FHFA STAFF WORKING PAPER SERIES

The Daily Microstructure of the Housing Market

Peter Chinloy
William D. Larson

February 2017

FEDERAL HOUSING FINANCE AGENCY
Office of Policy Analysis & Research
400 7th Street SW
Washington, DC 20024, USA

Please address correspondences to Peter Chinloy (peter.chinloy@temple.edu). Many thanks to Alex Bogin, Bob Collender, Susan Cooper, Don Corin, Debra Fuller, Catherine Jouet, Andy Leventis, Eric Maribojoc, Doug McManus, Brian Wiksell, Joe Williams, and Tony Yezer, who provided helpful questions, comments, and discussion. No outside funding was received for this research.

The Daily Microstructure of the Housing Market

Peter Chinloy, William D. Larson

February 2017

Abstract

The microstructure of the housing market includes periodic buyer liquidity constraints, high transaction costs, and bilateral negotiations on price and timing. These separately introduce daily price volatility and negative serial correlation that is suppressed at a monthly frequency. In a daily U.S. house price index, the annualized standard deviation of returns is 27 percent, versus 3 percent for monthly data. We attribute the daily volatility to repeating calendar-based liquidity price premiums (8 percentage points), transaction costs (7 pp), estimation and composition error (2 pp), and idiosyncratic shocks (10 pp). Monthly house price indices suggest housing has exceptionally high risk-adjusted returns. A daily index brings Sharpe ratios in line with other assets.

Keywords: liquidity, market microstructure, daily house price index, mortgages, volatility

JEL Classification: G21 · G23

Peter Chinloy
Department of Finance
Fox School of Business
Temple University
1801 Liacouras Walk
Philadelphia, PA 19122, USA
peter.chinloy@temple.edu

William D. Larson
Federal Housing Finance Agency
Office of Policy Analysis & Research
400 7th Street SW
Washington, DC 20024, USA
william.larson@fhfa.gov

1 Introduction

The microstructure of the market for housing contains many incentives that drive pricing and transaction timing for individual units. Transaction costs are high because houses are indivisible, expensive, heterogeneous, and are traded by individuals rather than large institutions. Negotiations are bilateral and opaque, with only the list price known to parties outside of the negotiations, which along with transaction costs, lead to high transaction spreads. Buyer liquidity is highest on paydays, and standard mortgage contracts offer liquidity that varies within a month.[1] Both liquidity factors provide opportunities for sellers to negotiate a relatively higher price on particular dates, leading buyers to potentially pay different prices for the same house. As a consequence, buyer and seller types vary by calendar date, and daily house prices may exhibit periodic, as opposed to ongoing, volatility.

Volatility of existing house price indices behaves contrary to this conventional wisdom about housing markets, let alone the personal experience of individual homeowners. When using standard house price indices at a monthly frequency, houses appear to be assets with exceptionally high risk-adjusted returns and little volatility. Table 1 shows between 2000 and 2014, housing returns have a Sharpe ratio of 1.3 and an annualized standard deviation of 3.3 percent. Over the same period with comparable monthly data, the S&P 500 has a Sharpe ratio of 0.18 and an annualized standard deviation of 13.8 percent.

Either housing offers returns an order of magnitude greater than stocks, or volatility measures for housing calculated using monthly house price indices do not represent the true volatility of house prices. Lending support to the notion that monthly house price index volatility is misleading are the findings of Case and Shiller (1989), Flavin and Yamashita (2002), and others, who show using micro-level data the annualized volatility of housing is closer to 15 percent, near that of stocks. Thus, a puzzle exists regarding the returns to housing: why is monthly house price index volatility so misleading, and is there a link to the microstructure of the housing market?

In this paper, we conceptualize daily house price movements as originating from order flow-based liquidity factors (e.g. Pastor and Stambaugh, 2003), market frictions (Roll, 1984), and idiosyncratic shocks. Liquidity factors, such as paydays and mortgage contract timing,

[1] In a typical mortgage contract, daily interest accrued between closing and the end of the month ("per-diem interest") is paid at closing.

depend on the calendar day and repeat each month, contributing to saw-toothed price behavior at a daily frequency, but with no effect on monthly measures. Separately, ongoing transaction costs cause small mean-reverting daily price fluctuations that have little effect when aggregated over time. Idiosyncratic innovations have no daily correlation, but when aggregated, introduce positive autocorrelation in the series (Working, 1960).

Due to liquidity and transaction cost frictions, daily house prices are negatively autocorrelated. But when aggregated to a monthly frequency, the within-month liquidity and transaction cost factors offset and diminish, and the measure is dominated by price innovations. This model is therefore able to explain the microstructure behind the temporal aggregation bias observed in house price indices. The negatively correlated factors are suppressed at a monthly frequency, but the underlying transactions reflect these high-frequency dynamics. Homeowners, mortgage holders, and investors in asset-backed securities face default risk that is obscured by standard low-frequency (i.e. monthly) indices.[2]

Study of the microstructure of the housing market, with its liquidity, frictions, and high-frequency dynamics, requires a daily house price measure. Using a sample of 33 million purchase-mortgage transactions, we construct a national, daily repeat-sales house price index from 2000 through 2014. This index has a high degree of autocorrelation and numerous large and significant calendar effects. In particular, transactions on the 1st business day of the month, the 15th calendar day of the month, and Mondays each give 1 to 2 percent price premiums for identical homes. Even controlling for these effects, there is still a large amount of suppressed volatility when moving from daily to monthly frequency house price indices. Of the raw 27.5 percent annualized standard deviation in daily returns, 2.5 percentage points are due to error related to sample composition and estimation error, 7.6 are due to the repeating calendar liquidity effects, 6.7 are due to large spreads vis-a-vis transaction costs, and 10.5 are due to idiosyncratic shocks.

Liquidity and transaction frictions separately generate negative autocorrelation at high frequencies, facilitating construction of standard financial time series indicators. The characteristics of the price and volume series resemble classic order flow dynamics, leading us

[2]While we make no argument about the efficiency of the housing and primary mortgage markets, as the forces that cause prices and transaction timing is endogenous, there may be important effects on the secondary mortgage market. Because houses transact at different prices for identical units, differential default and prepayment rates for houses purchased on different days of the month may result.

to construct classic Amihud (2002) and Pastor and Stambaugh (2003) measures of housing market liquidity. These measures are highly correlated and correctly signed throughout the sample. Volume-return relationships on the 1st and 15th of the month are particularly strong.

We also construct the Roll (1984) market friction measure. Transaction costs for housing are known to be high, and contribute to substantial price volatility. This measure, calculated using the daily house price index, is approximately five times an equivalent measure for the S&P 500. While the measure for stocks is often set to zero due to positive autocorrelation, there is always negative autocorrelation in the house price series, resulting in a well-behaved friction measure.

2 Background

The microstructure of financial markets is well-developed, highly efficient, and has been the subject of intense research. Some key findings from this literature are that surges in liquidity predict returns (Pastor and Stambaugh, 2003), and market frictions introduce volatility into financial time series (Roll, 1984). These factors introduce negative autocorrelation in returns—little information is conveyed by an individual trade despite the presence of fluctuations. This has been demonstrated using a variety of measures, including the ratio of leverage or assets to book value (Adrian, Etula, and Muir, 2014), the ratio of returns to volume (Amihud, 2002; Acharya and Pedersen, 2005), and measures of the bid-ask spread or transaction costs (Roll, 1984; Hasbrouck, 2009; Goyenko, Holden, and Trzcinka, 2009).

However, despite research suggesting links between liquidity and returns, these findings are not always empirically robust (Hou, Xue, and Zhang, 2014 and 2015). For instance, in the time series for Pastor and Stambaugh's (2003) liquidity measure in Pastor (2016), the liquidity factor has the opposite of the predicted sign in 1/3 of observations. In addition, in the time series for Hasbrouck's (2009) transaction friction indicator in Hasbrouck (2016), the value is negative and set to zero in 1/5 of the sample. Housing, on the other hand, has a number of characteristics that suggest high transaction frictions and liquidity constraints, making it an ideal market in which to measure and test these textbook financial theories.

There has been little research on daily financial indicators for housing, including daily prices, trading volumes, liquidity, and transaction costs.[3] Rather, research on the microstructure of housing markets tends to focus on the search process between buyers and sellers, buyer versus seller bargaining power, and incentives of real estate agents and brokers.[4] The housing market has many features similar to other physical and financial assets, such as transaction costs, liquidity constraints, and potential principal-agent dynamics. But certain characteristics unique to housing transactions may influence trade timing, pricing, and other outcomes relevant to financial markets.

First, houses are exceptionally difficult to price. Housing is a physical, immobile, rarely transacted, heterogeneous, composite commodity, whose value consists of the implicit value of land, structures, schools, safety, and other attributes. An individual home typically has a holding period measured in years, leading to the possibility of large changes in transaction prices for the same unit resulting from location- and attribute-specific supply and demand dynamics. This leads to the use of "comparable" transactions for appraisals, essentially a reliance on past sales of imperfect substitutes in somewhat different locations and time periods. All of which leads to substantial uncertainty regarding the current fair market price for a housing unit, contributing to substantial bid-ask spreads.

Housing is also expensive and indivisible, making it highly illiquid. The physical nature of housing requires deeds, recordation, and physical insurance of collateral when leveraged purchases are made. The illiquidity of the asset, combined with its heterogenous nature, makes information costly to acquire and transaction costs high. A typical housing transaction also involves representation by agents who charge a commission as a function of the sale price of the unit, typically 3 percent each for the buyer's and seller's agent. Overall, transaction costs for housing are estimated to be in the 7 to 10 percent range by Harris et al. (2013), including agent fees, mortgage origination fees, recordation taxes, but excluding moving and search costs. In the Roll (1984) model this would suggest a bid-ask spread nearly 20 percent of the price of the asset, not to mention the above uncertainty regarding pricing.[5]

[3]One major exception is the work of Bollerslev, Patton, and Wang (2015), which introduces daily house price indices for 10 cities. This paper will be discussed in greater detail in the following sections.

[4]For an excellent overview of recent literature on housing market microstructure, see Han and Strange (2015).

[5]Note that the "bid-ask spread" for housing is different than in financial markets. In financial markets, the bid and ask are known and transparent. In housing markets, the list price is not necessarily an ask price because it is non-binding, and bids are subject to negotiation.

It is also important to consider the temporal dimension of housing transactions. When an offer is agreed upon, it includes a price and a transaction date at some point (typically 30-45 days) in the future. The signing of a contract launches a further series of bilateral negotiations that may include additional seller concessions and potential changes to the closing date and effective price. At the eventual closing date, the transaction occurs. The result of this sequence may lead volumes and prices to be concentrated at specific times of the week or month in accordance to both observed and unobserved heterogeneity in the buyer, seller, agents, or housing units.

To fully examine each of these factors requires a search-and-matching model of buyers and sellers, such as Novy-Marx (2009) and Ngai and Tenreyro (2014). We do not attempt to perform such an exercise here, and instead focus on the time-series behavior of transaction volumes and prices. Rather, we posit some potential reasons why we might expect temporal clusters of outcomes, which we examine in a reduced-form model in the following sections.

It has been documented that liquidity-constrained sellers transact at lower prices, as documented by Hendel, Nevo and Ortalo-Magne (2009), who find that similar houses sell for different prices depending on the trading platform, and Campbell, Giglio and Pathak (2011), who find foreclosures and estate sales sell for 20 percent less than when the seller is less constrained. Sellers may also need to sell a house in order to close on a contingent purchase, and do so at a slightly lower price. While these factors may influence seller liquidity, to our knowledge there is no clear reason to expect there to be temporal clustering for these reasons.

However, there are some clear reasons for temporal clustering of transactions based on duelling incentives between buyers and sellers. Those who borrow pay daily interest between the closing date and the end of the month at closing, necessitating greater out-of-pocket-costs the earlier a buyer closes. Liquidity is thus higher for buyers later in the month, indicating they may be willing to pay slightly higher prices. On the seller side, a typical mortgage payment has a grace period until the 15th of the month, increasing seller liquidity early in the month. A seller may therefore be more willing to sell at a slightly lower price in order to close earlier in the month. This tension is balanced most clearly on the 15th, when a homeowner can execute a sale and a purchase on the same day and maximize liquidity.[6]

Some non-liquidity factors may influence timing as well. Both buyers and sellers have an incentive to close before the end of a calendar month due to monthly billing cycles for mortgages and rental contracts.[7] Operational considerations include the myriad of extensive paperwork and timed disclosure requirements for a transaction to take place. This includes mortgage loan disclosure statements, deed and title preparation, and a physical meeting to sign paperwork and complete the transaction. These are often tied to a weekly cycle which results in transaction volumes rising throughout the week to a peak on Fridays. Insofar as buyers versus sellers have different incentives to make these Friday deadlines, price differentials may arise.

Overall, the microstructure of the housing market is one with a variety of factors potentially leading to high bid-ask spreads, high transaction costs, and transaction volumes and prices that are dependent on the calendar day of the week or month. While some of these factors are symmetric, such as the desire to close early in the month to give time to move, others may not be, suggesting the possibility price premiums that either exist at the time of contract offer and ratification, or between contract signing and closing. Substantial within-month fluctuations in the prices of identical homes may therefore exist, with several implications to be explored in the next section.

3 Empirical Framework

In our stylized model, daily house price dynamics originate from three main sources: daily liquidity or "order flow" differences, bid-ask spread and transaction cost dynamics, and idiosyncratic shocks. Combined, this model gives a rich representation of daily house prices that allows us to explore both within-month dynamics and consequences of aggregation. The structure incorporates volume surges around liquidity events, such as paycheck timing or requirements that mortgages and rents be paid on specific days. It also takes into ac-

[6] Other mortgage contract provisions that may influence timing are interest "holidays" if a closing occurs on the 1st of a month, and a tendency of per-diem interest for units closing on the 15th of the month or later to be lumped into the balance of the mortgage. Finally, buyers may wish to close several weeks before the end of the month if there are modifications to the house that a buyer wishes to undertake before move-in, which for renters, typically coincides with the end of the month.

[7] Not all buyers are liquidity constrained. Some make cash purchases and thus rely on neither paycheck timing nor mortgage contract stipulations. However, a large fraction of all home buyers finance part of their purchase through debt, meaning they are potentially susceptible to small fluctuations in cash-on-hand and are beholden to standard mortgage contract timing.

count decentralized and bilateral trading, allowing for transactions within a bid-ask spread for a particular day. Each of these two processes has negative serial correlation at a daily frequency. However, when aggregated to a monthly frequency, this structure gives a series with positive autocorrelation under very modest assumptions. As the aggregation horizon expands, liquidity effects and bid-ask bounces are smoothed and dominated by the accumulation of idiosyncratic shocks.

The change in house prices at time t on day of the month d, $\Delta y_{t,d}$, is a function of a repeating sequence of calendar-based liquidity effects, z_d, frictional costs cq_t, and an independently and identically distributed (IID) innovation w_t. Each of these variables is described more fully in turn.

$$\Delta y_{t,d} = cq_t + z_d + w_t \tag{1}$$

The frictional shock is based on the bid-ask dynamics of Roll (1984) and Hasbrouck (2009). Transaction costs are c giving a bid-ask spread of $2c$, within which trades occur. In this model, trades occur at the bid or the ask, where q_t is a binary Markov variable set to -1, indicating a transaction at the bid, or +1, indicating a transaction at the ask, with equal, independent probabilities. While specified as model based on transaction-level data, it has also been used in a daily context by Roll (1984).

$$cq_t \in \{-c, c\};\ E[cq_t] = 0;\ var(cq_t) = -c^2/2;\ cov(cq_t, cq_{t-j}) = -c^2/4\ \forall\ j > 0 \tag{2}$$

The autocovarance of cq_t is $cov(cq_t, cq_{t-1}) = -c^2/4$ and the variance of cq_t is $var(q_t) = c^2/2$. The negative autocovariance is due to transactions occurring either at the bid or the ask, causing successive trades to occur with a difference of zero or $-2cq_{t-1}$. Because q is a Markov variable, the autocovariance is identical, regardless of the time period between observations, because the state of any particular day contains no information from the prior day.

Liquidity volatility is based on sign-directed volume, or order flow (Pastor and Stambaugh, 2003). While there are other liquidity models available, including leverage from financial intermediaries (Adrian, Etula, and Muir, 2014; Acharya and Pedersen, 2005), an order flow volume representation is preferred because it describes empirical house price and volume dynamics, as later sections show. In this model, when buyers are liquidity constrained, they are willing to pay in excess of the market return to gain liquidity. In such periods, trading volumes rise reflecting the introduction of additional sellers into the market.

When this liquidity shock fully recedes in the following period, there is negative autocorrelation in returns. This order flow variable is defined as a repeating sequence of fixed effects, z_d, with subscript d indicating the day of the month, with variance σ_z^2. These fixed effects are equal to $sign(\Delta y_{t-1}) \times N_{t-1}$, in the classic rendition of the model, where $sign(\Delta y_{t-1})$ is the sign of the prior period's price change, and N_{t-1} is the prior period's volume. In this model, price innovations affect current period prices though new information, and the next period through mean reversion at rate $\gamma \times N_t$, where γ is the order flow parameter of Pastor and Stambaugh (2003).

$$E[z_d] = 0;\ var(z_d) = \sigma_z^2;\ cov(z_d, z_{d-j}) = -\sigma_z^2/(D-1)\ \forall\ 0 < j < D-1 \qquad (3)$$

Because of our repeating sequence rendition, we specify the fixed effects as mean zero with autocorrelation equal to $-1/(D-1)$ where D is the number of days in a month.

In order to understand the dynamics of $\Delta y_{t,d}$, we must also make the additional assumption that q_t, z_d, w_t are independent. The daily autocorrelation for housing returns is then

$$corr(\Delta y_t, \Delta y_{t-1}) = -\frac{c^2/4 + \sigma_z^2/(D-1)}{c^2/2 + \sigma_z^2 + \sigma^2} < 0 \qquad (4)$$

This equation shows the daily autocorrelation of housing returns is negative, and due to both transaction costs (bid-ask spreads) and liquidity constraints (order flow dynamics). Repeat-sales residuals from Case and Shiller (1989) suggest $\sigma = 0.15$, transaction cost estimates assign $c = 0.1$, and $D = 19$ trading days in a typical month. Under a standard deviation of liquidity effects of $\sigma_z = 0.02$, the daily autocorrelation of housing returns is about -0.09.

Temporal aggregation removes most of the volatility from liquidity and friction sources. Aggregation has further consequences for the IID innovation term, introducing positive autocorrelation into the series. From the derivations of Working (1960) on the correlation of differences in a random walk with IID shocks of variance σ^2, if m is an aggregation period of variable X, let $\Delta X_{t,m} = \frac{1}{m}(\Delta X_t + \Delta X_{t-1} + ... + \Delta X_{t-m})$. Then, the variance of $\Delta X_{t,m}$ is $var(\Delta X_{t,m}) = \frac{2m^2+1}{3m}\sigma^2$, the covariance between $\Delta X_{t,m}$ and $\Delta X_{t-m,m}$ is $cov(\Delta X_{t,m}, \Delta X_{t-m,m}) = \frac{m^2-1}{6m}\sigma^2$, and the correlation between $\Delta X_{t,m}$ and $\Delta X_{t-m,m}$ is $corr(\Delta X_{t,m}, \Delta X_{t-m,m}) = \frac{m^2-1}{2(2m^2+1)}$. As m rises, the autocorrelation converges to 1/4.

Averaging over a month assigns $m = D$. Because z_d is a repeating sequence of mean zero, the mean and variance of the sum is zero when aggregated over a month. The variance and autocovariance of the bid-ask term cq_t is averaged over D periods. The effect on monthly autocorrelation from spreads is therefore falling with D, as both the variance and autocovariance fall at the same rate, while the effect from idiosyncratic shocks is rising. This leaves the following as the correlation of housing returns at a monthly frequency,

$$corr(\Delta y_{t,D}, \Delta y_{t-D,D}) = \left(-c^2\frac{1}{4D} + \sigma^2\frac{D^2-1}{6D})\right) \Bigg/ \left(c^2\frac{1}{2D} + \sigma^2\frac{2D^2+1}{3D}\right) \tag{5}$$

In the equation above, if m were at a daily frequency ($D = 1$), then ignoring the order flow effects, the returns autocorrelation would reduce to the daily autocorrelation for the transaction cost component, $-1/2$. Setting $D = 19$; $c = 0.1$; $\sigma = 0.15$ gives a monthly autocorrelation of about 0.25. This is both twice as large and opposite sign of the daily autocorrelation, and near the convergence point for the autocorrelation of idiosyncratic shock alone. In sum, as D rises, the effects of the microstructure of the housing market fall to zero, and the idiosyncratic shocks dominate. Monthly aggregation creates artificial momentum in the series that does not exist at a daily frequency.

The model shows how market microstructure can give negative autocorrelation in high-frequency series, but positive autocorrelation in low-frequency series. We stochastically specify this model using a GARCH specification. In practice, daily house prices require additional time series modeling considerations. To our knowledge, the only previous research on daily house price indices is Bollerslev, Patton, and Wang (2015). They consider ten U.S. cities, and begin their analysis by constructing a noise filter to uncover the latent signal in what they believe to be noisy daily estimates. After recovering the filtered appreciation rates from the raw series using the Kalman MA(1) filter, they fit this series using a heterogenous autoregressive model (HAR) with generalized autoregressive conditional heteroskedasticity (GARCH) terms.[8] The conditional mean portion of their HAR-GARCH model includes the

[8]In the literature on stocks, microstructure noise is often removed by smoothing with a Kalman filter as well (Diebold, 2006).

1st and 5th lags of the series in order to account for the prior day and the same day of the prior week, the lagged ten-city average level of daily appreciation, and the lagged monthly moving average of the series.

In order to simultaneously smooth the series while taking into account these rich deterministic terms and dynamics, we model the conditional mean of daily house prices with an MA(1) term, mimicking Bollerslev, Patton, and Wang's (2015) specification, and the possibility of higher-order lags up to a length of 19 business days (approximately one month). We also specify the order flow term using a vector of calendar-based fixed effects for the day of the week (Monday omitted), day of the month (the 1st omitted), and the first and last 3 business days of the month. The lags are chosen based on a combination of likelihood ratio tests and visual inspection of the autocorrelogram of the series.

The error w_t has lagged squared residual w_{t-1}^2 and the standard deviation of this error is s_t. The conditional variance equation also includes the log of the number of transactions used to construct the house price index N, summed over t and $t+1$.

$$\begin{cases} \Delta y_t = \Sigma_{k=1}^{K} \mu_k \Delta y_{t-k} + \mu_{avg} \frac{1}{19} \Sigma_{k=1}^{19} \Delta y_{t-k} + \Sigma_{j=1}^{5} \phi_{1j} DOW_j + \Sigma_{j=1}^{31} \phi_{2j} DOM_j + \\ \quad + \Sigma_{j=1}^{5} \phi_{3j} FIRSTBUSDAY_j + \Sigma_{j=1}^{5} \phi_{4j} LASTBUSDAY_j + \theta w_{t-1} + w_t \\ s_t^2 = \omega + \gamma w_{t-1}^2 + \lambda s_{t-1}^2 + \theta \ln N_t \end{cases} \tag{6}$$

In this model, $\mu_1 < 0$ implies negative autocovariance which is a necessary condition for spread-related volatility. When $\phi > 0$, the positive sign implies buyer liquidity constraints exist in the market and price premiums cause an increase in demand-driven volume.

4 Daily House Prices

This section implements a daily, repeat-sales house price index for the United States.[9] A national house price index is appropriate in this context because capital flows freely within the United States, and barriers to arbitrage across cities are due to high transaction costs, which we wish to measure. House price indices are constructed from purchase transactions

[9] The method is outlined in Appendix A, but omitted from the main text for brevity.

where a mortgage is acquired by Fannie Mae and Freddie Mac between 1975 and 2014. This dataset includes more than 33 million mortgages across the two enterprises, including the date of the closing, the loan-to-value ratio (LTV) at origination, and the origination loan amount, which allows us to calculate the purchase price. The model fitting the data includes annual dummy variables as controls for years between 1975 and 1999, and daily dummy variables from January 1, 2010 through December 31, 2014. Overall, the sample of repeat sales on days where the daily index is calculated includes 7.3 million transactions, giving estimates for 3,773 trading days between 2000 and 2014, or about 2,000 transactions per trading day.[10] A monthly house price index is estimated using identical methods for comparison purposes.

There are two main issues with the estimation of a daily house price index. The first is that the transaction price is essentially a smoothed, lagged spot price of housing because we observe transaction prices instead of contract prices. This introduces autocorrelation in the series up to a substantial lag length, as contracts signed on a particular date close on a number of separate but autocorrelated dates in the future. Complicating this further are the buyer and seller options to break the contract, potentially leading to bias in the observed transaction prices. A contract price or list price index, such as the one described by Anenberg and Laufer (2016), would potentially alleviate this issue while introducing others such as the changing relationship between the list and sale prices in hot versus cold housing markets (Carrillo et al., 2015).

The second issue is that transactions for a given time period may not be representative of the nation as a whole. Were two areas to appreciate at different rates, and at the same time, be sampled at different rates, a resulting price series could fluctuate due to changes in composition rather than changes in prices. We examine composition-related fluctuations across several dimensions in Section 7.3 and in the Appendix, and estimate this effect to be small.

Figure 4 shows transaction counts by month, day of the month, and day of the week. There is a clear annual cycle for housing, with volumes highest in June and lowest in January. Within months, we see evidence of borrowers shifting transactions from the early days of the

[10] Our method is similar to Bollerslev, Patton, and Wang (2015), who use an iterative procedure to estimate the index due to computing constraints. The daily index is estimated one month at a time, using daily dummy variables in the relevant month, plus monthly dummy variables as controls.

month to the 15th, the date at which mortgage interest is lumped into the balance of the loan and added to the next month's payment, as opposed to being due at closing. Transactions also have a within-month increasing trend, as both buyers and sellers attempt to close out transactions before the end of the month. Transactions on the first are also common, and are often the result of delays in transactions from the prior month, with exceptions often given by the lender regarding interest due at closing.

The extensive paperwork for contracts on rents, mortgages, payroll and structural conditions ties closings to business days (Monday through Friday). Within a week, there is an incentive to close later rather than earlier resulting from to the desire to close before a weekend. Three business days before closing, a "closing disclosure form" must be presented. If this spans a weekend, there is greater uncertainty regarding the terms and the possibility an item will change, requiring resubmission and the restarting of the 3 business day clock. In addition, on the weekend, the buyer and seller can move without losing payroll from work. Combined, these incentivize closing on Thursdays or Fridays rather than Mondays, Tuesdays, or Wednesdays.

Figure 2 shows monthly and daily house and S&P 500 closing prices between 2000 and 2014. The house price series rises during 2000 to 2005 while the S&P 500 falls after the "dot-com" crash in 2000 through 2002. After 2005, housing and stock prices appear to move in tandem. Volatility for house prices is relatively low in the 2000 to 2006 period, but rises after 2006 though the end of the sample. The average and standard deviations for capital gains are shown in Table 1. These show similar gross capital gains for housing and stocks, and for standard deviations at a daily frequency. However the standard deviation for returns to housing is much smaller at a monthly frequency, suggesting downward temporal aggregation bias.

When averaged over calendar dates, there appear to be patterns in the house price series that do not exist for the S&P 500. While month effects are small and statistically indistinguishable from zero, calendar day within the month and day within the week effects are large and significant. Figure 3 shows average daily capital gains for housing and for stocks over the sample for various calendar indicators. These simple statistics suggest calendar effects for housing but not for stocks.

These results appear to show effects on the 1st and 15th of the month for housing, with an identical house transacting with a daily appreciation rate of 1.25 percent if it happens to transact on the 1st, and 0.6 percent if it transacts on the 15th. The day-of-month calendar effects are correlated with volumes, suggesting liquidity effects on the price. There are also positive and robust but individually insignificant values for dates at the end of the month, suggesting the possibility of end-of-month effects as well. There also appear to be effects related to the day of the week. A house transacting on a Monday transacts with 0.8% daily appreciation, compared to Wednesdays and Fridays, which transact at -0.3 and -0.5 percent.

It is remarkable how large these statistics are, in magnitude. While transaction patterns fit what one would typically expect given what we know about housing and mortgage markets—with the exception of the spike around the 15th, which defies easy explanation—price patterns are not so easy to explain. For instance, while price and volumes tend to rise throughout the month due, presumably, to the desire to minimize per-diem interest, spikes in prices on the 1st and 15th are both large and surprising to the point where they suggest substantial heterogeneity is the cause.

Lending support to this heterogeneity hypothesis is a figure of average LTVs by calendar day of the month. This figure shows LTVs slowly rising within the month, consistent with the per-diem interest explanation. There are also negative spikes in LTVs on the first and near the middle of each month. These are a puzzle that we are unable to resolve within the present research. However, as our empirical framework hypothesizes and later sections show, because prices and volumes are correlated and tend to exhibit negative autocorrelation, we are able to attribute this within-month heterogeneity to liquidity and spread-related dynamics.

5 Model Results

This section presents comparisons between the time series properties for both housing and the S&P 500 at daily and monthly frequencies. In general, housing has a rich data generating process with long lags, calendar fixed effects, negative autocorrelation at short horizons, and positive autocorrelation at long horizons. Each of these factors lends support to our conceptual model of the microstructure of the housing market, which is characterized by high transaction costs and periods of illiquidity. In contrast, returns to the S&P tend to follow a random walk, with a slight negative autocorrelation at a daily frequency.

5.1 Daily Housing Capital Gains Models

Estimates from the sequence of models are presented in Table 2. Column 1 shows the first-order partial correlation of the series is about -0.36, suggesting a high degree of saw-toothed volatility. The Bollerslev, Patton, and Wang (2015) specification reported in column 2 gives a partial autocorrelation of 0.19 at the 1st lag and 0.32 at the 5th lag of the series. The trailing monthly average is -0.19 with the MA(1) parameter of -0.70. In the conditional variance equation, the lagged residual is 0.04 and the lagged variance parameter is 0.96. The log of the number of transactions used to estimate the index is -1.45, suggesting estimation-error related volatility.

Model 3 includes more lags than models 1 or 2. The lags are selected based on likelihood ratio (LR) tests and visual inspection of autocorrelograms. LR tests reject the null of no difference in the addition of lags past 25 lags, but explanatory power drops off substantially after $t-14$, with the exception of the lag at $t-19$. The 19th lag is important because there are often 19 trading days in a month, making this value analogous to a month-on-month partial difference. There are substantial effects of the 1st, 5th, 9th, 10th, 13th, 14th, and 19th lags. These tend to coincide with week-on-week effects in the case where the lag is a multiple of 4 or 5, depending on the number of trading days in a week.

The fourth model introduces the concept of calendar effects, specifying daily appreciation as a function of only a vector of fixed effects, with estimation using ordinary least squares. Fixed effects are mean-reverting if prices move upwards at the same time as volume spikes, consistent with order flow liquidity concepts. Estimated effects include those related to days-of-week, days-of-month, months, and first and last business days-of-month. The reference category for the day-of-week is Monday; for day-of-month it is the 1st; for month it is January; and for the business day of the month it is any business day that is not one of the first three or last three. Overall, these estimates suggest large day-of-week effects, a 15th day of the month effect, and 1st and last business day effects. These effects are significant, ranging between -0.6 and 1.9 percent for an identical house transacting on a different day.

The final estimation combines models 3 and 4 to include full dynamics along with calendar effects. The estimated calendar effects tend to shrink with the introduction of the time series variables (AR, MA terms), but are robust in sign. Autocorrelograms of the raw series and for residuals estimated in model 5 are shown in Figure 5. The panel for the raw housing series shows a series with five day-of-week effects (a 5-day cycle) with a long necessary lag order. The panel for the residual housing series shows that nearly all of the autocorrelation has been removed, indicating model 5 to be a well-specified model.

It may be tempting to conclude that, because housing transacts at higher prices on Mondays, the 1st business day, and the 15th of the month, that a homeowner should aim to transact a sale on one of these dates, and similarly, a buyer should avoid them. However, it should be noted that the transaction date, along with the price, are the result of a complex bilateral negotiation whose outcome results in these dates having higher observed transaction prices.

On the other hand, it may be possible to determine the collateral risk implications of these calendar effects. Mortgage default and prepayment differentials may arise either due to unobserved borrower characteristics or the higher transaction prices, all else equal, on dates with price differentials. In this case, there may be opportunity for arbitrage in the secondary mortgage market.

5.2 Daily S&P 500 Capital Gains Models

We also estimate identical models for daily S&P 500 capital gains over the same sample. The S&P has generally been shown to be efficient, with minimal autocorrelation or calendar effects. Models mimicking those in Table 2 are found in Table 3.

The partial autocorrelation estimate in column 1 is -0.06, about 1/7 that of the daily house price series, suggesting much more substantial saw-tooth-related volatility in measured house prices compared to stocks. The fit of each model is low, with R-squared statistics ranging from between 0.006 in model 1 to 0.2 in model 5. This compares to R-squared statistics between 0.10 and 0.48 for housing. Overall, it appears stock prices are much less explainable than house price dynamics. Liquidity and frictional dynamics are minimal, with lags of S&P 500 capital gains failing to be useful in predicting future gains. Similarly, there are no discernable liquidity surges leading to higher volumes and prices at particular dates. The autocorrelograms present a similar story, with very little series autocorrelation, and what little was there to begin with is eliminated using model 3.

Overall, these results suggest that while there may be liquidity effects by chance, they do not appear consistently. Results are inconclusive, similar to findings of Hou, Xue, and Zhang (2014). This is not the case for houses, which have complex dynamics and calendar effects, indicating both liquidity and frictional dynamics.

6 Volatility Decomposition

The time series models in the prior section suggest the presence of liquidity and friction-related volatility. In this section, we parse the raw volatilities in the housing and S&P 500 returns series in order to more fully understand the reasons for differences across assets and level of temporal aggregation. Sources considered are composition volatility arising from changes in geographic sampling; estimation error from the construction of the daily repeat-sales house price index; calendar liquidity effects, which we measure using calendar fixed effects; transaction costs, measured using autocorrelations in the data; and idiosyncratic shocks, which we attribute to residual volatility. Volatility from each source is assumed mutually exclusive, though there are potentially interactions between contributing factors. Table 4 shows the decomposition of the raw volatilities from these sources.

Volatility from geographic sample composition has two necessary conditions to exist: transactions must be sampled from different locations at changing frequencies, and the locations with changing sampling frequency must have different appreciation rates. In order to estimate the extent of composition volatility, we construct a composition index p^c_{md}, where subscripts m and d indicate the month and the day within the month. This index is constructed as a weighted average price change using de-meaned monthly house price appreciation rates by state, p_m, along with daily variation in state-level transaction shares, λ_{md}. This index gives daily variation in prices within a month due to a pure composition effect.

$$p^c_{md} = \sum_{s=1}^{51} (p_m - \bar{p}_m)\lambda_{md} \tag{7}$$

This composition index is shown in Figure 6. According to this measure, composition-related volatility is low throughout the sample, with peaks in 2008 through 2011. The mean of the series is zero by construction. We interpret the standard deviation of this series as the composition volatility in the raw daily index. This value is approximately 0.001, on average, suggesting composition contributes about 0.1 percentage points (5.74 percent of the total) to the raw series standard deviation. By construction, it is zero for the monthly series.

While this composition index considers state-level variation, composition differences may also arise due to other factors such as center-city versus suburban differentials, high versus low price areas, or factors related to weather or industrial composition. We consider an alternative approach to measuring the effects of changing composition on our estimates in Appendix 2 using annual ZIP code level indices. This approach assumes a representative ZIP code-level housing unit, with within-year variation in calendar effects as residuals. There may be further dimensions over which composition may vary in material ways even within ZIP codes, including foreclosures or distressed sales, but we believe these to have comparatively small effects on the index–it requires both changes in composition, and simultaneously, changes in appreciation rates that are not offsetting, in order to affect the daily index.

Next, as seen from the estimate on the transactions parameter in the conditional variance equation in Table 2, there is estimation error-related volatility in the index. We determine the influence of estimation error by comparing the standard deviation of capital gains to what it would have been if it was estimated with infinite underlying transactions. Because each price index value is estimated, this exercise is performed using statistical formulas for the sampling distribution of the difference between two means.

In the expression below, e_t is the fraction of the standard deviation that is estimation error at time t. The first term is a divisor, and is the average standard deviation s. The second is the difference between the standard errors of change in the index at t and $t-1$ using N_t and N_{t-1} observations, and the standard error of the difference with infinite observations.

$$e_t = \frac{1}{s}(s_{N_t,N_{t-1}} - s_{\infty,\infty}) \tag{8}$$

We assume that the variance of the index at t and $t-1$ is identical and equal to the square of the average sample standard deviation, giving $s_{N_t,N_{t-1}} = s\sqrt{(N_t + N_{t-1})/(N_t N_{t-1})}$. The standard error of the difference is zero when the number of transactions in both periods is infinite. When we calculate this metric over the sample, we estimate estimation error to contribute an additional 0.07 percentage points (3.7 percent) to the standard deviation of the daily series. For the monthly house price series, estimation error contributes 0.007 percentage points to the standard deviation (0.75 percent).

The influence of calendar-based liquidity factors on volatility is measured using a regression-based approach. First, we estimate a model relating daily returns to a vector of calendar fixed effects. The standard deviation of the residuals from this model, presented in column

4 of Tables 2 and 3, are interpreted as the standard deviation of the series absent calendar-based liquidity effects. Because liquidity potentially varies over time, we estimate this model separately for each year. The standard deviation of this residual series is 0.0124 for housing and 0.0127 for the S&P 500. The difference between the raw series and this residual series is the contribution of liquidity to volatility, and is 0.48 percentage points (27.9 percent of the total) for housing and 0.01 percentage points (0.7 percent) for stocks.[11]

Transaction friction-related volatility is also estimated using regression residuals. Recall from Roll (1984) that transaction costs are only measurable if the autocorrelation is negative, with a magnitude related to the negative autocovariance in the series. Accordingly, when serial correlation has been eliminated from the series, so to has the influence of transaction costs. We therefore proceed by first measuring the sign of the autocorrelation in the series, and then if it is negative, eliminating serial correlation from the housing and S&P 500 daily and monthly series by capturing the residuals from a well-specified time-series model. If the series has zero or positive autocorrelation, the influence of transaction frictions is set to zero. If it is negative, the differential standard deviation of these residuals with the raw series is interpreted as the influence of transaction costs on the volatility of the raw series.

The autocorrelation terms in column 1 of Tables 2 and 3 are negative, allowing us to proceed with the residual-based approach for both daily series. However, the autocorrelation for monthly house prices is positive, and for stocks it is zero, giving zero influence of transaction costs for monthly data. For the daily housing series, column 3 in Table 2 is 0.013, and the difference with the raw series is 0.423 percentage points (24 percent) compared to the raw series. For the daily S&P 500 series, the residual standard deviation is 0.01279, just 0.004 percentage points (0.33 percent) lower than the raw series, though it is statistically significant.

The contribution of idiosyncratic shocks is taken as the remainder of the raw standard deviation minus the sum of the four factors. Recall from Equation 5 that the standard deviation of the accumulation of idiosyncratic shocks is $\sqrt{(2m^2+1)/(3m)}$. Setting $m = 19$ gives a multiplier of 3.56. The monthly standard deviation for the S&P 500 is about 3 times that of the daily shock, suggesting stock volatility to consist of mostly IID innovations.

[11] It is worth noting that F-tests do not reject the null of no effect of the vector of calendar effects on returns for stocks.

For housing, the monthly multiple is only about 1.5, indicating daily idiosyncratic shock volatility is suppressed. We posit that this is due to the fact that, as noted in Section 4, the daily appreciation rates are a smoothed weighted average of contract prices at different points in the recent past. This has the effect of smoothing the idiosyncratic shock. When the raw 1.73 percent daily standard deviation for housing returns is annualized to 27.5 percent, we gain a picture of the shares of each factor to annualized volatility. Of this total, 1.6 percentage points are related to composition, 1.0 are related to measurement error, 7.6 to calendar liquidity effects, 6.7% to transaction frictions, and 10.5 to indiosyncratic shocks.

7 Financial Indicators for Housing

7.1 Transaction Costs

In Roll (1984) and Hasbrouck (2009), the midpoint of the bid-ask spread is the price. Market efficiency results in a spread equal to twice the transaction cost. In this simple framework, when transactions occur, they occur at either the bid or the ask, giving rise to negatively correlated transaction prices. This gives an estimate of transaction costs equal to c at time t as

$$c_t = \sqrt{(-cov(\Delta y_t, \Delta y_{t-1}))} \tag{9}$$

The measure is complicated by two main factors. First, the price index is a smoothed average of historical contract prices, causing any negative autocovariance in daily prices to be severely understated, attenuating the measure. Second, daily price changes that are due to measurement error or composition effects may bias the autocovariance upward. These issues are examined in robustness checks and found to be minimal.

Transaction costs are calculated as a rolling 252 trading day (one year) average. A depiction of this monthly transaction cost measure is presented in Figure 8. In general, frictional costs were low in the early 2000s at 0.5 percent, rising sharply in 2004 to a high of 1.25 percent in 2008, falling in 2010, and remaining at about 1.1 percent through 2014. This frictional cost is the difference in prices at execution by buyers of the same house. The cost does not include more conventional transaction costs of selling a house, including brokerage and fees, which are capitalized into the price. In the Roll (1984) model, the spread is twice the transaction cost, so buyers at virtually the same time are paying a range of 2.2 percent in price differential for the same house.

An equivalent transaction cost measure is calculated using daily S&P 500 data, also shown in Figure 8. This suggests stocks to be an asset with much lower transaction costs than housing, with costs ranging between 0 and 0.5 percent, and settling at about 0.2 percent after 2009. The exception is a surge during the financial crisis of 2008-2009 and the U.S. debt rating downgrade in August 2011, when it reaches 1 percent. These estimates are similar to Hasbrouck (2009), who shows transaction costs for stocks at or below 0.5 percent in the early 2000s. Overall, frictional costs for housing are about five times that of stocks for most of the sample.

7.2 Liquidity

It is also possible to construct measures of housing market liquidity. There are a number of liquidity measurement methods available, but here we focus on those related to returns and volume, as outlined in Section 3. These measures show the effects of order flow changes on returns under the logic that periods of high buyer demand, sellers are able to extract a premium when buyers are liquidity constrained. Two of the most common are the Amihud (2002) measure and the Pastor and Stambaugh (2003) volume adjustment. Both of these seek to track the relation between returns and volume, with strong associations indicating illiquidity.

The Amihud (2002) illiquidity measure (A_t) is the absolute value of the return divided by the volume.

$$A_t = \frac{|\Delta y_t|}{N_t} \tag{10}$$

Periods with constant volatility yet declining volumes are therefore observationally equivalent to periods with increasing volatility with constant trading volumes, as long as the ratio is unchanging.

The Pastor and Stambaugh (2003) measure is a parameter λ estimated using an OLS regression of the responsiveness of prior volume changes and return sign reversals on returns.

$$\Delta p_t = \alpha + \lambda sign(\Delta p_{t-1}) \ln N_{t-1} + e_t \tag{11}$$

In this model, λ measures the strength of volume-related saw-toothed volatility in returns. When prices change due to volumes, they are reversed to a greater degree in the following period when λ is both negative and large.

Both of these measures are calculated using a rolling 252 (annual) business day window, giving us daily values for both liquidity measures between 2001 and 2014. The Amihud measure is calculated as the rolling average, and the Pastor and Stambaugh measure is estimated with using an OLS regression for each day, holding constant the trailing one-month average appreciation rate.

These liquidity measures are presented in Figure 9, with higher values indicating less liquidity. The daily house price index facilitates the construction of these two separate measures, neither of which exists in the prior housing finance literature. Both series are highly correlated, providing useful cross-validation. Both indicate the U.S. housing market began the 2000s with relatively high levels of liquidity. As the boom period in house prices reached maturity and fell into decline, liquidity steadily receded until the end of 2009. Overall for housing, both liquidity and frictions measures are correctly signed throughout, a notable result because as Pastor (2016) and Hasbrouck (2016) indicate, the measures for stocks are incorrectly signed in 1/3 and 1/5 of the cases, respectively.

7.3 Volatility and Risk Adjusted Returns

Quarterly and monthly house price indices have traditionally produced remarkably low measures of house price volatility and therefore high measures of risk-adjusted returns. According to Ghysels, Plazzi, Valkanov and Torous (2013), between 1991 and 2010, the national Case-Shiller index at a quarterly frequency has an annualized standard deviation of 5 percent, and the national, monthly index produced by the Federal Housing Finance Agency has a standard deviation of 2 percent, both on capital gains near 3 percent per year.

Household or loan-level measures present a vastly different story, however. In contrast to the 2 percent to 5 percent annual standard deviations estimated using price indices, Case and Shiller (1989) use residuals from repeat-sales regressions to estimate the number to be near 15 percent, and Flavin and Yamashita (2002), using the Panel Study of Income Dynamics (PSID) between 1969 and 1992, produce an estimate of 14 percent.[12] Existing studies, including Calhoun et al. (1995), argue low measured volatility in house price indices is due to temporal aggregation bias. This issue has made house price indices of limited use in measuring house price volatility and risk-adjusted returns for housing. Investors and policymakers are aware of this disconnect between aggregate and disaggregate measures

[12] Case and Shiller (1989) go on to conclude "...there is so much noise in individual houses, the standard deviation of annual price changes is comparable to that on the aggregate stock market" (page 134).

of volatility, and prefer loan-level simulations of house price changes for stress testing, for instance, in the methodology used by the FHFA and described in Smith et al. (2016).

The academic literature, however, frequently measures house price volatility using house price indices at a monthly or quarterly frequency. For instance, Han (2010) and Peng and Thibodeau (2013) use the standard deviation of the change in house price indices as measures of house price risk studies linking house price risk with housing demand. Han (2013) uses the same approach to argue that the risk-return relationship for housing can be negative, in contrast to standard theories which suggest greater returns are correlated with greater risk. While we do not contest the findings of these papers, it is notable that it is still common practice to use variation in an index with temporal aggregation bias as a measure of price volatility.

The issue of risk-adjusted returns is highlighted by estimated Sharpe ratios for housing estimated using house price indices versus other assets. In the monthly index estimated in Section 2, the annualized capital gain is 2.7 percent with a standard deviation of 3.3 percent between 2000 and 2014. If one assumes a risk-free rate of return of 1.9 percent (average 10-year Treasury yield between 2000 and 2014), a depreciation rate of 2.5 percent (Harding, Rosenthal, and Sirmans, 2007), and imputed rental dividend of 6 percent (Poterba and Sinai, 2008), the Sharpe ratio for housing is nearly 1.3. This Sharpe ratio puts houses well above any risk-adjusted returns frontier.

The individual stock with the highest Sharpe ratio with 30 years of data during 1929-2010 is Berkshire Hathaway at 0.76, according to Frazzini, Kabiller and Pedersen (2013) and Frazzini and Pedersen (2014). Non-traded commercial real estate has a Sharpe ratio of 0.47 between 1978 and 2015, using reported data from the National Council of Real Estate Investment Fiduciaries. Between 2000 and 2014, the monthly S&P 500 composite has a Sharpe ratio of 0.18. If one assumes housing has a standard deviation of 14 percent, such as what is observed in the PSID by Flavin and Yamashita (2002), the Sharpe ratio for housing falls to 0.24, in line with stocks and other asset classes. The question becomes, is there a way to measure volatility using a price index, and can a daily house price estimate be part of the measure?

We posit that a daily house price index can be used as an accurate measure of house price volatility. Our daily house price index produces volatility and risk-adjusted return measures in line with those estimated using micro-level datasets. While monthly Sharpe ratios for housing and the S&P 500 are very different, at 1.31 and 0.18, respectively, ratios calculated with unconditional daily data are remarkably similar, at 0.15 and 0.12 between 2000 and 2014 (see Table 1). When the series are modeled using residual standard deviations from the time series models in Section 3, Sharpe ratios fall because some of the variation in the index is explained. However, the indices are still similar, at 0.20 for housing and 0.12 for the S&P 500.

These estimates suggest that at best, daily house price index almost fully accounts for the temporal aggregation bias found in monthly or quarterly indices. At worst, the increased error variance from estimation error and other sources exactly matches the remaining negative bias in the volatility of the index to the point where the indices appear to produce reliable measures. In either case, the volatility measures produced by the standard deviations of both the series itself, and the reduced-form residuals from a well-specified time series model, result in volatilities and risk-adjusted returns in line with those observed in micro data and with other asset classes.

8 Conclusion

In this paper, we have introduced a conceptual framework for the analysis of the microstructure of the housing market. This simple model presents daily house price appreciation as originating from three factors: liquidity effects, bid-ask spreads, and idiosyncratic shocks. Changes from both liquidity effects and bid-ask spreads are negatively correlated, meaning they offset over time. As a consequence, time-aggregated metrics smooth over these effects, giving a false representation of the volatility of house prices.

While price changes arising from bid-ask spreads are zero in expectation, we predict and observe calendar-related price and volume effects. When expected appreciation rates for a particular date or day of the week are non-zero, volumes are higher. These price changes are mean-reverting over the course of the month, a classic characteristic of liquidity-based price changes. Furthermore, we observe negative autocorrelation in excess of that which is predicted by the measurement and composition-related volatility, suggesting the existence of large transaction costs and spreads for housing.

Despite this evidence, we believe there to be other drivers of daily price changes. The magnitude of the observed calendar effects are exceptionally large, suggesting additional contributing factors related to seller and borrower heterogeneity. So, while liquidity and transaction costs play an important role in price dynamics, this series is ripe for further research.

Housing transactions have numerous discrete steps and many agents with potentially differing incentives, many of which could contribute to the observed data generating process. For instance, liquidity constrained borrowers may be willing to pay more to close on dates just after paychecks clear, or first-time homebuyers, who are less experienced with housing transactions, may be led to simultaneously pay more and close on particular dates that are advantageous for the agent. Additionally, while we cannot determine the efficiency of these outcomes because our models are reduced-form, the secondary mortgage market may present arbitrage opportunities were these observed daily price changes to result in unpriced collateral risk.

References

Acharya, V. V. and Pedersen, L. H. (2005). Asset pricing with liquidity risk. *Journal of Financial Economics*, 77(2):375–410.

Adrian, T., Etula, E., and Muir, T. (2014). Financial intermediaries and the cross-section of asset returns. *The Journal of Finance*, 69(6):2557–2596.

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56.

Anenberg, E. and Laufer, S. (2016). A more timely house price index. *Review of Economics and Statistics*.

Calhoun, C. A., Chinloy, P., and Megbolugbe, I. F. (1995). Temporal aggregation and house price index construction. *Journal of Housing Research*, 6(3):419.

Campbell, J. Y., Giglio, S., and Pathak, P. (2011). Forced sales and house prices. *The American Economic Review*, 101(5):2108–2131.

Carrillo, P. E., de Wit, E. R., and 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(3):609–651.

Case, K. E. and Shiller, R. J. (1989). The efficiency of the market for single-family homes. *The American Economic Review*, pages 125–137.

Diebold, F. X. (2006). On market microstructure noise and realized volatility. *Journal of Business and Economic Statistics*, 24:181–183.

Flavin, M. and Yamashita, T. (2002). Owner-occupied housing and the composition of the household portfolio. *The American Economic Review*, 92(1):345–362.

Frazzini, A., Kabiller, D., and Pedersen, L. H. (2013). Buffett's alpha. Working Paper 19681, National Bureau of Economic Research.

Frazzini, A. and Pedersen, L. H. (2014). Betting against beta. *Journal of Financial Economics*, 111(1):1–25.

Ghysels, E., Plazzi, A., Valkanov, R., and Torous, W. (2013). Forecasting real estate prices. *Handbook of Economic Forecasting*, 2:509–580.

Goyenko, R. Y., Holden, C. W., and Trzcinka, C. A. (2009). Do liquidity measures measure liquidity? *Journal of financial Economics*, 92(2):153–181.

Han, L. (2010). The effects of price risk on housing demand: empirical evidence from US markets. *Review of Financial Studies*, 23(11):3889–3928.

Han, L. (2013). Understanding the puzzling risk-return relationship for housing. *Review of Financial Studies*, 26(4):877–928.

Han, L. and Strange, W. C. (2015). Chapter 13 - the microstructure of housing markets: Search, bargaining, and brokerage. In Gilles Duranton, J. V. H. and Strange, W. C., editors, *Handbook of Regional and Urban Economics*, volume 5 of *Handbook of Regional and Urban Economics*, pages 813 – 886. Elsevier.

Harding, J. P., Rosenthal, S. S., and Sirmans, C. F. (2007). Depreciation of housing capital, maintenance, and house price inflation: Estimates from a repeat sales model. *Journal of Urban Economics*, 61(2):193–217.

Harris, B. H. et al. (2013). Tax reform, transaction costs, and metropolitan housing in the united states. *Washington, DC: The Urban Institute*.

Hasbrouck, J. Updated gibbs estimates of trading costs. `http://people.stern.nyu.edu/jhasbrou/Research`. Accessed: 2016-12-01.

Hasbrouck, J. (2009). Trading costs and returns for us equities: Estimating effective costs from daily data. *The Journal of Finance*, 64(3):1445–1477.

Hendel, I., Nevo, A., and Ortalo-Magne, F. (2009). The Relative Performance of Real Estate Marketing Platforms: MLS versus FSBOMadison.com. *American Economic Review*, 99(5):1878–98.

Hou, K., Xue, C., and Zhang, L. (2014). A comparison of new factor models. Working Paper 20682, National Bureau of Economic Research.

Ngai, L. R. and Tenreyro, S. (2014). Hot and cold seasons in the housing market. *American Economic Review*, 104(12):3991–4026.

Novy-Marx, R. (2009). Hot and cold markets. *Real Estate Economics*, 37(1):1–22.

Pastor, L. Liquidity factor data. `http://faculty.chicagobooth.edu/lubos.pastor/research`. Accessed: 2016-12-01.

Pastor, L. and Stambaugh, R. (2003). Liquidity risk and expected stock returns. *Journal of Political Economy*, 111(3):642–685.

Peng, L. and Thibodeau, T. (2013). Risk Segmentation of American Homes: Evidence from Denver. *Real Estate Economics*, 41(3):569–599.

Poterba, J. and Sinai, T. (2008). Tax expenditures for owner-occupied housing: Deductions for property taxes and mortgage interest and the exclusion of imputed rental income. *The American Economic Review*, 98(2):84–89.

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of Finance*, 39(4):1127–1139.

Sharpe, W. F. (1966). Mutual fund performance. *The Journal of Business*, 39(1):119–138.

Smith, S., Fuller, D., Bogin, A., Polkovnichenko, N., and Weiher, J. (2016). Countercyclical capital regime revisited: Tests of robustness. *Journal of Economics and Business*, 84:50–78.

Working, H. (1960). Note on the correlation of first differences of averages in a random chain. *Econometrica: Journal of the Econometric Society*, pages 916–918.

Figure 1: Housing Transactions by Calendar Indicator

A. Month

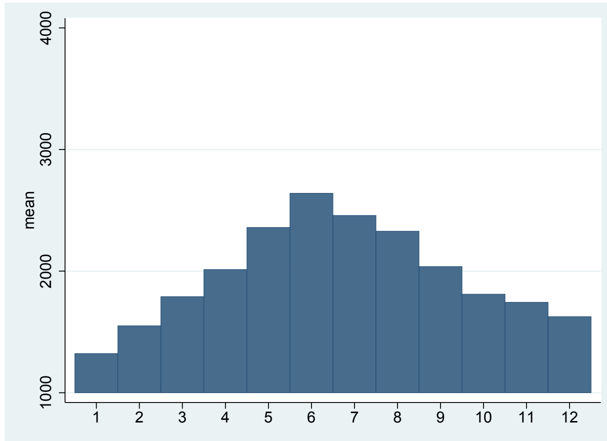

B. Day of Month

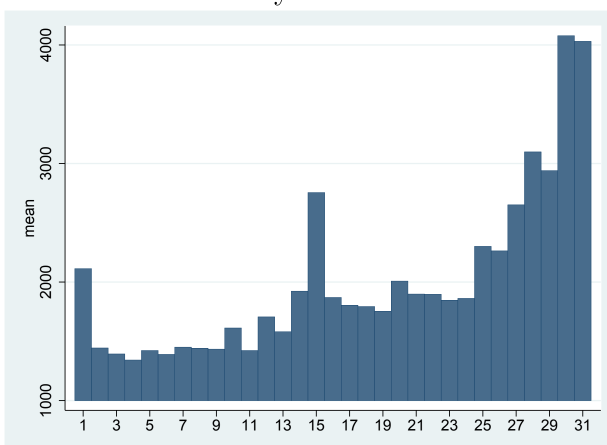

C. Day of Week

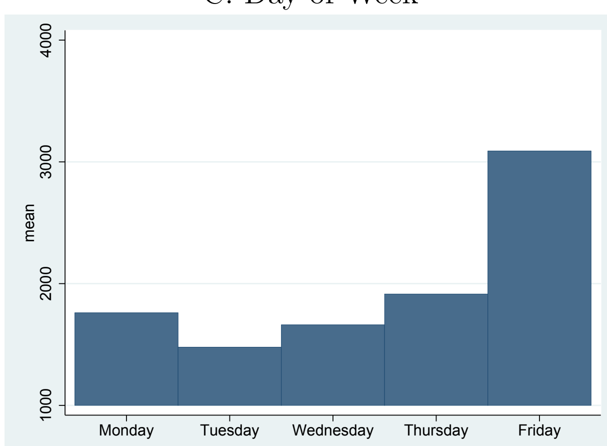

Note: The figures show average housing transactions by calendar indicator between 2000 and 2014. Transaction volumes are highest mid-year, rise throughout each month with spikes on the 1st and 15th, and are highest on Fridays.

Figure 2: Housing and S&P Price Indices - Monthly vs Daily Measures

A. Monthly Price Index

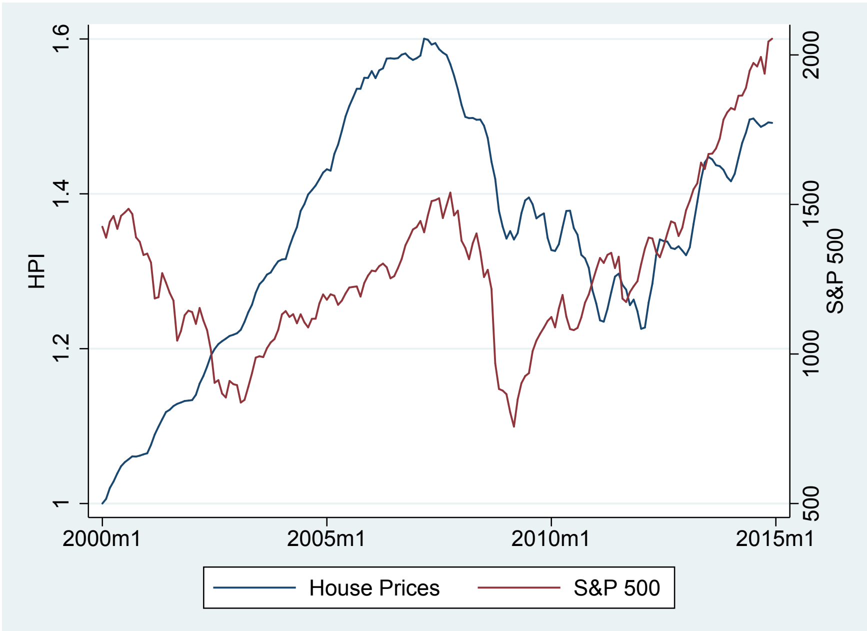

B. Daily Price Index

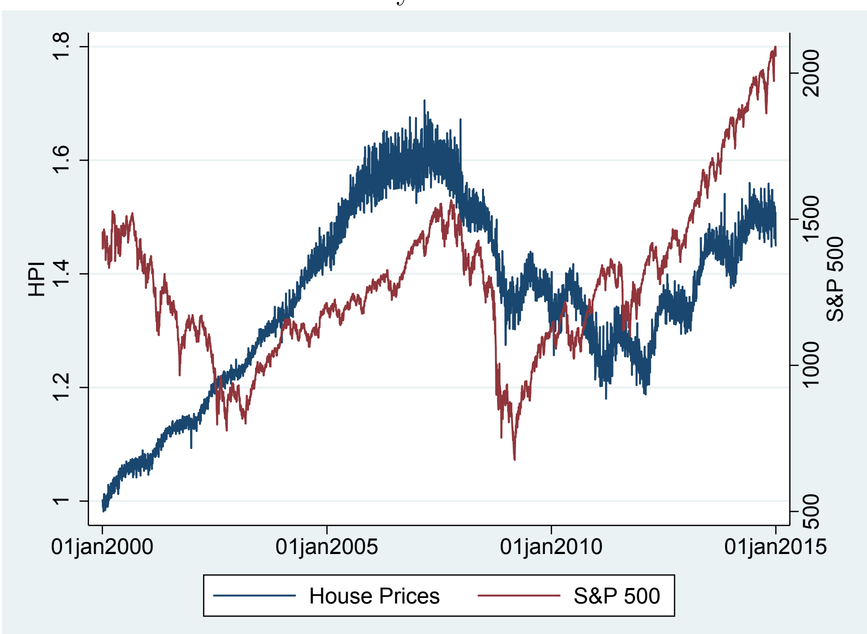

Note: The figures show monthly versus daily house prices for a repeat-sales index against closes for the S&P 500. The house price series rise during 2000 to 2005, while the S&P 500 falls after the dot-com crash in 2000 through 2002. House and stock prices after 2005 appear to be move together. House prices rise to a peak in early 2006 with rising volatility. Stocks reach their peak later, and have the sharp decline during the 2008-2009 financial crisis. The subsequent stock price rise since 2009, virtually doubling between 2008 and 2014. House prices increase by 30 percent from their trough, but with increased volatility.

Figure 3: Housing and S&P Capital Gains - Daily Averages by Calendar Indicator

A. Month

i. Housing | ii. S&P 500

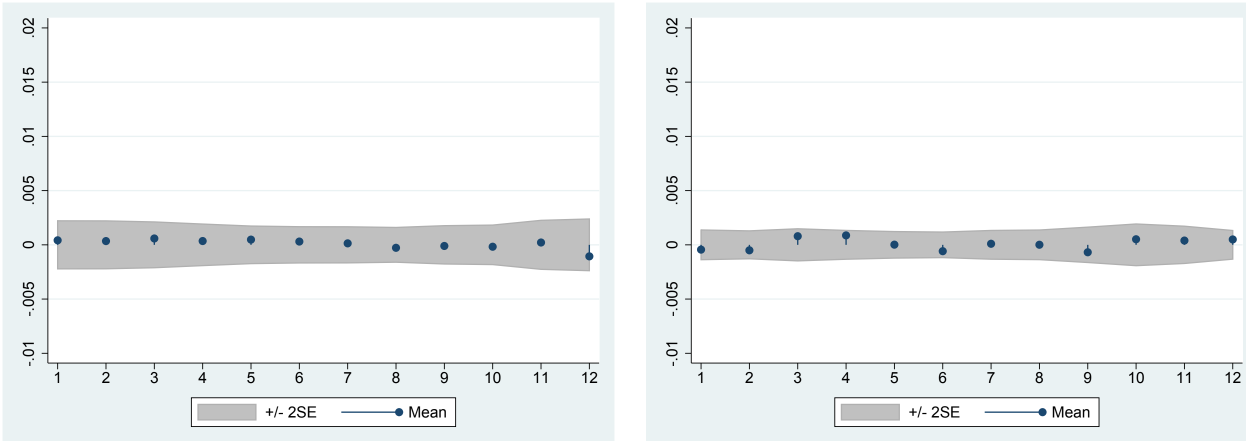

B. Day of Month

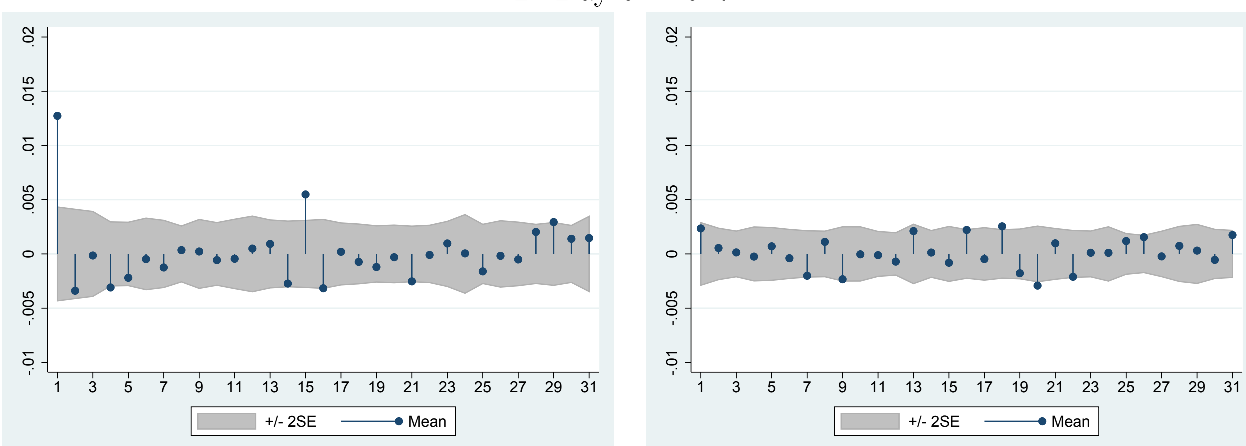

C. Day of Week

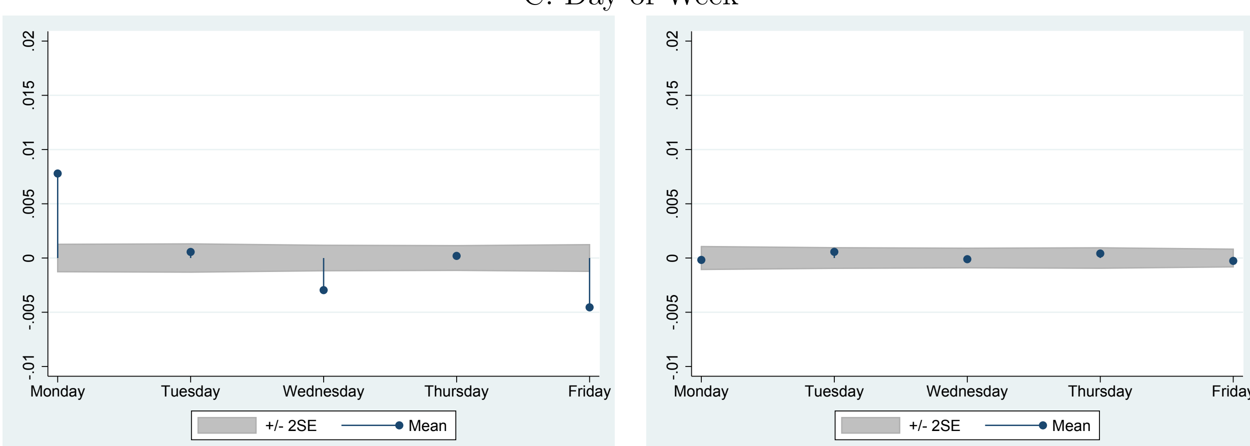

Note: The figures show average capital gains for housing and stocks between 2000 and 2014. Prices are highest for housing on the 1st and the 15th of the month, and on Mondays, and are lower on Fridays. Prices for stocks exhibit no clear differences by indicator.

Figure 4: Mortgage LTVs at Origination by Calendar Indicator

A. Month

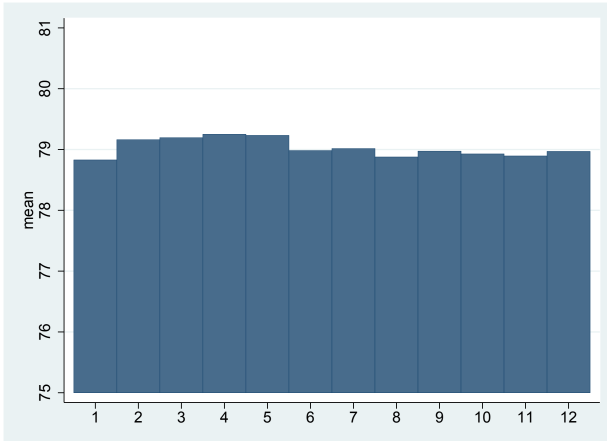

B. Day of Month

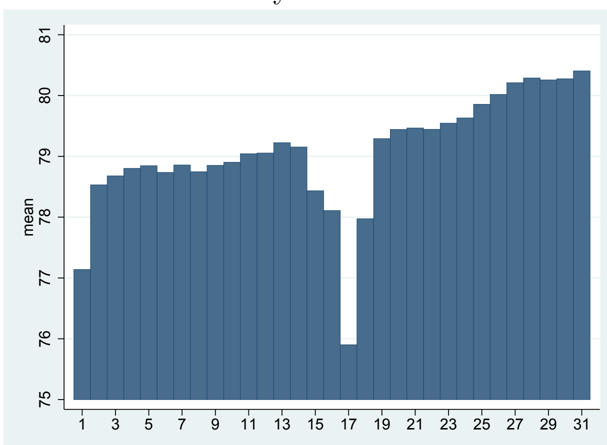

C. Day of Week

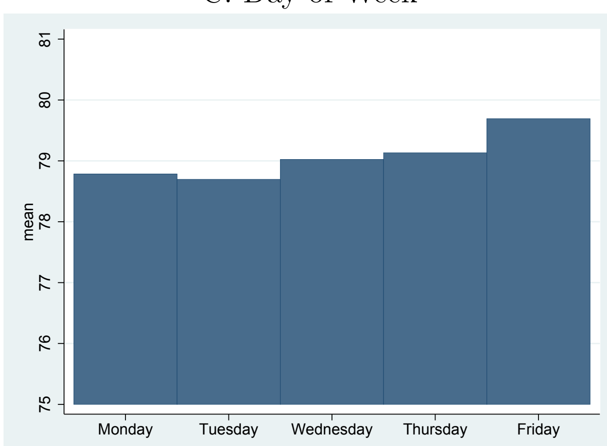

Note: The figures show average loan-to-value (LTV) ratios by calendar indicator between 2000 and 2014. LTVs ratios are roughly constant across months, rise throughout each month with negative spikes on the 1st and 15th, and are highest on Fridays.

Figure 5: Housing and S&P Capital Gains - Autocorrelograms

A. Raw Series

i. Housing

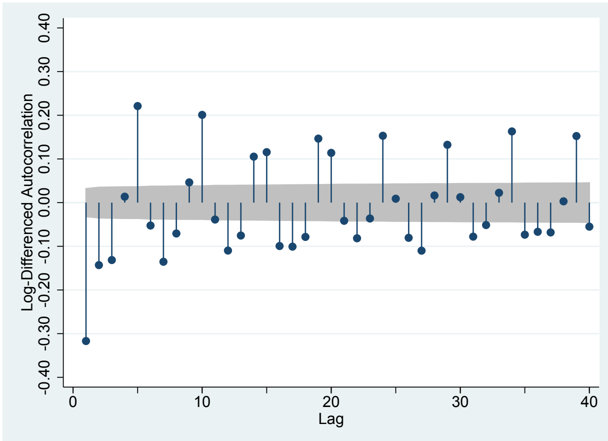

ii. S&P 500

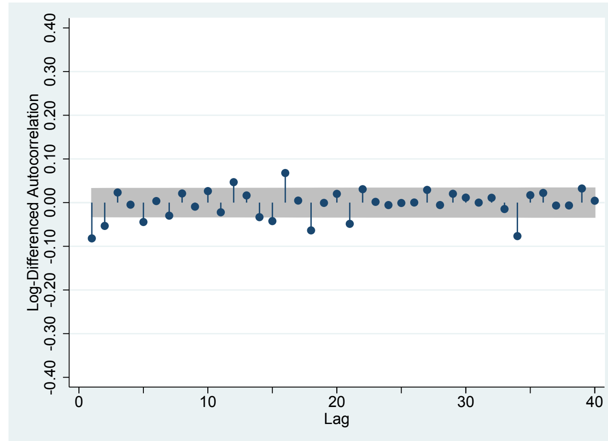

B. Residuals from Time Series Model

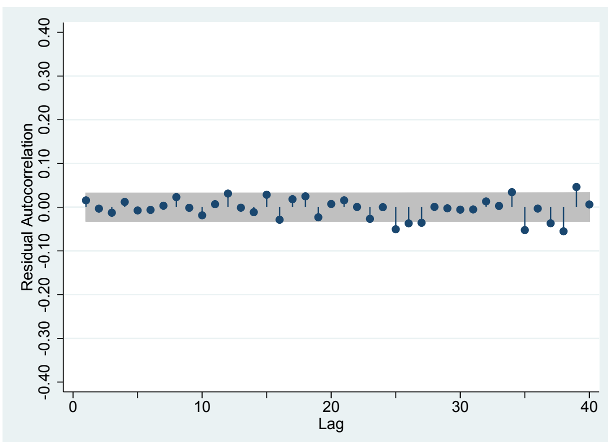

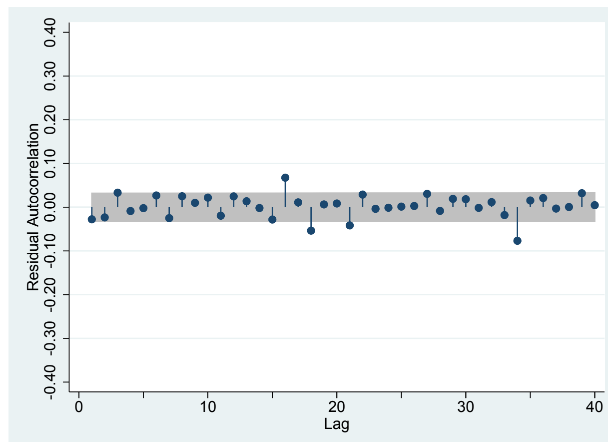

Note: The time series model for housing is presented in column 5 of Table 2 and column 3 of Table 3 for the S&P 500.

Figure 6: Daily Capital Gains to Housing

Daily Capital Gains

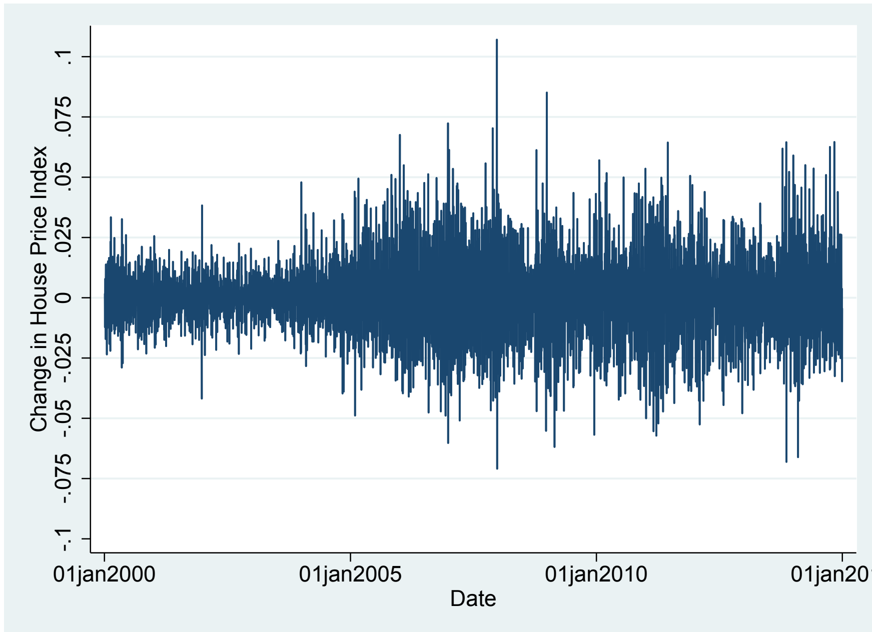

Note: The figure presents the daily log-difference in the daily price index for housing.

Figure 7: Composition Bias

Daily Change in Composition Index

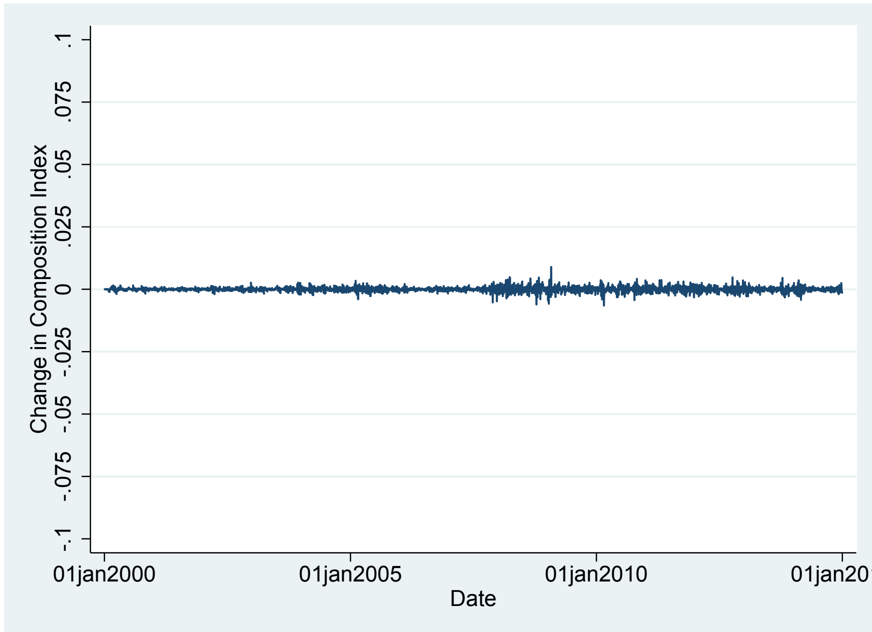

Note: The figure presents the daily composition index, calculated using monthly appreciation rates by state and daily state sample composition. See Equation 7.

Figure 8: Housing and S&P Transaction Cost Measures

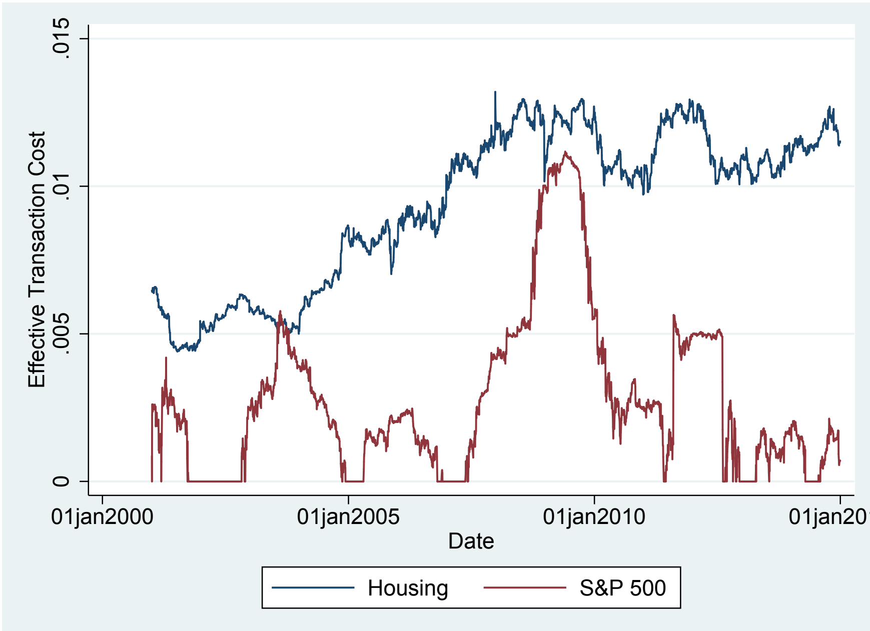

Note: Implied transaction costs are based on Roll (1984) measures, which require the first order autocovariance to be negative. When this is positive, the transaction cost measure is set to zero by convention. The discrete jump in implied transaction costs for the S&P 500 in 2011 is on August 8, 2011, when Standard and Poor's downgraded U.S. sovereign debt from AAA to AA+.

Figure 9: Housing Liquidity Measures

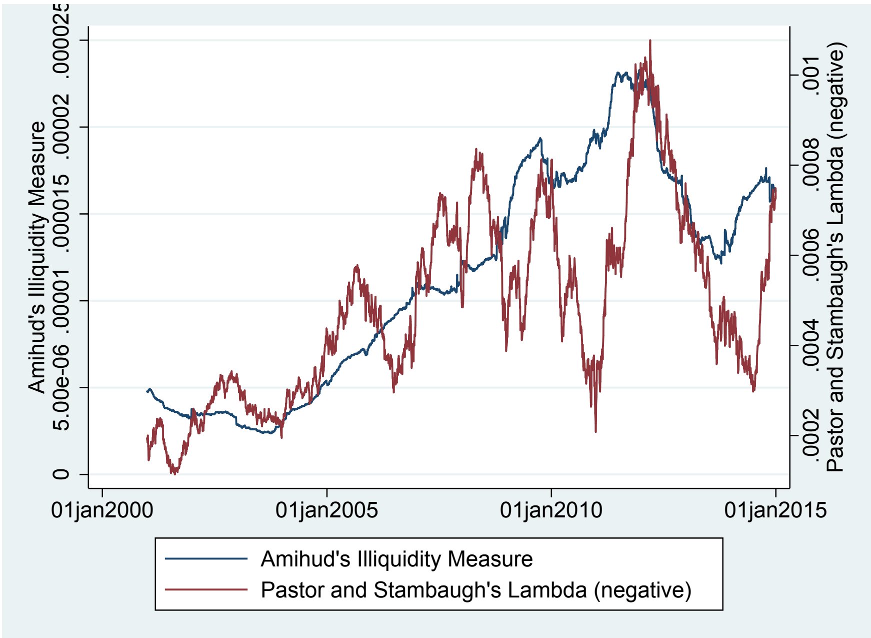

Note: The Amihud (2002) measure is the ratio of absolute returns to trading volume. The Pastor and Stambaugh (2003) λ is is estimated based on a regression model described in the text. Both are constructed based on a trailing 252 trading day (annual) average.

Table 1: Summary Statistics

	Obs per Period	Non-Annualized Mean	Non-Annualized SD	Annualized Mean	Annualized SD	Annualized Sharpe Ratio
Monthly:						
HPI (raw)	42,021	0.22%	0.95%	2.68%	3.31%	1.31
HPI (conditional)	42,021	0.22%	0.56%	2.68%	1.95%	2.22
S&P 500 (raw)	-	0.20%	3.99%	2.45%	13.83%	0.18
S&P 500 (conditional)	-	0.20%	3.71%	2.45%	12.84%	0.19
Daily:						
HPI (raw)	1,909	0.01%	1.73%	2.48%	27.48%	0.15
HPI (conditional)	1,909	0.01%	1.29%	2.48%	20.43%	0.20
S&P 500 (raw)	-	0.01%	1.28%	2.32%	20.38%	0.12
S&P 500 (conditional)	-	0.01%	1.27%	2.32%	20.08%	0.12

Note: The mean and standard deviations are for capital gains. Sharpe ratios include imputed rental income from housing = 6%; Depreciation = 2.5%; S&P dividend yield = 1.91%; Risk-free rate = 1.86%. The conditional standard deviations are calculated using residuals from time series models described in Sections 3 and 5. Annualized values for means are calculated by taking the non-annualized value and multiplying it by the number of periods in a year (12 months and 252 trading days); for standard deviations they are calculated by multiplying the non-annualized value by the square root of the periods in the year.

Table 2: Time Series Model Estimates, House Prices

	Dependent Variable: $\Delta \ln HPI$				
	[1]	[2]	[3]	[4]	[5]
Conditional Mean Equation					
y_{t-1}	-0.359***	0.192***	0.0830*		0.179***
y_{t-2}			-0.0498		-0.00502
y_{t-3}			-0.0419		-0.0174
y_{t-4}			0.0117		0.0424**
y_{t-5}		0.318***	0.168***		0.156***
y_{t-6}			0.00402		-0.00663
y_{t-7}			-0.0250		-0.0446**
y_{t-8}			0.00924		0.00141
y_{t-9}			0.0631***		0.0601***
y_{t-10}			0.179***		0.152***
y_{t-11}			0.0536**		0.0329*
y_{t-12}			0.0111		-0.0178
y_{t-13}			0.0403**		0.0179
y_{t-14}			0.0868***		0.0764***
y_{t-19}			0.167***		0.166***
w_{t-1}		-0.698***	-0.654***		-0.802***
$\bar{y}_{t-1}$		-0.187***	-0.171***		-0.129***
Tuesday				-0.00636***	-0.00169**
Wednesday				-0.00984***	-0.00306***
Thursday				-0.00680***	-0.00141
Friday				-0.0119***	-0.00482***
2nd of Month				-0.00239	-0.00284**
15th of Month				0.00918**	0.00492***
February				-0.000413	0.000893***
March				-0.000235	0.00182***
April				-0.000814	0.00139***
May				1.99e-05	0.00179***
June				-0.000648	0.00158***
1st Business Day				0.0186***	0.0117***
Last Business Day				0.00530**	0.00118
Constant	0.000240*	0.00101***	0.000789***	0.00410	0.000971
R-squared	0.100	0.354	0.429	0.103	0.478
Observations	3,772	3,771	3,771	3,772	3,771
Conditional Variance Equation					
w^2_{t-1}	0.0405***	0.0389***	0.0437***		0.0438***
s^2_{t-1}	0.955***	0.954***	0.948***		0.946***
ln Transactions	-1.449***	-1.412***	-1.397***		-1.356***
Constant	-3.115	-3.291**	-3.278**		-3.534**
Observations	3,772	3,771	3,771	3,772	3,771

Note: ***, **, and * denote significance at the 1%, 5%, and 10% level of significance. The dependent variable is the log-difference of house prices by trading day. Months, days of month, and business days that are not shown are estimated but not statistically different than zero in any model in which they are considered, and thus are omitted for brevity.

Table 3: Time Series Model Estimates, S&P 500

	Dependent Variable: Δ ln S&P 500 Close				
	[1]	[2]	[3]	[4]	[5]
Conditional Mean Equation					
y_{t-1}	-0.0556***	0.419**	0.875***		0.298
y_{t-2}			0.0238		-0.00861
y_{t-3}			0.0183		0.00364
y_{t-4}			0.00612		0.00420
y_{t-5}		-0.0401**	-0.0366		-0.0399**
y_{t-6}			0.00534		-0.00887
y_{t-7}			0.0228		0.00701
y_{t-8}			0.00378		8.55e-05
y_{t-9}			-0.0158		-0.0157
y_{t-10}			0.0267		0.0136
y_{t-11}			-0.00138		0.00427
y_{t-12}			0.0131		0.0231
y_{t-13}			-0.0157		-0.00616
y_{t-14}			-0.0129		-0.0279
y_{t-19}			-0.00313		-0.00692
w_{t-1}		-0.796***	-0.931***		-0.356
$\bar{y}_{t-1}$		-0.474**	-0.00352		-0.00475
Tuesday				0.000786	-4.43e-06
Wednesday				0.000100	7.53e-05
Thursday				0.000628	-0.000139
Friday				3.61e-05	-0.000315
2nd of Month				-0.00127	-0.000770
15th of Month				-0.00243	-0.00284
February				-1.30e-05	-0.000231
March				0.00125	0.000430
April				0.00137	0.000168
May				0.000468	7.70e-06
June				-6.59e-05	-0.000487
1st Business Day				0.000637	0.00133
Last Business Day				-0.00147	-0.00147
Constant	0.000489***	0.000542***	0.000545***	0.000717	0.00218
R-squared	0.006	0.010	0.011	0.014	0.021
Conditional Variance Equation					
w^2_{t-1}	0.0924***	0.0914***	0.0915***		0.0986***
s^2_{t-1}	0.896***	0.897***	0.897***		0.891***
Constant	1.65e-06***	1.62e-06***	1.62e-06***		1.56e-06***
Observations	3,772	3,771	3,771	3,772	3,771

Note: ***, **, and * denote significance at the 1%, 5%, and 10% level of significance. The dependent variable is the log-difference of house prices by trading day. Months, days of month, and business days that are not shown are estimated but not statistically different than zero in any model in which they are considered, and thus are omitted for brevity.

Table 4: Volatility Decomposition

	Housing		S&P 500	
	Daily	Monthly	Daily	Monthly
Standard Deviation	1.731%	0.954%	1.284%	3.993%
Decomposition:				
Composition Volatility	0.099%	-	-	-
Estimation Error	0.065%	0.007%	-	-
Calendar Liquidity Effects	0.483%	-	0.009%	-
Transaction Costs	0.423%	-	0.004%	-
Idiosyncratic Shocks	0.661%	0.947%	1.270%	3.993%

Note: The reported standard deviation is the raw series. Composition volatility is measured as the standard deviation of the composition index described in Equation 15. Estimation error is measured as $\sqrt{(N_t + N_{t-1})/(N_t N_{t-1})}$, where N_t is the the number of observations in period t. Calendar liquidity effects are calculated as the standard deviation of the raw series minus the standard deviation of residuals from an OLS regression of returns on calendar fixed effects—model 4 in Tables 2 and 3. Transaction costs are estimated as the standard deviation of the raw series minus the standard deviation of the residuals from ARMA-GARCH models—model 3 in Table 2 for housing, and model 1 in Table 3 for the S&P 500—estimated separately for each year and then averaged. Transaction costs are set to zero in periods where the autocovariance of the series is positive, following Hasbrouck (2009).

A Appendix 1: Repeat Sales Method

We begin by expressing the value of a house i at time t as a vector of characteristics X with time-invariant relative prices β, and a vector of price price levels δ multiplied by a vector of dummy variables D set to 1 in the relevant period and 0 in all others.

$$p_{it} = X_{it}'\beta + D_t'\delta + e_{it} \tag{12}$$

Under the assumption X_i does not vary over t, the difference between prices in period t and τ, defined as $dp_{t,\tau} \equiv p_t - p_\tau$, is

$$dp_{it\tau} = D_{t\tau}'\delta + e_{it\tau} \tag{13}$$

where $D_{t\tau} \equiv D_t - D_\tau$, and $e_{t\tau} \equiv e_t - e_\tau$. It is assumed that house price growth rates have an additional variance term related to the holding period $t-\tau$, with estimated residuals from equation 13 parameterized as follows

$$\hat{e}_{it\tau}^2 = \alpha_1 + \alpha_2(t-\tau) + \alpha_3(t-\tau)^2 + v_{it\tau} \tag{14}$$

These two expressions lead to the standard generalized least squares estimator,

$$\frac{dp_{it\tau}}{\sqrt{\tilde{e}_{it\tau}^2}} = \left(\frac{D_{t\tau}}{\sqrt{\tilde{e}_{it\tau}^2}}\right)'\delta + \frac{e_{it\tau}}{\sqrt{\tilde{e}_{it\tau}^2}} \tag{15}$$

where $\tilde{e}_{it\tau}^2$ is the estimated variance of the estimated 1st-stage residuals and $\hat{\delta}$ is the final estimated (log) price level. The final price index is calculated as $\hat{I}_t = \exp\hat{\delta}_t$ and normalized so that the first day of the index in 2000 is set to 1.

B Appendix 2: Further Composition Robustness

The composition-related volatility estimates in Section 6 use state-level variation in monthly prices and daily composition. There may be other sources of composition-related variation in prices, including variation in appreciation rates within states (cities vs suburbs vs rural areas) or price tiers (high vs low-priced units). Insofar as these different categories have different appreciation rates and are also sampled at different frequencies, composition-related volatility may exist.

The difficulty with extending the prior state-level approach to other categories is that it requires estimation of a monthly index by area. For most areas within states, this is infeasible. Furthermore, due to conforming loan limits, price tiers are unreliable as the sample is censored on the price dimension by an unknown amount. We therefore pursue a falsification robustness test as opposed to other direct measures of composition-related volatility.

Our strategy is to take each pair of transactions in the house price index construction section and use the ZIP code price indices developed by Bogin, Doerner, and Larson (2016) to predict the transaction price for the second sale. The difference between the actual and predicted sale price is then estimated as a function of the vector of calendar effects in the empirical sections in the main text. These calendar effects should be of similar sign and significance level to the estimates in Table 2.

This approach controls for a variety of differences in composition, including price tiers, borrower characteristics, schools, environmental quality, industry composition, and other factors related to proximity, assuming that each attribute is representative at a ZIP code level of geography. By conditioning appreciation on a local measure such as a ZIP code index, any residual variation may be reliably interpreted as a calendar effect that is untainted by composition. Of the 7.3 million transactions between 2000 and 2014, 4.5 million are the second or later observed sale of a particular housing unit. We use this sample to test for the existence of calendar effects.

Table 5 presents three columns. Columns 1 and 2 are from the "conditional mean" portions of columns 4 and 5 in Table 2. Column 3 presents estimates of 2nd-sale residuals from ZIP code-level house price index predictions regressed on the vector of calendar effects from columns 1 and 2. Estimates from this model are consistent with estimates in columns 1 and 2, indicating that this alternative approach also measures large and significant calendar effects on house prices. These include day-of-week effects that are largest on Monday, effects on 15th of the month, and the first and last business day of the month.

Overall, this model suggests that conditional on ZIP code and time period-related composition, there is still substantial calendar-related variation in the observed price changes of identical units. Moreover, this variation is consistent with the variation observed in the daily house price index. This approach therefore supports the findings in the main text of the paper.

Table 5: Time Series Model Estimates, Robustness

	[1] Daily HPI Fixed Effects Model	[2] Daily HPI ARMA-GARCH w/FEs Model	[3] ZIP code Residual Model
Tuesday	-0.00636***	-0.00169**	-0.00188***
Wednesday	-0.00984***	-0.00306***	-0.00229***
Thursday	-0.00680***	-0.00141	-0.00176***
Friday	-0.0119***	-0.00482***	-0.00269***
2nd of Month	-0.00239	-0.00284**	-0.000271
15th of Month	0.00918**	0.00492***	0.00347***
February	-0.000413	0.000893***	0.00208***
March	-0.000235	0.00182***	0.0112***
April	-0.000814	0.00139***	0.0222***
May	1.99e-05	0.00179***	0.0335***
June	-0.000648	0.00158***	0.0417***
1st Business Day	0.0186***	0.0117***	0.0396**
Last Business Day	0.00530**	0.00118	0.0610***
R-squared	0.10	0.48	0.007
Observations	3,772	3,771	4,530,119

Note: ***, **, and * denote significance at the 1%, 5%, and 10% level of significance. Columns 1 and 2 are identical to the conditional mean portion of columns 4 and 5 in Table 2. The dependent variable in these models is the log-difference of house prices by trading day. Months, days of month, and business days that are not shown are estimated but not statistically different than zero in any model in which they are considered, and thus are omitted for brevity. Column 3 gives estimates of residuals from 2nd-sale predictions based on ZIP code house price indices.