05
				Full sample					
RMSE	1.87	1.96	1.92	1.90	1.82	**1.57**	1.95	1.91	1.89
$(90\%)\mathcal{M}$						$*$			
p-value	0.01	0.01	0.01	0.01	0.01	1.00	0.01	0.01	0.01

Table 4 shows the prediction accuracy comparisons across the models considered. The rolling window estimation begins in 2016:Q1 with window length of 32 quarters. the full sample is 2016:Q1–2023:Q4. Bold values denote best model(s) for each horizon. The MCS p-value is computed by stationary bootstrap with bloc length 8 and 5,000 replications. $(90\%)\mathcal{M}$ refers to the 90% Model Confidence Set (MCS). MCS inclusion denoted by $*$. 'FC' is the forecast made at the end of the previous quarter. 'M1' is the nowcast in month 1 of the current quarter. 'M2' is the nowcast in month 2 of the current quarter. 'M3' is the nowcast in month 3 of the current quarter. 'QA' is the quarter average nowcast for the current quarter.

The results for both Warsaw and Poznan indicate that incorporating timely list-price data via our MIDAS framework improves predictive accuracy. At every horizon and under both rolling window and recursive estimation schemes, at least one MIDAS model achieves a smaller RMSE than standard time series models based solely on quarterly data. The RMSE differences across

[12]We did investigate the MIDAS-AR with common factor specification but the RMSEs were significantly worse than all other models. Furthermore, we also found that the MIDAS-AR common factor specification performed relatively poorly in the rolling window and recursive out-of-sample (OOS) exercises.

models narrow at longer horizons, suggesting that timely monthly data are particularly valuable during volatile periods, but less so in more stable times. This is clearly visible during the COVID-19 shock, where MIDAS models incorporating early-quarter list-price data (e.g., M1) were able to capture the sudden market contraction and rebound more effectively than quarterly models. This result suggests that MIDAS models are capable of capturing emerging trends quickly, before they are visible in transaction data.

In Warsaw, MIDAS models–particularly M1–consistently outperform standard time series benchmarks across horizons and estimation types. Notably, using list-price data from the first or second month yields the best predictions, while incorporating third-month data reduces accuracy.

Although it may seem counterintuitive that incorporating more recent data leads to lower accuracy, this aligns with our earlier findings (Section 3): list-price indices tend to lead transaction prices by about one month in Warsaw, so later list-price data may reflect transactions from the following quarter rather than the current one.

Results differ somewhat under the recursive estimation scheme. Although the MIDAS models still outperform standard time series models for Warsaw, the ordering of models changes: the quarter-average (QA) model is the most accurate across all horizons, tied with M3 at the 5-year horizon. Unlike in the rolling window exercise, predictive accuracy improves as more monthly data on the current quarter are incorporated, as one would expect. Regarding model confidence, only M2, M3, and QA are included in the 90 percent pseudo-MCS at the 3-year horizon. For the 5-year horizon, all models except MA(1) are included in the 90 percent pseudo-MCS, while for the full sample, all except AR(1) and MA(1) are included. This pattern suggests that, at longer evaluation horizons, predictive differences between MIDAS models and simpler benchmarks narrow to the point that statistical tests can no longer reliably distinguish between them. The added value of MIDAS models is clearest in shorter, more volatile periods where timely data matter most.[13]

[13] There are a few reasons for this result: (i) as horizon length increases, market conditions may stabilize, so simpler models catch up, (ii) the benefit of timely monthly data diminishes in quieter periods, (iii) the recursive estimation scheme has a longer sample period so that temporarily large misses do not have as much of an impact as they would do with the rolling window scheme with a shorter estimation window.

Table 5: Warsaw: Model Prediction Accuracy Comparison – Recursive Estimation

	AR(1)	MA(1)	AR(2)	ARMA(1,1)	FC	M1	M2	M3	QA
				Past 3-years					
RMSE	2.60	2.86	2.63	2.64	2.50	2.22	2.00	2.00	**1.99**
(90%)$\mathcal{M}$							$*$	$*$	$*$
p-value	0.01	0.00	0.01	0.01	0.01	0.01	0.21	0.21	1.00
				Past 5-years					
RMSE	2.25	2.50	2.22	2.22	2.16	1.98	2.12	**1.80**	**1.80**
(90%)$\mathcal{M}$	$*$		$*$	$*$	$*$	$*$	$*$	$*$	$*$
p-value	0.14	0.00	0.14	0.14	0.14	0.19	0.14	0.76	1.00
				Full sample					
RMSE	1.96	2.22	1.92	1.92	1.82	1.62	1.74	1.55	**1.51**
(90%)$\mathcal{M}$			$*$	$*$	$*$	$*$	$*$	$*$	$*$
p-value	0.03	0.00	0.26	0.27	0.27	0.48	0.31	0.57	1.00

Table 5 shows the prediction accuracy comparisons across the models considered. The recursive estimation begins in 2016:Q1 with initial sample length of 32 quarters. The full sample is 2016:Q1–2023:Q4. Bold values denote best model(s) for each horizon. The MCS p-value is computed by stationary bootstrap with bloc length 8 and 5,000 replications. (90%)$\mathcal{M}$ refers to the 90% Model Confidence Set (MCS). MCS inclusion denoted by $*$. 'FC' is the forecast made at the end of the previous quarter. 'M1' is the nowcast in month 1 of the current quarter. 'M2' is the nowcast in month 2 of the current quarter. 'M3' is the nowcast in month 3 of the current quarter. 'QA' is the quarter average nowcast for the current quarter.

We present the results for Poznan in Table A4 (rolling window estimation scheme) and Table A5 (recursive estimation scheme) in the Appendix and are similar to Warsaw. The MIDAS-based models generally outperform the standard time series models for all horizons and both estimation methods, and the recursive estimation scheme tends to be more accurate than the rolling window estimation scheme.

Model M1 is also best in the rolling window estimation scheme, but unlike with Warsaw, more models are included in the (90%) $\mathcal{M}$ than just M1 for the 3-year and full sample horizons. Similar to Warsaw, models M2–M3 and QA have smaller RMSEs in the recursive estimation scheme than in the rolling window estimation scheme. Still, unlike Warsaw, model M1 remains competitive in the recursive estimation scheme and achieves the lowest RMSE in both the 5-year and full sample horizons.

We interpret the smaller RMSE reductions in Poznan as an illustration that the usefulness of MIDAS depends on market conditions: in thinner, less liquid markets where list prices are less tightly linked to transactions, the MIDAS approach is less successful. Indeed, during the COVID-19 period, the advantage of MIDAS models in Poznan was modest, reflecting the weaker connection between list prices and transactions in this smaller market.

These results for Poznan are consistent with our earlier analysis in Section 3, where we found that list-price indices align less closely with transaction-price indices in Poznan compared to Warsaw, both in terms of timing and magnitude of growth rates (see Table 2). The weaker lead-lag relationship in Poznan likely reflects the smaller, less liquid market, in which list prices are a noisier signal for transaction prices.

This helps explain why the RMSE reductions from incorporating timely list-price data are more modest for Poznan, and why MIDAS models deliver less pronounced gains relative to standard time series models. The COVID-19 period provides further illustration of this point. While in Warsaw, the volatility of the pandemic years highlighted the advantage of using high-frequency list-price data, in Poznan, the benefits of MIDAS were less clear-cut. The smaller market and weaker alignment between list and transaction prices meant that timely data did not translate into equally strong predictive improvements during this volatile period.

To summarize, our results demonstrate a clear advantage of using higher-frequency (monthly) data to predict lower-frequency (quarterly) transaction prices. Both rolling window and recursive estimation schemes show that MIDAS models incorporating list-price data outperform standard time series models based solely on quarterly data. In particular, our preferred MIDAS specifications reduce RMSE by about 16–23 percent for Warsaw and 5–15 percent for Poznan relative to an AR(1) benchmark using either the rolling window or recursive estiamiton scheme.

Unlike many macroeconomic nowcasting applications, we find that predictive accuracy peaks when incorporating list-price data from the first or second month of the current quarter, and declines when third-month data are included. This reflects the typical six-week gap between a property's last listing and its transaction date: later list-price data often correspond to transactions that complete in the following quarter, introducing forward-looking noise rather than improving current-quarter predictions.

Ultimately, our results show MIDAS to be a viable method for producing accurate nowcasts of quarterly transaction prices using timely list-prices.

5 Conclusion

Transaction-based indices are widely regarded as the most reliable indicators of housing market trends because they reflect actual prices agreed upon by buyers and sellers (Eurostat, 2013; Silver, 2018; International Monetary Fund, 2020). However, transaction data are typically recorded with considerable delays, creating a challenge for real-time monitoring of housing market dynamics. This is especially problematic for macroprudential policy makers who need timely indicators to make informed decisions. Our study shows that list-price data from online platforms can help bridge this information gap caused by the delayed availability of transaction prices.

Drawing on 16 years of micro-level data from Warsaw and Poznan, we constructed quality-adjusted list- and transaction-price indices using the hedonic rolling-time-dummy (RTD) method, following international guidelines. This methodological choice ensured comparability between indices and enabled a robust analysis of their co-movement and lead-lag dynamics. We find that list-price and transaction-price indices co-move closely and that list-price indices lead transaction-price indices by one to two months. This lead-lag relationship is strongest in Warsaw, where the larger and more liquid market facilitates a tighter alignment between asking and realized prices.

Building on this relationship, we adapted the Mixed Data Sampling (MIDAS) regression framework to nowcast quarterly transaction-price growth using monthly list-price indices. Our preferred MIDAS specifications reduce one-quarter-ahead RMSE by approximately 16–23 percent for Warsaw and 5–15 percent for Poznan (rolling window and recursive estimation schemes respectively, full sample) relative to an AR(1) benchmark. The predictive advantage is greatest when incorporating list-price data from the first or second month of the quarter. Including third-month data tends to reduce accuracy, as these later list prices often correspond to transactions finalized in the following quarter, introducing forward-looking noise rather than improving current-quarter nowcasts.

These results have important implications. First, they confirm that list-price indices, when properly constructed and quality-adjusted, provide a valuable complementary signal to transaction-price indices. Second, MIDAS models based on list-price indices can provide early nowcasts of the direction and levels of the transaction price index, potentially improving the timeliness and effectiveness of macroprudential or monetary interventions. While these nowcasts are not substitutes for the official transaction price indices, they offer a means of detecting emerging housing market pressures earlier than traditional transaction-price indices allow, especially in places where transaction data release are subject to additional administrative delays. This is particularly useful during periods of higher-than-normal volatility, such as during the COVID-19 period.

Nonetheless, our study also illustrates that the predictive value of list-price indices depends on market structure and conditions. In Poznan, we found smaller advantages for the MIDAS approach compared to traditional forecasting methods than in Warsaw. This suggests that in less liquid or smaller markets, list-price data provide a weaker signal of future transaction prices. Moreover, even though we used 16 years of list and transaction data, our study period was marked by sustained price growth, which limited our ability to assess how well the MIDAS approach performs in a downturn.

Future research could usefully test this approach in other countries, smaller markets, or during periods of market downturns. Another possible future exploration could combine list-price data with other timely indicators such as mortgage approvals, search activity, or sentiment indices to further improve predictive performance.

References

Agarwal, S., Ben-David, I., and Yao, V. (2014). Collateral Valuation and Borrower Financial Constraints: Evidence from the Residential Real Estate Market. Management Science, 61(9):2220–2240.

Anenberg, E. and Laufer, S. (2017). A more timely house price index. Review of Economics and Statistics, 99(4):722–734.

Bajari, P., Cen, Z., Chernozhukov, V., Manukonda, M., Vijaykumar, S., Wang, J., Huerta, R., Li, J., Leng, L., Monokroussos, G., et al. (2023). Hedonic prices and quality adjusted price indices powered by ai. arXiv preprint arXiv:2305.00044.

Banti, C. and Phylaktis, K. (2019). Global liquidity, house prices and policy responses. Journal of Financial Stability, 43:79–96.

Beracha, E. and Seiler, M. J. (2014). The effect of listing price strategy on transaction selling prices. Journal of Real Estate Finance and Economics, 49:237–255.

Brockwell, P. J. and Davis, R. A. (1991). Time Series: Theory and Methods. Springer: New York, NY.

Cavallo, A. (2013). Online and official price indexes: Measuring argentina's inflation. Journal of Monetary Economics, 60(2):152–165.

Cavallo, A., Diewert, W. E., Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2018). Using online prices for measuring real consumption across countries. American Economic Review Papers and Proceedings, 108:483–487.

Cavallo, A. and Rigabon, R. (2016). The billion prices project: Using online prices for measurement and research. Journal of Economic Perspectives, 30(2):151–178.

Choi, C.-Y. and Hansz, J. A. (2021). From banking integration to housing market integration - evidence from the comovement of U.S. metropolitan house prices. Journal of Financial Stability, 54:100883.

Clements, M. P. and Galvão, A. B. (2008). Macroeconomic forecasting with mixed-frequency data: Forecasting output growth in the united states. Journal of Business & Economic Statistics, 26(4):546–554.

Court, A. T. (1939). Hedonic price indexes. In The Dynamics of Automobile Demand, pages 99–119. General Motors Corporation.

de Haan, J. (2010). Hedonic price indexes: A comparison of imputation, time dummy and 're-pricing' methods. Journal of Economics and Statistics (Jahrbuecher fuer Nationaloekonomie und Statistik), 230(6):772–791.

Diewert, W. E. (2002). Similarity and Dissimilarity Indexes: An Axiomatic Approach. Discussion Paper 02-10, Department of Economics, University of British Columbia, Vancouver, Canada.

Diewert, W. E. (2009). Durables and owner-occupied housing in a consumer price index. In Price Index Concepts and Measurement, pages 445–500. University of Chicago Press.

European Commission (2018). Report from the Commission to the European Parliament and the Council on the suitability of the owner-occupied housing (OOH) price index for integration into the harmonised index of consumer prices (HICP) coverage. European Commission: Brussels, 29.11.2018 COM(2018) 768 final.

European Systemic Risk Board (2019). Vulnerabilities in the residential real estate sectors of the EEA countries. European System of Financial Supervision.

Eurostat (2013). Handbook on residential property prices indices (RPPIs). Eurostat Methodologies and Working Papers.

Eurostat (2020). Practical guidelines on web scraping for the HICP. European Commission: Brussels.

Foroni, C., Marcellino, M., and Schumacher, C. (2015). Unrestricted Mixed Data Sampling (MIDAS): MIDAS Regressions with Unrestricted Lag Polynomials. Journal of the Royal Statistical Society Series A: Statistics in Society, 178(1):57–82.

Garzoli, M., Plazzi, A., and Valkanov, R. I. (2021). Backcasting, Nowcasting, and Forecasting Residential Repeat-Sales Returns: Big Data meets Mixed Frequency. Swiss Finance Institute Research Paper Series 21-21, Swiss Finance Institute.

Genesove, D. and Mayer, C. (2001). Loss aversion and seller behavior: Evidence from the housing market. Quarterly Journal of Economics, 116(4):1233–1260.

Ghysels, E., Kvedaras, V., and Zemlys, V. (2016). Mixed frequency data sampling regression models: The r package midasr. Journal of Statistical Software, 72(4):1–35.

Ghysels, E., Santa-Clara, P., and Valkanov, R. (2004). The midas touch: Mixed data sampling regression models.

Ghysels, E., Sinko, A., and Valkanov, R. (2007). Midas regressions: Further results and new directions. Econometric reviews, 26(1):53–90.

Hansen, P., Lunde, A., and Nason, J. (2011). The Model Confidence Set. Econometrica, 79(2):453–497.

Haurin, D., McGreal, S., Adair, A., Brown, L., and Webb, J. R. (2013). List price and sales prices of residential properties during booms and busts. Journal of Housing Economics, 22(1):1–10.

Hill, R., Scholz, M., Shimizu, C., and Steurer, M. (2018). An evaluation of the methods used by european countries to compute their official house price indexes. Economie et Statistique, 500-502:221–238.

Hill, R. J. (2013). Hedonic price indexes for housing: A survey, evaluation and taxonomy. Journal of Economic Surveys, 27(5):879–914.

Hill, R. J., Pfeifer, N., Steurer, M., and Trojanek, R. (2024). Warning: Some transaction prices can be detrimental to your house price index. Review of Income and Wealth, 70(2):320–344.

Hill, R. J. and Trojanek, R. (2022). An evaluation of competing methods for constructing house price indexes: The case of warsaw. Land Use Policy, 120:106226.

Hoekstra, R., Ten Bosch, O., and Harteveld, F. (2012). Automated data collection from web sources for official statistics: First experiences. Statistical Journal of the IAOS, 28(3-4):99–111.

Kolbe, J., Schulz, R., Wersing, M., and Werwatz, A. (2021). Real estate listings and their usefulness for hedonic regressions. Empirical Economics, 61(6):3239–3269.

Leamer, E. E. (2007). Housing is the business cycle. In Housing, Housing Finance and Monetary Policy. A Symposium Sponsored by the Federal Reserve of Kansas City.

Lyons, R. C. (2019). Can list prices accurately capture housing price trends? Insights from extreme markets conditions. Finance Research Letters, 30:228–232.

National Academies of Sciences, Engineering, and Medicine (2022). Modernizing the Consumer Price Index for the 21st Century. Washington, DC: Academies Press.

Pfeifer, N. and Steurer, M. (2022). Early real estate indicators during the covid-19 crisis. Journal of Official Statistics, forthcoming.

Pollakowski, H. O. (1995). Data Sources for Measuring House Price Changes. Journal of Housing Research, 6(3):377–387.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. Journal of Political Economy, 82:34–55.

Shimizu, C., Nishimura, K., and Watanabe, T. (2010). Housing prices in tokyo: A comparison of hedonic and repeat-sales measures. Journal of Economics and Statistics, 230:792–813.

Shimizu, C., Nishimura, K. G., and Watanabe, T. (2016). House prices at different stages of the buying/selling process. Regional Science and Urban Economics, 59(C):37–53.

Trojanek, R. (2018). Teoretyczne i metodyczne aspekty wyznaczania indeksów cen na rynku mieszkaniowym. Poznań University of Economics and Business Press, Poznań.

Trojanek, R. (2021). Housing price cycles in Poland – the case of 18 provincial capital cities in 2000–2020. International Journal of Strategic Property Management, 25(4):332–345.

Trojanek, R. (2025). From print advertisements to web portals: Hedonic house price and rental indices for Poland using three decades of listings. Working Paper.

Trojanek, R. and Huderek-Glapska, S. (2018). Measuring the noise cost of aviation – the association between the limited use area around warsaw chopin airport and property values. Journal of Air Transport Management, 67:103–114.

Wang, X., Li, K., and Wu, J. (2020). House price index based on online listing information: The case of China. Journal of Housing Economics, 50(July):101715.

Appendices

Appendix A List- and Transaction Price Indices

Figure A1: List and Transaction Price Indices for Warsaw, 2008 to 2023

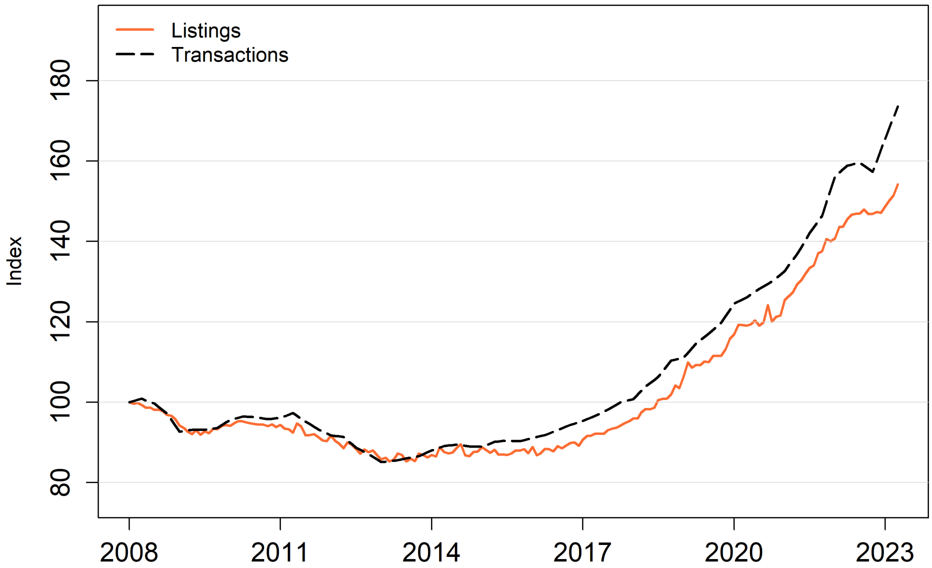

Figure A1 illustrates the monthly listing price index series vs the quarterly transaction price index series for Warsaw between 2008 and 2023. Index = 100 in Q1 2008.

Figure A2: List and Transaction Price Indices for Poznan, 2008 to 2023

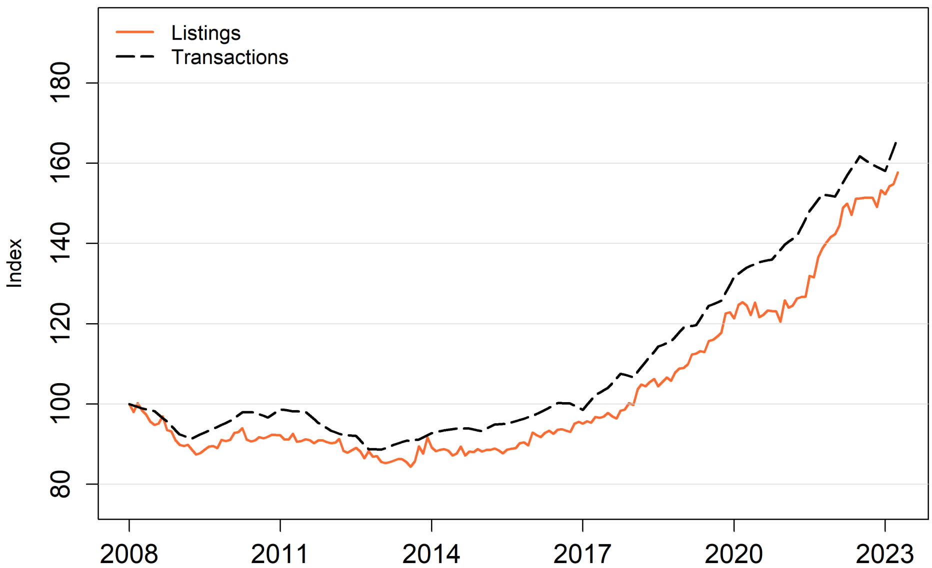

Figure A2 illustrates the monthly listing price index series vs the quarterly transaction price index series for Poznan between 2008 and 2023.

Figure A3: Monthly List- and Quarterly Transaction Price Changes for Poznan

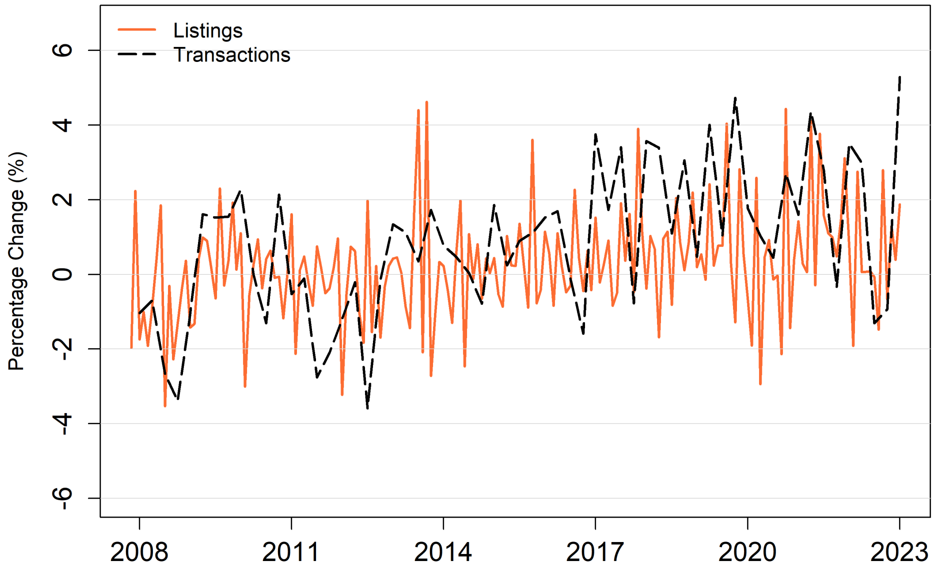

Figure A3 illustrates the monthly listing price and the quarterly transaction price growth rates for Poznan between 2008 and 2023.

Figure A4: Poznan: best list price alignment series according to DWM criterion

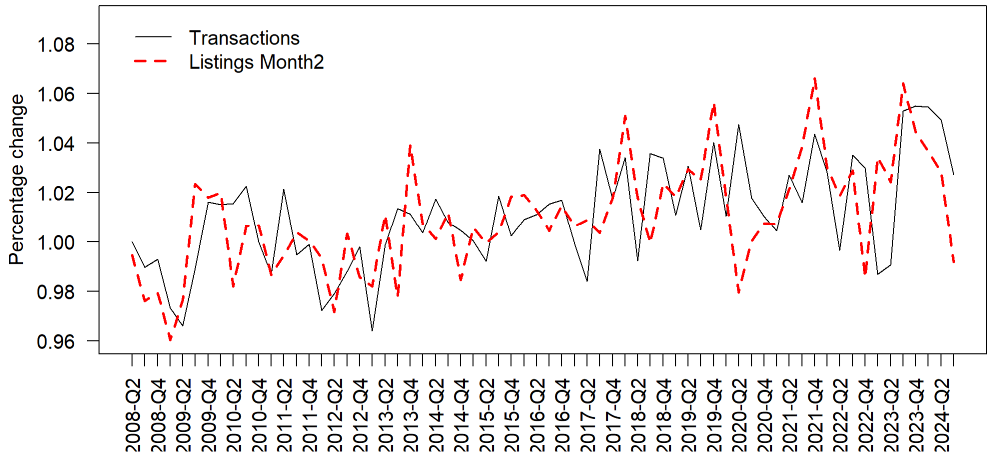

The results presented in Figure A4 illustrate the strong alignment between the quarterly transaction price index for Poznan and *month*1 list price index.

Appendix B List and Transaction Price Indices

Appendix B.1 Regression Outputs for Warsaw - List and Transaction Price Indices

Table A1: Summary Statistics – Warsaw

Variable	Transactions	Listings
const	12.3490*** (0.000)	12.1347*** (0.000)
district dummies	Yes	Yes
Time effects		
Transactions (Quarterly)		
2023-Q1	0.0369*** (0.001)	
2023-Q2	0.0773*** (0.000)	
2023-Q3	0.1251*** (0.000)	
2023-Q4	0.1702*** (0.000)	
Listings (Monthly)		
2023-02		0.0101*** (0.004)
2023-03		0.0185*** (0.000)
2023-04		0.0361*** (0.000)
2023-05		0.0481*** (0.000)
2023-06		0.0552*** (0.000)
2023-07		0.0863*** (0.000)
2023-08		0.1146*** (0.000)
2023-09		0.1371*** (0.000)
2023-10		0.1704*** (0.000)
2023-11		0.1769*** (0.000)
2023-12		0.2012*** (0.000)
Property characteristics		
area	0.0218*** (0.000)	0.0208*** (0.000)
area^2	-4.617e-05*** (0.000)	-4.446e-05*** (0.000)
age	-0.0031*** (0.000)	-0.0013*** (0.000)
construction technology	0.1433*** (0.000)	0.0601*** (0.000)
quality of apartment	–	0.0464*** (0.000)
garage	0.0526*** (0.000)	–
distance to city center	-3.924e-05*** (0.000)	–
R^2	0.807	0.881
N	16,451	49,573
Sample period	2022:Q4–2023:Q4	2023:01–2023:12

Note: p-values less than 0.05 are denoted by (*), p-values less than 0.01 are denoted by (**), and p-values less than 0.001 are denoted by (***). The district dummies (not shown) are all significant.

Appendix B.2 Regression Outputs for Poznan - List and Transaction Price Indices

Table A2: Summary Statistics – Warsaw

Variable	Transactions	Listings
const	12.1084*** (0.000)	11.9066*** (0.000)
district dummies	Yes	Yes
Time effects		
Transactions (Quarterly)		
2023-Q1	-0.0113 (0.229)	
2023-Q2	0.0385*** (0.000)	
2023-Q3	0.0898*** (0.000)	
2023-Q4	0.1390*** (0.000)	
Listings (Monthly)		
2023-02		0.0124 (0.155)
2023-03		0.0168* (0.038)
2023-04		0.0359*** (0.000)
2023-05		0.0359*** (0.000)
2023-06		0.0480*** (0.000)
2023-07		0.0745*** (0.000)
2023-08		0.0978*** (0.000)
2023-09		0.1043*** (0.000)
2023-10		0.1366*** (0.000)
2023-11		0.1418*** (0.000)
2023-12		0.1613*** (0.000)
Property characteristics		
area	0.0224*** (0.000)	0.0203*** (0.000)
area^2	-6.922e-05*** (0.000)	-5.603e-05*** (0.000)
age	-0.0019*** (0.000)	-0.0015*** (0.000)
construction technology	0.1148*** (0.000)	0.0618*** (0.000)
quality of apartment	–	0.0513*** (0.000)
garage	0.1359*** (0.000)	–
distance to city center	-3.076-05*** (0.000)	–
R^2	0.707	0.809
N	4,190	9,380
Sample period	2022:Q4–2023:Q4	2023:01–2023:12

Note: p-values less than 0.05 are denoted by (*), p-values less than 0.01 are denoted by (**), and p-values less than 0.001 are denoted by (***). The district dummies (not shown) are all significant.

Appendix C Nowcasting Results for Poznan

Table A3: Poznan: MIDAS Model Comparison

	Normalised Exponential Almon				Normalised Beta		Unrestricted	
	$j=2$		$j=3$		$j=3$		MIDAS	
Lag	BIC	p-value	BIC	p-value	BIC	p-value	BIC	p-value
0:2	**285.50**	**0.87**	289.66	0.00	293.32	0.00	289.66	–
0:3	**277.10**	**0.54**	280.35	0.54	282.53	0.11	284.28	–
0:4	277.10	0.01	292.16	0.00	**275.13**	**0.00**	276.86	–
0:5	274.11	0.00	**<u>268.68</u>**	**<u>0.64</u>**	269.18	0.34	280.04	–

Table A3 shows the results of the in-sample model selection using the BIC. Bold values denote best model per lag. A bold and underline value denote best overall model. The best model is an normalized exponential Almon MIDAS with up to six lags. Even though we only have monthly updates, the unrestricted MIDAS is never optimal in our case. The p-values shows the results of the test of the null hypothesis that the restrictions on the MIDAS regression coefficients implied by the polynomial weighting function are valid and supported by the data. Failure to reject the null implies the functional restrictions are supported by the data. We find these restrictions are valid for our preferred model.

Figure A5: Poznan: Restricted MIDAS Model Fit

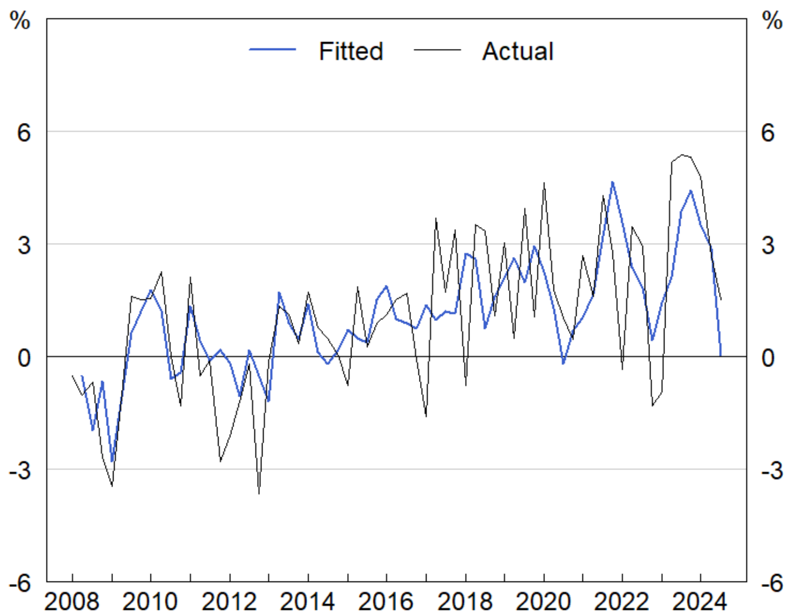

Figure A5 shows the fitted values from the optimal R-MIDAS models for Poznan between 2008:Q1 to 2024:Q3 (blue line) compared to the actual transaction price series (black line).

Table A4: Poznan: Model Prediction Accuracy Comparison – Rolling Estimation

	AR(1)	MA(1)	AR(2)	ARMA(1,1)	FC	M1	M2	M3	QA
Past 3-years									
RMSE	2.61	2.63	2.84	2.69	2.63	**2.43**	2.52	2.74	2.68
(90%)$\mathcal{M}$	$*$	$*$				$*$	$*$		$*$
p-value	0.21	0.27	0.01	0.03	0.09	1.00	0.27	0.09	0.27
Past 5-years									
RMSE	2.33	2.35	2.42	2.34	2.30	**2.09**	2.24	2.50	2.34
(90%)$\mathcal{M}$						$*$			
p-value	0.05	0.03	0.03	0.03	0.03	1.00	0.05	0.00	0.05
Full sample									
RMSE	2.33	2.34	2.34	2.27	2.29	**2.20**	2.33	2.47	2.33
(90%)$\mathcal{M}$	$*$	$*$	$*$	$*$	$*$	$*$	$*$		$*$
p-value	0.55	0.24	0.38	0.55	0.55	1.00	0.24	0.04	0.24

Table A4 shows the prediction accuracy comparisons across the models considered. The rolling window estimation begins in 2016:Q3 with window length of 34 quarters. the full sample is 2016:Q3–2023:Q3. Bold values denote best model(s) for each horizon. The MCS p-value is computed by stationary bootstrap with bloc length 8 and 5,000 replications. (90%)$\mathcal{M}$ refers to the 90% Model Confidence Set (MCS). MCS inclusion denoted by $*$. 'FC' is the forecast made at the end of the previous quarter. 'M1' is the nowcast in month 1 of the current quarter. 'M2' is the nowcast in month 2 of the current quarter. 'M3' is the nowcast in month 3 of the current quarter. 'QA' is the quarter average nowcast for the current quarter.

Table A5: Poznan: Model Prediction Accuracy Comparison – Recursive Estimation

	AR(1)	MA(1)	AR(2)	ARMA(1,1)	FC	M1	M2	M3	QA
				Past 3-years					
RMSE	2.45	2.54	2.67	2.58	2.28	2.13	2.20	**2.03**	2.08
(90%)$\mathcal{M}$	$*$	$*$		$*$	$*$	$*$	$*$	$*$	$*$
p-value	0.29	0.29	0.05	0.16	0.29	0.64	0.46	1.00	0.64
				Past 5-years					
RMSE	2.28	2.38	2.34	2.29	2.11	**1.88**	2.05	1.89	1.89
(90%)$\mathcal{M}$			$*$	$*$	$*$	$*$	$*$	$*$	$*$
p-value	0.05	0.06	0.16	0.16	0.16	1.00	0.16	0.90	0.90
				Full sample					
RMSE	2.32	2.40	2.32	2.29	2.19	**1.98**	2.21	2.09	2.03
(90%)$\mathcal{M}$			$*$	$*$	$*$	$*$	$*$	$*$	$*$
p-value	0.03	0.02	0.15	0.15	0.15	1.00	0.15	0.56	0.56

Table A5 shows the prediction accuracy comparisons across the models considered. The recursive estimation begins in 2016:Q3 with initial sample length of 34 quarters. The full sample is 2016:Q3–2024:Q4. Bold values denote best model(s) for each horizon. The MCS p-value is computed by stationary bootstrap with bloc length 8 and 5,000 replications. (90%)$\mathcal{M}$ refers to the 90% Model Confidence Set (MCS). MCS inclusion denoted by $*$. 'FC' is the forecast made at the end of the previous quarter. 'M1' is the nowcast in month 1 of the current quarter. 'M2' is the nowcast in month 2 of the current quarter. 'M3' is the nowcast in month 3 of the current quarter. 'QA' is the quarter average nowcast for the current quarter.