

Search and Predictability of Prices in the Housing Market*

Stig V. Møller[†]

Thomas Q. Pedersen[‡]

Christian M. Schütte[§]

Allan Timmermann[¶]

January 2022

Abstract

We develop a new housing search index (*HSI*) extracted from online search activity on a limited set of keywords related to the house buying process. We show that the *HSI* has strong predictive power over subsequent changes in house prices, both in-sample and out-of-sample and after controlling for the effect of commonly used predictors, and relate our findings to models of search-induced frictions. Our results imply that search data can be used as an early indicator of where the market is going.

Keywords: Internet search, housing markets, housing demand, forecasting, frictions, inelastic housing supply.

JEL codes: C10, E17, G10, R3.

*We thank two anonymous reviewers, an Associate Editor, and the Department Editor, David Sraer, for insightful and constructive comments. In addition, we thank Anil Kumar, participants at the 2019 SoFiE Conference in Shanghai, the IAAE 2019 Annual Conference in Nicosia, the 2019 Nordic Econometric Meeting in Stockholm, the 2019 International Symposium on Forecasting in Thessaloniki, the 2019 CFE in London and seminar participants at Aarhus University for comments. We acknowledge support from the Danish Finance Institute and from the Danish Council of Independent Research and from the Independent Research Fund Denmark (7024-00020B).

[†]CREATES, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus and the Danish Finance Institute. E-mail: svm@econ.au.dk.

[‡]CREATES, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus and the Danish Finance Institute. E-mail: tq-pedersen@econ.au.dk.

[§]CREATES, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus. E-mail: christianms@econ.au.dk.

[¶]UCSD, CEPR, and CREATES, University of California, San Diego, La Jolla, CA 92093. E-mail: atimmermann@ucsd.edu.

1 Introduction

Fluctuations in house prices have a profound impact on household welfare, financial stability, and the broader economy. For example, Case et al. (2012) estimate that the decline in U.S. housing wealth during 2005-2009 implied a decline in consumption of about \$350 billion per year. Further, in an analysis covering more than 60 countries, Reinhardt and Rogoff (2009) show that house price bubbles have historically been among the best predictors of banking crises across both advanced and emerging market economies. In response to the importance of variation in house prices for macroeconomic stability, the European Commission recently included house prices in its early warning system for macroeconomic imbalances (the “MIP Scoreboard”). Reliable and accurate predictions of house prices are evidently of great importance for policy makers as they can be used to predict future costs of living and revenues from real estate taxes or provide an early warning of an incipient price boom—or, potentially a weakening—in the housing market. Accurate forecasts are also valuable to households planning to buy or sell in the residential real estate market, particularly if available at the local market level.

The housing market is characterized by a highly heterogeneous and complex product, local segmentation, and a slow price discovery process caused by a variety of frictions. Buying a house is, therefore, a search intensive process involving a lengthy review of homes for sale and price comparisons across the inventory of homes listed for sale at a given point in time. Much of this search process is conducted online. A recent report by the National Association of Realtors (NAR, 2020) shows that home buyers use the internet as their main source of information about the housing market, with as many as 93% of home buyers using the internet to search for a home.

This paper develops and tests a set of hypotheses about the relation between online housing search volume and changes in house prices. Our first and main hypothesis is that search activity, which tracks peoples’ intentions of buying a house and thereby proxies for housing demand, should have a positive relation with house prices. Given various frictions in the housing market, an increase in search activity is propagated into future periods, implying sluggish price adjustment in response to an increase in demand such that search activity should hold predictive power for future variation in house prices – an insight that follows directly from theoretical search-based models (e.g. Berkovec and Goodman, 1996, Genesove and Han, 2012, and Carrillo et al., 2015). Because the house search process tends to be lengthy, our second hypothesis is that internet search volume has predictive

power at both short and long horizons, but also that its predictive power declines at longer horizons where frictions are less likely to be binding. Our third hypothesis is that the predictive power of housing search, being a proxy for housing demand, is particularly strong in housing markets with low supply elasticity. Since housing markets are inherently local and segmented, our fourth hypothesis is that local search activity contains important information about local house prices beyond what is captured in national search activity.

The intense and lengthy search process involved in buying a house coupled with large frictions in the housing market means that it is natural to expect internet search volume for housing to have predictive power for future house prices. Using Google Trends search data, we start out with the keyword “buying a house” and add related search terms supplied by Google, all of which are related to the search process of future home buyers. To capture common variation across search volume indices, we define the Housing Search Index (*HSI*) as the first principal component of the search volume indices. This provides us with a simple and intuitive measure of housing demand. We validate our search-based measure of demand by comparing it to data on home tours and writing offers.

We show that demand for housing as measured through online search activity predicts future house prices at both short and long horizons. At the one-month horizon, the *HSI* explains more than 50% of the variation in national house price growth, while at the one-year horizon the explanatory power is close to 65%. The predictive power of the *HSI* peaks at horizons around 3-8 months, which is consistent with the time buyers typically spend finding a home from the initial search process to closing the deal. Across horizons, the *HSI* produces far more accurate forecasts of future house prices than standard housing market determinants – a result that holds both in-sample and out-of-sample. Overall, the *HSI* tracks the housing market with relatively high accuracy. The index captures not only the turbulence surrounding the financial crisis and the more stable period the housing market has experienced in recent years, but also the unusual development in house prices following the outbreak of the Covid-19 pandemic.

Demand for housing is generally believed to be a function of key macroeconomic variables such as interest rates, employment and credit conditions. To better understand the mechanism behind housing search activity, we examine the relation between the *HSI* and a range of variables typically used to explain dynamics in the housing market. We find that internet search for housing has a

negative correlation with the level of the mortgage rate, indicating that households intensify search in times with low financing costs. Otherwise the *HSI* has a relatively low correlation with key housing market determinants as well as with various risk premium proxies.

Google Trends provide data also on local online search volume. This is a key advantage relative to macroeconomic data since housing markets tend to be local in nature (Del Negro and Otrok, 2007, Gyourko et al., 2013, Glaeser et al., 2014, and Hernández-Murillo et al., 2017). In regressions across 77 Metropolitan Statistical Areas (MSAs), we show that local housing search is a strong predictor also for local house prices, generally explaining more than 40% of the one-month-ahead variation in MSA-level house prices. Furthermore, controlling for national search activity, we show that local housing search remains a significant predictor of local house prices, which is direct evidence that housing markets are influenced by local search dynamics.

We next exploit cross-sectional variation in local housing markets to corroborate our interpretation that the *HSI* is a proxy for latent housing demand. We do so along two dimensions. First, our MSA-level regressions show a large dispersion in the economic effect on house prices from changes in search activity. Provided that the *HSI* captures variation in housing demand, we would expect to see a larger economic effect in local housing markets with a more constrained housing supply. Using the supply inelasticity measure of Saiz (2010), we show that this is indeed the case. Second, theoretical search-based models imply that the time it takes for a house to be sold should fall in response to an increase in demand. We test this using the *HSI* as proxy for demand and find supportive evidence of a negative relation between the *HSI* and time-on-market.

The Covid-19 pandemic caused a massive shock to the U.S. economy and housing market and we would not necessarily expect the relation between search activity and house prices to remain robust during this period. To explore the impact of the pandemic on our results, we estimate an MSA-level panel model that includes the *HSI* along with a measure of housing supply (for-sale inventory data from Zillow) and a stringency index for Covid-19 lockdown measures. We find that demand and supply effects along with Covid-19 restrictions combine to capture nearly two-thirds of the monthly variation in house prices across MSAs during the pandemic.

Other papers have studied the relation between online search and housing. Wu and Brynjolfsson (2015) find that search data are more effective for predicting house transactions than for predicting house prices and that online search has rather limited predictive power over house prices. This

contrasts with our findings, but the reason for the difference is easy to comprehend: Wu and Brynjolfsson use two broad, predefined search categories (real estate listings and real estate agencies) containing several individual search terms, complicating the economic interpretation of their search activity measures. Conversely, we explicitly use terms that capture search activity from potential house buyers and therefore are more strongly related to housing demand and have a highly significant predictive power over variation in house prices across several horizons. Beracha and Wintoki (2013) use search volume for "real estate i", where "i" is the name of a city. They show that abnormal search volume for a city lead to abnormal changes in house prices for that city. We find that our suggested procedure has considerable stronger predictive power over future house prices compared to the procedure used by Beracha and Wintoki (2013).¹

Our analysis is also related to the literature that exploits online search activity to measure peoples' attention and its impact on asset prices. For example, Da et al. (2011) construct a direct measure of investor attention through online search activity for individual stock tickers and show that an increase in attention predicts higher stock prices in the ensuing two weeks. At a more aggregate level, Da et al. (2015) use daily search activity to construct a Financial and Economic Attitudes Revealed by Search (FEARS) index using keywords such as recession, unemployment and bankruptcy. They find that the index predicts short-term return reversals as well as temporary increases in volatility.² Andrei and Hasler (2015) provide both a theoretical framework and empirical results which support attention as a key determinant of asset prices.

We contribute to this literature by showing that demand for housing as measured through online search activity is a strong predictor of house prices. The predictive ability of search activity for house prices follows naturally from the high search intensity involved in buying a house as well as the frictions present in the housing market. Consequently, search activity has a relatively large and long-lasting impact on future house prices – both in absolute terms and when compared to other asset classes.

Our paper is also directly related to the literature on predictability of house prices, including studies such as Rapach and Strauss (2009), Plazzi et al. (2010), Ghysels et al. (2013), Soo (2018),

¹We focus on the link between search behavior and *future* house prices. Gargano et al. (2021) study the reverse relation, namely how *past* price growth explains differences in search behavior across prospective home buyers. They find that prospective home buyers experiencing higher growth in their postcode of residence search more broadly across locations and house characteristics.

²Joseph et al. (2011) also find that the more difficult stocks are to arbitrage, the stronger the link between search intensity (as measured by online ticker search) and future returns.

Cox and Ludvigson (2019), and Bork et al. (2020). This literature typically uses either economic variables such as interest rates, employment and credit conditions or sentiment-based variables as predictors. The underlying intuition here is that supply and especially demand are largely driven by these variables which, consequently, contain important information about future house prices. We extend this literature by proposing a more direct measure of demand and show that it strongly outperforms standard variables used to predict future house prices.

In addition to the higher predictive accuracy of our *HSI* measure, there are several other advantages of using online search data in forecasting house prices compared to data gathered from government agencies. Many macroeconomic variables are often announced with a substantial time delay, only available at a low frequency, and subject to substantial data revisions, complicating real-time forecasting. In contrast, Google search data are readily available at a high frequency without time delay and are not subject to data revisions.³

The rest of the paper is structured as follows. Section 2 explains how we build on the theoretical insights from search-based models as well as how we measure housing demand and construct the national and local search indices. This section also contains an analysis of how the *HSI* relates to standard housing market determinants. Section 3 contains an empirical analysis of the predictive power of search activity in the housing market over future house prices. Section 4 explores variation in local housing markets and relates our findings to variation in local supply elasticities and speculative demand. Section 5 analyzes the housing market during the Covid-19 pandemic. Finally, Section 6 contains concluding remarks.

2 Search Activity in the Housing Market

Online search volume has been shown to track investor sentiment in stock and bond markets (Da et al., 2015). It is plausible to expect that search activity also contains valuable information for tracking and quantifying variation in the demand for housing – a highly complex and segmented market. Specifically, aggregate internet search volume for phrases such as “buying a house” is likely to reflect genuine interest in actually buying a house and should thereby provide a timely and observable signal that is correlated with the underlying (latent) variation in housing demand.

³Guo (2009) and Ghysels et al. (2017) show that asset return predictability from macroeconomic data tends to be considerably weaker when using real-time macroeconomic data as opposed to using revised macroeconomic data.

2.1 Search as a Leading Indicator for Housing Demand

We start by motivating our choice of housing search activity as a leading indicator for demand in the housing market, building on theoretical work from the search and matching literature. The idea behind these search models is that since no central clearing house exists, buyers and sellers look for each other until they are matched. Since search is a costly activity, agents will aim at optimizing the effort over time. Several models within this framework imply that positive (negative) demand shocks lead to subsequent positive (negative) house price changes, motivating why housing search as a proxy for demand should have predictive power over future house price changes.⁴

Piazzesi et al. (2020) point out that, although supply in the housing market can be proxied by the number of homes available for sale in a given market, demand (the number of potential buyers), remains unobserved. A similar observation is made by Han and Strange (2015) who argue that although we have measures for seller time-on-market, there is no parallel for buyer time-on-market as a proxy for buyer search. Since buyers are arguably more active than sellers, empirical research on buyer search intensity is essential for reaching a better understanding of housing markets. Our paper attempts to make up for this shortcoming, arguing that we can use internet search activity, segmented by local markets at the MSA level, as a proxy for the search behavior of home buyers across time.

Our study is related to Piazzesi et al. (2020) who document that search activity is positively correlated with house prices in the cross-section of U.S. cities. In contrast to their study, we characterize search intensity dynamics over time at an aggregate MSA level across the U.S., instead of focusing on cross-sectional search for individual houses at a single point in time. Our study confirms the positive relationship between search activity and prices but by analyzing the time series dimension, we can capture the effect of current search intensity on future price appreciation. In this respect, our study confirms the theoretical predictions of Berkovec and Goodman (1996), who present a model in which frictions in the search and matching process imply that current demand shocks impact not only current but also future house price changes. In their model, buyers and sellers have imperfect information about the underlying market conditions, implying that price expectations adjust gradually in response to a demand shock.⁵

⁴See Han and Strange (2015) for a detailed survey of the literature on housing search models.

⁵Kraimer (2001) and Novy-Marx (2009) also analyze frictions in the search and matching process of home buyers and sellers.

Carrillo et al. (2015) develop a search and matching model in which measures of market tightness, defined as the ratio of buyers and sellers in the market, predict future house price changes. More buyers entering the market during times of increasing demand leads to market tightness which in turn is followed by an increase in the bargaining power of sellers and higher likelihood of a sale. Since buyers and sellers do not hold perfect information about market conditions (e.g., the size of demand shocks), an increase in market tightness today leads to an increase in house prices in the future.⁶ Other search-based models can generate similar mechanisms of sluggish price adjustments. For instance, building on Wheaton (1990), Diaz and Jerez (2013) specify a search model that propagates the effect of aggregate shocks to future periods. A key element of their model is that search and matching frictions produce trading delays such that not all agents seeking to buy a new home can do so right away, implying that the effect of aggregate shocks is propagated to future periods. Genesove and Han (2012) develop a search and matching model in which lagged seller response, due to gradual adjustment of the seller’s reservation price, results in sluggish price adjustments after a demand shock. In a similar vein, Head et al. (2014) show that time-consuming search and matching generates sluggish price adjustments in response to a shock.

Taken together, the theoretical insights from search-based models imply that a shift in current demand will lead to future price changes, which is the main hypothesis of this paper.

2.2 Construction of the Housing Search Index

To quantify internet search activity, we use Google Trends data from which we obtain a time series index on the volume of queries for a given search term in a given geographic area.⁷ Google Trends provides a set of related queries for every main query. The list of related queries (or, equivalently, *related terms*) includes between 0 and 25 different terms, with the final number depending on the search volume of the main query, i.e. high volume series will usually have 25 related queries while lower volume series will feature fewer. Google does not disclose the methodology it uses to select related queries, but the resulting terms are usually intuitively related to the main query. From the perspective of quantifying housing demand, this feature is appealing for two reasons. First, each

⁶van Dijk and Francke (2018) create a proxy for tightness in the Dutch housing market which relates positively to changes in house prices.

⁷Google dominates the U.S. search engine market with a 63 percent market share as of October 2018 (Statista, 2018). Data on search volume are also available for other services owned by Google such as Image Search, News Search, Google Shopping and YouTube Search, but these account for far smaller volumes than general Google searches.

semantically related keyword can provide additional information about housing demand beyond that contained in the original query. Second, since related terms are likely to be correlated, this induces a natural factor structure which allows us to build an aggregate measure of housing demand.

Google Trends data are available from 2004 onwards. Our sample runs from 2004:1 to 2021:1 at the monthly frequency.⁸ To obtain a simple and clean measure of housing demand, we initially use “buying a house” as our main search term and subsequently obtain a list of 22 related terms: “when buying a house”, “buying a home”, "buy a house", "mortgage", "buying a new house", "before buying a house", "how to buy a house", "real estate", "steps to buying a house", "buying a house calculator", "first time buying a house", "buying a house process", "house buying process", "homes for sale", "building a house", "buying a house with bad credit", "cost of buying a house", "buying a house to rent", "mortgage calculator", "houses for sale", "buying a house tips", and "buying a foreclosure house". These search terms are all related to the home buying process and as such should proxy for housing demand. The three remaining related search terms are excluded either because they are unrelated to housing ("buying a car") or because the search volume is low. We define low volume series as those for which more than 10% of observations equal zero.⁹

Some of the related terms may be measured with more noise than others. To filter out the noise and more accurately estimate latent demand, we use a targeted PCA approach which ensures that only the most relevant search indices are included to compute the latent demand factor. Specifically, our implementation follows Bai and Ng (2008) as we use the elastic net estimator of Zou and Hastie (2005) to select the ten most relevant search indices and then apply principal component analysis to summarize the most important information from these indices into one common component. We interpret this principal component as a summary measure for housing search and refer to it as the Housing Search Index (*HSI*).¹⁰

Before extracting the first principal component, we transform the search indices as follows. Following Da et al. (2011, 2015) and Vozlyublenniaia (2014), we first convert the series to their natural

⁸As noted by D’Amuri and Marcucci (2017), Google Trends are created based on a sample of queries that change according to the time and IP address used to download the data. To account for sampling error, we compute the index for all Google Trends queries using an average over 15 different days. The correlation across different samples is always above 0.99. Hence, the results are, for all practical purposes, robust to this issue.

⁹The two excluded terms are "help buying a house" and "buying a house cash".

¹⁰Our main goal is to produce a simple and easy-to-interpret index of housing search, which is why we use a simple targeted PCA approach. However, the predictive results that we report below are generally highly robust to using more advanced machine learning techniques. We refer to Section A.7 in the Online Appendix for further details.

logarithm.¹¹ To account for the possibility that the individual Google Trends series could follow different trends, we adopt a sequential testing strategy in the spirit of Ayat and Burrige (2000) and similar to Borup and Schütte (2022).¹² We further remove seasonality by regressing each series on monthly dummy variables and study the residuals from this regression.

2.3 Housing Search and Prices

Panel A in Figure 1 displays a time series of the *HSI* along with the log growth rate in the seasonally adjusted monthly Federal Housing Finance Agency (FHFA) purchase-only house price index for the United States.¹³ Housing search and growth in house prices move closely together. In particular, we note that the *HSI* captures the negative growth rates in 2009-2010 that followed the collapse in the housing market, the subsequent recovery, as well as the more stable house price growth seen in recent years. The *HSI* also captures the unusual development in house prices following the outbreak of the Covid-19 pandemic. As an initial response to Covid-19, house prices dropped slightly but subsequently experienced large positive growth rates – a development mirrored in the *HSI*.

To explore the dynamic relation between the *HSI* and movements in house prices, Panel A of Figure 2 shows regression slope coefficients, associated *t*-statistics and *R*²-values of monthly price changes from $t - 1$ to t on lagging, contemporaneous and leading values of the *HSI*:

$$p_t - p_{t-1} = \alpha_j + \beta_j HSI_{t+j} + \varepsilon_t, \quad j = -12, \dots, 12, \quad (1)$$

where p_t is the log of the FHFA house price index in month t . We find much larger coefficients and

¹¹There is no consensus in the literature as to whether Google Trends data are best characterized by stationarity, trend stationarity or a unit root since this can be very sensitive to the query in question. Vozlyublennaia (2011), Choi and Varian (2012), Bijl et al. (2016) and D’Amuri and Marcucci (2017) do not perform any differencing or detrending of the series, which suggests that the Google Trends data they use are stationary. Da et al. (2015) study the log-differences (growth rates) of their data.

¹²The idea is to successively test for stationarity, linear trend stationarity and quadratic trend stationarity using an augmented Dickey-Fuller (ADF) test. Specifically, the first test computes an ADF test with a constant term. If the null of non-stationarity is rejected, we stop and use the series without any transformation; conversely, if the null is maintained, we use an ADF test that includes both a constant and a linear time trend. If the null of this second test is rejected, we linearly detrend the series by using the residuals of a regression of the series on a constant and a time trend; otherwise we compute a final ADF test that includes a constant, a linear trend and a quadratic trend. If we reject the null of this test, we detrend the series but include a quadratic trend in the regression.

¹³It is well-known that house prices display strong seasonal variation with high prices during spring and summer and low prices during fall and winter. Section A.13 in the Online Appendix shows that a non-seasonally adjusted housing search index to a large extent captures the seasonal component in house prices.

R^2 -values using lags rather than leads of the *HSI*, suggesting that movements in the *HSI* precede movements in the FHFA house price index. The strongest statistical relation between the *HSI* and changes in house prices occurs at lags of the *HSI* ranging from one through four months. At these lags, the predictive power of the *HSI* over monthly house price changes is more than 50%. Leads of the *HSI* are also significantly related to house price changes, but increasing the lead length substantially reduces the magnitude of the slope coefficient, the degree of statistical significance, and R^2 -values.

Table 1 shows results from tests of bi-directional Granger causality between the *HSI* and house price changes. Regardless of lag length, we generally find that the Granger causality runs from the *HSI* to house price changes and not the other way around, once two or more lags are included. Overall, the results indicate that the *HSI* is a leading indicator of subsequent changes in house prices – a point we explore more in-depth in Section 3.

2.4 Housing Search and Transactions

If online search activity provides an accurate signal about peoples’ intentions of buying a house, we should expect to find a positive relation between *HSI* and subsequent house sales. To explore this relation, Panel B of Figure 1 displays *HSI* along with monthly sales of existing housing units from the National Association of Realtors (NAR). The figure shows a strong positive relation between online search activity and house sales, which supports the conjecture that people only engage in a costly search process if they have true intentions of completing a transaction. The figure also shows that *HSI* tends to lead home sales, as we observe a substantial decrease in search activity prior to the large drop in house sales leading up to the financial crisis and likewise an increase in search activity prior to the increase in sales in 2009 and 2011-2012. Even in the unexpected event of the Covid-19 pandemic, we see how *HSI* leads house sales first with a small decrease in search activity as an initial response to the outbreak of the disease followed by a historically high degree of housing-related online search mirroring the development in transactions.

To evaluate the lead-lag relation between *HSI* and house sales, we undertake a similar regression analysis as that performed in equation (1):

$$sales_t = \alpha + \beta HSI_{t+j} + \varepsilon_t, \quad j = -12, \dots, 12, \quad (2)$$

where $sales_t$ is the sales of existing single-family housing units from NAR in month t . Panel B of Figure 2 shows the slope coefficients, associated t -statistics and R^2 -values as functions of j . The large values for $j < 0$ strongly suggest that search activity leads house sales. In contrast, we see no discernible relation between sales and future search activity, suggesting that increased sales activity does not prompt an increase in the volume of searches for buying a house. Consistent with this, the Granger causality tests in Table 1 imply that the HSI helps to forecast home sales, while the reverse does not hold.¹⁴

Taken together, Panels A and B in Figure 2 suggest that online housing search volume leads both house prices and home sales but that the lead times are very different, being notably shorter (1-4 months) for house prices than for actual home sale transactions (10-12 months).

2.5 Demand Interpretation

It is important to validate our interpretation of HSI as a measure of latent demand. As mentioned in Section 2.1, although measures of housing supply are readily available, direct measures of demand are much harder to obtain. The best alternative measure of demand we could find comes from Redfin, which is one of the largest real estate brokers in the U.S. (<https://www.redfin.com/>). Redfin's Housing Demand Index (HDI) is described by the company as "the industry's first and only measure of housing activity prior to purchase". The index is based on the number of customers requesting home tours and writing offers in major metro areas of the U.S. Tours are weighted by averaging the number of tours per written offer.¹⁵ This measure is arguably also closely related to the measure used by Piazzesi et al. (2020) who use user click data from the real estate website Trulia.com to measure search activity.

We believe Redfin's data is representative of the U.S. population of home buyers for two reasons. First, Redfin's website received more than 24 million unique visitors per month in 2020, and it is currently ranked as the 4th largest online brokerage in the U.S. by online market share. Second, in contrast to other competitors (such a Trulia.com) the company also operates offline using local brokers. This offline activity provides the company with a broader coverage of customers, making

¹⁴Home sales is highly persistent with an AR(1) coefficient of 0.97. As a robustness check, we also conducted the Granger causality tests using the first difference of home sales, which led to the same conclusion, namely that the Granger causality runs from the HSI to home sales and not the other way around.

¹⁵More details on the methodology can be found here: <https://www.redfin.com/news/redfin-housing-demand-index-methodology/>

it more representative of the U.S. population as a whole.

The *HDI* is constructed at a weekly frequency and starts in the first week of 2018. As a first validation we aggregate the *HDI* to a monthly frequency by taking the average value per month and plot it together with our *HSI*. As Panel A in Figure 3 illustrates, the two indices closely follow each other (correlation of 0.83) and both capture the large spike in demand that happened during the latter part of the Covid-19 period. To further validate our interpretation, we exploit the fact that Google Trends can be obtained at a weekly frequency and construct a weekly version of our *HSI* using the same keywords and methodology as we use for the monthly version. Since the *HDI* data at the time of writing spans until the first week of April 2021, we expand the sample for the weekly *HSI* to end at this date, giving us a total of 170 observations.¹⁶ The weekly *HSI* is plotted together with the Redfin’s *HDI* in Panel B of Figure 3. The correlation between the two weekly indices is 0.93.

One natural concern from the above validation exercise is that the high correlation between our *HSI* and Redfin’s *HDI* makes the *HSI* a superfluous measure of demand. However, *HSI* has several advantages over *HDI*. First, Redfin’s *HDI* is only available from the first week of 2018, whereas *HSI* starts in 2004, providing a longer historical sample to draw inference. Second, *HSI* can be constructed in real time for any MSA for which there is Google data, whereas *HDI* is only available (at the time of writing) at the national level. Finally, *HSI* tends to capture demand earlier than Redfin’s *HDI*. In Table 1, we show the results from tests of bi-directional Granger causality between *HSI* and *HDI*. The tests show that lagged values of *HSI* have predictive power over *HDI* and not the other way around. These results suggest that *HSI* is able to capture latent demand earlier than the *HDI*. The most likely explanation for this relation is that people often use Google to explore the housing market before actually committing to a specific realtor.

2.6 Housing Search and Other Housing Market Variables

Housing search activity is likely to be correlated with a variety of other economic variables. It is therefore important to address to what extent we can explain variation in housing search by means of macroeconomic fundamentals and other determinants of outcomes in the housing market. For example, does housing search increase in periods with low interest rates, high employment, good

¹⁶Results using a sample that ends in January 2021, as with the monthly series, are almost identical.

credit conditions, and high sentiment? Moreover, does housing search still predict movements in house prices after controlling for other economic variables?

To better understand the drivers behind housing search, we regress the *HSI* on a set of commonly used housing market determinants. Motivated by studies such as Rapach and Strauss (2009), Plazzi et al. (2010), Ghysels et al. (2013), Bork and Møller (2018), Cox and Ludvigson (2019) and Bork et al. (2020), we include the following set of variables in our analysis:

- Employment (*payrolls*): The year-over-year log employment growth rate (total nonfarm payrolls).
- Inflation (*infl*): The log difference in the Consumer Price Index for all urban consumers (all items).
- Building permits (*permits*): The log difference in new private housing units authorized by building permits.
- Housing starts (*starts*): The log difference in new privately owned housing units.
- Term spread (*term*): The 10-year treasury constant maturity rate minus federal funds rate.
- Mortgage rate (*mort*): The level of the 30-year fixed mortgage rate.
- Price-rent ratio (*pr*): The log ratio of the house price to the rent of primary residences.
- Loans outstanding (*loans*): The log change in commercial and industrial loans outstanding.
- Sentiment (*sent*): Fraction of respondents who answer that now is a "good time" to buy a house from the University of Michigan's Survey of Consumers.¹⁷

In addition, we include broad economic activity indices and risk premium variables:

- The Chicago Fed National Activity Index (*cfnai*).
- The Aruoba, Diebold, and Scotti (2009) Business Conditions Index (*ads*).
- The price-dividend ratio (*pd*): The log ratio of the S&P500 index and twelve month trailing dividends.¹⁸

¹⁷All other variables are from the Federal Reserve Bank of St. Louis (FRED) database.

¹⁸Data are obtained from Robert Shiller's website.

- Risk aversion (*ra*): Risk aversion index from Bekaert et al. (2021).
- Uncertainty (*unc*): Uncertainty index measured in annual volatility units from Bekaert et al. (2021).

Table 2 shows the results from the contemporaneous regression model

$$HSI_t = \alpha + x_t' \beta + \varepsilon_t, \quad (3)$$

where x_t contains the standard housing market determinants either individually in univariate regressions (left column) or combined in a multivariate regression (right column). In the univariate regressions, common house price predictors such as employment, inflation, the term spread, price-to-rent ratio, loans outstanding, and sentiment bear little-to-no relation to the volume of housing search. Building permits and housing starts are both significantly positively related to housing search volume. In addition, stock market based risk premium proxies such as the price-dividend ratio are significantly related to search activity. However, with R^2 values around 3-9%, these variables explain only a very small part of variation in the HSI . The most striking relation is between the level of the mortgage rate and search activity with a strong indication that periods of low mortgage rates coincide with periods of high search activity.

Combining our full list of standard housing market determinants in a multivariate regression (right column), we notice that some of the results in the univariate regressions are reversed. Search activity is now strongly positively related to the price-rent ratio, indicating that search volume tends to increase in times with high valuation ratios. However, even with the full list of standard housing market determinants, we can only explain around 70% of the variation in the HSI . With roughly 30% of the variation in the HSI left unexplained, a relatively large component of time-series movements in the volume of housing search is, thus, uncorrelated with standard activity measures from the housing market.¹⁹

¹⁹Some of the information contained in these variables might represent noise that just happens to be correlated with HSI . We analyze this possibility in Section A.3 in the Online Appendix where we consider a placebo test that generates artificial times series by resampling from the panel of regressors. When regressing HSI on the 14 artificial time series, the median R^2 across bootstrap replications is above 20%.

2.7 Local Housing Search

Online search activity can be used to quantify a local component in housing demand. Specifically, Google Trends can be used to extract search activity that occurs within smaller geographical areas, allowing us to study the importance of housing search in the cross-section of local housing markets. This is an important feature because existing evidence suggests that local market factors help explain movements in house prices across the U.S. (e.g. Del Negro and Otrok, 2007, and Hernández-Murillo et al., 2017).

We later analyze whether the effect from housing search activity on house prices depends on local housing supply. To do this, we use Saiz’s (2010) supply elasticity measures across Metropolitan Statistical Areas (MSAs). Saiz (2010) provides results for the 95 MSAs with a population over 500,000 in 2000. Google defines metropolitan areas slightly differently from the U.S. Office of Management and Budget (OMB) which leads us to exclude 18 MSAs from our analysis. For the remaining 77 MSAs there is a one-to-one mapping between the definitions of Google and OMB.

We define local housing search using the same keywords as for the aggregate U.S. housing market and exploit that Google Trends automatically includes geographical idiosyncrasies of home buyer search patterns in each MSA through the related terms. In this way, the search data will be heavily localized. While search activity for individuals residing in a given MSA counts in the overall search volume for that particular MSA, some individuals may also be interested in buying a home in one of the neighboring MSAs. To allow for such potential moves across MSA borders, we also include search activity in the state in which the MSA is located. Based on the local search activity, we construct local *HSI* using the same targeted PCA approach as for the national index.

To illustrate the differences across local housing markets, Figure 4 shows the local *HSI* along with the growth rate in the local Freddie Mac house price index for Miami and Wichita. Among the 77 MSAs included in our analysis, Miami and Wichita have the lowest and highest supply elasticity, respectively, cf. Saiz (2010). For Miami we see a very similar pattern in house prices as compared to the national market, although with a larger boom-bust cycle. We also observe a very strong relation between the local *HSI* and growth in house prices similar to that found for the national market. In contrast, house prices in Wichita did not experience a boom-bust cycle from 2004 to 2010 and monthly growth rates never stray far away from zero. We should therefore expect differences in the economic effect on local house prices from shocks to local *HSIs*, a point we explore in Section 4.

3 Search Volume and Predictability of House Prices

If online search activity tracks peoples’ intentions of buying a house – and thus proxies for the demand for housing – we would expect increases in the *HSI* to be associated with higher subsequent house prices. Given various frictions in the housing market, an increase in demand is propagated into future periods, which leads sluggish price adjustment in response to an increase in demand (Berkovec and Goodman, 1996, and Carrillo et al. 2015). Figures 1 and 2 support this conjecture by showing a strong positive relation between housing search and future growth in house prices.

To more formally explore the predictive power of housing search with respect to movements in house prices, we estimate predictive regressions

$$p_{t+h} - p_t = \alpha_h + \beta_h HSI_t + \theta'_h Z_t + \varepsilon_{t+h}, \quad (4)$$

where p_t is the log of the FHFA house price index, h is the forecast horizon, *HSI* is the predictive variable, and Z_t is a vector of control variables. We consider four different horizons, namely $h = 1, 3, 6$ and 12 months. To account for autocorrelation in house price growth (Case and Shiller, 1989) and overlaps in the dependent variable in (4) when $h > 1$, we compute bootstrap standard errors from a circular block bootstrap that resamples the data in blocks of consecutive observations, reproducing serial correlation and other dependencies in the data.^{20,21}

Panel A in Table 3 reports the estimate of β_h , the corresponding bootstrap t -statistic in parenthesis, and the R^2 in square brackets. The *HSI* is a strong predictor of future house prices with highly significantly positive slope estimates, consistent with future house prices rising when current search (demand) for housing is high. The predictive power of the *HSI* is high when measured by the R^2 , which ranges from 56% for $h = 1$ to around 70% for $h = 6$. The economic magnitude is also large, as a one standard deviation increase in the *HSI* is associated with a 0.4% increase in expected house price growth at the one-month horizon rising to 4.4% at the one-year horizon. For example, starting from May 2020 and until the last sample point in January 2021, *HSI* stays more than two

²⁰We resample the regressand and the regressor jointly in blocks with an average size of 24, which is close to the optimal block length according to the Politis and White (2001) automatic selection procedure. Section A.1 in the Online Appendix shows results for other choices of the average block length as well as results using Hodrick (1992) t -statistics aimed at circumventing the issue with overlapping data.

²¹We also examined the possibility that the predictive regressions suffer from small sample bias arising from cross-correlation in error terms as studied by Stambaugh (1999). We can rule this out as the innovations in *HSI* are only weakly correlated with those of the predictive regression.

standard deviations above its mean, which gives to rise to monthly predictions of more than 1% each month throughout this period. We see from these results that the effect of *HSI* on predicting house price changes is both statistically and economically significant.²²

Panel B of Table 3 controls for all 14 variables described in Section 2.6. We see that *HSI* retains its statistical significance across all horizons, although its slope coefficient is somewhat reduced.²³ These results suggest that housing search activity carries important information about future house prices over and above that embedded in standard housing market predictors. To further verify this claim, we use the residuals from the multivariate regression in Table 2 to construct a version of the housing search index, HSI^\perp , that is orthogonal to the standard predictors. Panel C in Table 3 shows that the slope coefficients for HSI^\perp remain positive and significant across horizons. This evidence suggests that *HSI* contains complementary information about future house prices that is not subsumed by standard housing market variables. In terms of predictive power, the R^2 generated by HSI^\perp ranges from around 7% to 11% across horizons, which is still sizeable given that we have cleaned *HSI* from all information embedded in 14 control variables.²⁴ Furthermore, we analyze whether HSI^\perp granger causes house price growth rates. As shown in Table 1, lagged values of HSI^\perp predict future house price growth rates, while the reverse is not the case.

Finally, in Panel D of Table 3, we control for the AR(1) component in house price growth. A number of studies have documented that growth in house prices exhibit positive serial dependence, which can arise due to frictions and illiquidity and may also reflect the procedure used to construct the house price indices (Ghysels et al., 2013).²⁵ The results show that *HSI* stays strongly statistically significant when controlling for the AR(1) component. The coefficients on *HSI* are reduced across horizons, but the economic magnitude of the predictability remains substantial. The results suggest that *HSI* contains relevant information about future house price growth rates not already embedded in the AR(1) component.

²² Across the four horizons, *HSI* delivers an R^2 that is at least 10 percentage points higher than the best performing individual search terms such as "buying a house", "when buying a house", "buying a home", "buying a new house", "building a house", and "cost of buying a house" which produce R^2 values around 20-50% across the four horizons. This shows the value added by extracting common information from a broad set of search terms.

²³ Section A.2 in the Online Appendix reports estimation results for the control variables, which are often insignificant.

²⁴ Some of the lost explanatory power may reflect that we control for a large number of variables and some of these variables may by chance explain part of the variance of *HSI* without being relevant variables for predicting house prices. In that sense, the R^2 values reported in Panel C of Table 3 may be viewed as a conservative measures of the additional predictive power gained from using *HSI*.

²⁵ The FHFA calculates their monthly repeat-sales house price index without the use of temporal aggregation. In contrast, the monthly Case-Shiller house price index is based on a three-month moving average window, implying that this index is substantially more autocorrelated than the FHFA index.

3.1 Comparison with Wu and Brynjolfsson (2015)

Our *HSI* focuses on the buying side of the housing market through the chosen keywords. Accordingly, we interpret the search index as a proxy for latent demand. Wu and Brynjolfsson (2015) also consider the use of online search activity to predict house prices and sales. Instead of using specific keywords, they consider predefined search categories supplied by Google Trends, namely “Real estate agencies” and “Real estate listings”. Google classifies search queries into categories using an undisclosed natural language classification engine (Choi and Varian, 2012) and it is unclear how we should interpret these categories other than that they relate to the topic given by the name of the category. Wu and Brynjolfsson (2015) find that these two search categories hold some predictive power for future house prices, but also that prices are more difficult to predict than house sales.²⁶

To facilitate a direct comparison with Wu and Brynjolfsson (2015), Table 4 explores the predictive power of the two search categories “Real estate agencies” and “Real estate listings” and compares these to our *HSI*.²⁷ Panel A shows that the two predefined categories hold some predictive power for growth in house prices with R^2 values ranging from 9% for $h = 1$ to 18% for $h = 12$. Compared to *HSI* this degree of predictive power is, however, of limited magnitude, which is also evident from Panel B, where we use all three search-based predictors simultaneously. These results strongly suggest that a more carefully chosen set of keywords with a clear economic interpretation is important for the predictive power of online search compared to broad search categories.

Panel C in Figure 1 plots the two predefined search categories along with the log growth rate in the FHFA house price index. Compared to Panel A in the same figure, “Real estate agencies” and “Real estate listings” do not capture movements in house prices to the same extent as the *HSI*. In particular, we notice that the predefined categories show an increase in search activity during the first part of the bust period and lag house prices in the second part of that period. They also do not spike after the outbreak of the Covid-19 pandemic. Regressing *HSI* on the two search categories “Real estate agencies” and “Real estate listings” produces an R^2 of only 4%, suggesting that these predefined search categories are only weakly correlated with the *HSI*.

In summary, we confirm Wu and Brynjolfssons (2015) finding that the predefined search categories “Real estate agencies” and “Real estate listings” to some degree can predict future house prices,

²⁶Dietzel (2016) uses a similar approach for real estate subcategories to analyze turning points in housing markets.

²⁷We detrend and deseasonalize the predefined search categories similar to the other search indices as described in Section 2.2.

but also find that their predictive power is limited. A likely explanation of this is that these broad categories reflect both the buying and selling sides of the housing market. Our much stronger prediction results based on the *HSI* suggest an additional explanation, namely that the predefined categories contain too much irrelevant information which reduces their predictive power.²⁸

3.2 Predictability at Longer Horizons

Table 3 covers forecast horizons up to 12 months. Searching for a house is often a lengthy process so it is not surprising that the *HSI* displays strong predictive power also over long horizons up to a year. However, we would also expect that its predictive power declines for very long horizons since home buyers have an incentive to limit the search period to avoid excessively large search costs. To visualize the predictive power over very long horizons, Figure 5 summarizes the slope coefficients, associated *t*-statistics and R^2 values for horizons up to five years ($h = 60$). The figure shows that the *HSI* is a significant predictor of house price growth up to a horizon of roughly five years, but also that the explanatory power steadily declines after its peak at horizons around 3 to 8 months. Our finding that short-horizon effects are larger than long-horizon effects is consistent with the search models of Berkovec and Goodman (1996) and Genesove and Han (2012).²⁹

Our 17-year sample from 2004-2021 means that we only have a limited number of independent observations at the longer horizons. Caution should therefore be exercised when interpreting these results, especially at the longest horizons. However, a decline in the predictive power at a horizon of roughly 8 months seems plausible given the time it typically takes to buy a home from the initial search process to closing the deal. NAR (2020) reports that the typical search time for a home is 10 weeks. Prior to searching for a home, buyers are likely to gather information about the house buying process itself. Once a buyer has found a house, the buyer and seller have to agree on a price, the house must be inspected, and the loan application must be approved, with the latter steps typically taking 40-50 days. As such, the predictive pattern of *HSI* is different from conventional

²⁸We also compared our procedure with that of Beracha and Wintoki (2013), who analyze search activity for a particular MSA by using the search term "real estate i", where "i" is the given MSA. We find that our procedure has considerably stronger predictive power over future house prices than that of Beracha and Wintoki (2013). For example, in Miami, Toledo, and Houston, our local *HSIs* generate R^2 s of 69%, 64%, and 51% at the one-month horizon compared to 20%, 9%, and 0% when using "real estate Miami", "real estate Toledo", and "real estate Houston", respectively. In general, the series of Beracha and Wintoki (2013) show weak co-movement with *HSI*.

²⁹In the dynamic search model of Berkovec and Goodman (1996), a key mechanism is lagged price responses to demand shocks. In their simulations, the price adjustment takes place within a few months following a change in demand.

predictors such as the price-rent ratio for which the predictive power builds up over long horizons.

3.3 Alternative Measures of House Price Changes

So far we have focused on the FHFA purchase-only house price index, which is one of the primary indices used in the literature. To illustrate that the predictive power of *HSI* is not only restricted to the FHFA house price index, we next consider other commonly used house price indices, namely the Case-Shiller national home price index, the Freddie-Mac house price index, and the Zillow home value index for single-family homes. The FHFA, Case-Shiller, and Freddie-Mac indices are similar in the sense that they are all constructed using a repeat sales methodology. In contrast, the Zillow index instead uses a valuation model to estimate prices for individual homes. All four indices differ in geographic coverage, price information source, and weighting scheme to form aggregate indices. These differences could be important when analyzing house price predictability, so Table 5 shows predictive results for each of the four house price measures. Despite the different methodologies used to measure house prices, our results show that the strong predictive power of *HSI* holds for all four house price indices implying that the choice of house price index is less important.

Another question is whether the predictive ability of *HSI* extends to commercial properties which, unlike residential properties, are purchased entirely from an investment perspective.³⁰ Similar to residential real estate, commercial real estate is characterized by various frictions that may induce sluggish price adjustments. The *HSI* is designed to capture demand for residential properties, but may capture common variation in the residential and commercial real estate markets. To investigate this possibility, we use the CoStar composite value-weighted index of commercial properties across the U.S., which is constructed based on the repeat-sale methodology. From Panel E in Table 5, we see that *HSI* significantly predicts the growth rate in prices of commercial properties, although the extent of predictability is lower both in terms of statistical significance and predictive power.

3.4 Out-of-Sample Forecasting Tests

To be practically useful for policy makers and households, it is critical that our *HSI* could have been used to improve forecasts of house prices in real time. Moreover, full-sample predictive regressions

³⁰Plazzi et al. (2010) provide evidence of significant time-variation in expected returns on commercial properties.

such as those reported in Table 3 potentially overfit the data.

To address these issues, we next consider a set of out-of-sample forecasting experiments in which we recursively compute the *HSI* and estimate the coefficients of the predictive model using only information available at the time of the forecast. We use the first three years of our sample (2004-2006) as our initial estimation period and reserve the remaining sample (2007-2021) for out-of-sample testing.³¹

Table 6 reports Campbell and Thompson (2008) out-of-sample R^2 values (R_{OoS}^2) and Diebold and Mariano (1995) t -statistics (t_{DM}) for comparing predictive accuracy against a given benchmark. In each case, R^2 values are computed relative to a "historical average" benchmark that assumes constant growth rates in house prices. The null hypothesis is $R_{OoS}^2 \leq 0$, while the alternative hypothesis is $R_{OoS}^2 > 0$.

We find that the *HSI* is able to explain more than 50% of the out-of-sample variation in next month's growth in aggregate house prices. The predictive power increases with the forecast horizon and reaches its peak for $h = 3$ with $R_{OoS}^2 = 65\%$, declining to $R_{OoS}^2 = 54\%$ for $h = 12$. Diebold-Mariano tests strongly reject the null hypothesis that $R_{OoS}^2 \leq 0$ at all forecast horizons.

Further, the table shows that forecasts based on the *HSI* strongly outperform forecasts based on popular determinants of house prices across all horizons. In most cases, these variables generate negative R_{OoS}^2 statistics. Exceptions include building permits and housing starts, but in both cases the R_{OoS}^2 statistics are not significantly positive and never exceed 3%. For $h = 12$, sentiment generates the largest R_{OoS}^2 statistic of 14% among the standard predictors, but again it is not statistically significant according to the Diebold-Mariano test. Stock market based risk premium measures such as the price-dividend ratio do not outperform the historical mean benchmark. The same goes for commonly used business cycle indicators such as *cfnai*.³² These results make it less likely that the predictive ability of *HSI* stems from a typical risk compensation channel.

To assess if the strong predictive power of *HSI* is restricted to certain periods in time, we follow Welch and Goyal (2008) and plot the difference in the cumulative sum of squared forecast errors

³¹We use an expanding estimation window but obtain similar results with rolling windows. When generating the out-of-sample forecasts, we account for a two-month publication lag in house prices.

³²These variables sometimes generate extremely negative R_{OoS}^2 statistics, which typically arises from very substantial movements in the variables during the Covid-19 period, but with predictions in the opposite direction of the movements in the housing market.

(CSSFE) for $h = 1$ in Panel A of Figure 6. The benchmark is again constant growth rates in house prices. An upward sloping CSSFE implies that *HSI* delivers better forecasts than the benchmark and vice versa if the CSSFE is downward sloping. Figure 6 shows that *HSI* holds important information about future house prices irrespective of the market conditions, but also that online search activity is especially useful in turbulent times as witnessed under the financial crisis in 2007-2009 and the Covid-19 pandemic.

Finally, we compare the predictive ability of *HSI* to that of an AR(1) model, which captures the positive persistence in house price growth rates. As shown in Panel A of Table 6, *HSI* generates higher R_{OoS}^2 values than the AR(1) model across all horizons. To more formally compare *HSI* with the AR(1) model, we use forecast encompassing tests (Chong and Hendry, 1986):

$$y_{t+h} = \varpi_h^{HSI} \hat{y}_{t+h}^{HSI} + \varpi_h^{ar1} \hat{y}_{t+h}^{ar1} + \varepsilon_{t+h}, \quad (5)$$

where $y_{t+h} = p_{t+h} - p_t$ is the realized h -month ahead log price growth rate and \hat{y}_{t+h} is the forecasted value using either the *HSI* or the AR(1) component. We implement the test by estimating:

$$e_{t+h}^{ar1} = \varpi_h^{HSI} (e_{t+h}^{ar1} - e_{t+h}^{HSI}) + u_{t+h}, \quad (6)$$

where $e_{t+h} = y_{t+h} - \hat{y}_{t+h}$ is the forecast error. We test the null hypothesis that $\varpi_h^{HSI} = 0$, which would imply that the AR(1) forecast encompasses the forecast of the *HSI*. We also estimate the reverse regression and test whether $\varpi_h^{ar1} = 0$. Panel B of Table 6 reports the results. The estimates of ϖ_h^{HSI} are strongly statistically significant across all horizons, implying that *HSI* contains relevant information beyond what is contained in the AR(1) forecast. Moreover, the weight on forecasts from the *HSI* model exceeds the weight on forecasts from the AR(1) model at all horizons.

In conclusion, our out-of-sample analysis confirms the strong in-sample predictive ability of the *HSI* and shows that online search activity is a consistently strong predictor of future house prices in turbulent as well as in calmer periods. The analysis also emphasizes the strong predictive power of *HSI* compared both to standard house price determinants that generally have difficulties predicting future house prices out-of-sample as well as forecasts from an AR(1) model.³³

³³Section A.6 in the Online Appendix reports results from a bootstrap analysis, which further supports the strong out-of-sample predictive power of *HSI* over future movements on house prices.

4 Variation in Search across Local Housing Markets

National accounts data are often limited in geographic scope, and a key advantage of Google Trends data is that they have few geographical restrictions. This fact is particularly important for our analysis because housing markets are local in nature and we would not expect nationally aggregated data to capture all the complexities of local housing market dynamics.

To explore the predictive power of local versions of the *HSI*, we estimate MSA-level regressions,

$$p_{it+h} - p_{it} = \alpha_i + \beta_i HSI_{it} + \varepsilon_{it+h}, \quad (7)$$

where p_{it} is the log of the Freddie Mac house price index and HSI_{it} is the housing search index, both for MSA i in month t . Figure 7 summarizes the results through a scatter plot of the estimated slope coefficients (β_i) versus R_i^2 values across the 77 MSAs introduced in Section 2.7. To ease comparisons across MSAs all search indices are standardized and the slope coefficients are multiplied by 1,200 to measure the annualized change in house prices after a one standard deviation change in search activity. For brevity, we only present results for $h = 1$, but the conclusion is robust across longer forecast horizons as we will verify in a panel setting in Section 4.1. The strong predictive power of the *HSI* at the national level reappears in individual local housing markets with slope coefficients that are significantly positive for all but one MSA and with 54 MSAs generating R^2 values above 40%.

Across the 77 MSAs, the estimated slope coefficients range from 0.06 (Scranton) to almost 15 (Stockton) on an annualized basis. This implies a large dispersion in the economic effect on local house prices from shocks to demand as proxied by search activity. For example, a one standard deviation increase in the local *HSI* leads to an annualized 11.9% increase in expected house price growth in Miami the following month, while the corresponding response is only 0.1% in Wichita.

4.1 Local Variation in Supply Elasticities

The effect of changes in demand on prices should, in theory, be stronger in markets where the supply response is more inelastic compared to markets with a relatively flat supply curve where the supply response is more elastic. For example, in the search-model of Novy-Marx (2009), the

amplification effect of a shock to demand is stronger in markets where agents are less responsive. Following Saiz (2010), we therefore analyze whether the effect of HSI on house prices is stronger in MSAs with a more inelastic housing supply. We start by estimating predictive panel regressions, which allow us to analyze the average predictive relationship across all MSAs. In particular, we regress the h -month-ahead log house price growth in MSA i on the lagged housing search index in MSA i , constraining the slope coefficients to be identical across MSAs but allowing for individual MSA-specific fixed effects, i.e., imposing $\beta_i = \beta_j$ in (7). Following Thompson (2011), we compute standard errors that are robust to heteroskedasticity as well as correlation along both the time and MSA dimensions. Panel A in Table 7 shows the results. Local HSI significantly predicts local house price growth rates across all horizons. The predictive power of the local HSI as measured by the within- R^2 continues to be very large and is roughly 35% across all four horizons. Moreover, consistent with the national evidence, increased local housing search activity is associated with positive future growth rates in local house prices.

We next interact the local HSI with the degree of housing supply elasticity as computed by Saiz (2010). This allows us to analyze whether house prices in MSAs with a more inelastic housing supply react stronger to changes in housing demand as measured by search activity. To test this effect, we estimate

$$p_{it+h} - p_{it} = \alpha_i + (\beta + \beta_E \times Elasticity_i) HSI_{it} + \varepsilon_{it+h}, \quad (8)$$

where $Elasticity_i$ is the supply elasticity measure of Saiz (2010). Panel B in Table 7 shows the results. The results show that the relation between variation in local search and house prices is significantly stronger in MSAs with low supply elasticity compared to those with high supply elasticity. To visualize these results, Figure 7 shows the ten most supply-constrained MSAs in red and the ten least supply constrained MSAs in green. We see that there is some degree of clustering of the MSAs in accordance with the panel results in Table 7.

In Panel C of Table 7, we include local control variables. In general, access to fundamental variables at the MSA-level is quite limited, especially at the monthly frequency. We include local employment growth (*payrolls*), the local price-rent ratio (*pr*), and local realized volatility as a measure of uncertainty (*unc*).³⁴ We compute the realized volatility measure using a rolling window of three

³⁴Data on employment and rent of primary residences are available from the Bureau of Labor Statistics. While employment data are available at the MSA-level, rental data are not available across all MSAs. We therefore use

months based on squared de-measured returns ³⁵ From the results in Panel C, we see that the inclusion of these control variables leads to only very small changes in slope coefficients on *HSI*. These results suggest that *HSI* contains predictive content beyond that contained in typical risk-compensation variables.

In conclusion, our results suggest that variation in local housing demand as proxied by our search index possesses strong predictive power over growth rates in local house prices. Moreover, changes in local housing demand have a larger economic impact on house prices in MSAs with a more constrained supply of housing.

4.2 National-Level versus MSA-Level Search

To analyze the extent to which housing markets are influenced by local search dynamics relative to national search activity, we next augment the panel regression model with the national-level *HSI*. As Panel A in Table 8 shows, local housing search stays statistically significant across all forecast horizons after controlling for national-level housing search. These results imply that housing markets are strongly influenced by local search dynamics, consistent with findings in the literature that housing markets are local in nature (Del Negro and Otrok, 2007, Gyourko et al., 2013, Glaeser et al., 2014, and Hernández-Murillo et al., 2017).

We also evaluate if the effect of local supply elasticity is affected by variation in the national-level housing search. Panel B of Table 8 shows that the effect of local-level search remains stronger in MSAs with low supply elasticity. In addition, Panel C of Table 8 shows that these results are largely unaffected by including control variables.

4.3 Local Variation in Time-on-Market

To shed further light on the economic channel that generates the predictive ability of *HSI*, we analyze the relation between *HSI* and the time it takes for a home to be sold. According to the search-based models of e.g. Stein (1995), Krainer (2001), Novy-Marx (2009), and Diaz and Jerez (2013), there is a negative correlation between price and time-on-market (*TOM*). Likewise, the

rental data at the regional level (West, North Central, Northeast, and South) and map the MSAs to the region in which the MSA is located.

³⁵This way of computing realized volatility follows from high-frequency measures of e.g. Andersen et al. (2001).

search model of Genesove and Han (2012) in which sellers react to demand shocks with a delay implies that TOM decreases following a shock to demand. Thus, given that HSI proxies for demand, we expect to see a negative relation between HSI and TOM .

To examine if this holds, we estimate panel regressions,

$$TOM_{it+h} = \alpha_i + \beta HSI_{it} + \varepsilon_{it+h}, \quad (9)$$

where TOM_{it} is the average time-on-market measured in days in MSA i in month $t + h$.³⁶

Consistent with predictions from theoretical search models, Table 9 shows that there is a negative relation between TOM and HSI . The effect of HSI on TOM as measured by the magnitude of the predictive coefficient and the within- R^2 is strongest at the one and three month horizons. At these horizons, HSI is strongly statistically significant with two-way clustered t -statistics of -3.66 and -3.72 , respectively. HSI remains statistically significant at the six month horizon but turns insignificant at the 12-month horizon. These results imply that an increase in search activity leads to a decrease in time-on-market with an effect that peaks at relatively short horizons. This evidence appears in line with the analysis Genesove and Han (2012) who find that short-run effects are stronger than long-run effects.

4.4 Out-of-Sample Forecasts

We next examine out-of-sample forecasts of house prices in each of the 77 MSAs with the first forecast made for 2007:1 and the last for 2021:1. We compute out-of-sample forecasts using both the MSA-level and national-level HSI s as predictive variables and use recursive estimation with an expanding window. Panel B in Figure 6 gives an overview of the out-of-sample R^2 values across MSAs. R_{OoS}^2 values are systematically high across MSAs and forecast horizons. At the one-month horizon, the median R_{OoS}^2 is 57% while the first and third quartile values are 49% and 62% across MSAs. The extent of predictability reaches its peak at the three-month horizon with a median R_{OoS}^2 of 57%, which is only slightly reduced at longer horizons. For the annual horizon, the median R_{OoS}^2 is 43%, while the first and third quartile values are 24% and 55%.

³⁶From the Zillow database, we have obtained time-on-market data for 72 out of 77 MSAs in our cross section over the period 2018:1 to 2020:1.

Panel C in Figure 6 visualizes the out-of-sample performance at the one-month horizon by plotting the median value of the cumulative squared forecast error for the no-predictability benchmark and that of the online search model. The out-of-sample performance of online search has been stable over time as *HSI* succeeds in consistently outperforming the no-predictability benchmark. Of special interest is the period around the outbreak of the Covid-19 pandemic which triggered an initial drop in house prices, quickly followed by rapidly rising house prices. Consistent with the findings for national-level house prices, Panel C illustrates that the MSA-level predictive power of online search actually strengthens during the pandemic.

The predictive power of online search is consistently strong over time, but gets even stronger during periods of turmoil such as the financial crisis in 2007-2009 and the Covid-19 pandemic. Theoretical models of search imply that the effect of shocks to demand is amplified and can lead to excess volatility (e.g. Novy-Marx, 2009 and Anenberg and Bayer, 2020). This mechanism of search models may explain why we find the largest gains in predictability relative to the historical mean benchmark during times of high price volatility.

Because *HSI* has the potential to act as an early indicator of where the market is going, it is important to examine how *HSI* performs under different market conditions and price paths. We therefore separately analyze the out-of-sample predictive power of *HSI* during periods of upturns and downturns in the housing market. Following Burnside et al. (2016), we define \bar{p}_{it} as the centered moving average of the log house price in MSA i at time t :

$$\bar{p}_{it} = \frac{1}{2n + 1} \sum_{j=-n}^n p_{it+j}.$$

An upturn is then given as an interval of time for which $\Delta\bar{p}_{it} > 0$ for all t while a downturn is an interval of time for which $\Delta\bar{p}_{it} < 0$ for all t .³⁷ In Panel D of Figure 6, we show the median R_{OoS}^2 across MSAs during periods of upturns and downturns. While *HSI* generates strong out-of-sample predictability in both upturns and downturns, R_{OoS}^2 increases during periods of downturns. Thus, the predictive power of *HSI* is robust across different market conditions but is strengthened when the housing market is in a downturn.

We next analyze whether the predictability mainly is concentrated in MSAs with either high or low levels of housing market volatility. We sort the 77 MSAs in two groups based on local standard

³⁷We set n to equal 5 months but obtain similar results with other reasonable choices of n .

deviations of house price growth rates. Panel E of Figure 6 shows that out-of-sample predictability is slightly higher in housing markets with high volatility but is generally strong across forecast horizons in both low and high volatility markets. Overall, there is strong and robust evidence that *HSI* is a useful out-of-sample predictor of house prices across MSAs.

4.5 Economic Sources of Predictability of House Prices

Our results show that there is substantial variation in search activity over time and across MSAs and that search activity precedes movements in future house prices – a finding that is in line with the theoretical search-and-matching modelling framework (e.g. Carillo et al., 2015).

To shed further light on the degree of time-variation in expected house price changes, we identify episodes of intense housing search activity based on a threshold of one standard deviation of the local *HSI*. Across MSAs, we identify 2,321 months with housing search activity more than one standard deviation above the (local) mean. Panel A in Figure 8 shows the median house price change following these episodes as well as the 1st and 3rd quartiles. From the figure, we see that the median growth rate in house prices following periods with high search is 0.8% at the one-month horizon, 2.1% at the three-month horizon, 4.1% at the six-month horizon and 7.0% at the annual horizon. These realized house price changes suggest that significant economic gains can potentially be achieved from being an early buyer in a market with increasing demand.

Panel A also illustrates that the potential savings from buying a house h months early in an increasing market varies strongly across MSAs. When search is one standard deviation above the mean, the 25th percentile growth rate is 0.6%, 1.7%, 3.3%, and 5.9% at the one-month, three-month, six-month, and one-year horizons, while for the 75th percentile, the price changes are 1.0%, 2.8%, 5.2%, and 8.7%, respectively. MSAs in the 1st quartile have an average supply elasticity of about 2.5, while those in the 4th quartile typically are more inelastic markets with an average supply elasticity of about 1.2.

Because the *HSI* captures local changes in housing demand, we would also expect episodes with low search activity to coincide with subsequent downward pressure on house prices. We identify periods with low search activity as months in which the local housing search falls one standard deviation below its mean. Across MSAs, we find 2,792 events with low search activity. Panel B

shows that episodes with low *HSI* are associated with a subsequent median decline in house prices of -0.2% , -0.7% , -1.2% , and -2.2% at the one-month, three-month, six-month, and one-year horizons. This suggests that it is risky to buy early during times when the *HSI* is low. That is, there is a strong incentive to suspend or reduce search efforts because it is likely it will become possible to buy a house at a lower price, the longer the buyer waits.

Overall, our results suggest that the potential economic gains from exploiting predictability in house prices as identified by the *HSI* can be quite large. However, frictions can generate trading delays, implying that not all agents seeking to buy a new home can do so right away. Furthermore, we cannot entirely rule out that search intensity comove with a time-varying risk-premium component such that episodes of intense search reflect a high required premium for buying a house at the relevant point in time. However, we offer two pieces of evidence that run counter to the time-varying risk premium interpretation.

First, Table 2 shows that proxies for time-varying risk premia such as the price-dividend ratio and the risk aversion and uncertainty indexes of Bekaert et al. (2021) only explain a very small portion of the variation in the *HSI*. Moreover, Table 6 shows that, in contrast with the *HSI*, these risk premium proxies produce very poor out-of-sample forecasts of changes in house prices.

Second, if variation in the *HSI* reflects a time-varying risk premium, we would expect its predictive power over residential house prices to carry over to the REIT market. In contrast, if the predictive power of the *HSI* arises from search frictions, it should hold little or no predictive power for publicly traded REITs whose prices are largely unaffected by search frictions. In Section A.12 in the Online Appendix, we show that the *HSI* has very limited predictive ability over REIT returns.

These findings suggest that the predictive power of the *HSI* over future house prices does not arise from a risk compensation channel, but is more likely to reflect sluggish price adjustments in the residential real estate market due to frictions.

5 Search and House Prices During Covid-19

The Covid-19 pandemic triggered the sharpest reduction in economic activity recorded in modern times. Despite this massive contraction, in 2020 U.S. house prices experienced their largest gains

since 2005 (Mahertz, 2021).³⁸ Recessions are traditionally accompanied by stagnant or declining housing markets, so the increased house prices has led to wide speculation by experts and the media about its possible causes (Demsas, 2021, Friedman, 2021a, Passy, 2021). Some authors suggest that the primary cause of higher prices is supply constraints since the number of homes for sale and new houses built across metropolitan areas in the U.S. plummeted during 2020 (Badget and Bui, 2021, Friedman, 2021b). Others have pointed towards factors affecting demand (Romem, 2020), including falling mortgage rates, increased preference for more space and suburban housing resulting from a shift towards working from home and the increased adoption of new technologies that facilitate and accelerate the buying process due to social distancing norms.

From the perspective of theoretical models of search, a possible interpretation is that the Covid-19 pandemic caused a shock to demand due to a shift in housing preferences, which implied an increase in buyers entering the market and hence an increase in the bargaining power of sellers (e.g. Carillo et al., 2015). Due to sluggish price adjustments arising from frictions, the demand shock influences prices not only on initial impact but also in subsequent periods. In addition, the feedback effect described by Novy-Marx (2009) may have substantially magnified the initial shock, especially due to the tight supply constraints.

A combination of supply and demand shocks thus appear to have affected house prices during the pandemic. From a policy perspective it is important to quantify their respective effects since the optimal policy response depends heavily on their relative importance.³⁹ To address this point, we estimate the following panel regression

$$p_{it+1} - p_{it} = \alpha_i + \beta_D HSI_{it} + \beta_S S_{it} + \gamma' Z_{it} + \varepsilon_{it+1}, \quad (10)$$

where $p_{it+1} - p_{it}$ is the one-month change in the log house price index for MSA i in month $t + 1$, HSI_{it} is our housing search index, S_{it} is a proxy for the housing supply given by the Zillow for-sale-inventory of houses, and Z_{it} is a vector of controls including the Covid-19 stringency index of Hale et al. (2021) and the number of Covid-19 cases, all measured for MSA i in month t .

Table 10 shows results from estimating (10) over the period from February 2020 to January 2021,

³⁸At 11.3% in January 2021, trailing 12-month returns on the national FHFA index recorded their highest value in our sample. This compares with a pre-Covid-19 maximum of 10.1% for September 2005.

³⁹A recent report from OECD (2020) on the housing market during Covid-19 notes that policy responses to curb housing demand can affect the long-run supply.

the end of our sample. While the short time span is a concern, the cross-sectional dimension of our data helps to achieve more powerful tests and reliable estimates.⁴⁰ For comparison with the full sample results (see Table 7, Panel A for $h = 1$), we first estimate (10) using only *HSI* as a proxy for demand (column 1). The slope coefficient during the Covid-19 period is slightly lower than for the full sample (0.33 vs. 0.42), while the R^2 is slightly higher (40.9% vs. 37.2%), demonstrating that the *HSI* remained a robust predictor of growth in house prices during the pandemic. Next, we include supply (column 2) and finally also the lockdown stringency measure of Hale et al. (2021) (column 3) and the number of Covid-19 cases (column 4) as control variables. Across all specifications, the *HSI* and supply measures are highly statistically significant with the expected signs. For example, in the full specification, the coefficient estimates on the *HSI* and supply measures are 0.13 and -0.16. A one standard deviation change in the *HSI* is thus associated with a 1.6% predicted (annualized) change in house prices, while the corresponding impact from a supply change is 1.9%.

Interestingly, Covid-19 restrictions on their own were associated with large negative changes in house prices: the estimated slope coefficient on the stringency index of Hale et al. (2021) is -0.25 with a t -statistic exceeding ten. Controlling for Covid-19 restrictions is clearly important as their introduction leads to a reduction in the estimated slope coefficient on the *HSI* index from 0.33 to 0.13. These results suggest that tight supply constraints and increased demand for houses combined to lead to higher house prices during the pandemic.⁴¹

6 Concluding Remarks

In this paper, we show that online data on search for housing can be used to accurately quantify variation in the demand for housing both at the national (U.S.) and regional (metropolitan) level. Moreover, such data can be used to robustly predict changes in house prices, both in-sample and out-of-sample, at short and long-term horizons, and in periods with rapidly or slowly changing house prices.

Our housing search index produces significantly more accurate forecasts of house prices than con-

⁴⁰The cross-sectional dimension consists of 72 MSAs for which we have data on both prices, supply and demand. The control variables in this regression are available only at the state level, so we map the state-level data to the individual MSAs for these variables.

⁴¹Section A.10 in the Online Appendix provides additional results on the Covid-19 period.

ventional measures of variation in housing demand such as employment, interest rates, sentiment, or proxies for risk. These variables only provide a partial account of housing demand and embed much less information about future house prices than our more direct measure obtained from search activity which reflects peoples' interest in buying a house regardless of whether the motive is based on fundamentals or is of a more speculative nature.

Our findings of strong predictability of future changes in house prices do not suggest arbitrage opportunities and also do not appear to be driven by time-varying risk premia. Instead, our results are more consistent with search-based models with frictions which imply that shocks affecting the housing market will only be reflected in future house prices through a gradual adjustment process.

References

- Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P. (2001). The distribution of realized exchange rate volatility. *Journal of American Statistical Association* 96, 42-55.
- Andrei, D., Hasler, M. (2015). Investor attention and stock market volatility. *Review of Financial Studies* 28, 33-72.
- Anenberg, E., Bayer, P. (2020). Endogenous sources of volatility in housing markets: the joint buyer-seller problem. *International Economic Review* 61, 1195-1228.
- Ayat, L., Burridge, P. (2000). Unit root tests in the presence of uncertainty about the non-stochastic trend. *Journal of Econometrics* 95, 71-96.
- Badgett, E., Bui, Q. (2021). Where have all the houses gone? *The New York Times*, May 14.
- Bai, J., Ng, S. (2008). Forecasting economic time series using targeted predictors. *Journal of Econometrics* 146, 304-317.
- Bekaert, G., Engstrom, E.C., Xu, N.R. (2021). The time variation in risk appetite and uncertainty. *Management Science*, forthcoming.
- Beracha, E., Wintoki, M. (2013). Forecasting residential real estate price changes from online search activity. *Journal of Real Estate Research* 35, 283-312.
- Berkovec, J.A., Goodman, J.L. (1996). Turnover as a measure of demand for existing homes. *Real Estate Economics* 24, 421-440.
- Bijl, L., Kringhaug, G., Molnár, P., Sandvik, E. (2016). Google searches and stock returns. *International Review of Financial Analysis* 45, 150-156.
- Bork, L., Møller, S.V. (2018). Housing price forecastability: A factor analysis. *Real Estate Economics* 46, 582-611.
- Bork, L., Møller, S.V., Pedersen, T.Q. (2020). A new index of housing sentiment. *Management Science* 66, 1563-1583.
- Borup, D., Schütte, E. (2022). In search of a job: Forecasting employment growth using Google trends. *Journal of Business & Economic Statistics* 40, 186-200.
- Burnside, C., Eichenbaum, M., Rebelo, S. (2016). Understanding booms and busts in housing markets. *Journal of Political Economy* 124, 1088-1147.
- Campbell, J., Thompson, S. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies* 21, 1509-1531.
- Carrillo, P. E., de Wit, E. R., Larson, W. (2015). Can tightness in the housing market help predict subsequent home price appreciation? Evidence from the United States and the Netherlands. *Real Estate Economics* 43, 609-651.
- Case, K. E., Shiller, R. J. (1989). The efficiency of the market for single family homes. *American Economic Review* 79, 125-137.

- Case, K.E., Quigley, J.M., Shiller, R.J. (2012). Wealth effects revisited 1975-2012. Cowles Foundation Discussion Paper no. 1884. Yale University.
- Choi, H., Varian, H. (2012). Predicting the present with Google Trends. *Economic Record* 88, 2-9.
- Chong, Y., Hendry, D. (1986). Econometric evaluation of linear macroeconomic models. *Review of Economic Studies* 53, 671-690.
- Cox, J., Ludvigson, S.C. (2019). Drivers of the great housing boom-bust: Credit conditions, beliefs, or both? *Real Estate Economics* 2019, 1-33.
- Da, Z., Engelberg, J., Gao, P. (2011). In search of attention. *Journal of Finance* 66, 1461-1499.
- Da, Z., Engelberg, J., Gao, P. (2015). The sum of all FEARS: Investor sentiment and asset prices. *Review of Financial Studies* 28, 1-32.
- D'Amuri, F., Marcucci, J. (2017). The predictive power of Google searches in forecasting US unemployment. *International Journal of Forecasting* 33, 801-816.
- Del Negro, M., Otrok, C. (2007). 99 luftballons: Monetary policy and the house price boom across US states. *Journal of Monetary Economics* 54, 1962-1985.
- Demsas, J. (2021). Covid-19 caused a recession. So why did the housing market boom? *Vox*, Feb 5.
- Diaz, A., Jerez, B. (2013). House prices, sales, and time on the market: A search-theoretic framework. *International Economic Review* 54, 837-872.
- van Dijk, D.W., Francke, M.K. (2018). Internet search behavior, liquidity and prices in the housing market. *Real Estate Economics* 46, 368-403.
- Diebold, F.X., Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 253-263.
- Dietzel, M. (2016). Sentiment-based predictions of housing market turning points with google trends. *International Journal of Housing Markets* 9, 108-136.
- Friedman, N. (2021a). The pandemic ignited a housing boom – but it is different from the last one. *The Wall Street Journal*, March 15.
- Friedman, N. (2021b). U.S. existing-home prices climb on tight supply. *The Wall Street Journal*, April 23.
- Gargano, A., Giacoletti, M., Jarnecic, E. (2021). Local experiences, attention and spillovers in the housing market. *Journal of Finance*, forthcoming.
- Genesove, D., Han, L. (2012). Search and matching in the housing market. *Journal of Urban Economics* 72, 31-45.
- Ghysels, E., Horan, C., Moench, E. (2017). Forecasting through the rearview mirror: Data revisions and bond return predictability. *Review of Financial Studies* 31, 678-714.
- Ghysels, E., Plazzi, A., Valkanov, R., Torous, W. (2013). Forecasting real estate prices. *Handbook of Economic Forecasting*, G. Elliot and A. Timmermann (Eds.), 509-580.
- Glaeser, E.L., Gyourko, J., Morales, E., Nathanson, C.G. (2014). Housing dynamics: An urban approach. *Journal of Urban Economics* 81, 45-56.

- Guo, H. (2009). Data revisions and out-of-sample stock return predictability. *Economic Inquiry* 47, 81-97.
- Gyourko, J., Mayer, C., Sinai, T. (2013). Superstar cities. *American Economic Journal: Economic Policy* 5, 167-199.
- Hale, T., Angrist, N., Goldszmidt, R., Kira, B. Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S., Tatlow, H. (2021). A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nature Human Behaviour* 5, 529-538.
- Han, L., Strange, W. (2015). The microstructure of housing markets: search, bargaining, and brokerage. *Handbook of Regional and Urban Economics*, Edited by Duranton, G., Henderson J.V., Strange, W.C., 5, 813-886.
- Head, A., Lloyd-Ellis, H., Sun, H. (2014). Search, liquidity, and the dynamics of house prices and construction. *American Economic Review* 104, 1172-1210.
- Hernández-Murillo, R., Owyang, M.T., Rubio, M. (2017). Clustered housing cycles. *Regional Science and Urban Economics* 66, 185-197.
- Hodrick, R. (1992). Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial Studies* 5, 357-386.
- Huang, D., Jiang, D., Tu, J., Zhou, G. (2015). Investor sentiment aligned: A powerful predictor of stock returns. *Review of Financial Studies* 28, 791-837.
- Joseph, K., Wintoki, M.B., Zhang, Z. (2011). Forecasting abnormal stock returns and trading volume using investor sentiment: Evidence from online search. *International Journal of Forecasting* 27, 1116-1127.
- Krainer, J. (2001). A theory of liquidity in residential real estate markets. *Journal of Urban Economics* 49, 32-53.
- Mahertz, T. (2021). The U.S. housing market gained more value in 2020 than in any year since 2005. *Zillow*, Jan 26.
- National Association of Realtors (NAR) (2020). Home Buyers and Sellers Generational Trends Report. Technical Report, National Association of Realtors.
- Newey, W., West, K. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703-708.
- Novy-Marx, R. (2009). Hot and cold markets. *Real Estate Economics* 37, 1-22.
- OECD (2020). Housing amid covid-19: Policy responses and challenges. *OECD contributions*.
- Passy, J. (2021). Is the U.S. housing market heading for a crash? Here is what the experts say. *MarketWatch*, May 3.
- Piazzesi, M., Schneider, M., Stroebel, J. (2020). Segmented housing search. *American Economic Review* 110, 720-759.
- Plazzi, A., Torous, W., Valkanov, R. (2010). Expected returns and expected growth in rents of commercial real estate. *Review of Financial Studies* 23, 3469-3519.
- Politis, D.N., White, H. (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews* 23, 53-70.

- Rapach, D., Strauss, J. (2009). Differences in housing price forecastability across US states. *International Journal of Forecasting* 25, 351-372.
- Reinhardt, C.M., Rogoff, K.S. (2009). *This Time is Different: Eight Countries of Financial Folly*. Princeton University Press.
- Romem, I. (2020). Explaining the frenzy in the housing market. *The New York Times*, Dec, 3.
- Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics* 125, 1253-1296.
- Soo, C. K. (2018). Quantifying animal spirits: News media and sentiment in the housing market. *Review of Financial Studies* 31, 3689-3719.
- Stambaugh, R. F. (1999). Predictive regressions. *Journal of Financial Economics* 54, 375-421.
- Statista (2018). Market share of search engines in the United States 2008-2018, available from www.statista.com.
- Stein, J.C. (1995). Prices and trading volume in the housing market: A model with down-payment effects. *Quarterly Journal of Economics* 110, 379-406.
- Thompson, S.B. (2011). Simple formulas for standard errors that cluster by both firm and time. *Journal of Financial Economics* 99, 1-10.
- Vozlyublennaia, N. (2014). Investor attention, index performance and return predictability. *Journal of Banking and Finance* 41, 17-35.
- Welch, I., Goyal, A. (2008). A comprehensive look at the performance of equity premium prediction. *Review of Financial Studies* 21, 1455-1508.
- Wheaton, W.C. (1990). Vacancy, search, and prices in a housing market matching model. *Journal of Political Economy* 98, 1270-1292.
- Wu, L., Brynjolfsson, E. (2015). The future of prediction: How Google searches foreshadow housing prices and sales. In: A. Goldfarb, S. M. Greenstein, and C. E. Tucker (Eds.), *Economic analysis of the digital economy*, Chicago: University of Chicago Press.
- Zou, H., Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of Royal Statistical Society, Series B*, 67, 301-320.

Table 1. Granger Causality Tests. The table reports results from standard Granger causality tests:

$$HSI_t = \delta + \sum_{i=1}^p \theta_i HSI_{t-i} + \sum_{i=1}^p \gamma_i y_{t-i} + \varepsilon_t$$

$$y_t = \delta + \sum_{i=1}^p \theta_i y_{t-i} + \sum_{i=1}^p \gamma_i HSI_{t-i} + \varepsilon_t$$

where HSI is the housing search index and y is either house price changes, home sales, or Redfin’s housing demand index (HDI). HSI^\perp is the part of HSI that is orthogonal to the other predictive variables. The table shows p -values from the joint test that $\gamma_1 = \gamma_2 = \dots = \gamma_p = 0$. We use the Newey and West (1987) estimator with automatic lag selection. For house prices and home sales, the sample period is 2004:1 to 2021:1 with monthly observations, while the sample frequency is weekly for HDI and the sample runs from the first week of 2018 until the first week of April 2021.

Null hypothesis	$p = 1$	$p = 2$	$p = 3$	$p = 4$
HSI does not Granger cause house price changes	0.000	0.000	0.000	0.000
House price changes do not Granger cause HSI	0.022	0.209	0.719	0.669
HSI does not Granger cause home sales	0.002	0.005	0.003	0.004
Home sales do not Granger cause HSI	0.916	0.411	0.226	0.204
HSI does not Granger cause HDI	0.000	0.000	0.000	0.000
HDI does not Granger cause HSI	0.616	0.667	0.661	0.144
HSI^\perp does not Granger cause house price changes	0.002	0.016	0.052	0.066
House price changes do not Granger cause HSI^\perp	0.267	0.485	0.285	0.449

Table 2. The Relation Between Housing Search and Alternative Variables. The table reports results from regressions, $HSI_t = \alpha + x_t'\beta + \varepsilon_t$, where HSI_t is the housing search index and x_t contains standard house price determinants. We show results from univariate regressions using one variable at a time as well as from a multivariate regression. For each regression, the table reports estimates of β , corresponding t -statistics in parenthesis, and the R^2 in square brackets. We compute standard errors using the Newey and West (1987) estimator with automatic lag selection. All variables are standardized to facilitate comparison of the β estimates. The sample period is 2004:1 to 2021:1.

	Univariate	Multivariate
<i>payrolls</i>	-0.22 (-1.17) [4.86]	-0.02 (-0.28)
<i>infl</i>	0.03 (0.26) [0.10]	0.01 (0.22)
<i>permits</i>	0.30 (3.24) [9.13]	0.05 (1.34)
<i>starts</i>	0.16 (2.40) [2.63]	-0.01 (-0.47)
<i>term</i>	-0.04 (-0.29) [0.19]	0.51 (5.43)
<i>mort</i>	-0.57 (-4.57) [32.05]	-1.18 (-12.57)
<i>pr</i>	-0.00 (-0.01) [0.00]	1.09 (6.93)
<i>loans</i>	-0.10 (-0.66) [1.10]	0.15 (2.53)
<i>sent</i>	0.17 (1.20) [2.92]	0.10 (1.25)
<i>cfnai</i>	0.18 (1.01) [3.27]	0.06 (1.09)
<i>ads</i>	0.17 (1.06) [2.89]	0.12 (2.53)
<i>pd</i>	0.28 (2.66) [7.87]	0.04 (0.26)
<i>ra</i>	-0.20 (-2.21) [4.20]	-0.11 (-1.43)
<i>unc</i>	-0.20 (-1.73) [3.89]	0.10 (0.61) [70.99]

Table 3. Predicting House Prices With Housing Search and Alternative Variables.

The table reports results from predictive regressions, $p_{t+h} - p_t = \alpha_h + \beta_h HSI_t + \theta_h' Z_t + \varepsilon_{t+h}$, where p is the log of the FHFA house price index, HSI is the housing search index, Z is a vector of control variables, and h is the forecasting horizon in months. Panel A reports results using HSI on its own (i.e. $\theta_h = 0$). Panel B controls for the 14 predictive variables defined in Section 2.6. Panel C uses the part of HSI that is orthogonal to the other predictive variables (HSI^\perp). Panel D controls for an AR(1) component. For each regression, the table reports the estimate of β , the corresponding t -statistic in parenthesis, and the R^2 in square brackets. We compute standard errors using a circular block bootstrap. All predictive variables are standardized and slope coefficients are multiplied by 100 to facilitate comparison across variables. The sample period is 2004:1 to 2021:1.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A: HSI alone				
HSI	0.43 (6.39) [56.91]	1.26 (6.74) [67.37]	2.44 (6.19) [70.41]	4.37 (5.35) [64.42]
Panel B: HSI and control variables				
HSI	0.27 (4.66) [69.47]	0.77 (3.84) [78.73]	1.56 (3.36) [80.45]	2.31 (3.27) [80.77]
Panel C: Orthogonalized HSI				
HSI^\perp	0.15 (3.02) [7.12]	0.47 (3.01) [9.39]	0.96 (2.90) [10.93]	1.55 (2.89) [8.15]
Panel D: Controlling for AR(1) component				
HSI	0.27 (5.74) [64.35]	0.84 (6.21) [74.71]	1.69 (6.41) [77.88]	3.01 (5.18) [70.65]

Table 4. Predicting House Prices With Alternative Search Indices. The table reports results from predictive regressions, $p_{t+h} - p_t = \alpha + \beta'x_t + \varepsilon_{t+h}$, where p_t is the log of the FHFA house price index, x_t is a vector of predictive variables, and h is the forecast horizon in months. For each regression, the table reports slope estimates, the corresponding t -statistics in parenthesis, and the R^2 in square brackets. We compute standard errors using a circular block bootstrap. All predictive variables are standardized to facilitate comparison of the β estimates and the log price change is multiplied by 100. Panel A shows results for the predefined search categories used by Wu and Brynjolfsson (2015), while Panel B includes the *HSI* jointly with the predefined search categories. The sample period is 2004:1 to 2021:1.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A: Predefined search categories				
Real estate agencies	0.25 (1.86)	0.79 (2.11)	1.75 (2.71)	3.60 (3.23)
Real estate listings	-0.14 (-1.31) [8.97]	-0.44 (-1.34) [11.74]	-1.07 (-1.62) [15.37]	-2.35 (-1.88) [18.21]
Panel B: <i>HSI</i> joint with predefined search categories				
<i>HSI</i>	0.41 (6.40)	1.21 (6.43)	2.33 (5.77)	4.15 (4.60)
Real estate agencies	0.10 (1.61)	0.24 (1.43)	0.46 (1.46)	0.92 (1.16)
Real estate listings	-0.02 (-0.33) [59.16]	-0.06 (-0.41) [69.02]	-0.25 (-0.98) [71.38]	-0.87 (-1.96) [65.49]

Table 5. Alternative House Price Indices. The table reports results from predictive regressions, $p_{t+h} - p_t = \alpha + \beta HSI_t + \varepsilon_{t+h}$, where p_t is the log house price measured using either the FHFA index (Panel A), the Case-Shiller index (Panel B), the Freddie-Mac index (Panel C), the Zillow index (Panel D), or the CoStar commercial property index (Panel E). For each regression, the table reports the estimate of β , the corresponding t -statistic in parenthesis, and the R^2 in square brackets. We compute standard errors using a circular block bootstrap. The sample period is 2004:1 to 2021:1.

$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A: FHFA			
0.43	1.26	2.44	4.37
(6.39)	(6.74)	(6.19)	(5.35)
[56.91]	[67.37]	[70.41]	[64.42]
Panel B: Case-Shiller			
0.46	1.40	2.73	5.12
(4.75)	(5.10)	(5.16)	(4.86)
[53.71]	[59.36]	[62.18]	[62.90]
Panel C: Freddie-Mac			
0.49	1.49	2.88	5.26
(4.89)	(5.29)	(5.26)	(4.92)
[65.38]	[68.34]	[67.92]	[63.74]
Panel D: Zillow			
0.39	1.22	2.51	5.02
(4.39)	(4.74)	(5.20)	(5.71)
[55.96]	[60.97]	[66.38]	[71.19]
Panel E: CoStar			
0.34	1.12	2.56	5.73
(1.96)	(2.20)	(2.29)	(2.18)
[6.26]	[10.37]	[18.84]	[31.32]

Table 6. Out-of-Sample Tests. Panel A reports the Campbell and Thompson (2008) out-of-sample R^2 (R_{OoS}^2) and in parenthesis the p-value from the Diebold and Mariano (1995) t -statistic, computed using the Newey and West (1987) estimator with h lags, where h is the forecast horizon in months. The null hypothesis is that the R_{OoS}^2 is equal to zero or negative and the alternative hypothesis is that it is positive. Panel B reports coefficient estimates from forecast encompassing tests for whether the weights in (5) are equal to zero with p -values shown in parentheses.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A: R_{OoS}^2 statistics				
<i>HSI</i>	51.37 (0.00)	64.54 (0.00)	63.91 (0.00)	54.25 (0.02)
<i>payrolls</i>	-44.66 (0.93)	-94.68 (0.96)	-124.68 (0.95)	-61.46 (0.91)
<i>infl</i>	-4.22 (0.87)	-1.52 (0.77)	-1.15 (0.77)	-0.81 (0.86)
<i>permits</i>	2.91 (0.39)	2.58 (0.40)	0.29 (0.48)	2.36 (0.33)
<i>starts</i>	1.02 (0.33)	1.47 (0.21)	1.65 (0.16)	1.14 (0.22)
<i>term</i>	-5.84 (0.80)	-10.44 (0.81)	-14.92 (0.78)	-21.92 (0.74)
<i>mort</i>	2.35 (0.43)	-5.20 (0.59)	-8.93 (0.64)	-33.21 (0.78)
<i>pr</i>	-12.97 (0.93)	-26.63 (0.98)	-42.68 (0.96)	-70.34 (0.92)
<i>loans</i>	-7.53 (0.79)	-13.91 (0.90)	-23.66 (0.88)	-49.02 (0.87)
<i>sent</i>	-0.49 (0.53)	-2.47 (0.59)	-0.17 (0.50)	13.79 (0.26)
<i>cfnai</i>	-65.29 (0.84)	-109.85 (0.86)	-111.84 (0.88)	-7.45 (0.67)
<i>ads</i>	-91.49 (0.86)	-141.42 (0.88)	-139.03 (0.89)	-4.91 (0.65)
<i>pd</i>	-32.02 (0.77)	-96.27 (0.86)	-179.35 (0.87)	-288.30 (0.87)
<i>ra</i>	-189.95 (0.88)	-392.86 (0.90)	-511.52 (0.93)	-274.89 (0.93)
<i>unc</i>	-23.38 (0.76)	-69.09 (0.83)	-120.67 (0.85)	-253.75 (0.85)
<i>ar1</i>	44.13 (0.00)	55.54 (0.00)	54.95 (0.00)	31.15 (0.14)
Panel B: Encompassing tests				
<i>HSI</i>	0.59 (0.00)	0.60 (0.00)	0.60 (0.00)	0.79 (0.00)
<i>ar1</i>	0.41 (0.00)	0.40 (0.00)	0.40 (0.00)	0.21 (0.13)

Table 7. Predicting Local House Prices With Local Housing Search: Evidence From Panel Regressions. The table reports results from fixed effects panel regressions of the form, $p_{it+h} - p_{it} = \alpha_i + \beta HSI_{it} + \beta_E HSI_{it} \times Elasticity_i + \phi' Z_{it} + \varepsilon_{it+h}$, where p_{it} is the log of the Freddie-Mac house price index in MSA i , HSI_{it} is the housing search index in MSA i , $Elasticity_i$ is supply elasticity in MSA i , Z_{it} contains control variables, and h is the forecast horizon in months. The control variables are local employment growth, the local price-rent ratio, and local realized volatility. In Panels A and B, we set $\beta_E = \phi = 0$ and $\phi = 0$, respectively. For each regression, the table reports the estimate of β , the corresponding t -statistic in parenthesis, and the within R^2 in square brackets. We compute standard errors using Thompson (2011) two-way clustered robust-statistics with h lags. HSI is standardized to facilitate interpretation of the β estimates. The sample period is 2004:1 to 2021:1.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A: Local HSI				
HSI	0.42 (9.43) [37.20]	1.22 (7.59) [37.89]	2.32 (6.26) [36.89]	4.19 (5.58) [33.81]
Panel B: Effect of Supply Elasticity				
HSI	0.65 (7.49)	1.92 (6.32)	3.67 (5.23)	6.58 (4.49)
$HSI \times Elasticity$	-0.12 (-4.02) [40.06]	-0.35 (-3.65) [40.93]	-0.67 (-3.22) [39.99]	-1.20 (-2.79) [36.57]
Panel C: Adding Control Variables				
HSI	0.61 (7.98)	1.81 (6.94)	3.41 (5.70)	5.77 (4.19)
$HSI \times Elasticity$	-0.11 (-4.03) [43.40]	-0.34 (-3.73) [44.05]	-0.67 (-3.30) [44.15]	-1.18 (-2.75) [45.50]

Table 8. Predicting Local House Prices With National and Local Housing Search: Evidence From Panel Regressions. The table reports results from fixed effects panel regressions of the form, $p_{it+h} - p_{it} = \alpha_i + \beta_{US}HSI_{US_t} + \beta HSI_{it} + \beta_E HSI_{it} \times Elasticity_i + \phi' Z_{it} + \varepsilon_{it+h}$, where p_{it} is the log of the Freddie-Mac house price index in MSA i , HSI_{US_t} is the national-level housing search index, HSI_{it} is the housing search index in MSA i , $Elasticity_i$ is supply elasticity in MSA i , Z_{it} contains control variables, and h is the forecast horizon in months. The control variables are local employment growth, the local price-rent ratio, and local realized volatility. In Panels A and B, we set $\beta_E = \phi = 0$ and $\phi = 0$, respectively. For each regression, the table reports the estimate of β , the corresponding t -statistic in parenthesis, and the within R^2 in square brackets. We compute standard errors using Thompson (2011) two-way clustered robust-statistics with h lags. HSI is standardized to facilitate interpretation of the β estimates. The sample period is 2004:1 to 2021:1.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A: U.S. vs. Local HSI				
U.S. HSI	0.28 (10.35)	0.86 (9.78)	1.66 (8.23)	3.01 (5.56)
Local HSI	0.25 (7.95) [48.52]	0.74 (6.90) [50.36]	1.41 (6.30) [49.90]	2.58 (6.86) [46.27]
Panel B: Effect of supply elasticity				
U.S. HSI	0.28 (10.16)	0.85 (9.64)	1.64 (8.20)	2.97 (5.57)
Local HSI	0.48 (6.28)	1.40 (5.41)	2.70 (4.80)	4.86 (4.62)
Local $HSI \times Elasticity$	-0.11 (-4.07) [51.10]	-0.33 (-3.60) [53.11]	-0.64 (-3.18) [52.69]	-1.12 (-2.78) [48.71]
Panel C: Adding Control Variables				
U.S. HSI	0.30 (10.51)	0.88 (9.27)	1.63 (7.89)	2.61 (5.24)
Local HSI	0.43 (6.52)	1.32 (5.61)	2.59 (4.87)	4.62 (4.04)
Local $HSI \times Elasticity$	-0.11 (-4.27) [54.27]	-0.32 (-3.75) [54.80]	-0.63 (-3.23) [53.93]	-1.12 (-2.71) [52.82]

Table 9. Predicting Time-on-Market With Local Housing Search. The table reports results from fixed effects panel regressions of the form, $TOM_{it+h} = \alpha_i + \beta HSI_{it} + \varepsilon_{it+h}$, where TOM_{it} is the time-on-market measured in days in MSA i , HSI_{it} is the housing search index in MSA i , and h is the forecast horizon in months. For each regression, the table reports the estimate of β , the corresponding t -statistic in parenthesis, and the within R^2 in square brackets. We compute standard errors using Thompson (2011) two-way clustered robust-statistics with one lag. HSI is standardized to facilitate interpretation of the β estimates. The sample period is 2004:1 to 2021:1.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
HSI	-3.49	-3.57	-2.82	-0.72
	(-3.66)	(-3.72)	(-2.70)	(-0.82)
	[19.96]	[21.05]	[13.36]	[0.84]

Table 10. The Housing Market During the Covid-19 Pandemic. The table reports results from fixed effects panel regressions of the form, $p_{it+1} - p_{it} = \alpha_i + \beta_D HSI_{it} + \beta_S S_{it} + \gamma' Z_{it} + \varepsilon_{it+1}$, where p_{it} is the log of the Freddie-Mac house price index in MSA i in month t , HSI_{it} is the housing search index, S_{it} is housing supply as measured by the for-sale-inventory, and Z_{it} is a vector of control variables, including the monthly change in Covid-19 restrictions and Covid-19 cases. For each regression, the table reports slope estimates, the corresponding t -statistic in parenthesis, and the within R^2 in square brackets. We compute standard errors using Thompson (2011) two-way clustered robust-statistics with one lag. To facilitate interpretation of the estimates, we standardize all regressors. The sample period is 2020:2 to 2021:1.

<i>HSI</i>	0.33 (9.38)	0.23 (5.99)	0.13 (4.35)	0.13 (4.31)
Supply		-0.16 (-2.05)	-0.11 (-1.92)	-0.16 (-2.22)
Covid-19 restrictions			-0.25 (-9.42)	-0.24 (-7.14)
Covid-19 cases				-0.07 (-1.71)
R^2_{within}	[40.92]	[46.79]	[62.43]	[63.24]

Figure 1. Housing Search Index. Panel A shows the housing search index (HSI) along with the log growth rate in the seasonally adjusted Federal Housing Finance Agency (FHFA) purchase-only house price index. Panel B shows the HSI along with the monthly sales of existing single-family housing units from the National Association of Realtors. Panel C shows search volume for the predefined search categories "Real estate agencies" and "Real estate listings" along with the log growth rate in the FHFA House Price Index. The sample period is 2004:1 to 2021:1.

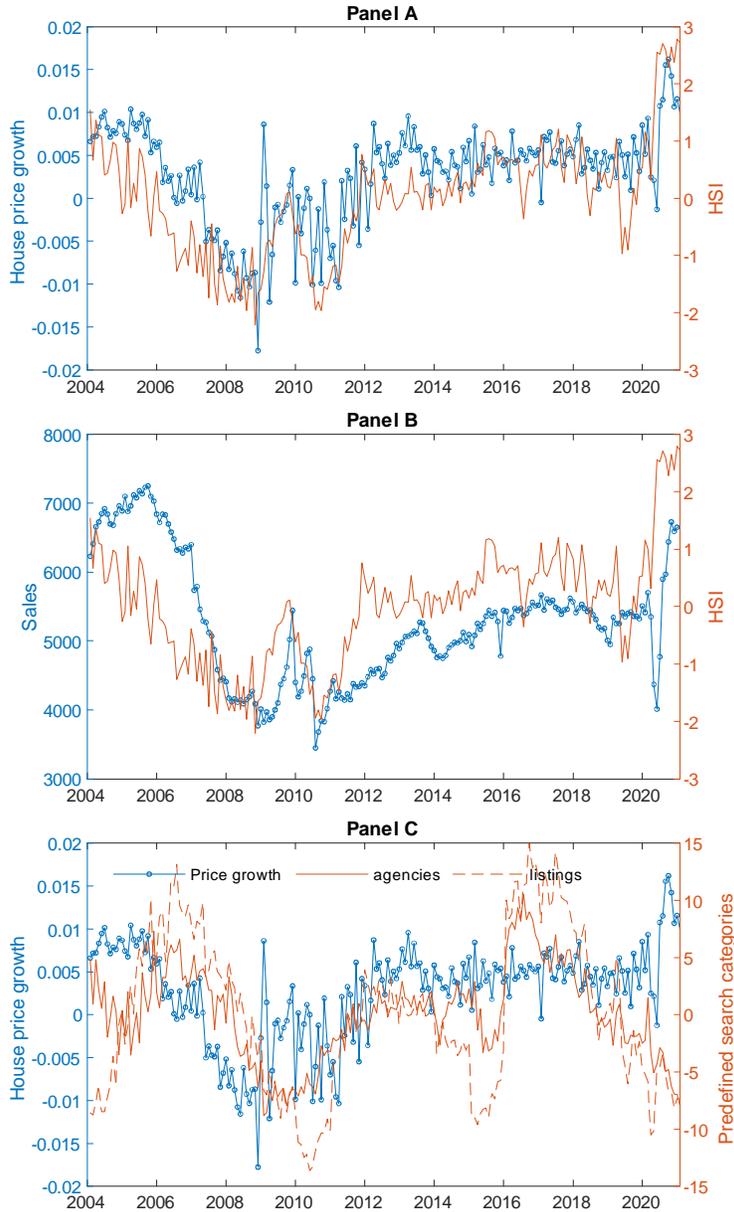


Figure 2. Lead-Lag Relations. Panel A shows regression slope coefficients, associated t -statistics and R^2 values of monthly price changes from $t-1$ to t on HSI_{t+j} for $j \in \{-12, 12\}$. Panel B shows the results from regressing monthly house sales at time t on HSI_{t+j} for $j \in \{-12, 12\}$. Standard errors are calculated using the Newey and West (1987) procedure with 12 lags.

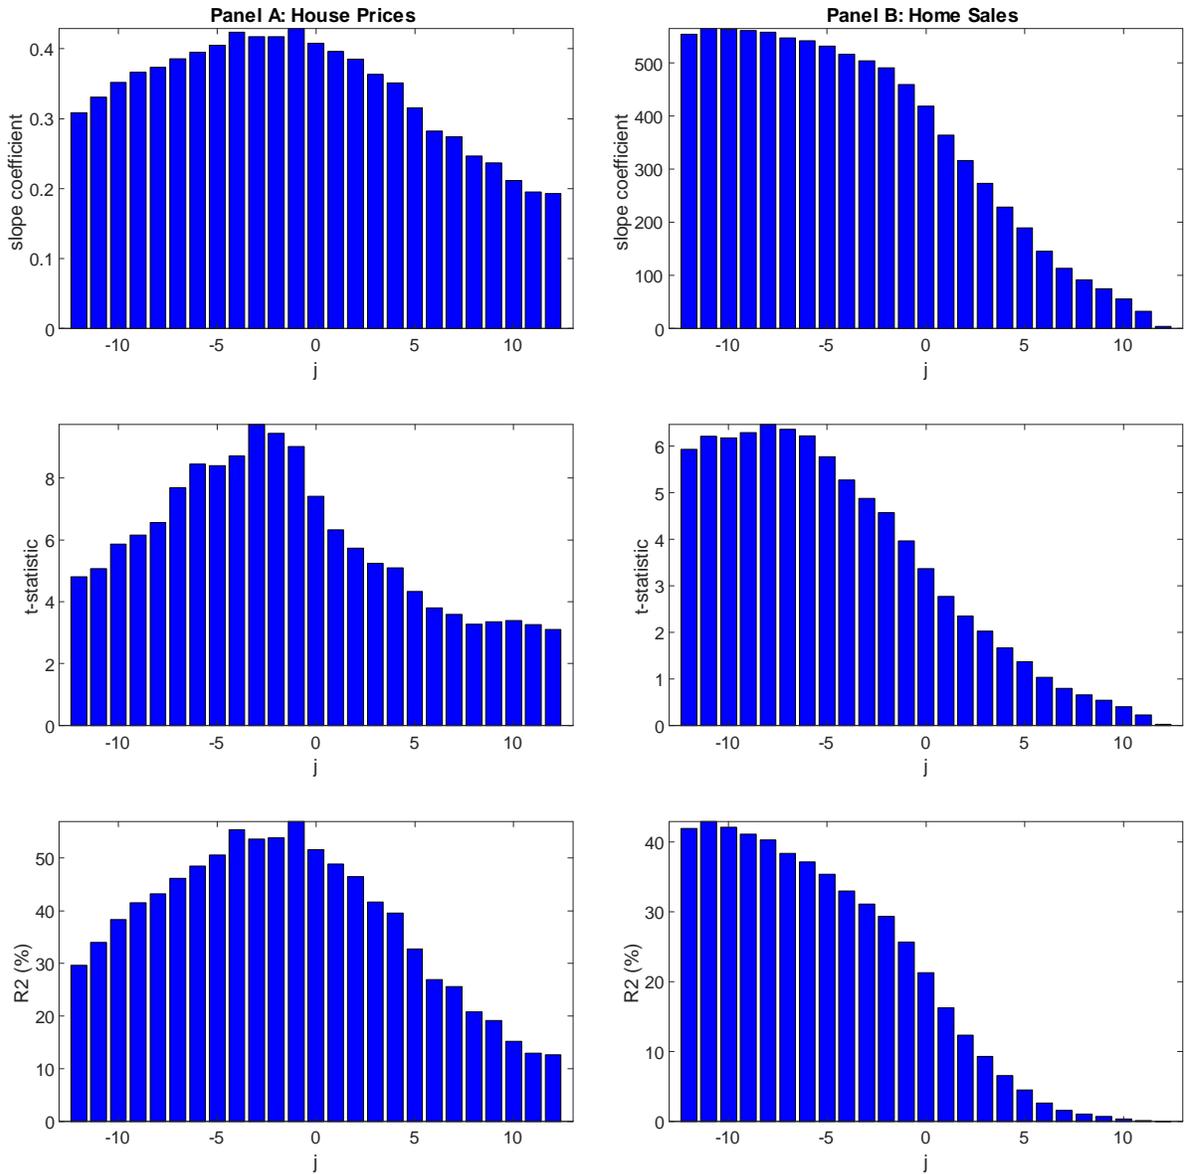


Figure 3. Housing Demand. The figure shows the Housing Search Index (HSI) along with the Housing Demand Index (HDI) constructed by Redfin. In Panel A, the sample frequency is monthly and covers the period 2018:1 to 2021:1. In Panel B, the sample frequency is weekly and the sample period runs from the first week of 2018 until the first week of April 2021.

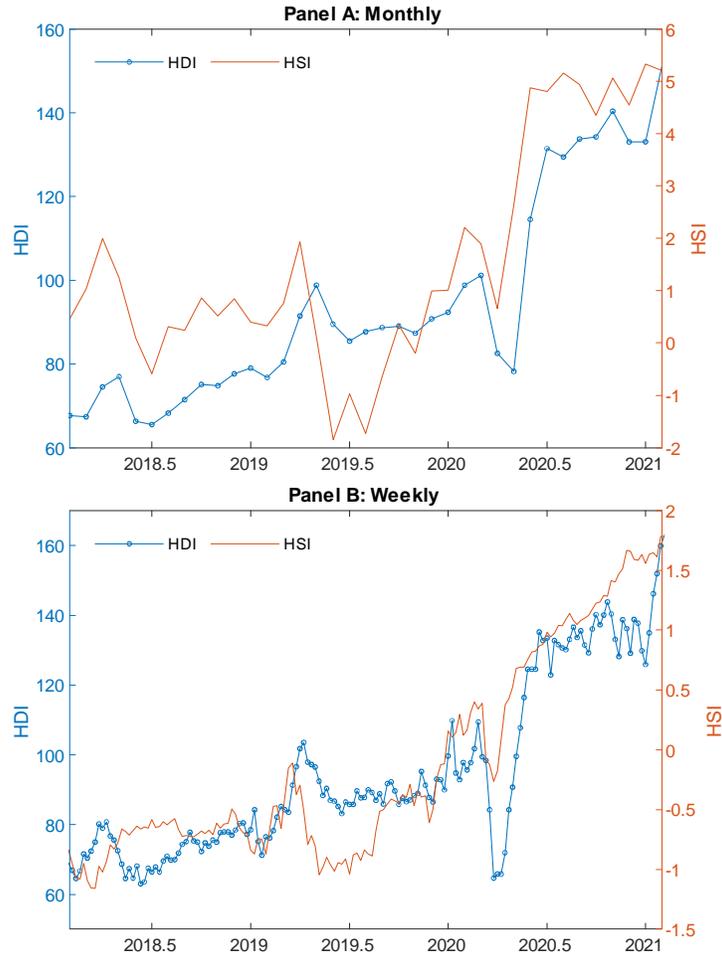


Figure 4. Local Housing Search. Panel A and B show the local *HSI* and log growth rate in house prices in Miami (FL) and Wichita (KS), respectively. The sample period is 2004:1-2021:1.

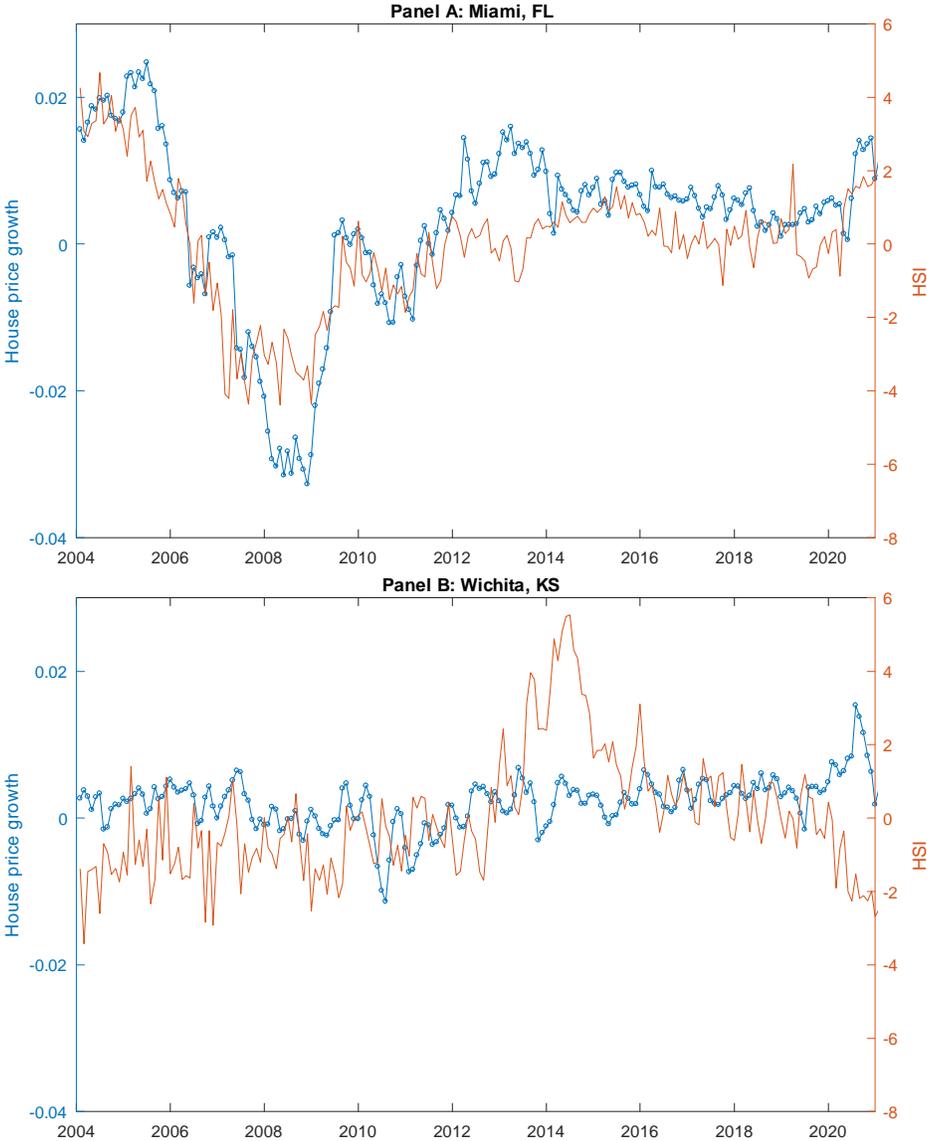


Figure 5. Long-Horizon Predictability. The figure shows estimated slope coefficients, associated t -statistics and R^2 values from the regression, $p_{t+h} - p_t = \alpha + \beta HSI_t + \varepsilon_{t+h}$, as a function of h . We compute standard errors using a circular block bootstrap. The sample period is 2004:1-2021:1.

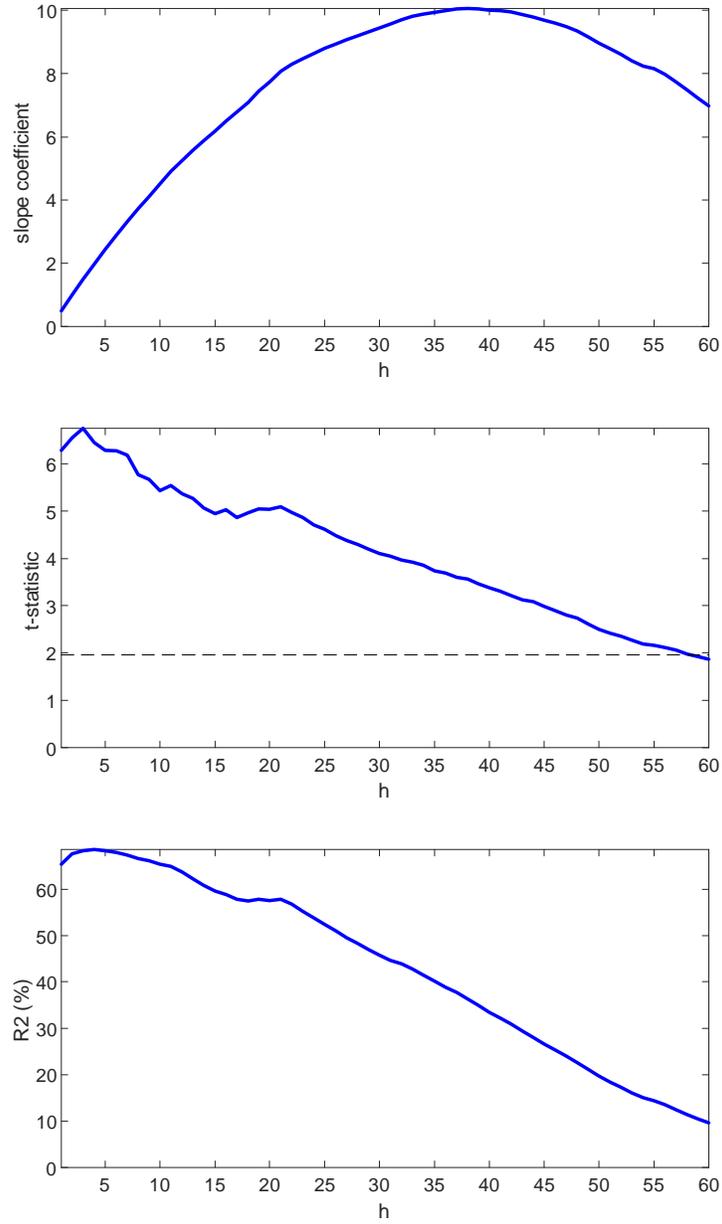


Figure 6. Out-of-Sample Forecast Errors. Panel A shows the cumulative sum of squared forecast errors of the no-predictability benchmark minus the cumulative sum of squared forecast errors of a model based on the national-level *HSI*. The forecast horizon is $h = 1$ month and the out-of-sample period runs from 2007:1 to 2021:1. Panel B shows the 1st quartile, median and 3rd quartile out-of-sample R^2 values at horizons of $h = 1, 3, 6,$ and 12 months across MSAs. Panel C shows the median of the cumulative sum of squared forecast errors of the no-predictability benchmark minus the cumulative sum of squared forecast errors of *HSI* across MSAs. Panel D plots the median out-of-sample R^2 at horizons of $h = 1, 3, 6,$ and 12 months across MSAs in downturns and upturns, whereas Panel E shows results for MSAs with high and low volatility.

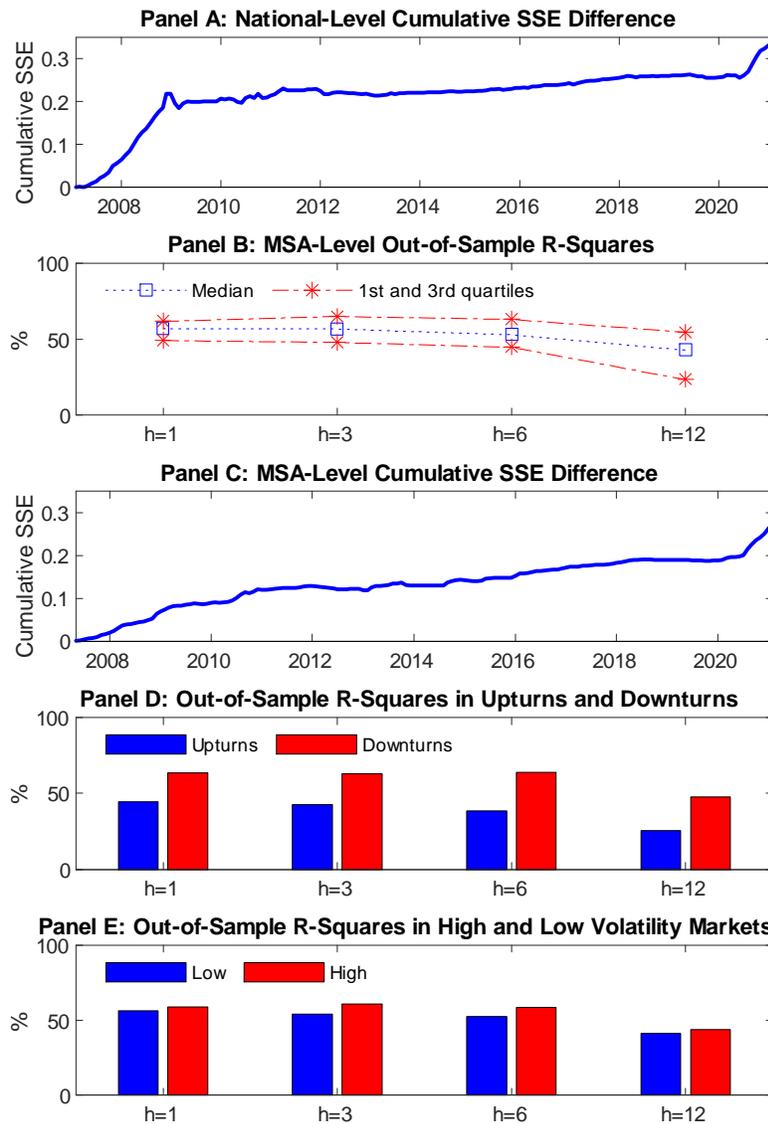
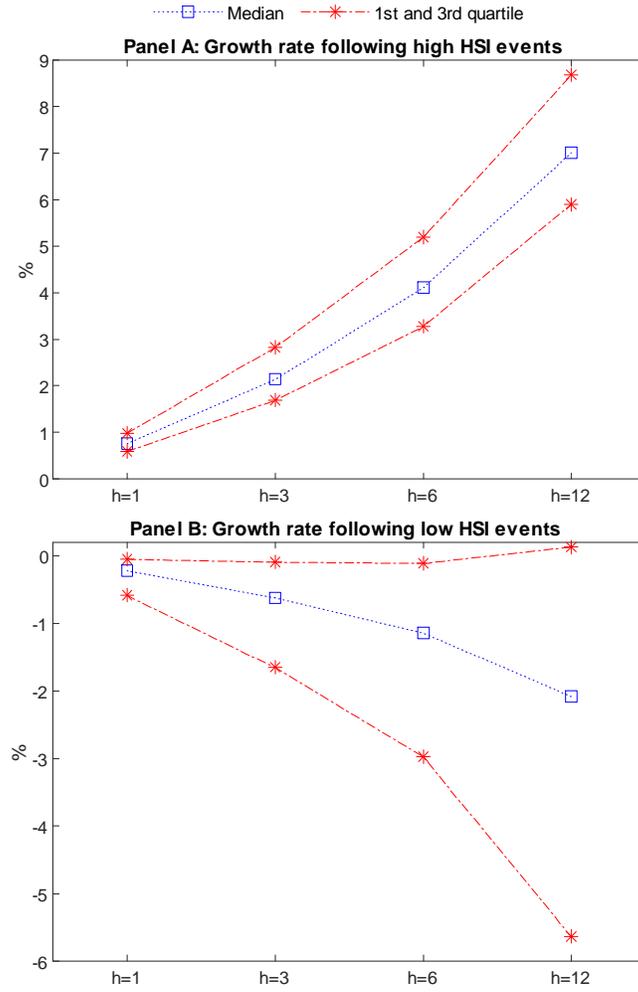


Figure 8. Economic Value of house price forecasts. Panel A shows 1st quartile, median and 3rd quartile growth rates in house prices after episodes where *HSI* is one standard deviation above the local mean. Panel B shows the results following events where *HSI* is one standard deviation below the local mean. Forecast horizons range from one-month ahead ($h = 1$) to one-year ahead ($h = 12$).



Online Appendix for

"Search and Predictability of Prices in the Housing Market"

January 2022

This appendix provides additional results and robustness checks of the analysis reported in the main paper. Below we provide a brief description of the robustness checks:

- Section A.1 illustrates robustness to the method of inference.
- Section A.2 reports in-sample results for conventional determinants of outcomes in housing markets.
- Section A.3 reports results from a placebo test.
- Section A.4 reports results from the selling side of market.
- Section A.5 shows out-of-sample results using alternative search indices.
- Section A.6 reports bootstrap results from a "useless" factor test.
- Section A.7 reports prediction results from more advanced machine learning procedures.
- Section A.8 analyzes the relation between *HSI* and speculative behavior proxies.
- Section A.9 reports results for one-way and two-way clustered standard errors in panel regressions.
- Section A.10 provides details of additional results for the Covid-19 period.
- Section A.11 reports results on the relation between search and price expectations.
- Section A.12 analyzes the REITs market.
- Section A.13 analyzes non-seasonally adjusted data.
- Section A.14 provides additional results on the predictive ability of local search.

A.1 Inference

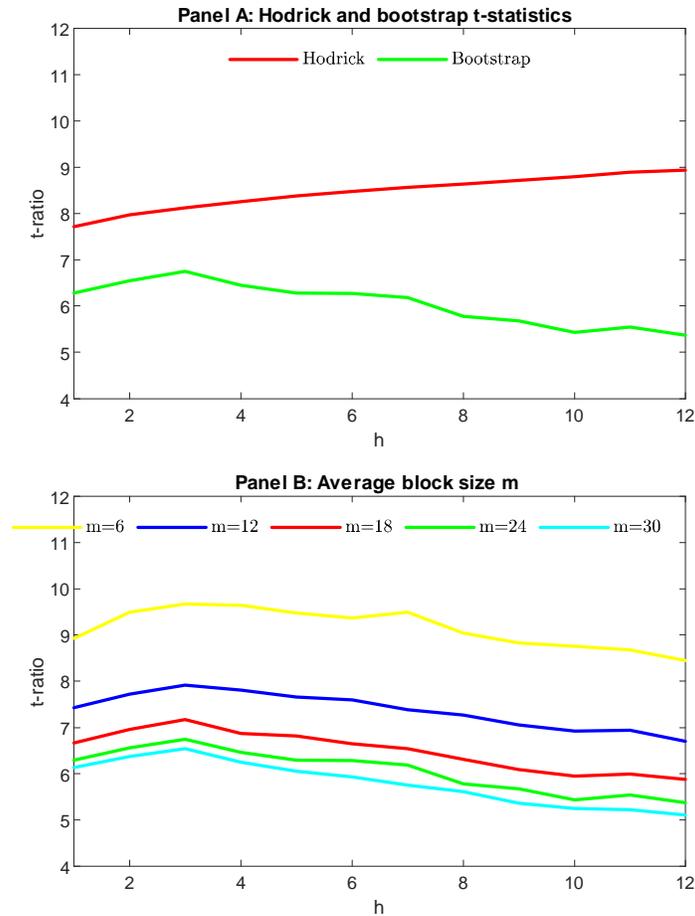
In the main paper, we use a circular block bootstrap to conduct inference when estimating overlapping multi-period ahead forecast regressions of the form, $p_{t+h} - p_t = \alpha + \beta'x_t + \varepsilon_{t+h}$, where p_t is the log of the FHFA house price index, x_t is a vector of predictors, and h is the forecast horizon in months. We resample the regressand and the regressor jointly in blocks with an average block size of $m = 24$, which is close to the optimal block length according to Politis and White's (2001) automatic selection procedure (the optimum is slightly below 24). We use 20,000 replications.

Figure A.1 shows results for the robustness of the method used for inference. In Panel A of Figure A.1, we compare block bootstrap t -statistics with those of the Hodrick (1992) procedure, which is aimed at circumventing issues with overlapping data.¹ We use HSI as the predictor and consider horizons from one through 12 months. We find that the Hodrick t -statistics confirm the statistical significance of HSI across horizons, but that these t -statistics are less conservative compared to the block bootstrap t -statistics.

In Panel B of Figure A.1, we consider various choices of the average block length: $m = 6, 12, 18, 24, 30$. Following Politis and White's (2001) automatic selection procedure, we set the average block size to $m = 24$ when generating the results reported in the main paper. However, the statistical significance of HSI is robust towards other reasonable choices of m as shown in Panel B of Figure A.1. From the figure, we see that using $m = 18$ or $m = 30$ gives bootstrap t -statistics close to those obtained with $m = 24$, whereas $m = 12$ and especially $m = 6$ generates somewhat higher bootstrap t -statistics. In any case, HSI remains statistically significant across the various specifications.

¹Hodrick's (1992) procedure is based on reverse regressions. Wei and Wright (2013) show that the reverse regression methodology can help in circumventing size distortions in long-horizon forecasting regressions with near-unit roots.

Figure A.1. Robustness to the method of inference.



A.2 Commonly Used Housing Market Determinants

In the main paper, Table 3 provides in-sample evidence that HSI retains its predictive ability when controlling for a long list of predictive variables. Table A.1 shows predictive results from each of the 14 control variables. Among all variables, HSI generates by far the strongest predictive results in terms of statistical significance as well as R^2 values.

Table A.1. Predicting House Prices With Alternative Variables. The table reports results from predictive regressions, $p_{t+h} - p_t = \alpha + \beta x_t + \varepsilon_{t+h}$, where p_t is the log of the FHFA house price index, x_t is a predictive variable, and h is the forecast horizon in months. For each regression, the table reports the estimate of β , the corresponding t -statistic in parenthesis, and the R^2 in square brackets. We compute standard errors using a circular block bootstrap. All predictive variables are standardized and slope coefficients are multiplied by 100 to facilitate comparison across variables. The sample period is 2004:1 to 2021:1.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
<i>HSI</i>	0.43 (6.39) [56.91]	1.26 (6.74) [67.37]	2.44 (6.19) [70.41]	4.37 (5.35) [64.42]
<i>payrolls</i>	0.03 (0.12) [0.25]	0.06 (0.08) [0.15]	0.32 (0.23) [1.22]	2.72 (1.59) [24.98]
<i>infl</i>	0.00 (0.03) [0.00]	-0.12 (-0.82) [0.58]	-0.25 (-0.97) [0.77]	-0.34 (-0.69) [0.40]
<i>permits</i>	0.17 (4.13) [9.27]	0.38 (3.01) [6.05]	0.62 (2.73) [4.61]	1.34 (2.75) [6.07]
<i>starts</i>	0.08 (2.19) [2.06]	0.23 (2.49) [2.16]	0.38 (2.22) [1.74]	0.85 (2.41) [2.42]
<i>term</i>	-0.06 (-0.51) [1.17]	-0.10 (-0.27) [0.42]	-0.04 (-0.05) [0.02]	0.59 (0.35) [1.17]
<i>mort</i>	-0.24 (-1.89) [18.00]	-0.72 (-1.95) [21.82]	-1.44 (-2.03) [24.55]	-2.95 (-2.16) [29.37]
<i>pr</i>	0.04 (0.28) [0.45]	0.03 (0.08) [0.05]	-0.16 (-0.19) [0.30]	-1.09 (-0.67) [3.99]
<i>loans</i>	-0.07 (-0.62) [1.36]	-0.00 (-0.02) [0.00]	0.36 (0.62) [1.56]	0.34 (0.22) [0.40]
<i>sent</i>	0.12 (1.28) [4.65]	0.35 (1.22) [5.36]	0.82 (1.36) [7.97]	2.36 (2.28) [18.77]
<i>cfnai</i>	0.15 (0.77) [6.70]	0.31 (0.49) [4.15]	0.31 (0.26) [1.13]	1.78 (2.53) [10.69]
<i>ads</i>	0.13 (0.53) [4.90]	0.41 (1.50) [3.96]	0.16 (0.50) [0.59]	1.64 (1.75) [9.10]
<i>pd</i>	0.20 (2.16) [11.98]	0.47 (1.78) [9.43]	0.76 (1.61) [6.78]	1.27 (1.34) [5.43]
<i>ra</i>	-0.22 (-2.68) [15.36]	-0.48 (-2.46) [9.75]	-0.85 (-1.97) [8.60]	-1.75 (-2.73) [10.37]
<i>unc</i>	-0.19 (-2.11) [11.18]	-0.46 (-1.90) [8.88]	-0.75 (-1.58) [6.67]	-1.64 (-1.95) [9.11]

A.3 Placebo Tests

In the main paper, we analyze the relation between HSI and a large set of commonly used housing market determinants by estimating the contemporaneous regression, $HSI_t = \alpha + x_t'\beta + \varepsilon_t$, where x_t contains a list of 14 variables. Using the full list of housing market determinants leads to an R^2 of about 70%, meaning that about 30% of the variation in HSI is left unexplained. Some of the time-series variation in these 14 variables might be correlated with HSI by chance. We analyze this possibility by running a placebo test that generates artificial times series by resampling from the panel of regressors. We use a circular block bootstrap to resample the panel of housing market determinants and use 20,000 replications. For each replication, we regress the non-resampled HSI on the randomly resampled panel of regressors and save the R^2 . We set the average block size to $m = 6, 12, 18, 24$, and 30. Table A.2 reports the median R^2 across the different block sizes. We see that the placebo regressors generate a median R^2 in the range from 19.4% to 30.5% across block sizes. This suggests that the estimated $R^2 = 0.70$ obtained from regressing HSI on all 14 variables is inflated. There is therefore a risk of excluding valuable information when orthogonalizing HSI against all 14 housing market determinants.

Table A.2. Results of placebo test. We use a circular block bootstrap to resample the panel of control variables. For each replication, we regress the non-resampled HSI on the resampled panel of variables and save the R^2 from this regression. The table reports the median R^2 value across different block sizes.

	$m = 6$	$m = 12$	$m = 18$	$m = 24$	$m = 30$
R^2	19.4%	25.0%	27.8%	29.2%	30.5%

A.4 Buying versus Selling Side of the Housing Market

Wu and Brynjolfsson (2015) conjecture that house prices are difficult to predict using two predefined real estate categories since these categories can reflect both the buying and selling sides of the housing market. This motivates us to explore the predictive power of a search index based on the main search term “selling a house” instead of “buying a house”. We follow the approach used in constructing the HSI described in Section 2.2, but now use a keyword intended to capture the selling side of the housing market. The related search terms are: "when selling a house", "selling

a home", "selling your house", "selling my house", "selling a house taxes", "how to sell a house", "selling your home", "tax on selling a house", "selling house by owner", "cost of selling a house", "capital gains", "taxes on selling a house", "closing costs", "capital gains tax", and "selling a house tips". These search terms are all directly related to the home selling process. We denote the index based on these search terms by HSI^{sell} .

Panel A in Table A.3 shows that this search index based on the selling side of the housing market holds limited predictive power over future house prices. The slope coefficients are only significantly different from zero for $h = 1$ and the R^2 -values never exceed 8%. The slope coefficients for HSI^{sell} have the same positive sign as for HSI , i.e. an increase in search activity on the selling side of the housing market is associated with an increase in future house prices. Accordingly, we are careful not to interpret HSI^{sell} as a measure of housing supply. However, when we control for HSI , the slope coefficients on HSI^{sell} turn negative except for $h = 12$ as shown in Panel B of Table A.3.

In conclusion, search activity on the buying side of the housing market appears to dominate search activity on the selling side in terms of predictive power over movements in future house prices.

Table A.3. Predicting House Prices With Alternative Search Indices. The table reports results from predictive regressions, $p_{t+h} - p_t = \alpha + \beta'x_t + \varepsilon_{t+h}$, where p_t is the log of the FHFA house price index, x_t is a vector of predictors, and h is the forecast horizon in months. For each regression, the table reports slope estimates, the corresponding t -statistics in parenthesis, and the R^2 in square brackets. We compute standard errors using a circular block bootstrap. All predictors are standardized to facilitate comparison of the β estimates and the log price change is multiplied by 100. Panel A shows the results for HSI^{sell} , which is an alternative search index based on the selling side of the housing market, and Panel B includes HSI and HSI^{sell} in a joint regression. The sample period is 2004:1 to 2021:1.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A: Search index based on the selling side of the housing market				
HSI^{sell}	0.11 (0.79) [3.68]	0.34 (0.82) [4.97]	0.68 (0.81) [5.43]	1.51 (0.93) [7.72]
Panel B: HSI joint with the selling side of the housing market				
HSI	0.46 (6.35)	1.31 (6.64)	2.48 (6.09)	4.27 (5.02)
HSI^{sell}	-0.07 (-1.04) [58.05]	-0.15 (-0.75) [68.22]	-0.15 (-0.38) [70.64]	0.33 (0.42) [64.76]

A.5 Out-of-Sample Results With Alternative Search Indices

Rather than using specific keywords, Wu and Brynjolfsson (2015) employ two predefined search categories supplied by Google Trends, namely “Real estate agencies” and “Real estate listings”. In the main paper, we show that the *HSI* contains stronger in-sample predictive power than these two alternative search indices. As a further robustness check, we analyze the out-of-sample evidence. In Table A.4, we compare the out-of-sample predictive power of the *HSI* against the two predefined search categories. Panel A reports results for the *HSI*, while Panel B reports results from the predefined search categories. Both “Real estate agencies” and “Real estate listings” generate negative out-of-sample R^2 s when used on their own as well as in a joint specification. In contrast, *HSI* is able to outperform the historical mean benchmark. Thus, the out-of-sample evidence confirms the results obtained from the in-sample analysis.

Table A.4. Out-of-Sample Tests With Alternative Search Indices. The table reports the Campbell and Thompson (2008) out-of-sample R^2 (R_{OoS}^2) and in parenthesis the p-value from the Diebold and Mariano (1995) t -statistic, computed using the Newey and West (1987) estimator with h lags, where h is the forecast horizon in months. The null hypothesis is that the R_{OoS}^2 is equal to zero or negative and the alternative hypothesis is that it is positive. Panel A reports results for *HSI*. Panel B reports results from the predefined search categories used by Wu and Brynjolfsson (2015).

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A				
<i>HSI</i>	51.37 (0.00)	64.54 (0.00)	63.91 (0.00)	54.25 (0.02)
Panel B				
Real estate agencies	-1.08 (0.57)	-4.45 (0.68)	-6.68 (0.70)	-17.85 (0.82)
Real estate listings	-4.36 (0.90)	-7.63 (0.90)	-12.21 (0.89)	-20.89 (0.83)
Joint	-0.79 (0.56)	-4.04 (0.68)	-5.15 (0.67)	-12.97 (0.71)

A.6 Bootstrap Analysis

To further validate the statistical significance of the *HSI* in forecasting house prices, we consider a simulation experiment that is comparable to the "useless" factor tests of Kan and Zhang (1999a,b). In particular, we generate 10,000 bootstrap samples by row-wise resampling from the panel of search

indices (with replacement). The resampled panels have the same length as the original panel of search indices. For each bootstrap sample, we recursively estimate the HSI and generate out-of-sample forecasts, then save the R_{OoS}^2 statistic. Because the resampled placebo Google search data should bear no relation to the realized house price growth rates, the HSI should not be useful in forecasting growth in house prices. Basically, the resampled search indices represent random noise and so are “useless”.

Table A.5. Bootstrapped p-values. Panel A reports the Campbell and Thompson (2008) out-of-sample R^2 (R_{OoS}^2). In Panel B, we generate 10,000 bootstrap samples of the search indices used to compute the HSI . For each bootstrap sample, we recursively estimate the HSI based on the resampled search indices, generate out-of-sample forecasts, and then compute the R_{OoS}^2 . The table reports the empirical p -value, which is the fraction of artificial R_{OoS}^2 statistics that exceed the actual R_{OoS}^2 statistic. Results are shown for three different bootstrap techniques: row resampling, a parametric bootstrap, and a block bootstrap.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A: Out-of-sample R^2				
HSI	51.37	64.54	63.91	54.25
Panel B: Bootstrapped p-values for HSI				
Row resampling	0.00	0.00	0.00	0.00
Parametric AR(1)	0.00	0.00	0.00	0.00
Block bootstrap	0.00	0.00	0.00	0.00

We analyze the empirical distribution of the R_{OoS}^2 statistic by computing empirical p -values. The simulations in Panel B in Table A.5 show that the fraction of bootstrapped R_{OoS}^2 statistics that exceed their empirical counterparts from Panel A equals zero across all horizons. Hence, the chance of obtaining the same goodness-of-fit with random Google data as we find with the actual data is virtually zero. As further robustness checks, we consider two alternative bootstrap procedures that take into account the persistence in the data. The first procedure uses a parametric bootstrap in which we estimate an AR(1) model for each Google series and retain the estimated coefficients along with the residuals from each regression to construct a panel of placebo series that have the same autoregressive coefficient and variance as the series in the Google Trends panel. In addition, we consider a non-parametric circular block bootstrap procedure similar to the row-wise resampling above. However, instead of drawing one row at a time, for each series we build the placebo series from blocks of size m . For each series, we select the optimal value of m , using the automatic selection procedure developed by Politis and White (2004). Results from these two alternative bootstrap procedures are also shown in Panel B. The findings are identical to those obtained from

the row-resampling bootstrap, implying that it is extremely unlikely that the observed R_{OoS}^2 were due to chance. These robustness tests corroborate the robustness of our findings on the highly significant out-of-sample predictive power of the *HSI* over future movements in house prices.

A.7 Machine Learning

We construct the *HSI* using a simple targeted PCA approach. However, it may be possible to achieve more accurate forecasts using more advanced machine learning techniques. To test this possibility, we perform an out-of-sample forecasting exercise using three popular machine learning techniques that have been shown to achieve strong predictive performance even on difficult forecasting problems such as asset risk premiums (Gu et al., 2020). The three methods we consider are Random Forest (Breiman, 2001), Gradient Boosted Trees (GBT)², and an artificial neural network (ANN) with a single hidden layer.³ All of these methods incorporate the possibility of complex non-parametric nonlinear relations between the predictors and house prices but also run the risk of overfitting the data. To avoid this problem, we follow the most common approach in the literature and select *tuning* parameters adaptively from the data using a recursive cross-validation scheme. Subsection A.7.1 provides further details on our cross validation sample splitting scheme, choice of hyper-parameters and choice of neural network architectures. As in the main paper, we use the first three years as our initial estimation period and reserve the remaining sample for out-of-sample evaluation.

Machine learning methods are inherently designed to take advantage of high-dimensional data so it is possible that they require a larger set of predictors to achieve their forecasting potential. For this reason, when utilizing these methods we consider an extended set of predictors where we add the top 25 related search terms for each of our original set of 23 predictors used to form the *HSI*. After removing duplicates, this expanded set contains a total of 292 Google Trends (GT) terms. We note that, although it would be possible to construct a factor like the *HSI* using this expanded set of predictors, this factor would not have the simple demand search interpretation as the *HSI* since many of the additional terms are not necessarily reflecting housing demand. Thus, our high-

²In particular, we consider the powerful XGboost algorithm of Chen and Guestrin (2016), which builds on the non-parametric procedures developed by Breiman (1997) and Friedman (2001).

³In unreported results, we also considered a a deep artificial neural network with three hidden layers, but this ANN underperformed the network with a single layer.

dimensional panel does not offer the same interpretability as the smaller data used to extract the *HSI* from.⁴

Table A.6 reports the Campbell and Thompson (2008) out-of-sample R^2 values (R_{OoS}^2) and Diebold and Mariano (1995) t -statistics (t_{DM}) against the historical average. Overall, the forecasting performance of these methods using a high-dimensional panel is similar to what we obtained for the *HSI*. For example, the GBT model which is the best performing machine learning approach, obtains an average R_{OoS}^2 across horizons of 58.6%, which is only marginally above that obtained using the *HSI* (58.5%).

The fact that an expanded set of GT predictors in combination with models that allow for nonlinearities does not result in significantly improved predictive performance implies that most of the information embedded in Google search data is likely to be already captured by the *HSI* and that the search terms are probably characterized by a strong factor structure. The exception to this rule appears to be forecasts at the one-month horizon, where the ML methods outperform the *HSI*.

The panel we use for this forecasting experiment is based on terms that are somehow related to the original set of 23 GT terms used to form the *HSI*. If the purpose is to obtain the best possible forecast, it is possible that a larger set of terms with broader coverage, in combination with ML methods, will result in better forecasting performance. Such an exercise would, however, result in even more complexity when trying to interpret the information captured by the forecasting model.

Table A.6. Out-of-Sample Tests. The table reports the Campbell and Thompson (2008) out-of-sample R^2 (R_{OoS}^2) and in parenthesis the p-value from the Diebold and Mariano (1995) t -statistic, computed using the Newey and West (1987) estimator with h lags, where h is the forecast horizon in months. The null hypothesis is that the R_{OoS}^2 is equal to zero or negative and the alternative hypothesis is that it is positive.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
RF	60.80 (0.00)	60.47 (0.00)	52.93 (0.00)	35.55 (0.04)
GBT	58.85 (0.00)	64.83 (0.00)	60.55 (0.00)	50.29 (0.00)
ANN	53.18 (0.00)	56.26 (0.00)	61.80 (0.00)	49.15 (0.00)

⁴For example, the expanded set of predictor terms include terms like “average cost of building a house” and “real estate news”, which can contain predictive information for housing returns but are not necessarily reflecting demand.

A.7.1 Machine Learning Model Specifications

All models are implemented in `Python` with the choice of hyperparameters shown below. We use the recursive *walk forward* cross validation scheme implemented in the `TimeSeriesSplit` function of the `sklearn` package. We use 12 validation periods (splits) and set the gap equal to the forecast horizon. Using K-fold cross validation with $K = 5$ folds leads to similar results.

Gradient Boosted Trees The number of trees considered are $B = 500$ with a shrinkage parameter (or learning rate) set to $\tau = 0.1$. We validate the depth of the trees using the following grid of values: maximum tree depth ($L \in [1, 2, 3, 5]$), the percentage of observations used in each tree in: $[60\%, 80\%, 100\%]$, the percentage of predictors used in each tree in: $[60\%, 80\%, 100\%]$, and the across-trees regularization parameter $\gamma \in [0, 0.1]$. The implementation is conducted with the `XGBRegressor` from the `XGboost` package. All other settings are set to the default of the package.

Random Forest The number of trees considered is $B = 500$, and the trees are allowed to grow fully deep (L can be large) consistent with Breiman (2001). With the number of predictors given by $p = 292$, we validate the number of predictors to be considered at each split using the following grid: $[p, 2/3p, 1/3p, \sqrt{p}, \log_2 p]$. The implementation is conducted with the `RandomForestRegressor` from the `sklearn` package. All other settings are set to the default of the package.

Artificial Neural Network We use a batch size of 12 in the stochastic gradient descent optimizer, set the learning rate to a (by default) constant value of 0.01, and use 1,000 epochs. The activation function is the rectified linear activation function (ReLU). We use the popular square root rule to select to number of neurons, hence with 292 predictors, the hidden layer has $\lfloor \sqrt{292} \rfloor = 17$ neurons. We cross-validate a single parameter which is the penalty parameter for the ℓ_2 regularization using the following grid of values: $[0.01, 0.001, 0.0001]$. The implementation is conducted using `MLPRegressor` from the `sklearn` package, using default settings for the ADAM stochastic gradient descent optimizer. The results are robust to changes in the number of epochs and learning rate.

A.8 Speculative Demand

During the early 2000s house prices in many MSAs increased dramatically and reached record high levels, which was followed by a collapse in house prices and a severe crisis in the U.S. economy. A growing literature suggests that speculation in the housing market was an important driver of the boom and argues that economic fundamentals accounted for just a small fraction of the changes in prices during the housing boom (e.g. Akerlof and Shiller, 2009, Chinco and Mayer, 2016, and Nathanson and Zwick, 2018).⁵ Given that the *HSI* is a direct measure of peoples' intention to buy a house and hence captures the demand side of the market, we may expect that the predictive power of *HSI* reflects both fundamental and non-fundamental sources of demand for housing.

As an attempt to analyze to what extent *HSI* captures speculation-driven demand, we construct a search-based measure of the interest in "flipping houses", which is likely to reflect speculative motives, and examine its relation to *HSI*.⁶ We construct a "flipping houses" search index by using "flipping houses" as our main search term and then obtain the following list of related terms: "flipping homes", "flip houses", "house flipping", "house flip", "flip the house", "how to flip houses", and "real estate flipping". Next, we construct a house flipping search index by extracting the first principal component of these search indices.

We start by plotting *HSI* together with the house flipping index in Figure A.2. From the figure, we see that the house flipping index increases substantially in the mid 2000s and reaches its highest value at the beginning of 2007, in line with the literature that has shown that speculation in the housing market was an important driver of the boom in the early 2000s. *HSI* also takes on high values in the 2004-2005 period, but does not show strong comovement with the flipping houses index. Furthermore, while *HSI* increases substantially during the Covid-19 pandemic, the house flipping index shows a more modest increase. Over the full sample, the correlation between the two series is only 0.10. Thus, when orthogonalizing *HSI* with respect to the flipping houses index, we find that it retains basically all of its predictive ability. These results indicate that the predictive ability of *HSI* is not driven by speculative activity.

⁵Other factors contributing to the boom and bust in house prices have been put forward in the literature, including credit conditions (Mian and Sufi, 2009, and Favilukis et al., 2017) and low interest rates resulting from excessively loose monetary policy (Taylor, 2014).

⁶Recent papers on house flipping include Goldstein (2018) and Bayer et al. (2020).

Figure A.2. Flipping Houses. The figure plots the flipping houses index together with the housing search index (*HSI*). The sample period is 2004:1-2021:1.



As another measure of speculation, we follow Gao et al. (2020) and compute the fraction of non-owner-occupied home purchases using the Home Mortgage Disclosure Act (HMDA) data from which it is possible to map individual mortgage level data to the 77 MSAs in our sample. As Gao et al. (2020) point out, decisions to buy a non-owner-occupied home are to a greater extent driven by speculative motives than by decisions to buy a primary home. The HMDA data is only available annually, but we can still compare general movements in the housing speculation measure of Gao et al. (2020) with those of the *HSI*. For the aggregate U.S. housing market, Gao et al. (2020) find that the degree of speculation peaked in 2005 with a share of non-owner occupied home purchases of about 15%. This is also the period where the house pricing boom was close to reaching its peak. Interestingly, the fraction of investment properties purchased is now substantially lower than it was during the first part of the 2000s.

Across the 77 MSAs in our sample, the median value of the share of investment properties purchased was 6.8% in 2020, a drop from 7.9% in 2019 and 8.8% in 2018. This relatively low level of housing speculation during the pandemic holds across the 77 MSAs in our sample. For example, Gao et al.

(2020) show that the share of non-owner occupied home purchases was almost 30% in Las Vegas in 2005. We find that this number has dropped to about 9% in 2020. These strong movements in speculation stand in stark contrast to the time-series movements in the *HSI*, which reached an all-time high during the pandemic.

In conclusion, these results suggest that the *HSI* captures general movements in demand and is not dominated by speculative activity.

A.9 Clustered Standard Errors

We use two-way clustered *t*-statistics when estimating fixed-effects panel regressions in the main paper. In Table A.7, we show results from both one-way and two-way clustered *t*-statistics (clustered by time and MSA).

Table A.7. Predicting Local House Prices With Local Housing Search: Evidence From Panel Regressions. The table reports results from fixed effects panel regressions of the form, $p_{it+h} - p_{it} = \alpha_i + \beta HSI_{it} + \varepsilon_{it+h}$, where p_{it} is the log of the Freddie-Mac house price index in MSA i , HSI_{it} is the housing search index in MSA i , and h is the forecast horizon in months. For each regression, the table reports the estimate of β , the corresponding *t*-statistic in parenthesis, and the within R^2 in square brackets. We compute standard errors using Thompson (2011) clustered robust-statistics with h lags. The table shows results using both single-clustered and double-clustered *t*-statistics. *HSI* is standardized to facilitate interpretation of the β estimates. The sample period is 2004:1 to 2021:1.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
β	0.42	1.22	2.32	4.19
t (cluster by time)	(13.26)	(9.12)	(6.98)	(5.98)
t (cluster by MSA)	(12.88)	(12.56)	(12.16)	(11.93)
t (cluster by both time and MSA)	(9.43)	(7.59)	(6.26)	(5.58)
R^2	[37.20]	[37.89]	[36.89]	[33.81]

A.10 Additional Results on Covid-19 Period

In the main paper, we reports results from the following panel regression

$$p_{it+1} - p_{it} = \alpha_i + \beta_D HSI_{it} + \beta_S S_{it} + \gamma' Z_{it} + \varepsilon_{it+1}, \quad (\text{A1})$$

where $p_{it+1} - p_{it}$ is the one-month change in the log house price index for MSA i in month $t + 1$, HSI_{it} is our housing search index, S_{it} is a proxy for the housing supply given by the Zillow for-sale-inventory of houses, and Z_{it} is a vector of controls including the Covid-19 stringency index of Hale et al. (2021) and the number of Covid-19 cases, all measured for MSA i in month t . To visualize the model fit, Figure A.3 plots realized month-on-month house price growth rates during the Covid-19 pandemic together with predicted values from the full specification in (A1). We focus on the four largest MSAs as measured by population. Across MSAs, growth in house prices were in general low (even negative) in March and April of 2020 but bounced back sharply between May and August and remained high subsequently. The figure illustrates that demand and supply effects along with Covid-19 restrictions combine to capture a substantial part of the variation in house prices across MSAs during the pandemic. If we only use supply as regressors, the model fit deteriorates, but once we account for HSI the predicted values are much closer to the actual values.

To give a sense of how search activity changed during the Covid-19 period in the various MSAs, Figure A.4 shows scatter plots of one-month-ahead price changes ($h = 1$) against search, marking Covid-19 points in red and pre-Covid-19 points in blue. The strong positive relation between one-month lagged search and house price changes is present both in the pre-Covid-19 period and during the pandemic. The relation remains remarkably stable both for MSAs in which the pandemic triggered the highest search activity and price growth observed in the entire sample (New York and Dallas) as well as for MSAs with a more modest pandemic impact (Los Angeles and Chicago).

While our data on housing supply from Zillow does not allow us to go as far back in time as the HSI data, it allows us to compare the period from November 2017 to January 2021. Figure A.5 shows that the pandemic was associated with a notably lower housing supply across the top-four MSAs with particularly strong reductions for New York and Chicago. Notice also how the housing supply constraint gradually tightens from month to month during 2020 – a pattern quite different from that seen in the HSI during the pandemic.

Figure A.3. House Price Movements During the Covid-19 Pandemic. The figure plots realized month-on-month house price growth rates during the Covid-19 pandemic together with predicted values from equation (A1). As regressors, we use either supply on its own, demand (HSI) on its own, or the full specification equation (A1). The figure shows results for the four largest MSAs as measured by population.

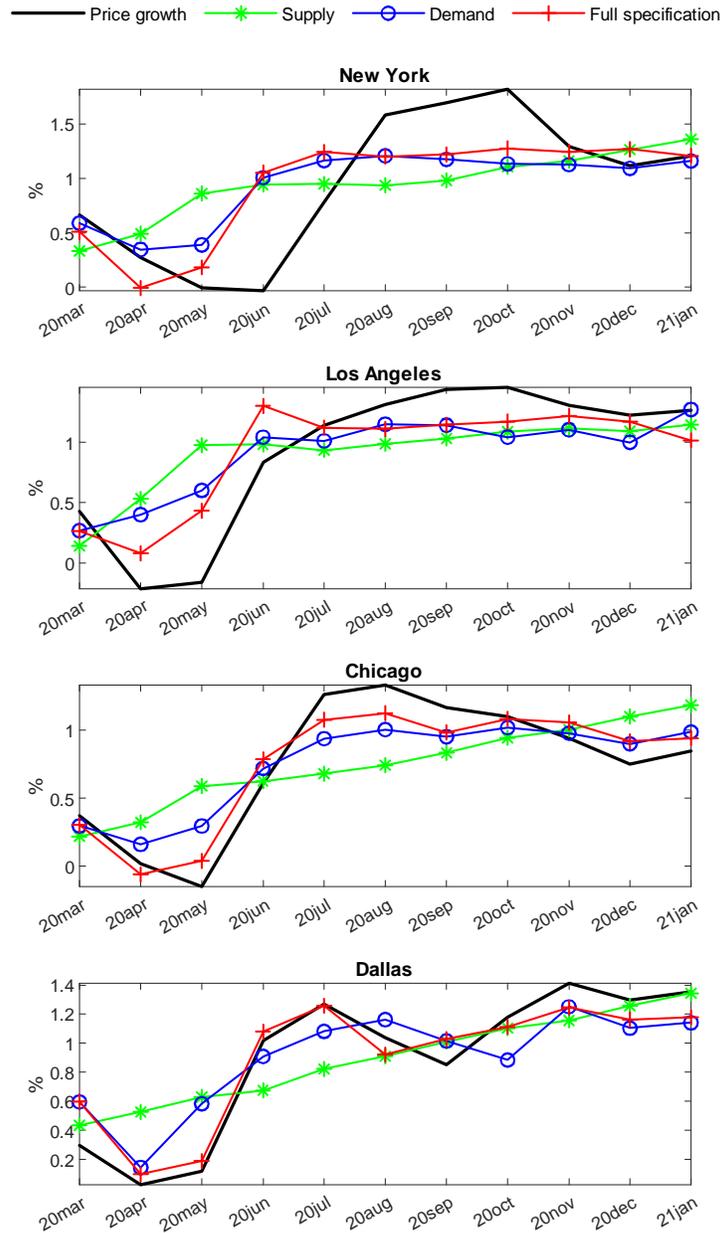


Figure A.4. Price Changes vs. Search. The figure plots realized house price growth rates (in percent) against one-month lagged values of *HSI* (standardized). Covid-19 points are marked in red with pre-Covid-19 points marked in blue. The figure shows results for the four largest MSAs as measured by population.

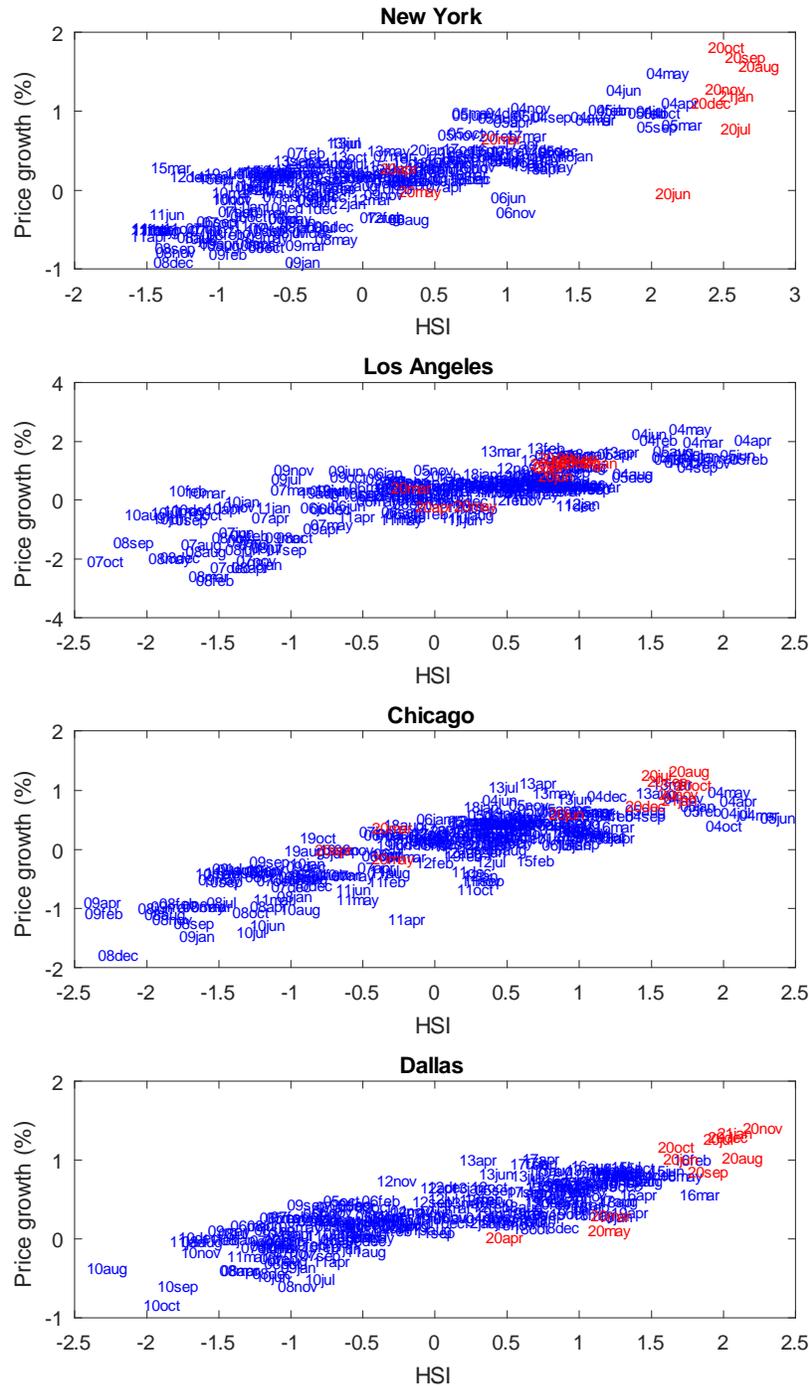
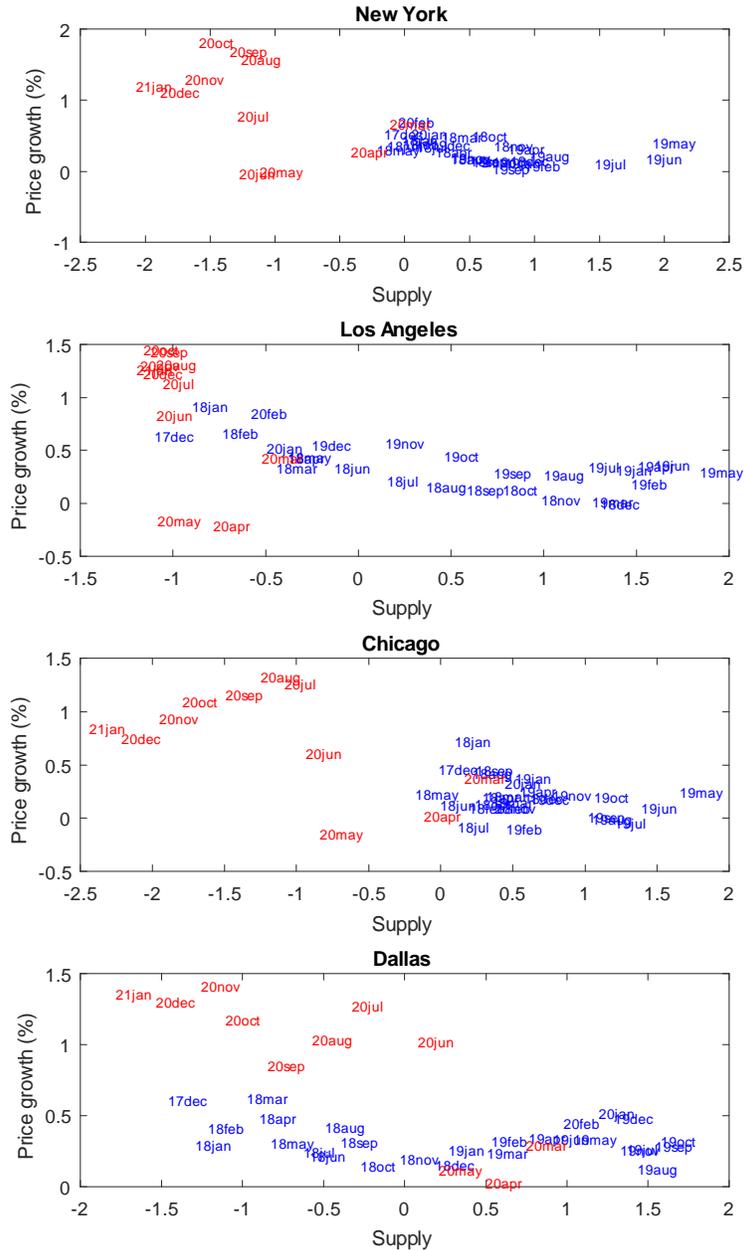


Figure A.5. Price Changes vs. Supply. The figure plots realized house price growth rates (in percent) against one-month lagged values of supply as measured by the inventory of houses for sale (standardized). Covid-19 points are marked in red with pre-Covid-19 points marked in blue. The figure shows results for the four largest MSAs as measured by population.



A.11 Search and Price Expectations

Our estimates of changes in house prices based on search activity are model-based and so may be subject to model misspecification biases. As we next discuss, however, we can obtain more direct measures of households' expectations about future house prices.

Following the analysis in Shimer (2004), analogous to search in the labor market we can think of home buyers' optimal search intensity as depending on three factors, namely (i) the sensitivity of the likelihood of finding a desirable home with respect to variation in search intensity. If the chance of finding a suitable home is highly sensitive to the amount of search, home buyers should be more willing to vary their search efforts in response to shifts in the housing market. Conversely, if the probability of finding a home is either very low (due to a tight housing market) or very high (due to an excess of supply), home buyers are unlikely to vary their search by much due to such shifts; (ii) the expected present value of rents or user benefits from owning a home, including shifts in expectations of future house prices. If home buyers expect future prices to be much higher, they should increase their search intensity, expecting to benefit from such price increases; (iii) the marginal cost of searching. This may change, e.g., as a result of new online search tools being launched (decreasing search costs) but could also simply reflect variation in the marginal value of time across economic recessions and expansions.

The third factor is likely to vary less over time than the first two. Provided that the cost of search is relatively constant, variation over time in search intensity should predominantly be driven by movements in the returns to search, i.e., the first two factors.

While we do not directly observe the likelihood of finding a house, we can construct a measure of house price expectations. Since 2007 the University of Michigan Surveys of Consumers has asked homeowners: "what do you think will happen to the prices of homes like yours in your community over the next 12 months?". To get a measure of house price expectations, we use the monthly time series of the fraction of people saying that house prices will increase minus the fraction responding that they will decrease. The survey data are available at the regional level (West, North Central, Northeast, and South), which we match with the MSA-level house search indices by taking averages within each region. We then analyze the relation between search activity and house price expectations by computing correlation coefficients across regions. The results indicate that housing

search intensity is strongly linked to house price expectations as the correlation coefficients range from 0.56 (Northeast) to 0.78 (North Central). These results are consistent with the notion that home buyers increase their search intensity and, thus, their demand, in part because of higher expected future house prices.

A.12 REITs

Prices of residential real estate investment trusts (REITs) provide a direct market-based view of real estate conditions. Unlike residential house price indices such as the purchase-only FHFA index, it is possible for individuals to invest and trade in REITs. Because REIT prices should reflect investor expectations about future fundamentals, they may potentially also contain relevant forward looking information about future FHFA house price growth rates.

To analyze this possibility, we obtain information about publicly traded residential REITs from the S&P Global SNL Real Estate Database and merge it with stock return data obtained from CRSP. The residential REITs are those that have primary property type specified as "Multifamily Apartments." We use data on 21 residential REITs that have been traded during our sample period from 2004:1 to 2020:12. Seven out of these are traded during the whole period, while the rest have missing observations in certain parts of the sample. To get an aggregate measure of the residential REIT market, we compute an equal-weighted average of the firm-specific REIT returns. Next, we estimate predictive regressions of the form, $p_{t+h} - p_t = \alpha + \beta R_t^{REIT} + \varepsilon_{t+h}$, where p_t is the log of the FHFA house price index, R_t^{REIT} is the average return on the residential REIT market, and h is the forecast horizon in months.

In Panel A of Table A.8, we see that there is a positive relation between R_t^{REIT} and future growth rates on the FHFA index, but it is not statistically significant according to the bootstrap-based t -statistics. These results imply that although REIT returns represent a market-based view of real estate conditions, they are not useful in forecasting residential house prices.

Table A.8. REITs. Panel A reports results from predictive regressions of h -step-ahead log growth rates on the FHFA house price index using residential REIT returns as predictors. Panel B reports results from predictive regressions of h -step-ahead log excess returns on residential REITs using HSI as the predictor. For each regression, the table reports the estimate of β , the corresponding t -statistic in parenthesis, and the R^2 in square brackets. We compute standard errors using a circular block bootstrap. All predictive variables are standardized and slope coefficients are multiplied by 100. The sample period is 2004:1-2020:12.

$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A			
0.11	0.17	0.28	0.52
(1.88)	(1.18)	(0.88)	(0.78)
[3.67]	[1.17]	[0.96]	[0.92]
Panel B			
0.77	2.19	4.53	9.42
(0.98)	(0.95)	(1.18)	(1.37)
[1.23]	[3.18]	[5.51]	[11.23]

The REIT market also provides an opportunity to analyze the source of predictability emanating from the HSI . If HSI captures a time-varying risk premium component due to time-varying risk or risk aversion, we would expect that its predictive power for residential house price indices carries over to the REIT market. In contrast, if the predictive power from the HSI stems from search frictions, it should hold no predictive power for REITs since returns on REITs are largely unaffected by such frictions. In Panel B, we show results from forecasting h -step-ahead log excess returns on the residential REIT market using HSI as our predictive variable. We see that the HSI has very limited predictive power over REIT returns. The sign on HSI is positive in accordance with the predictive results obtained when forecasting house prices on the residential real estate market. However, HSI is insignificant across all horizons. We view these results as indicative evidence that the predictive power of HSI over future house prices does not arise from a risk compensation channel, but is more likely to reflect sluggish price adjustments in the residential real estate market due to frictions.

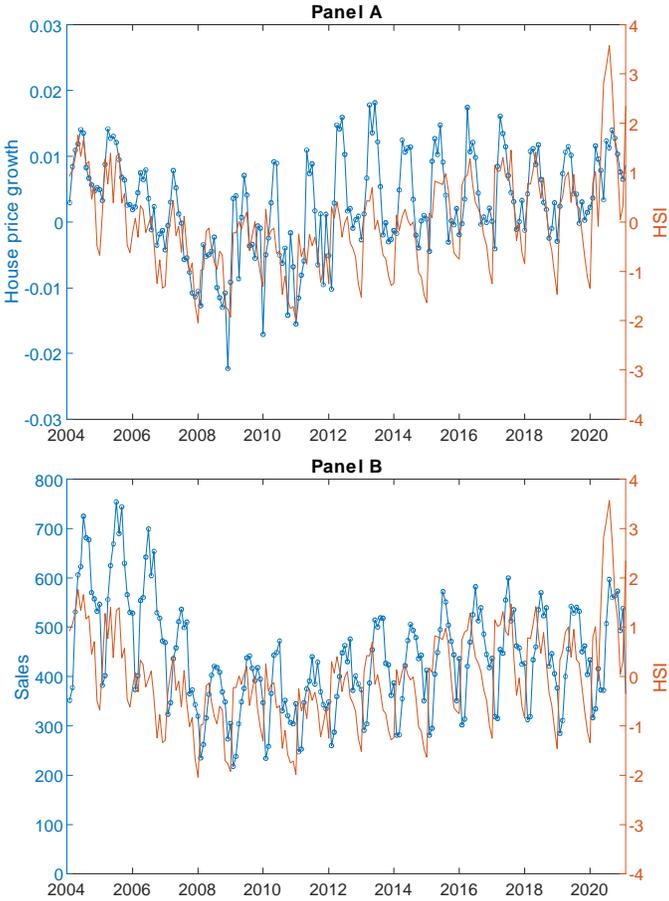
A.13 Non-Seasonally Adjusted Data

In the main paper, we analyze seasonally-adjusted houses prices, but it is well-known that house prices contain a strong seasonal component with high prices during spring and summer and low

prices during fall and winter (e.g. Ngai and Tenreyro, 2014). Fundamentals such as rental flows, the level of the mortgage rate, or credit variables do not follow such a seasonal cycle. Growth in output and related economic variables typically have a strong seasonal cycle, but with a boom in the fourth quarter (e.g. Barsky and Miron, 1989), which does not match the seasonal pattern on the housing market. In addition, it seems unlikely that sentiment/beliefs about the outlook for the housing market should move up and down each quarter according to the season. In the paper, we measure sentiment as the fraction of respondents who answer that now is a "good time" to buy a house from the University of Michigan's Survey of Consumers (Cox and Ludvigson, 2019). This series does not show any pronounced seasonality. If *HSI* captures seasonal variation in house prices, it can therefore be taken as evidence that it contains information beyond what is contained in typically used housing market determinants. To check this, we therefore reconstructed *HSI* based on seasonally unadjusted data using the approach outlined in Section 2.2 of the paper.

In panel A of Figure A.6, we plot the non-seasonally-adjusted search index together with non-seasonally adjusted house price growth. As is evident from the figure, the two series tend to move closely together with a strong seasonal cycle in both search and house prices. Search activity tends to reach its lowest points in November and December, while the peak points vary more from year to year but always occur in the spring or early summer months. In comparison, house prices tend to be lowest in December and often peak in May or June. Panel B illustrates the seasonal variation in sales, which typically peaks in June and reaches its lowest point in January, which also implies that search tends to move before sales.

Figure A.6. Non-Seasonally Adjusted Housing Search Index. Panel A shows the non-seasonally adjusted housing search index (HSI) along with the log growth rate in the non-seasonally adjusted purchase-only FHFA house price index. Panel B shows the non-seasonally adjusted HSI along with the non-seasonally adjusted sales of existing single-family housing units from the National Association of Realtors. The sample period is 2004:1-2021:1.



In Panel A of the Table A.9, we use the seasonally unadjusted HSI to predict seasonally unadjusted house price growth rates. We see that the predictability results are strong for $h = 1$ but weaken for $h > 1$. By increasing the forecast horizon, we often predict from one season into another and, in addition, the use of overlapping growth rates removes the extent of seasonality. In Panel B, we therefore use both the non-seasonally and seasonally adjusted HSI as predictive variables. We find that the non-seasonally and seasonally adjusted HSI jointly have strong predictive power for non-seasonally adjusted house prices across all forecast horizons.

Table A.9. Predicting Seasonally Unadjusted House Prices. The table reports results from predictive regressions of the h -step-ahead growth rate in house prices, $p_{t+h}^{NSA} - p_t^{NSA}$, where p^{NSA} is the log of the non-seasonally adjusted FHFA house price index. We use both the seasonally adjusted and unadjusted housing search index as predictors. For each regression, the table reports the slope estimates, the corresponding t -statistics in parenthesis, and the R^2 in square brackets. We compute standard errors using a circular block bootstrap. All predictors are standardized and slope coefficients are multiplied by 100 to facilitate comparison across variables. The sample period is 2004:1-2021:1.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
Panel A: Seasonally unadjusted HSI				
NSA HSI	0.52 (6.12) [47.23]	1.18 (5.47) [34.93]	1.29 (2.81) [14.50]	2.48 (3.31) [20.73]
Panel B: Seasonally unadjusted and adjusted HSI				
NSA HSI	0.42 (5.48)	0.61 (3.27)	-0.50 (-3.07)	-0.00 (-0.00)
SA HSI	0.15 (1.76) [49.51]	0.85 (3.55) [44.97]	2.78 (6.88) [53.88]	4.35 (5.06) [64.23]

A.14 Local Search

There is widespread evidence that housing markets are local in nature and segmented (see, e.g., Del Negro and Otrok, 2007, Gyourko et al., 2013, Glaeser et al., 2014, and Hernández-Murillo et al., 2017). Consistent with this evidence, in Table 8 of the main paper, we show that local housing search stays statistically significant across all forecast horizons when controlling for national-level housing search. To further verify that local search, unrelated to the aggregate search patterns of the economy, plays a significant role in predicting local house price changes, we orthogonalize local search with respect to US-level aggregate search:

$$HSI_{it} = b_i HSI_{US,t} + HSI_{it}^{Local} \quad (\text{A2})$$

where HSI_{it}^{Local} is the part of MSA-level search, HSI_{it} , that is unrelated to the aggregate US-level search behavior, $HSI_{US,t}$. Next, we examine whether expected local house price changes have both a local and national component by estimating

$$p_{it+h} - p_{it} = \alpha_i + \beta_{Local} HSI_{it}^{Local} + \beta_{US} HSI_{US,t} + \varepsilon_{it+h}, \quad (\text{A3})$$

where β_{Local} measures the impact of the part of local search that is orthogonal to aggregate national-level search behavior. Table A.10 shows the results of estimating (A3) on our sample of MSAs. Across horizons, we see that both national-level search and the part of local search that is unrelated to aggregate search patterns are strongly statistically significant. These results confirm that local-specific search, orthogonal to aggregate search patterns of the economy, plays a significant role in the prediction of local house price changes.

Table A.10. Predicting Local House Prices With Local Search Orthogonal to US search: Evidence From Panel Regressions. The table reports results from fixed effects panel regressions of the form, $p_{it+h} - p_{it} = \alpha_i + \beta_{Local} HSI_{it}^{Local} + \beta_{US} HSI_{US,t} + \varepsilon_{it+h}$, where p_{it} is the log of the Freddie-Mac house price index in MSA i , $HSI_{US,t}$ is the national-level housing search index, HSI_{it}^{Local} is the part of MSA-level search that is orthogonal to $HSI_{US,t}$, and h is the forecast horizon in months. We compute standard errors using Thompson (2011) two-way clustered robust-statistics with h lags. HSI is standardized to facilitate interpretation of the β estimates. The sample period is 2004:1 to 2021:1.

	$h = 1$	$h = 3$	$h = 6$	$h = 12$
U.S. HSI	0.43 (11.80)	1.28 (10.36)	2.44 (8.51)	4.38 (6.37)
Local HSI^\perp	0.20 (8.29) [48.11]	0.59 (7.07) [49.86]	1.13 (6.45) [49.26]	2.08 (7.28) [45.43]

References

- Akerlof, G., Shiller, R. (2009). *Animal spirits: How human psychology drives the economy and why it matters for global capitalism*. Princeton, NJ: Princeton University Press.
- Barsky, R.B., Miron, J.A. (1989). The seasonal cycle and the business cycle. *Journal of Political Economy* 97, 503-534.
- Bayer, P., Geissler, C., Mangum, K., Roberts, J. (2020). Speculators and middlemen: The strategy and performance of investors in the housing market. *Review of Financial Studies* 33, 5212–5247.
- Breiman L. (1997). *Arcing the edge*. Technical report 486, Department of Statistics, University of California Berkeley.
- Breiman, L. (2001). Random forests. *Machine learning* 45, 5-32.
- Campbell, J., Thompson, S. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies* 21, 1509-1531.
- Chen, T., Guestrin, C. (2016). Xgboost: A scalable tree boosting system. Working paper.
- Chinco, A., Mayer, C. (2016). Misinformed speculators and mispricing in the housing market. *Review of Financial Studies* 29, 486-522.
- Cox, J., Ludvigson, S.C. (2019). Drivers of the great housing boom-bust: Credit conditions, beliefs, or both? *Real Estate Economics* 2019, 1-33.
- Del Negro, M., Otrok, C. (2007). 99 luftballons: Monetary policy and the house price boom across US states. *Journal of Monetary Economics* 54, 1962-1985.
- Diebold, F.X., Mariano, R.S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 253-263.
- Favilukis, J., Ludvigson, S., Van Nieuwerburgh, S. (2017). The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy* 125, 140-223.
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics* 29, 1189-1232.
- Gao, Z., Sockin, M., Xiong, W. (2020). Economic consequences of housing speculation. *Review of Financial Studies* 33, 5248-5287.
- Glaeser, E.L., Gyourko, J., Morales, E., Nathanson, C.G. (2014). Housing dynamics: An urban approach. *Journal of Urban Economics* 81, 45-56.
- Guo, H. (2009). Data revisions and out-of-sample stock return predictability. *Economic Inquiry* 47, 81-97.
- Gyourko, J., Mayer, C., Sinai, T. (2013). Superstar cities. *American Economic Journal: Economic Policy* 5, 167-199.
- Hale, T., Angrist, N., Goldszmidt, R., Kira, B. Petherick, A., Phillips, T., Webster, S., Cameron-

- Blake, E., Hallas, L., Majumdar, S., Tatlow, H. (2021). A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nature Human Behaviour* 5, 529-538.
- Han, L., Strange, W. (2015). The microstructure of housing markets: search, bargaining, and brokerage. *Handbook of Regional and Urban Economics*, Edited by Duranton, G., Henderson J.V., Strange, W.C., 5, 813-886.
- Head, A., Lloyd-Ellis, H., Sun, H. (2014). Search, liquidity, and the dynamics of house prices and construction. *American Economic Review* 104, 1172-1210.
- Hernández-Murillo, R., Owyang, M.T., Rubio, M. (2017). Clustered housing cycles. *Regional Science and Urban Economics* 66, 185-197.
- Hale, T., Angrist, N., Goldszmidt, R., Kira, B. Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S., Tatlow, H. (2021). A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nature Human Behaviour* 5, 529-538.
- Hodrick, R. (1992). Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial Studies* 5, 357-386.
- Kan, R., Zhang, C. (1999a). GMM tests of stochastic discount factor models with useless factors. *Journal of Financial Economics* 54, 103-127.
- Kan, R., Zhang, C. (1999b). Two-pass tests of asset pricing models with useless factors. *Journal of Finance* 54, 203-235.
- Mian, A, Sufi, A. (2009). The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis. *Quarterly Journal of Economics* 124, 1449-1496.
- Nathanson, C.G., Zwick, E. (2018). Arrested development: Theory and evidence of supply-side speculation in the housing market. *Journal of Finance* 73, 2587-2633.
- Ngai, L. R., Tenreyno, S. (2014). Hot and cold seasons in the housing market. *American Economic Review* 104, 3991-4026.
- Newey, W., West, K. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703-708.
- Politis, D.N., White, H. (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews* 23, 53-70.
- Shimer, R. (2004). Search intensity. Working Paper, University of Chicago.
- Taylor, J.B. (2014). Causes of the financial crisis and the slow recovery: A 10-year perspective. Maily M, Taylor JB, eds. *Across the Great Divide: New Perspectives on the Financial Crisis* (Hoover Press, Stanford, CA), 51-65.
- Thompson, S.B. (2011). Simple formulas for standard errors that cluster by both firm and time. *Journal of Financial Economics* 99, 1-10.
- Wei, M., Wright, J. (2013). Reverse regressions and long-horizon forecasting. *Journal of Applied Econometrics* 28, 353-371.
- Wu, L., Brynjolfsson, E. (2015). The future of prediction: How Google searches foreshadow housing prices and sales. In: A. Goldfarb, S. M. Greenstein, and C. E. Tucker (Eds.), *Economic analysis of the digital economy*, Chicago: University of Chicago Press.