The Causes and Consequences of House Price Momentum

Adam M. Guren*

Boston University

July 7, 2015

First Version: November 10, 2013

Abstract

House price changes are positively autocorrelated over two to three years, a phenomenon known as momentum. This paper introduces, empirically grounds, and quantitatively analyzes an amplification mechanism that can generate substantial momentum from small frictions and demonstrates that the resulting momentum helps explain the short-run dynamics of housing markets. The amplification is due to a concave demand curve in relative price, which implies that increasing the quality-adjusted list price of a house priced above the market average rapidly reduces its probability of sale, but cutting the price of a below-average priced home only slightly improves its chance of selling. This creates a strategic complementarity that incentivizes sellers to set their list price close to others'. Consequently, frictions that cause slight insensitivities to changes in fundamentals lead to prolonged adjustments because sellers gradually adjust their price to stay near the average. I provide new micro empirical evidence for the concavity of demand—which is often used in macro models with strategic complementarities—by instrumenting a house's relative list price with a proxy for the seller's equity. I find significant concavity, which I embed in an equilibrium housing search model. Quantitatively, the model can explain the momentum observed empirically with a small fraction of rule-of-thumb sellers. Strong house price momentum leads households to re-time their purchase or sale, thereby explaining the excess volatility of inventory and the strong negative relationship between price changes and inventory levels.

*Email: guren@bu.edu. I would like to thank John Campbell, Raj Chetty, Emmanuel Farhi, Edward Glaeser, Nathaniel Hendren, and Lawrence Katz for their advice and guidance. Gary Chamberlain, John Friedman, Simon Gilchrist, James Hamilton, Max Kasy, Alisdair McKay, Jonathan Parker, Alp Simsek, Andrei Shleifer, and Lawrence Summers also provided thoughtful discussions and comments. Nikhil Agarwal, Rebecca Diamond, Will Diamond, Peter Ganong, Ben Hebert, Michal Kolesar, Tim McQuade, Pascal Noel, Mikkel Plagborg-Moller, Ben Schoefer, and seminar participants at the Booth School of Business, Boston University, Brown, Duke, the Federal Reserve Bank of Boston, the Federal Reserve Bank of New York, the Federal Reserve Board of Governors, the Haas School of Business, Kellogg School of Management, the NYU Search Theory Workshop, the Stern School of Business, UC Berkeley, and the Wharton School provided numerous discussions and suggestions. Andrew Cramond at DataQuick, Mike Simonsen at Altos Research, T.J. Doyle at the National Association of Realtors, and Mark Fleming, Sam Khater, and Kathryn Dobbyn at CoreLogic assisted me in understanding their data. Shelby Lin and Christopher Palmer provided assistance with the DataQuick data.

1 Introduction

A puzzling and prominent feature of housing markets is that aggregate price changes are highly positively autocorrelated, with a one percent annual price change correlated with a 0.30 to 0.75 percent change in the subsequent year (Case and Shiller, 1989).[1] This price momentum lasts for two to three years before prices mean revert, a time horizon far greater than most other asset markets. Substantial momentum is surprising because predictable price changes should be arbitraged away by households, either by altering their bidding and bargaining or by re-timing their purchase or sale, and because most pricing frictions dissipate quickly. Rational models with frictions have not been able to account for momentum quantitatively, and most behavioral models require large departures from rationality.

This paper introduces, empirically grounds, and quantitatively analyzes an amplification mechanism that can generate substantial momentum from small frictions. The mechanism relies on a strategic complementarity among list-price-setting sellers that makes the optimal list price for a house depend positively on the prices set by others. Strategic complementarities of this sort are frequently used in macroeconomic models, but there is limited empirical evidence of their importance or strength. In analyzing momentum in the housing market, I provide micro empirical evidence for a prevalent strategic complementarity in the macroeconomics literature and, using a calibrated equilibrium search model, demonstrate that its ability to amplify underlying frictions is quantitatively significant.

I also show that momentum has important consequences that help explain several perplexing features of the dynamics of housing markets relating to sales and inventory in addition to price. These dynamics, which are analogous to several features of business cycles, matter for the macroeconomy because housing markets affect household balance sheets, the financial system, and business cycles and are a potential channel for monetary policy. House price momentum may also explain why recoveries from housing-triggered cycles are slow.

The propagation mechanism I introduce relies on two components: costly search and a demand curve that is concave in relative price. Search is inherent to housing because no two houses are alike and idiosyncratic taste can only be learned through costly inspection. Search and idiosyncratic taste also limit arbitrage by creating endogenous transaction costs and by making the market price for a house difficult to ascertain. Concave demand in relative price implies that the probability a house sells is more sensitive to list price for houses priced above the market average than below the market average. While concave demand may arise in housing markets for several reasons, I focus on the manner in which asking prices direct buyer search. The intuition is summarized by an advice column for sellers: "Put yourself in the shoes of buyers who are scanning the real estate ads...trying to decide which houses to visit in person. If your house is overpriced, that will be an immediate turnoff. The buyer will probably clue in pretty quickly to the fact that other houses look like better bargains and move on."[2] In other words, the probability that a house is visited by

[1]See also Cutler et al. (1991), Head et al. (2014), Glaeser et al. (2014), and Titman et al. (2014).

[2]"Settling On The Right List Price for Your House," Ilona Bray, http://www.nolo.com/legal-encyclopedia/listing-

buyers decreases rapidly as a home's list price rises relative to the market average. This generates a concave demand curve in relative price because at high relative prices buyers are on the margin of looking and purchasing, while at low relative prices they are mostly on the margin of purchasing.

Concave demand incentivizes list-price-setting sellers—who have market power due to search frictions—to set their list prices close to the mean. Intuitively, raising a house's relative list price reduces the probability of sale and profit dramatically, while lowering its relative price increases the probability of sale slightly and leaves money on the table. Modest frictions that generate initial insensitivities of prices to changes in fundamentals cause protracted price adjustments because sellers find it optimal to gradually adjust their price so that they do not stray too far from the market average.

To evaluate the concavity of the effect of unilaterally changing a house's relative quality-adjusted price on its probability of sale, I turn to micro data on listings for the San Francisco Bay, Los Angeles, and San Diego metropolitan areas from 2008 to 2013. I address bias caused by unobserved quality by instrumenting relative list price with the amount of aggregate price appreciation since the seller purchased. The identification strategy takes advantage of the fact that sellers with low appreciation since purchase set higher list prices because the equity they extract from the sale of their current home constrains their ability to make a down payment on their next home (Stein, 1995; Genesove and Mayer, 1997). Because I compare listings within a ZIP code and quarter, this supply-side variation identifies the curvature of demand if unobserved quality is independent of when a seller purchased their home. The instrumental variable estimates reveal a concave relationship that is statistically and economically significant. My findings about the concavity of demand are robust to other sources of relative price variation that are independent of appreciation since purchase.

To assess the strength of this propagation mechanism, I embed concave demand in a Diamond-Mortensen-Pissarides equilibrium search model. While the amplification mechanism would interact with several sources of price insensitivity such as infrequent price adjustment or information frictions, I focus on a particular source of initial insensitivity: a small fraction of backward-looking rule-of-thumb sellers as in Campbell and Mankiw (1989) and Gali and Gertler (1999). Backward-looking expectations are frequently discussed as a potential cause of momentum (*e.g.*, Case and Shiller, 1987; Case et al. 2012), but some observers have voiced skepticism about widespread non-rationality in housing markets given the financial importance of housing transactions for most households. With a strategic complementarity, far fewer backward-looking sellers are needed to explain momentum because the majority of forward-looking sellers adjust their prices gradually so they do not deviate too much from the backward-looking sellers (Haltiwanger and Waldman, 1989; Fehr and Tyran, 2005). This, in turn, causes the backward-looking sellers to observe more gradual price growth and change their price by less, creating a two-way feedback that amplifies momentum.

I calibrate the parameters of the model that control the shape of the demand curve to match the micro empirical estimates and the remainder of the model to match steady state and time series moments. The calibrated model generates substantial amplification of the underlying frictions: the

house-what-price-should-set-32336-2.html.

model generates three years of positively autocorrelated price changes as observed empirically if 37.5 percent of sellers are backward-looking. By contrast, without concave demand, 73 percent of sellers would have to be backward-looking to generate a three-year response. Furthermore, the loss backward-looking sellers experience due to their failure to optimize is quite small (a mean of 0.74 percent of the sale price).

The amplification mechanism adapts two ideas from the macro literature on goods price stickiness to frictional asset search. First, the concave demand curve is similar to "kinked" demand curves (Stiglitz, 1979) which, since the pioneering work of Ball and Romer (1990) has been frequently cited as a potential source of real rigidities. In particular, a "smoothed-out kink" extension of Dixit-Stiglitz preferences proposed by Kimball (1995) is frequently used to tractably introduce real rigidities through strategic complementarity in price setting. Second, the repeated partial price adjustment caused by the strategic complementarity is akin to Taylor's (1980) "contract multiplier." A lively literature has debated the importance of strategic complementarities and kinked demand in particular for propagating goods price stickiness by analyzing calibrated models, by assessing whether the ramifications of strategic complementarities are borne out in micro data (Klenow and Willis, 2006), and by examining exchange-rate pass through for imported goods (*e.g.,* Gopinath and Itshoki, 2010; Nakamura and Zerom, 2010). My analysis of housing markets adds to this literature by directly estimating a concave demand curve and assessing its ability to amplify frictions in a calibrated model.

Having established a propagation mechanism for house price momentum empirically and theoretically, I show that momentum can help explain two puzzling features of housing markets: the excess volatility of inventory and the strong negative relationship between price changes and inventory levels. The volatility of houses for sale is difficult to reconcile with most calibrations of housing search models in a direct analogue to Shimer's (2005) unemployment volatility puzzle for labor search models, while the negative "housing Phillips curve" relationship between price changes and inventory levels is surprising because in most asset pricing models, price changes are correlated with changes in fundamentals such as inventory (Caplin and Leahy, 2011). Momentum explains both features because forward-looking buyers and sellers re-time their purchase decisions due to expectations of predictable future price changes, creating sudden swings in inventory as prices adjust gradually. For instance, at a trough, marginal buyers rush to purchase before prices rise, while marginal sellers wait to obtain a better price for their home, leading inventory to plummet.[3]

The remainder of the paper proceeds as follows. Section 2 introduces facts about housing dynamics and describes the state of the literature that explains these dynamics. Section 3 analyzes micro data to assess whether housing demand curves are concave. Section 4 presents the model.

[3]Buyer and seller quotes in newspapers provide suggestive evidence of such re-timing. In 2013, when prices were rising, a buyer explained to the *Wall Street Journal* "if you don't get in now, things are going to skyrocket over the next year," while a seller who delayed putting their house on the market told the *Journal* that "the extra money—that was worth [waiting] for the year." This effect is part of the folk wisdom of housing markets, yet has not appeared in the academic literature. For instance, Calculated Risk Blog describes a conversation with a real estate agent who argues that "In a market with falling prices, sellers rush to list their homes, and inventory increases. But if sellers think prices have bottomed, then they believe they can be patient, and inventory declines."

Section 5 calibrates the model to the micro estimates and assesses the degree to which strategic complementarities amplify momentum. Section 6 relates momentum to the excess volatility of inventory and housing Phillps curve. Section 7 concludes.

2 Three Facts About Housing Dynamics

This section presents three facts about housing dynamics—price momentum, the excess volatility of inventory, and the "housing Phillips curve"—using price indices from CoreLogic, sales data from the National Association of Realtors, and data on the stock of houses for sale from the United States Census, all from 1976 to 2013. The robustness of these facts is shown in Appendix B.

2.1 Momentum

Since the pioneering work of Case and Shiller (1989), price momentum has been considered one of the most puzzling features of housing markets. While other financial markets exhibit momentum, the housing market is unusual for the strength of the effect and the horizon over which it persists.[4]

Fact 1: *Price changes are serially correlated for 8 to 14 quarterly lags.*

House price momentum has consistently been found across cities, countries, time periods, and price index measurement methodologies (Cho, 1996; Titman et al., 2014). Figure 1 shows two nation-wide measures of momentum for the CoreLogic repeat-sales house price index for 1976 to 2013 and a third measure for the same index across 103 cities.[5] Panel A shows that autocorrelations are positive for 11 quarterly lags of the quarterly change in the price index adjusted for inflation and seasonality. Panel B shows an impulse response in log levels to an initial one percent price shock estimated from an AR(5). In response to the shock, prices gradually rise for two to three years before mean reverting. Finally, panel C shows a histogram of AR(1) coefficients estimated separately for 103 metropolitan area repeat-sales house price indices from CoreLogic using a regression of the annual change in log price on a one-year lag of itself as in Case and Shiller (1989):

$$\Delta_{t,t-4} \ln p = \beta_0 + \beta_1 \Delta_{t-4,t-8} \ln p + \varepsilon. \tag{1}$$

β_1 is positive for all 103 cities, and the median city has an annual AR1 coefficient of 0.60. Appendix B replicates these facts for a number of countries, price series, and measures of autocorrelation and consistently finds two to three years of momentum.[6] Appendix B also shows that there is

[4]Note that the "momentum" I analyze refers to autocorrelation in aggregate price time series net of dividends, which is distinct from the short-term over-performance of stocks that recently performed best that is also called "momentum." Time-series momentum holds for a number of other asset classes but over over shorter horizons than in the housing market (Cutler et al., 1991; Moskowitz et al. 2012).

[5]As discussed in Appendix B, price indices that measure the median price of transacted homes display momentum over roughly two years as opposed to three years for repeat-sales indices.

[6]House prices appear to be sticky but not their rate of change, unlike goods prices (Fuhrer, 2011). In particular, neither the evidence presented here nor the structural panel VAR in Head et al. (2014) shows evidence of high autocorrelations or delayed "hump shaped" impulse responses for house price changes.

Figure 1: Momentum in Housing Prices

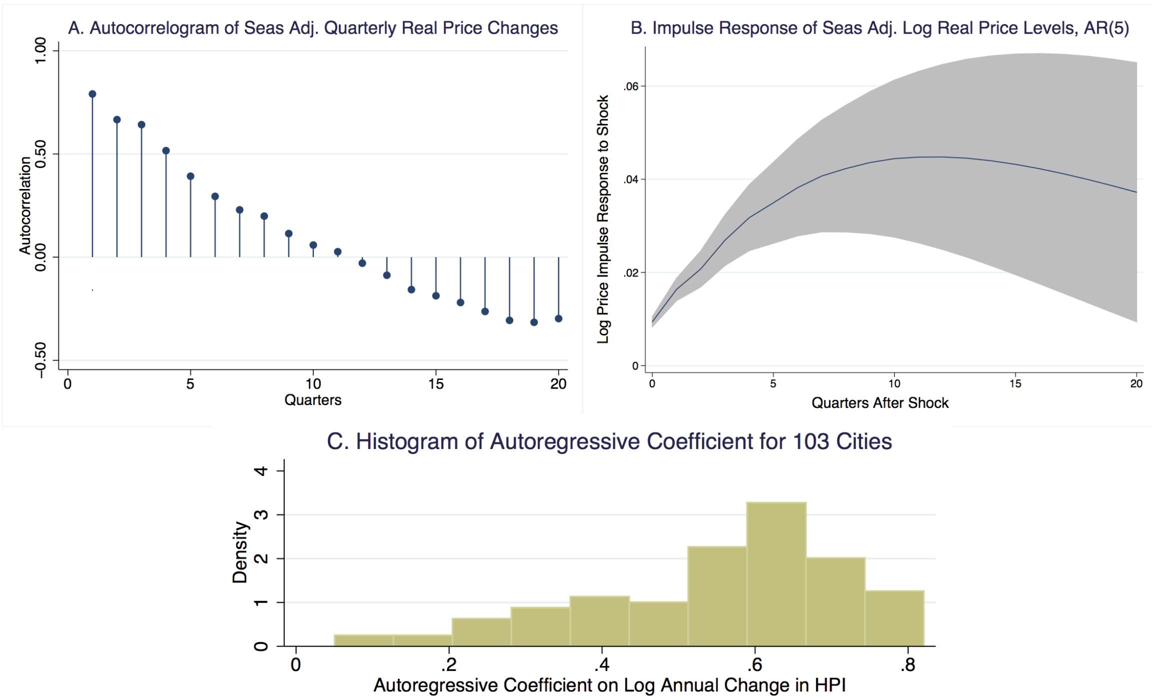

Notes: Panel A and B show the autocorrelation function for quarterly real price changes and an impulse response of log real price levels estimated from an AR(5) model, respectively. The IRF has 95% confidence intervals shown in grey. An AR(5) was chosen using a number of lag selection criteria, and the results are robust to altering the number of lags. Both are estimated using the CoreLogic national repeat-sales house price index from 1976-2013 collapsed to a quarterly level, adjusted for inflation using the CPI, and seasonally adjusted. Panel C shows a histogram of annual AR(1) coefficients of annual house price changes as in regression (1) estimated separately on 103 CBSA division repeat-sales house price indices provided by CoreLogic. The local HPIs are adjusted for inflation using the CPI. The 103 CBSAs and their time coverage, which ranges from 1976-2013 to 1995-2013, are listed in Appendix A.

no evidence of an asymmetry in falling markets relative to rising markets, although the test for asymmetry has limited power.

The existing evidence suggests that momentum cannot be explained by serially correlated changes in fundamentals. Case and Shiller (1989) argue that momentum cannot be explained by autocorrelation in interest rates, rents, or taxes. Glaeser et al. (2014) estimate a dynamic spatial equilibrium model and find that "there is no reasonable parameter set" consistent with short-run momentum. Capozza et al. (2004) find significant momentum after accounting for six comprehensive measures of fundamentals in a vector error correction model.

Five main explanations have been offered for momentum in asset markets and for the housing market more specifically. First, a behavioral finance literature hypothesizes that investors initially underreact to news due to behavioral biases (Barberis et al., 1998, Hong and Stein, 1999) or loss aversion (Frazzini, 2006) and then "chase returns" due to extrapolative expectations about price appreciation. Both extrapolative expectations and loss aversion are considered to be important forces in the housing market (Case and Shiller, 1987; Berkovec and Goodman, 1996; Glaeser et al., 2014; Genesove and Mayer, 2001). These models have not been calibrated quantitatively to the housing market.

Second, Glaeser and Nathanson (2015) provide a behavioral theory in which momentum arises from agents using a behavioral approximation whereby they neglect to account for the fact that previous buyers were learning from prices and instead take past prices as direct measures of demand. In doing so, they advocate for theories based on "small irrationality from the many rather than major irrationality from the few," such as the first class of explanations. As I describe below, my model is one of *small* irrationality from the few. Given the lack of evidence on heterogeneity in non-rationality, it is difficult to empirically distinguish Glaeser and Nathanson's approach, which can be thought of as an alternate micro foundation for my backward-looking sellers, from mine. I do, however, provide corroborating micro evidence for concave demand.

Third, Anenberg (2014) shows that gradual learning about market conditions by sellers can create momentum. Anenberg's structural model of learning can only explain an annual AR(1) coefficient of 0.124 relative to between 0.3 and 0.75 in the data. This reflects the fact that learning happens reasonably quickly in housing markets absent a strategic complementarity. The strategic complementary presented here would complement and amplifies learning-based explanations for momentum, as discussed below.

Fourth, Burnside et al. (2015) show that momentum could result from a gradual spread of optimism if sentiment drives house prices rather than fundamentals. It is difficult to quantitatively measure the speed of the spread of sentiment, but the strategic complementarity presented here complements sentiment-based explanations by reducing the degree of persistence in sentiment necessary to explain momentum.

Finally, in work most closely related to this paper, Head et al. (2014) demonstrate that search frictions, short sale constraints, housing with consumption value, and a gradual construction response can together cause the liquidity of houses to adjust slowly in response to a shock to local incomes, which creates momentum. While Head et al.'s calibrated model generates substantial autocorrelation over one year due to the gradual adjustment of the liquidity of houses, it generates no momentum over two years. Furthermore, 85 percent of the impulse response in price in occurs in the first quarter. My model with backward-looking sellers also features search frictions that make prices rise somewhat after an initial jump, but this momentum-generating effect is strengthened considerably by Head et al.'s assumptions about construction. Adding their mechanisms to my model would if anything reduce the number of backward-looking sellers needed to quantitatively explain momentum.

2.2 Excess Volatility of Inventory and Housing Phillips Curve

I relate two other facts and puzzles about the short-run dynamics of housing cycles to momentum.

Fact 2: *At an annual frequency, the volatility of inventory as measured by months of supply is three times that of real price and the volatility of sales volume is twice that of real price.*

Table 1: Cyclical Summary Statistics for Income, House Price, Sales, and Inventory

	$\sigma_{\log x_q - \log x_{q-4}}$	ρ_x, Real HPI	ρ_x, Sales Volume	ρ_x, Months of Supply
Real Disposable Pers. Income	0.016	0.819	0.668	0.497
Real House Price Index	0.065		0.726	0.305
Sales Volume	0.143			-0.263
Inventory: Months of Supply	0.207			

Notes: All series are for 1976-2013 at a quarterly frequency. The first column shows the standard deviation of annual log changes, while the other columns show correlation coefficients of log levels at a quarterly frequency. Real disposable personal income is BEA series DPIC96. Real price is the CoreLogic national repeat-sales house price index adjusted for inflation using the CPI. Sales volume is from the National Association of Realtors single-family existing home series. Months of supply is created by dividing homes listed for sale from the Census Vacancy Survey by the NAR sales series. The volume, income, and months of supply series are all seasonally adjusted.

Figure 2: Price Changes Correlated With Inventory Levels

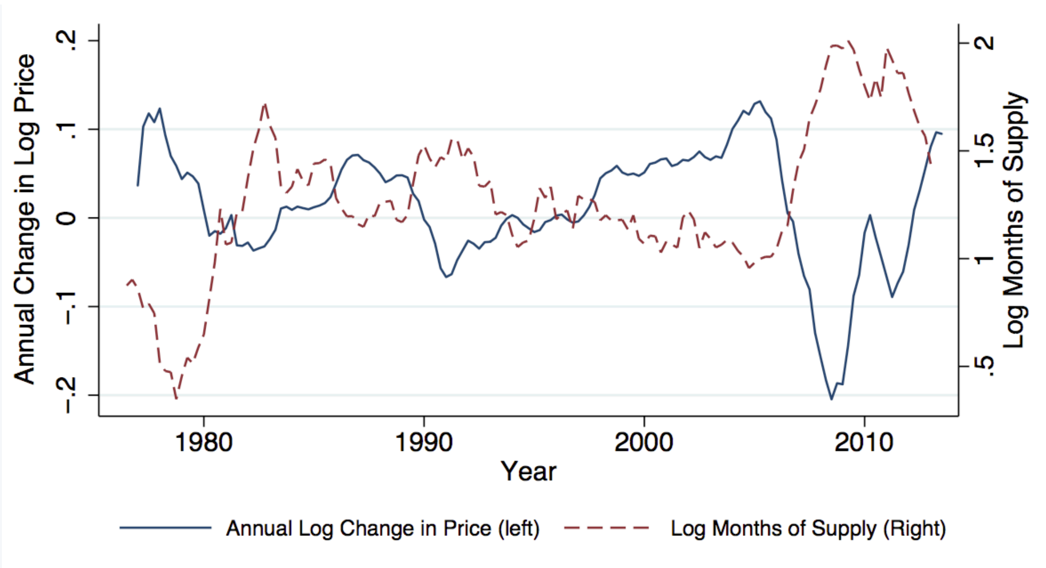

Notes: The figure shows the time series of the annual change in the log CoreLogic national repeat-sales house price index plotted against log months of supply. The latter is calculated by dividing homes vacant for sale from the Census Vacancy Survey by sales of existing single-family homes from the National Association of Realtors (NAR), measured at the midpoint of the yearlong period over which the change in price is computed.

Table 1 shows the standard deviation of annual log changes for four series: real disposable personal income, real house prices, sales volume, and "for sale" inventory measured as months of supply (a common metric in the housing market). Price is four times more volatile than income, and volume and inventory are, in turn, two and three times more volatile than price. The excess volatility of inventory is something many models have a difficult time explaining, as I discuss in Section 6.

Fact 3: *(Housing Phillips Curve) Price changes are negatively correlated with inventory levels, with a one log point increase in months of supply correlated with a 0.14 log point decrease in annual* *price growth (Peach, 1983; Caplin and Leahy, 2011).*

Figure 2 shows the time series of the annual change in the log price index and of log inventory, measured as months of supply at the midpoint of the year over which the change in price is calculated. This relationship is reminiscent of the Phillips curve as it relates for sale inventory—the equivalent of unemployment in the housing market—to price appreciation. The visible inverse co-movement in the series is confirmed by a regression: a one log point increase in months of supply is associated with a 0.14 log point decrease in the annual change in prices with an R-squared of 0.53. Appendix B shows that this relationship holds within nearly all metropolitan areas. This relationship is puzzling because most asset pricing models imply price changes should be correlated with *changes* in variables that reflect fundamentals, such as inventory, rather than with their levels. With mean reverting shocks, such models imply a positive correlation between price changes and inventory levels because when inventory levels are high, inventories tend to fall and prices tend to rise. Caplin and Leahy (2011) show that this effect can be eliminated but not reversed if prices are posted before shocks are realized.

3 Are Housing Demand Curves Concave?

I propose an amplification channel for momentum based on search and a concave demand curve in relative price. Search is a natural assumption for housing markets, but the relevance of concave demand requires further explanation.

A literature in macroeconomics argues that strategic complementarities among goods producers can amplify small pricing frictions into substantial price sluggishness by incentivizing firms to set prices close to one another. Strategic complementarities operate either through a monopolistic firm's marginal cost or its markup, which pushes a firm to price close to the market average if demand is concave in relative price. "Kinked demand" was introduced by Stiglitz (1979), who hypothesized that firms that increase their price induce consumers to search for a new firm, but firms that cut their price only gain a few active searchers. Ball and Romer (1990) show that this can create real rigidities and possibly explain why prices are so sticky despite small menu costs. Kimball (1995) generalizes Dixit-Stiglitz-style aggregator to allow for concave demand, which is used as an important real rigidity in several popular New Keynesian models (*e.g.*, Smets and Wouters, 2007). Despite the frequency with which it is used, there is little direct evidence for concave demand.[7]

Because momentum is similar to price stickiness in goods markets, I hypothesize that a similar strategic complementarity may amplify house price momentum. There are several reasons why concave demand may arise in housing markets. First, buyers may avoid visiting homes that appear to be overpriced. Second, buyers may infer that underpriced homes are lemons. Third, a house's relative list price may be a signal of seller type, such as an unwillingness to negotiate (Albrecht et al.,

[7] The most direct evidence to date comes from Nakamura and Zerom (2010), who directly estimate the "super elasticity" (elasticity of change of the elasticity) of demand for coffee using a random coefficients structural model and find evidence for concave demand. By contrast, Klenow and Willis (2006) argue that price changes in goods markets are too large to be consistent with concave demand.

2014). Fourth, homes with high list prices may be less likely to sell quickly and may consequently be more exposed to the tail risk of becoming a “stale” listing that sits on the market without selling (Taylor, 1999). Fifth, buyers may infer that underpriced homes have a higher effective price than their list price because their price is likely to be increased in a bidding war (Han and Strange, 2012b). Sixth, the law of one price—which would create a step-function demand curve—may be smoothed into a concave demand curve by uncertainty about what a house is worth.

Nonetheless, concrete evidence is needed for the existence of concave demand in housing markets before it is adopted as an explanation for momentum. Consequently, this section assesses whether demand is concave by analyzing micro data on listings matched to sales outcomes for the San Francisco Bay, Los Angeles, and San Diego metropolitan areas from April 2008 to February 2013.[8]

The relevant demand curve for list-price-setting sellers is the effect of unilaterally changing a house's relative quality-adjusted list price on its probability of sale. Detecting a nonlinear effect is challenging because quality is poorly measured, list prices are endogenously correlated with unobserved quality, and market conditions vary. The principal econometric challenge is that quality differences unobserved to the econometrician are likely positively correlated with price leading to an estimated demand curve that is far more inelastic than the true demand curve. The analysis is also complicated by the high number of foreclosures and short sales during the period that I analyze. Short sales, which occur when a home is sold for less than the outstanding mortgage balance, are especially worrisome because they often involve lengthy negotiations between the seller and their mortgage servicer which artificially decrease the probability of sale.

To surmount these challenges, I use a non-linear instrumental variable approach that traces out the demand curve using plausibly exogenous supply-side variation in seller pricing behavior. Before explaining the econometric strategy and presenting my main estimates, I first discuss the data.

3.1 Data

I combine data on listings with data on housing characteristics and transactions. The details of data construction can be found in Appendix A. The listings data come from Altos Research, which every Friday records a snapshot of homes listed for sale on multiple listing services (MLS) from several publicly available web sites and records the address, MLS identifier, and list price. The housing characteristics and transactions data come from DataQuick, which collects and digitizes public records from county register of deeds and assessor offices. This data provides a rich one-time snapshot of housing characteristics from 2013 along with a detailed transaction history of each property from 1988 to 2013 that includes transaction prices, loans, buyer and seller names and characteristics, and seller distress. I limit my analysis to non-partial transactions of single-family existing homes as categorized by DataQuick.

I match the listings data to a unique DataQuick property ID. To account for homes being delisted and re-listed, listings are counted as contiguous if the same house is re-listed within 90 days

[8]These areas were selected because both the listings and transactions data providers are based in California, so the matched dataset for these areas is of high quality and spans a longer time period.

and there is not an intervening foreclosure. If a matched home sells within 12 months of the final listing date, it is counted as a sale, and otherwise it is a withdrawal. The matched data includes 83 percent of single-family transactions in the Los Angeles area and 73 percent in the San Diego and San Francisco Bay areas. It does not account for all transactions due to three factors: a small fraction of homes (under 10%) are not listed on the MLS, some homes that are listed in the MLS contain typos or incomplete addresses that preclude matching to the transactions data, and Altos Research's coverage is incomplete in a few peripheral parts of each metropolitan area.

I limit the data to homes listed between April 2008 and February 2013.[9] I drop cases in which a home has been rebuilt or significantly improved since the transaction, the transaction price is below \$10,000, or a previous sale occurred within 90 days. I exclude ZIP codes with fewer than 500 repeat sales between 1988 and 2013 because my empirical approach requires that I calculate a local house price index. These restrictions eliminate approximately five percent of listings.

The final data set consists of 665,770 listings leading to 467,806 transactions. I focus on the 432,311 listings leading to 319,121 transactions with an observed prior transaction, and my IV procedure is limited to a more restricted sample described below. Table 2 provides summary statistics for several different subsamples that I use in the analysis.

3.2 Empirical Approach

3.2.1 Econometric Model

Before presenting the empirical approach, I introduce an econometric framework for how changes in list price around a quality-adjusted average price affect probability of sale. Each possible sequence of list prices is associated with a distribution of time to sale. To simplify the analysis, the unit of observation is a listing associated with an initial log list price, p. I work with a summary statistic of the time to sale distribution, d, which in the main text is an indicator for whether the house sells within 13 weeks, with a withdrawal counting as a non-sale. I vary the horizon and use time to sale for the subset of listings that sell in robustness checks. The data consist of homes, denoted with a subscript h, from markets defined by a location ℓ (a ZIP code in the data) and time period t (a quarter in the data).

I am interested in the impact of quality-adjusted list price relative to the average quality-adjusted list price in the market on probability of sale.[10] The quality-adjusted average list price $\tilde{p}_{h\ell t}$ has two additive components: the average log list price in location ℓ at time t, represented by a fixed effect

[9] The Altos data begins in October 2007 and ends in May 2013. I allow a six month burn-in so I can properly identify new listings, although the results are not substantially changed by including October 2007 to March 2008 listings. I drop listings that are still active on May 17, 2013, the last day for which I have data. I also drop listings that begin less than 90 days before the listing data ends so I can properly identify whether a home is re-listed within 90 days and whether a home is sold within six months. The Altos data for San Diego is missing addresses until August 2008, so listings that begin prior to that date are dropped. The match rate for the San Francisco Bay area falls substantially beginning in June 2012, so I drop Bay area listings that begin subsequent to that point.

[10] While I focus on list prices, it is important to test the robustness of the results to using transaction prices to ensure that bargaining or price wars that occur after a list price is chosen do not undo any concavity in list price. Appendix C shows all results are robust to using transaction prices.

Table 2: Summary Statistics For Listings Micro Data

Sample	All	Prior Trans	IV	All	Prior Trans	IV
Only houses that sold?	No	No	No	Yes	Yes	Yes
Transaction	70.30%	73.80%	66.70%	100%	100%	100%
Prior Transaction	64.90%	100%	100%	68.20%	100%	100%
REO	20.50%	24.90%	0%	26.70%	31.90%	0%
Short Sales	20.60%	24.20%	0%	20.20%	23.70%	0%
Positive Appreciation Since Purchase	63.00%	43.00%	100%	60.60%	42.30%	100%
Initial List Price	\$642,384	\$586,071	\$818,282	\$581,007	\$541,616	\$790,112
Transaction Price				\$534,826	\$497,828	\$731,906
Weeks on Market				15.06	15.69	12.39
Sold Within 13 Wks	43.40%	44.10%	46.80%	61.70%	59.70%	70.10%
Beds	3.283	3.236	3.309	3.266	3.226	3.298
Baths	2.185	2.119	2.275	2.145	2.097	2.256
Square Feet	1,810	1,722	1,911	1,763	1,694	1,888
N	665,770	432,311	111,358	467,806	319,121	74,320

Notes: Each column shows summary statistics for a different sample. The first three columns show the full sample, houses with an observed prior transaction, and the IV sample for all listed homes regardless of whether they sell. The second three columns show the same three samples but limited to houses that sell. The data covers listings between April 2008 and February 2013 in the San Francisco Bay, Los Angeles, and San Diego areas as described in Appendix A. REOs are sales of foreclosed homes and foreclosure auctions. Short sales include cases in which the transaction price is less than the amount outstanding on the loan and withdrawals that are subsequently foreclosed on in the next two years. Appreciation since purchase is based on the ZIP code repeat-sales price index described in Appendix A.

$\xi_{\ell t}$, and quality $q_{h\ell t}$ that is only partially observable to the econometrician:

$$\tilde{p}_{h\ell t} = \xi_{\ell t} + q_{h\ell t}. \tag{2}$$

In a Walrasian world, there would be no variation in $p_{h\ell t} - \tilde{p}_{h\ell t}$ because sellers would all price homes at $\tilde{p}_{h\ell t}$ understanding that homes priced above $\tilde{p}_{h\ell t}$ would not sell and that pricing below $\tilde{p}_{h\ell t}$ leaves money on the table. In the housing market, however, there are search frictions and substantial amounts of idiosyncratic preference that cause demand to be a downward-sloping function of $p_{h\ell t} - \tilde{p}_{h\ell t}$, which can be thought of as the seller's relative markup. Variation in the relative markup represents differences in sellers' outside options due to factors like liquidity.

Formally, I model the probability of sale $d_{h\ell t}$ as:

$$d_{h\ell t} = g\left(p_{h\ell t} - \tilde{p}_{h\ell t}\right) + \psi_{\ell t} + \varepsilon_{h\ell t}. \tag{3}$$

The demand curve in relative price $g(\cdot)$ is assumed to be invariant across markets defined by a location and time net of an additive fixed effect $\psi_{\ell t}$ that represents local market conditions. $\varepsilon_{h\ell t}$

is an error term that represents luck in finding a buyer and is assumed to be independent of the relative markup $p_{h\ell t} - \tilde{p}_{h\ell t}$.[11]

If $\tilde{p}_{h\ell t}$ were observable, one could directly estimate (3). However, observable measures of quality are imperfect, so quality $q_{h\ell t}$ likely has a component that is unobserved to the econometrician. I consequently model quality as a linear function of observed measures of quality $X_{h\ell t}$ and quality unobserved by the econometrician $u_{h\ell t}$:

$$q_{h\ell t} = \beta X_{h\ell t} + u_{h\ell t}. \tag{4}$$

Building on Yavas and Yang (1995), I include two measures of each house's value at listing as quality measures in $X_{h\ell t}$: a repeat-sales predicted price equal to the price the last time the house sold converted to today's prices using a repeat-sales house price index and a predicted price from a hedonic index that values the house based on its characteristics. To construct the repeat-sales predicted price, I first estimate interval-weighted geometric repeat-sales house price index for each ZIP code as in Case and Shiller (1989). The log index for a given time period is a time dummy in a regression of log house price on house and time fixed effects. The log predicted price $\hat{p}^{repeat}_{h\ell t}$ at time t for a house h in location ℓ that sold for $P_{h\ell\tau}$ at time τ is equal to $\log\left(P_{h\ell\tau}\frac{\phi_{\ell t}}{\phi_{\ell\tau}}\right)$, where $\phi_{\ell t}$ is the ZIP code repeat-sales index at time t. To construct the hedonic predicted price, I estimate a hedonic house price index for each ZIP code using a third order polynomial in age, log square feet, bedrooms, and bathrooms for the hedonic factor. The predicted log price $\hat{p}^{hedonic}_{t}$ is the sum of a house's hedonic value as implied by a regression and the fixed effect in the regression for a given time period. The construction of both indices follows practices common in the literature and is detailed in Appendix A. I include both predicted prices in $X_{h\ell t}$ because each approach has its virtues (Meese and Wallace, 1997). In Appendix C, I show the results are robust to modeling quality as a more flexible function of the predicted prices and to including other observables in $X_{h\ell t}$.

Combining (2) and (4), the reference price $\tilde{p}_{h\ell t}$ can be written as:

$$\tilde{p}_{h\ell t} = \xi_{\ell t} + \beta X_{h\ell t} + u_{h\ell t}, \tag{5}$$

where again $\xi_{\ell t}$ is a fixed effect that represents the average price in location ℓ at time t and $u_{h\ell t}$ is unobserved quality.

Unobserved quality creates two hurdles to identifying $g(\cdot)$. First, there is an endogeneity problem as unobserved quality and price are likely positively correlated. Second, unobserved quality creates a measurement error problem because the true $\tilde{p}_{h\ell t}$ is not observable. Both of these problems lead to bias in the estimated $g(\cdot)$.

[11] Demand shocks like $\varepsilon_{h\ell t}$ traditionally cause an endogeneity problem because they are correlated with price. However, here the variable of interest is relative price, so the effect of demand shocks on average price levels is absorbed into $\xi_{\ell t}$. Similarly, the effect of prices on aggregate demand is absorbed into $\psi_{\ell t}$. It is thus natural to assume that $\varepsilon_{h\ell t}$ is independent of the relative markup.

3.2.2 Instrument

To identify the demand curve $g(\cdot)$ in the presence of unobserved quality, I use plausibly exogenous supply-side variation in the list price due to the liquidity needs of sellers. Sellers face a trade-off between selling at a higher price and selling faster. Sellers with less liquidity and consequently a higher marginal utility of cash on hand choose a higher list price and longer time on the market. A proxy for liquidity that is orthogonal to unobserved quality and seller patience can thus serve as an instrument for list price.

The proxy for liquidity that I use is the equity a seller extracts from their sale. Housing is a large component of household wealth, and many sellers use the equity they extract from sale for the down payment on their next home (Stein, 1995). This increases the marginal utility of cash on hand for sellers who extract very little equity from their house because each additional dollar of equity they extract can be leveraged to buy a substantially better house. The marginal utility of cash is lower for sellers extracting substantial equity because their purchasing power is limited more by their creditworthiness and overall budget than the cash they have on hand. Consequently, homeowners with lower equity positions set higher list prices and sell their houses at higher prices (Genesove and Mayer, 1997; Genesove and Mayer, 2001).

Because financing and refinancing decisions make the equity of sellers endogenous, I use as my instrument the log of appreciation in the ZIP repeat-sales house price index since purchase $z_{h\ell t} = \log\left(\frac{\phi_{\ell t}}{\phi_{\ell\tau}}\right)$, where ϕ is the repeat-sales house price index, t is the period of listing, and τ is the period of previous sale.[12] This would be isomorphic to equity if all homeowners took out an identical mortgage and did not refinance. The instrument thus compares sellers who purchase identical homes with identical mortgages but who have different amounts of cash on hand to make their next down payment because one seller's home appreciated more in value than the other's.

If variation in seller liquidity represented by $z_{h\ell t}$ is independent of unobserved quality and is the only source of variation in price conditional on quality and average price, $z_{h\ell t}$ can be used as an instrument to trace out the demand curve $g(\cdot)$. Because existing evidence shows that the effect of equity is non-linear and strongest for sellers with low equity (Genesove and Mayer, 1997), I let $z_{h\ell t}$ affect price through a flexible function $f(\cdot)$. Formally, $g(\cdot)$ is identified if:

Condition 1.

$$z_{h\ell t} \perp (u_{h\ell t}, \varepsilon_{h\ell t})$$

and

$$\begin{aligned} p_{h\ell t} &= f(z_{h\ell t}) + \tilde{p}_{h\ell t} \\ &= f(z_{h\ell t}) + \xi_{\ell t} + \beta X_{h\ell t} + u_{h\ell t}. \end{aligned} \tag{6}$$

The first half of Condition 1 is an exclusion restriction that requires that appreciation since

[12] Here $z_{h\ell t}$ is a measure of liquidity, whereas when multiplied by the previous price $P_{h\ell\tau}$ in $X_{h\ell t}$ it is used to convert the previous price to present values and constrained to have the same coefficient as the previous price $P_{h\ell t}$.

purchase have no direct effect on the outcome, either through fortune in finding a buyer $\varepsilon_{h\ell t}$ in equation (3) or through unobserved quality $u_{h\ell t}$. If this is the case, $z_{h\ell t}$ only affects probability of sale through the relative markup $p_{h\ell t} - \tilde{p}_{h\ell t}$. Because I use ZIP $\times$ quarter of listing fixed effects $\xi_{\ell t}$, the variation in $z_{h\ell t}$ comes from sellers who sell at the same time in the same market but purchased at different points in the cycle. Condition 1 can thus be interpreted as requiring that unobserved quality be independent of when the seller purchased.

This assumption is difficult to test because I only have a few years of listings data, so flexibly controlling for when a seller bought weakens the effect of the instrument on price in equation (6) and widens the confidence intervals to the point that any curvature is not statistically significant. Nonetheless, I evaluate the identification assumption in four ways as documented in Appendix C. First, I vary the observable measures of quality. Second, I include a linear time trend in date of purchase or time since purchase. Third, I limit the sample to sellers who purchased prior to 2004 and again include a linear time trend, eliminating variation from sellers who purchased near the peak of the bubble or during the bust. In all three cases, the results remain robust, although standard errors widen with smaller sample sizes. Finally, I show that the shape of the estimated demand curve is similar for IV and OLS, although OLS results in a more inelastic demand curve due to bias created by the positive correlation of price with unobserved quality. While these tests assuage some concerns, if homes with very low appreciation since purchase are of substantially lower unobserved quality despite their higher average list price, my identification strategy would overestimate the true amount of curvature in the data.[13]

I focus on sellers for whom the exogenous variation is cleanest and consequently exclude three groups. First, many individuals who have had negative appreciation since purchase are not the claimant on the residual equity in their homes—their mortgage lender is. For these individuals, appreciation since purchase is directly related to how far underwater they are, which in turn affects the foreclosure and short sale processes of the mortgage lender or servicer. Because I am interested in market processes, I exclude short sales, withdrawals that are subsequently foreclosed upon, and individuals who have had negative appreciation since purchase from the analysis. Second, mortgage servicers and government-sponsored enterprises selling foreclosed homes have no reason to be sensitive to the amount of appreciation since the foreclosed-upon homeowner purchased and are dropped. Finally, investors who purchase, improve, and flip homes typically have a low appreciation in their ZIP code since purchase but improve the quality of the house in unobservable ways. To minimize the effect of investors, I exclude sellers who previously purchased with all cash, a hallmark of investors.

The second part of Condition 1 requires that liquidity embodied in $z_{h\ell t}$ is the only reason for variation in $p_{h\ell t} - \tilde{p}_{h\ell t}$. This is a strong assumption because there may be components of liquidity that are unobserved or other reasons that homeowners list their house at a price different from $\tilde{p}_{h\ell t}$, such as heterogeneity in discount rates. If the second part of the condition did not hold, then the estimates would be biased because the true $p_{h\ell t} - \tilde{p}_{h\ell t}$ would equal $f\left(z_{h\ell t}\right) + \zeta_{h\ell t}$, and the

[13]One concern is that sellers with higher appreciation since purchase improve their house in unobservable ways with their home equity. However, this would create a positive relationship between price and appreciation since purchase while I find a strong negative relationship.

unobserved measurement error $\zeta_{h\ell t}$ enters $g(\cdot)$ nonlinearly. This is an issue because the measurement error induced by unobserved quality is non-classical. In (6), unobserved quality is a residual that is independent of observed quality $\beta X_{h\ell t}$, and so the measurement error induced by $u_{h\ell t}$ is Berkson measurement error, in which the measurement error is independent of the observed component, rather than classical measurement error, in which the measurement error is independent of the truth $\tilde{p}_{h\ell t}$. An instrument such as $z_{h\ell t}$ can address classical measurement error in a non-linear setting, but it cannot address Berkson measurement error, which is why an additional assumption is necessary. If this assumption were to fail, my results would be biased.

However, the bias created by measurement error does not cause significant spurious concavity, which I show using two strategies. First, I prove that if the measurement error created by other sources of variation in the relative markup $p_{h\ell t} - \tilde{p}_{h\ell t}$ is independent of the variation induced by the instrument, the measurement error would not cause spurious concavity. Intuitively, noise in $p_{h\ell t} - \tilde{p}_{h\ell t}$ would cause the observed probability of sale at each observed $p_{h\ell t} - \tilde{p}_{h\ell t}$ to be an average of the probabilities of sale at true $p_{h\ell t} - \tilde{p}_{h\ell t}$s that are on average evenly scrambled. Consequently, the curvature of a monotonically-decreasing demand curve is preserved. An analytical result can be obtained if the true $g(\cdot)$ is a polynomial regression function as in Hausman et al. (1991):

Lemma 1. *Consider the econometric model described by (3) and (5) and suppose that:*

$$z_{h\ell t} \perp (u_{h\ell t}, \varepsilon_{h\ell t}), \tag{7}$$

$$p_{h\ell t} = f(z_{h\ell t}) + \zeta_{h\ell t} + \tilde{p}_{h\ell t}, \tag{8}$$

$\zeta_{h\ell t} \perp f(z_{h\ell t})$, *and the true regression function* $g(\cdot)$ *is a third-order polynomial. Then estimating* $g(\cdot)$ *assuming that* $p_{h\ell t} = f(z_{h\ell t}) + \tilde{p}_{h\ell t}$ *yields the true coefficients of the second- and third-order terms in* $g(\cdot)$. *If* $g(\cdot)$ *is a second-order polynomial, the same procedure yields the true coefficients of the first- and second-order terms.*

Proof. See Appendix C. □

While a special case, Lemma 1 makes clear that the bias in the estimated concavity is minimal if $\zeta_{h\ell t} \perp f(z_{h\ell t})$.

Second, while spurious concavity is a possibility if the measurement error created by other sources of variation in the relative markup were correlated with the instrument, the amount of concavity generated would be far smaller than the concavity I observe in the data. Appendix C presents Monte Carlo simulations that show that if the instrument captures most of the variation in the relative markup $p_{h\ell t} - \tilde{p}_{h\ell t}$ at low levels of appreciation since purchase but very little of the variation at high levels of appreciation since purchase, spurious concavity arises because the slope of $g(\cdot)$ is attenuated for low relative markups but not high relative markups. However, to spuriously generate a statistically-significant amount of concavity, one would need a virtually perfect instrument at low levels of appreciation since purchase and essentially all of the variation in price at high levels of appreciation since purchase to be measurement error. Because this is implausible, I conclude that spurious concavity due to Berkson measurement error is not driving my findings on concavity.

3.2.3 Estimation

Under Condition 1, $p_{h\ell t} - \tilde{p}_{h\ell t} = f(z_{h\ell t})$, and $g(\cdot)$ can be estimated by a two-step procedure that first estimates equation (6) and then uses the predicted $f(z_{h\ell t})$ as $p_{h\ell t} - \tilde{p}_{h\ell t}$ to estimate equation (3). Both equations are estimated by OLS, and in the main text I weight the specifications by the inverse standard deviation of the error in the repeat-sales index to account for the reduced precision of the predicted prices in areas with fewer transactions. I use a third-order polynomial for $f(\cdot)$, and Appendix C shows that the results are robust to the order of the polynomial used.

I assess the degree of concavity in two ways. First, I use a quadratic polynomial in $p_{h\ell t} - \tilde{p}_{h\ell t}$ for $g(\cdot)$ and test whether the quadratic term is statistically distinguishable from zero. Based on Lemma 1 above, the linear and quadratic terms are both unbiased, and so testing whether the quadratic term is less than zero is the simplest way to assess concavity. To account for spatial correlation, I calculate standard errors by block bootstrapping the entire procedure and clustering on 35 units defined by the first three digits of the ZIP code (ZIP-3), and I use the bootstrapped 95 percent confidence interval as my preferred test for concavity. Second, to visualize the data, I construct a binned scatter plot, which bins the data into 25 equally-sized groups of the log list price relative to the reference price, $p_{h\ell t} - \tilde{p}_{h\ell t}$, and, for each bin, plots the mean of $p_{h\ell t} - \tilde{p}_{h\ell t}$ against the mean of the probability of sale net of the average probability of sale in the market, $d_{h\ell t} - \psi_{\ell t}$. This approximates $g(\cdot)$ using indicator variables for the 25 bins of $p_{h\ell t} - \tilde{p}_{h\ell t}$, as detailed in Appendix C. I overlay a third-order polynomial fit with pointwise 95 percent confidence bands with the binned scatter plot.

There may be small-sample bias introduced into the estimation if $g(\cdot)$ is non-linear and the fixed effects $\xi_{\ell t}$ are imprecisely estimated with a small number of homes in a ZIP-quarter cell.[14] Appendix C shows that the results are not substantially changed by limiting the sample to fixed effect cells with at least 15 homes. Because the error in the estimated fixed effects is likely minimal for these cells, this suggests that imprecision in the estimated fixed effects is not driving the results.

3.3 Results

Figure 3 shows the resulting first and second stage binned scatter plots. As shown in panel A, the instrument induces a small amount of variation in the list price set by sellers.[15] This is the variation I use to identify the shape of demand (and is the x-axis of panel B). The first stage is strong with a joint F statistic for the third order polynomial of the instrument in (6) of 128. Panel B shows a

[14] There are 9,200 fixed effects. 0.33 percent of the data is unused because there is only a single house sold in a ZIP-quarter cell.

[15] Genesove and Mayer (1997) find that a house with 100 percent loan-to-value ratio is on average listed at a price four percent higher than a home with an 80 percent loan-to-value ratio. Subsequent work (Genesove and Mayer, 2001) finds slightly smaller numbers conditioning on whether a seller has experienced a nominal loss. Nonetheless, the similarity between their four percent figure and the amount of variation induced by the instrument in my first stage is reassuring.

Table 3: The Effect of List Price on Probability of Sale: Regression Results

Panel A: Ordinary Least Squares				
Dependent Var:	Sell Within 13 Weeks			
Sample:	All Listings With Prior Observation (432,311 obs, 421,297 After Dropping 1st and 99th % and Cells With One Obs)			
Controls:	ZIP × Quarter × Distress FE, Repeat and Hedonic Predicted Price			
	Constant	Linear	Quadratic	Quadratic Bootsrapped 95% CI
Coefficient	0.471***	-0.249***	-0.605***	[-0.796,-0.472]
	(0.005)	-0.016	-0.087	
Panel B: Ordinary Least Squares, IV Sample				
Sample:	Listings With Prior Obs, excluding REO, Short Sales, Investors, Neg Appreciation (111,358 obs, 108,460 After Dropping 1st and 99th % and Cells With One Obs)			
Controls:	ZIP × Quarter FE, Repeat and Hedonic Predicted Price			
	Constant	Linear	Quadratic	Quadratic Bootsrapped 95% CI
Coefficient	0.511***	-0.352***	-0.715**	[-1.213,-0.436]
	(0.012)	(0.031)	(0.223)	
Panel C: Instrumental Variable				
Sample:	Same a Panel B			
Controls:	ZIP × Quarter FE, Repeat and Hedonic Predicted Price			
	Constant	Linear	Quadratic	Quadratic Bootsrapped 95% CI
Coefficient	0.491***	-1.138***	-33.466*	[-76.788,-21.130]
	(0.009)	(0.290)	(14.289)	

Notes: * $p < 0.05$, ** $p<0.01$, *** $p<0.001$. Each row shows regression coefficients when $g(\cdot)$ in equation (3) is approximated using a quadratic polynomial. This relationship represents the effect of the log relative markup on the probability of sale within 13 weeks. In the IV panel, a first stage regression of log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and log predicted price using both a repeat-sales and a hedonic methodology, as in (6) is estimated by OLS. The predicted value of the polynomial of the instrument is used as the relative markup in equation (3), which is estimated by OLS. The sample is restricted to non-REOs, non-short sales, properties with positive appreciation since purchase, and properties not previously purchased with all cash (investors). In the OLS panels, quality is assumed to be perfectly measured by the hedonic and repeat-sales predicted prices and have no unobserved component. Panel B thus regresses log list price on fixed effects and the predicted prices and uses the residual as the estimated relative markup into equation (3), as described in Appendix C. OLS uses the full set of listings with a previous observed transaction, so to prevent distressed sales from biasing the results, the fixed effects are at the quarter of initial listing x ZIP x distress status level. Distress status corresponds to three groups: normal sales, REOs (sales of foreclosed homes and foreclosure auctions), and short sales (cases where the transaction price is less than the amount outstanding on the loan and withdrawals that are subsequently foreclosed on in the next two years). Both procedures are weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. Before creating the spline, the 99th and 1st percentiles of the relative markup are dropped, as are any observations fully absorbed by fixed effects. Standard errors and the 95 percent confidence interval for the quadratic term are computed by block bootstrapping the entire procedure on 35 ZIP-3 clusters.

Figure 3: Instrumental Variable Estimates of the Effect of List Price on Probability of Sale

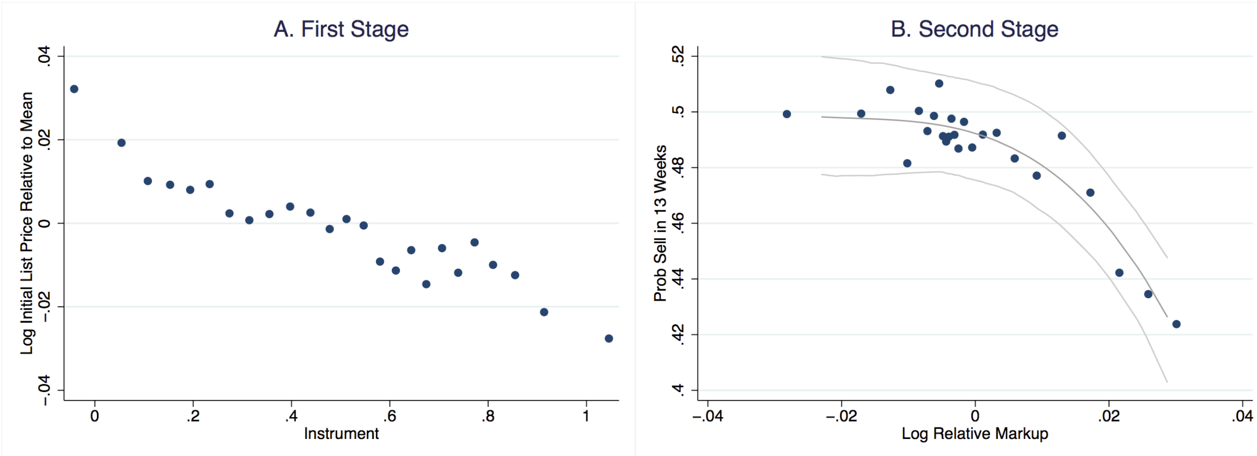

Notes: Panel B shows a binned scatter plot of the probability of sale within 13 weeks net of fixed effects (with the average probability of sale within 13 weeks added in) against the estimated log relative markup $p - \tilde{p}$. It also shows an overlaid cubic fit of the relationship, as in equation (3). To create the figure, a first stage regression of the log list price on a third-order polynomial in the instrument, fixed effects at the ZIP x first quarter of listing level, and repeat sales and hedonic log predicted prices, as in (6), is estimated by OLS. The predicted value of the polynomial of the instrument is used as the relative markup. The figure splits the data into 25 equally-sized bins of this estimated relative markup and plots the mean of the estimated relative markup against the mean of the probability of sale within 13 weeks net of fixed effects for each bin, as detailed in Appendix C. Before binning, the 1st and 99th percentiles of the log sale price residual and any observations fully absorbed by fixed effects are dropped. The entire procedure is weighted by the reciprocal of the standard deviation of the prediction error in the repeat-sales house price index in the observation's ZIP code from 1988 to 2013. The sample is limited to the IV subsample of homes that are not sales of foreclosures or short sales, sales of homes with negative appreciation since the seller purchased, or sales by investors who previously purchased with all cash. The grey bands indicate a pointwise 95-percent confidence interval for the cubic fit created by block bootstrapping the entire procedure on 35 ZIP-3 clusters. Panel A shows the first stage relationship between the instrument and log initial list price in equation (6) by residualizing the instrument and the log initial list price against the two predicted prices and fixed effects, binning the data into 25 equally-sized bins of the instrument residual, and plotting the mean of the instrument residual against the mean of the log initial list price residual for each bin. N = 111,293 observations prior to dropping the 1st and 99th percentiles and unique zip-quarter cells.

clear concave relationship in the second stage, with very inelastic demand for relatively low priced homes and elastic demand for relatively high priced homes. This curvature is also visible in the cubic polynomial fit.

Table 3 shows regression results when $g(\cdot)$ is approximated by quadratic polynomial. Panel C shows the IV results. The concavity visible in Figure 3 is apparent, with a large and statistically significant quadratic term with a bootstrapped 95 percent confidence interval that lies substantially below zero.

As a point of comparison, Panel A shows OLS results for the full sample of homes with a prior observed transaction and panel B shows OLS results for the IV sample. The fixed effects are at the ZIP $\times$ quarter $\times$ REO seller $\times$ short seller level to prevent distressed sales from biasing the results. OLS assumes away unobserved quality and should be positively biased if $\tilde{p}_{h\ell t}$ is positively correlated with $p_{h\ell t}$ due to omitted unobserved quality. This is the case: the estimated demand curve is more elastic for IV than OLS. In fact, the OLS bias is strong enough that the demand curve slopes significantly upward in the lowest tercile. Nonetheless, there is clear and significant concavity in the OLS results, and the similarity between the full data set and IV sample assuages concerns about sample selection.

At the mean price, the estimates imply that raising one's price by one percent reduces the probability of sale within 13 weeks by approximately 1.4 percentage points on a base of 46.8 percentage points, a reduction of 3 percent. This corresponds to a one percent price hike increasing the time to sale by three to five days. By contrast, increasing the list price by five percent reduces the probability of sale within 13 weeks by 14 percentage points, a reduction of 30 percent. These figures are slightly smaller than those found by Carrillo (2012), who estimates a structural search model of the steady state of the housing market with multiple dimensions of heterogeneity using data from Charlottesville, Virginia from 2000 to 2002. Although we use very different empirical approaches, in a counterfactual simulation, he finds that a one percent list price increase increases time on the market by a week, while a five percent list price increase increases time on the market by nearly a year. Carrillo also finds small reductions in time on the market from underpricing, consistent with the nonlinear relationship found here.

Appendix C shows the finding of concavity survives many robustness checks. The results are robust across geographies, time periods, and specifications, although in some cases restricting to a smaller sample leads to insignificant results. The results are also robust to time trends in date of purchase, limiting the sample to sellers who purchased before the bubble, controlling for other measures of quality, allowing for differential sorting by letting β vary across time and space, accounting for the uniqueness of a house in its neighborhood, and accounting for different price tiers within ZIP codes. Appendix C also shows that concavity is clearly visible in the reduced-form relationship between the instrument and probability of sale and that the results are robust to using transaction prices rather than using list prices. The instrumental variable results thus provide evidence of demand concave in relative price for these three MSAs from 2008 to 2013.[16]

4 A Model of House Price Momentum

This section introduces an equilibrium search model with concave demand. The model includes two additional ingredients new to the housing search literature. First, because concave demand only amplifies existing price insensitivity, I introduce small number of backward-looking rule-of-thumb sellers as in Haltiwanger and Waldman (1989) and Gali and Gertler (1999). Second, I include an endogenous entry decision for buyers and sellers because entry is a form of intertemporal arbitrage that works against the forces that generate momentum in the model. Furthermore, with an active but not overwhelming entry margin, the re-timing of purchases and sales in light of momentum can explain the excess volatility of inventory and the housing Phillips curve relationship.

[16] Aside from the tail end of my sample, this period was a depressed market. The similarity between my results and Carrillo's provide some reassurance that the results I find are not specific to the time period, but I cannot rule out that the nonlinearity would look different in a booming market.

Table 4: Notation in the Model

Variable	Description	Note
	Masses	
N	Total Population (Housing Stock Mass One)	
B_t	Endogenous Mass of Buyers	Value Fn V_t^b
S_t	Endogenous Mass of Sellers	Value Fn V_t^s
R_t	Endogenous Mass of Renters	Value Fn V_t^r
H_t	Endogenous Mass of Homeowners	Value Fn V_t^h
	Flow Utilities	
b	Flow Utility of Buyer (Includes search cost)	
s	Flow Utility of Seller (includes search cost)	
u_t	Flow Utility of Renter	Shocked Variable
h	Flow Utility of Homeowner	
	Moving Shock Probabilities	
λ^h	Prob Homeowner Gets Shock	
λ^r	Prob Renter Gets Shock	
	Costs	
c	Stochastic Cost for Homeowner to Stay in Home	$C(c) \sim U(\underline{c}, \bar{c})$
k	Stochastic Cost for Renter to Stay Renter (Negative)	$K(k) \sim U(\underline{k}, \bar{k})$
c_t^*	Threshold c Above Which Homeowners Enter	Endogenous
k_t^*	Threshold k Above Which Renters Enter	Endogenous
	Other Fixed Parameters	
β	Discount Factor	
L	Probability Seller Leaves Metro Area	
V^0	PDV of Exiting Metro Area	

The model builds on search models of the housing market, such as Wheaton (1990), Krainer (2001), Novy-Marx (2009), Piazzesi and Schneider (2009), Caplin and Leahy (2011), Genesove and Han (2012), Ngai and Tenreyro (2013), and Head et al. (2014). I first introduce a framework that models a metropolitan area with a fixed population and housing stock. I then describe the housing market component and show how sellers set list prices. I finally introduce rule-of-thumb sellers and define equilibrium. The notation used in the model is summarized in Tables 4 and 5.

4.1 Setting

Time is discrete and all agents are risk neutral. Agents have a discount factor of β. There is a fixed housing stock of mass one, no construction, and a fixed population of size N.[17]

[17] Construction is omitted for parsimony. While momentum is negatively correlated with the housing supply elasticity, momentum is important even in elastic MSAs. The model best applies to areas with inelastic housing supply in which momentum is stronger, although it is also relevant to the short run in areas with more elastic housing

There are four types of homogenous agents: a mass B_t of buyers, S_t of sellers, H_t of homeowners, and R_t of renters. These agents have flow utilities (inclusive of search costs) b, s, h, and u_t (u is time varying because it will be shocked), and value functions V_t^b, V_t^s, V_t^h, and V_t^r, respectively. Buyers and sellers are active in the housing market, which is described in the next section. The rental market, which serves as a reservoir of potential buyers, is unmodeled aside from the flow utility net of rents. Each agent can own only one home, which precludes short sales and investor-owners. Sellers and buyers are homogenous but sellers may differ in their list prices.

Each period with probability λ^h and λ^r, respectively, homeowners and renters receive shocks that cause them to separate from their current house or apartment, as in Wheaton (1990). However, rather than automatically entering the housing market, the shocks cause homeowners and renters to draw a one-time cost, $c \sim C(\cdot)$ for homeowners and $k \sim K(\cdot)$ (likely negative) for renters, that can be paid to stay in their current house or apartment and receive the same flow utility as before instead of moving.[18] The cost distributions are parameterized as uniform so that $c \sim U(\underline{c}, \bar{c})$ and $k \sim U\left(\underline{k}, \bar{k}\right)$, which gives an approximately constant seller entry elasticity appears to be constant as shown to hold in the data in Appendix E. This setup captures that potential movers have heterogeneous reasons to buy or sell and consequently differ in the ease with which they can re-time their transaction.

A renter who decides not to pay the cost k enters the market as a homogenous buyer. A homeowner who decides not to pay the cost c learns after making their entry decision whether they leave the MSA with probability L, in which case they become a seller and receive a net present value of V^0 for leaving, or whether they remain in the city with probability $1 - L$. If they remain in the city, they simultaneously become a buyer and a homogenous seller. These two roles are assumed to be quasi-independent so that the value functions do not interact and no structure is put on whether agents buy or sell first, as in Ngai and Tenreyro (2013) and Guren and McQuade (2015).

A homeowner who draws a cost c enters the market if:

$$c \geq V_t^h - V_t^s - LV^0 - (1 - L) V_t^b \equiv c_t^*. \tag{9}$$

Similarly a renter enters if

$$k \geq V_t^r - V_t^b \equiv k_t^*. \tag{10}$$

The cutoffs c_t^* and k_t^* determine the marginal entrant and control their flow into the market.[19] Because the population is constant, every time a seller leaves the city they are replaced by a new entrant. Entrants draw a cost of being a renter and decide whether to rent or buy in the same manner as a renter who just experienced a shock. The full closed system is illustrated diagrammatically

supply. See Head et al. (2014) for a model with a construction margin.

[18]The model would be isomorphic if homeowners and renters drew a cost every period and had a much lower probability of drawing a cost that leads to a move as long as the densities of the two distributions of costs were the same at the margin.

[19]This setup makes two implicit assumptions for tractability. First, although individuals are heterogeneous in their motivation to move, once they enter the market they are homogenous. Second, if an individual decides not to move today, they do not make another decision about moving until they get another shock.

Figure 4: Schematic Representation of Closed GE System

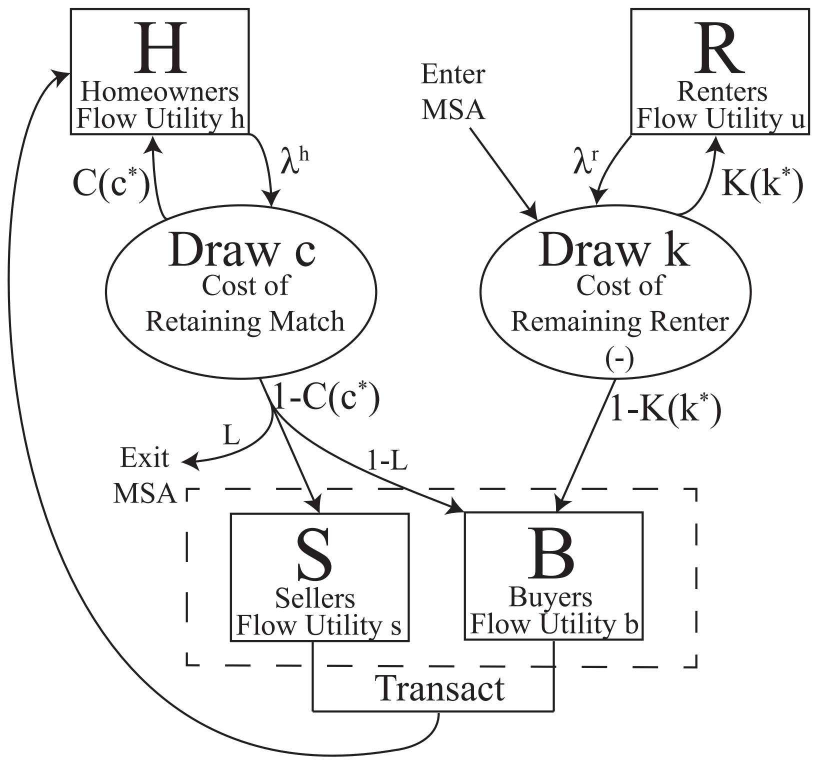

in Figure 4. I defer the laws of motion that formalize the system until after I have defined the probabilities of purchase and sale. The value functions of the homeowner and renter are:

$$V_t^h = h + \beta E_t \left[\begin{array}{c} \lambda^h \left(1 - C\left(c_{t+1}^*\right)\right)\left[V_{t+1}^s + LV^0 + (1-L)\, V_{t+1}^b\right] \\ -\lambda^h C\left(c_{t+1}^*\right) E\left[c|c < c_{t+1}^*\right] + \left(1 - \lambda^h\left(1 - C\left(c_{t+1}^*\right)\right)\right) V_{t+1}^h \end{array} \right] \tag{11}$$

$$V_t^r = u_t + \beta E_t \left[\begin{array}{c} \lambda^r \left(1 - K\left(k_{t+1}^*\right)\right) V_{t+1}^b - \lambda^r K\left(k_{t+1}^*\right) E\left[k|k < k_{t+1}^*\right] \\ + \left(1 - \lambda^r\left(1 - K\left(k_{t+1}^*\right)\right)\right) V_{t+1}^r \end{array} \right]. \tag{12}$$

The $E\left[c|c < c_t^*\right]$ and $E\left[k|k < k_t^*\right]$ terms reflect the expected costs if a homeowner or renter decide to not move and pay the cost of accommodating the shock.

4.2 The Housing Market

The search process occurs at the beginning of each period and unfolds in three stages. First, sellers post list prices $\hat{p}_t$. Second, buyers observe a noisy binary signal about each house's quality relative to its price. Buyers direct their search either towards houses that appear reasonably-priced for their quality or towards houses that appear to be overpriced for their quality, which defines two sub-markets: follow the signal (submarket f) or do not follow the signal (submarket d). After choosing a submarket, buyers search randomly within the submarket an stochastically find a house to inspect. Third, matched buyers inspect the house and decide whether to purchase it.

Table 5: Notation in Housing Market

Variable	Description	Note
	Utilities	
ε	Match-Specific One-Time Utility Benefit	$\sim F(\varepsilon)$
υ_h	Permanent House Quality	Mean Zero
	Stochastic Draws	
$\eta_{h,t}$	Noise in υ_h in Binary Signal About Quality, IID Common to All Buyers in Period t	$\sim G(\eta)$
	Parameters / Values	
θ_t^m	Market Tightness $= B/S$ in Submarket $m = \{f, d\}$	Endogenous
$\tilde{\theta}_t$	Vector of θ^m s	Endogenous
$q^m(\theta_t^m)$	Prob. Seller Meets Buyer in m (Matching Fn)	Endogenous
ξ^m	Constant in Matching Function in m	
γ	Matching Function Elasticity in all ms	
p_t	List Price Adjusted for Permanent Quality	Endogenous
Ω_t	Distribution of Prices	Endogenous
ε_t^*	Threshold ε for Purchase	Endogenous
μ	Threshold for Binary Signal	
	Distribution Parameters	
$\underline{\varepsilon}, \bar{\varepsilon}$	Bounds of Dist of $F(\varepsilon)$, Uniform With Mass Point	
χ	Size of Mass Point of Uniform Dist at $\bar{\varepsilon}$	
σ	Logistic Variance Param for $G(\eta)$	

At the inspection stage, buyers observe their idiosyncratic valuation for the house ε, which is match-specific, drawn from $F(\varepsilon)$ at inspection, and realized as utility at purchase. They also observe the house's permanent quality υ_h, which is common to all buyers, mean-zero, gained by a buyer at purchase, and lost by a seller at sale. I assume all sales occur at list price, or equivalently that risk neutral buyers and sellers expect that the average sale price will be an affine function of the list price.[20] This assumption is made for tractability and is not essential to the propagation mechanism. It is also empirically realistic: although many houses do sell above or below list price, Appendix A.3 shows that in the merged Altos-DataQuick micro data, the modal transaction price is the list price, and the average and median differences between the list and transaction price are less than 0.01 log points and do not vary much across years.[21] Letting $p_t \equiv \hat{p}_t - \upsilon_h$ be the quality-adjusted list price, the buyer purchases if his or her surplus from doing so $V_t^h + \varepsilon - p_t - b - \beta V_t^b$ is positive. This leads to a threshold rule to buy if $\varepsilon > p_t + b + \beta V_{t+1}^b - V_t^h \equiv \varepsilon_t^*(p_t)$ and a probability of purchase given

[20]This assumption restricts what can occur in bargaining or a price war. Several papers have considered the role of various types of bargaining in a framework with a list price in a steady state search model, including cases in which the list price is a price ceiling (Chen and Rosenthal, 1996; Haurin et al., 2010), price wars are possible (Han and Strange, 2013), and list price can signal seller type (Albrecht et al., 2014).

[21]An important feature of the housing market is that most price changes are decreases. Consequently, the difference between the initial list price and the sale price fluctuates substantially over the cycle as homes that do not sell cut their list price. I abstract from such duration dependence to maintain a tractable state space.

inspection of $1 - F(\varepsilon_t^*)$.

At the signal stage, buyers observe a binary signal from their real estate agent or from advertisements that reveals whether each house's quality-adjusted price relative to the average quality-adjusted price is above a threshold μ. However, quality υ_h (or equivalently the observation of the average price) is subject to mean zero noise $\eta_{h,t} \sim G(\cdot)$, where $G(\cdot)$ is assumed to be a fixed distribution.[22] This noise, which represents how well a house is marketed in a given period, is common to all buyers but independent and identically distributed across periods. The signal thus indicates a house is reasonably priced if,

$$p_t - E_{\Omega_t}[p_t] - \eta_{h,t} \leq \mu,$$

where Ω_t is the cumulative distribution function of list prices and the notation $E_{\Omega_t}[\cdot]$ represents an expectation with respect to Ω_t. Consequently, a house with quality-adjusted price p_t is searched by buyers in submarket f with probability $1 - G(p_t - E_{\Omega_t}[p_t] - \mu)$ and is searched by buyers in submarket d with probability $G(p_t - E_{\Omega_t}[p_t] - \mu)$. I assume that search is more efficient if buyers follow the signal than if they do not because they have the help of a realtor or are looking at better-marketed homes. In equilibrium, buyers follow a mixed strategy randomize whether they search submarket f or d so that the value of following the signal is equal to the value of not following it. I consider an equilibrium in which all buyers choose the same symmetric strategy with a probability of following the signal of ϕ_t.

After choosing a submarket, buyers search randomly within that sub-market and cannot direct their search to any particular type of home within that market. The probability a house in submarket m meets a buyer is determined according to a constant returns to scale matching function, $q^m(\theta_t^m)$, where θ_t^m is the ratio of buyers to sellers in submarket $m = \{f, d\}$. The probability a buyer meets a seller is then $q^m(\theta_t^m)/\theta_t^m$. The matching function captures frictions in the search process that prevent all reasonably-priced homes and all buyers from having an inspection each period. For instance, buyers randomly allocating themselves across houses may miss a few houses, or there may not be a mutually-agreeable time for a buyer to visit a house in a given period.

Given this setup, the mass of sellers in the f submarket is S_t times the weighted average probability that any given seller is in the f submarket $E_\Omega[1 - G(\cdot)]$, and the mass of sellers in the d submarket is similarly $S_t E_\Omega[G(\cdot)]$. Consequently, the market market tightness in the f and d submarkets is:

$$\begin{aligned}
\theta_t^f &= \frac{B_t^{follow}}{S_t^f} = \frac{B_t \phi_t}{S_t E_{\Omega_t}[1 - G(p_t - E_{\Omega_t}[p_t] - \mu)]} \\
\theta_t^d &= \frac{B_t^{do\ not\ follow}}{S_t^d} = \frac{B_t(1 - \phi_t)}{S_t E_{\Omega_t}[G(p_t - E_{\Omega_t}[p_t] - \mu)]}.
\end{aligned}$$

[22] Because the signal reveals no information about the house's permanent quality υ_h, posted price $\hat{p}_t$, or match quality ε_m, the search and inspection stages are independent.

The probability a buyer who follows the signal buys a house is then:

$$
\begin{aligned}
Pr\left[Buy|follow\right] &= \frac{q^f\left(\theta_t^f\right)}{\theta_t^f}\int \frac{1-G\left(p_t-E_{\Omega_t}\left[p_t\right]-\mu\right)}{E_{\Omega_t}\left[1-G\left(p_t-E_{\Omega_t}\left[p_t\right]-\mu\right)\right]}\left(1-F\left(\varepsilon_t^*\left(p_t\right)\right)\right)d\Omega_t\left(p_t\right) \\
&= \frac{q^f\left(\theta_t^f\right)}{\phi_t\theta_t}\int \left(1-G\left(p_t-E_{\Omega_t}\left[p_t\right]-\mu\right)\right)\left(1-F\left(\varepsilon_t^*\left(p_t\right)\right)\right)d\Omega_t\left(p_t\right) \\
&= \frac{1}{\phi_t\theta_t}E_{\Omega_t}\left[d^f\left(p_{t,}\Omega_t,\tilde{\theta}_t\right)\right],
\end{aligned}
$$

where $\theta_t = B_t/S_t$ is the aggregate market tightness and,

$$
d^f\left(p_t,\Omega_t,\tilde{\theta}_t\right) = q^f\left(q_t^f\right)\left(1-G\left(p_t-E_{\Omega_t}\left[p_t\right]-\mu\right)\right)\left(1-F\left(\varepsilon_t^*\left(p_t\right)\right)\right). \tag{13}
$$

Similarly, the probability a buyer buys if they do not follow the signal is:

$$
Pr\left[Buy|do\ not\ follow\right] = \frac{1}{\left(1-\phi_t\right)\theta_t}E_{\Omega_t}\left[d^d\left(p_{t,}\Omega_t,\theta_t,\phi_t\right)\right],
$$

where,

$$
d^d\left(p_t,\Omega_t,\tilde{\theta}_t\right) = q^d\left(\theta_t^d\right)G\left(p_t-E_{\Omega_t}\left[p_t\right]-\mu\right)\left(1-F\left(\varepsilon_t^*\left(p_t\right)\right)\right). \tag{14}
$$

Note that the demand curve faced by sellers, which is the *ex-ante* probability of sale for a house with a list price p_t, can be written as:

$$
\begin{aligned}
d\left(p_t,\Omega_t,\tilde{\theta}_t\right) &= Pr\left[Good\ Signal\right]Pr\left[Sell|Good\ Signal\right]+Pr\left[Bad\ Signal\right]Pr\left[Sell|Bad\ Signal\right] \\
&= d^f\left(p_t,\Omega_t,\tilde{\theta}_t\right)+d^d\left(p_t,\Omega_t,\tilde{\theta}_t\right).
\end{aligned} \tag{15}
$$

Similarly, the total probability a buyer buys is given the ϕ_t randomization strategy is:

$$
\begin{aligned}
Pr\left[Buy\right] &= \phi_t\frac{1}{\phi_t\theta_t}E_{\Omega_t}\left[d^f\left(p_{t,}\Omega_t,\tilde{\theta}_t\right)\right]+\left(1-\phi_t\right)\frac{1}{\left(1-\phi_t\right)\theta_t}E_{\Omega_t}\left[d^d\left(p_{t,}\Omega_t,\theta_t,\phi_t\right)\right] \\
&= \frac{1}{\theta_t}E_{\Omega_t}\left[d\left(p_t,\Omega_t,\tilde{\theta}_t\right)\right].
\end{aligned}
$$

Given these probabilities of purchase and sale, the stock of buyers is equal to the stock of buyers who failed to buy last period plus the stock of renters and flow of new entrants who decide to buy, while the stock of sellers is equal to those sellers who failed to sell last period plus homeowners who put their house up for sale. These are formalized by:

$$
B_t = \left(1-\frac{1}{\theta_{t-1}}E_{\Omega_{t-1}}\left[d\left(p_{t-1},\Omega_{t-1},\tilde{\theta}_{t-1}\right)\right]\right)B_{t-1} \tag{16}
$$

$$
+\lambda^r\left(1-K\left(k_{t-1}^*\right)\right)R_{t-1}+\left(1-LK\left(k_{t-1}^*\right)\right)\lambda^h\left(1-C\left(c_{t-1}^*\right)\right)H_{t-1}
$$

$$
S_t = \left(1-E_{\Omega_{t-1}}\left[d\left(p_{t-1},\Omega_{t-1},\tilde{\theta}_{t-1}\right)\right]\right)S_{t-1}+\lambda^h\left(1-C\left(c_{t-1}^*\right)\right)H_{t-1}. \tag{17}
$$

Because there are mass one of homes that can either be owned by a homeowner or up for sale and mass N of agents who can either be renters, homeowners, or buyers,

$$
\begin{aligned}
1 &= H_t + S_t && (18)\\
N &= R_t + B_t + H_t. && (19)
\end{aligned}
$$

These equations together with (16) and (17) implicitly define laws of motion for H_t and R_t.

When a buyer buys, they get the expected surplus $E\left[\varepsilon-\varepsilon_t^*|\varepsilon>\varepsilon_t^*\right]$ where ε_t^* is a function of p_t. The value of a buyer who follows, $V_t^{b,f}$, a buyer who does not follow $V_t^{b,d}$, and a buyer prior to choosing a submarket V_t^b, are then:

$$
\begin{aligned}
V_t^{b,f} &= b+\beta E_t V_{t+1}^b+\frac{1}{\phi_t\theta_t}E_{\Omega_t}\left[d^f\left(p_t,\Omega_t,\tilde{\theta}_t\right)E\left[\varepsilon-\varepsilon_t^*|\varepsilon>\varepsilon_t^*\right]\right] && (20)\\
V_t^{b,d} &= b+\beta E_t V_{t+1}^b+\frac{1}{\left(1-\phi_t\right)\theta_t}E_{\Omega_t}\left[d^d\left(p_t,\Omega_t,\tilde{\theta}_t\right)E\left[\varepsilon-\varepsilon_t^*|\varepsilon>\varepsilon_t^*\right]\right] && (21)\\
V_t^b &= \max\left\{V_t^{b,f},V_t^{b,d}\right\}.
\end{aligned}
$$

In equilibrium, buyers are indifferent between the two markets, so $V_t^{b,f}=V_t^{b,d}$ or:

$$
\frac{E_{\Omega_t}\left[d^f\left(p_t,\Omega_t,\tilde{\theta}_t\right)E\left[\varepsilon-\varepsilon_t^*|\varepsilon>\varepsilon_t^*\right]\right]}{E_{\Omega_t}\left[d^d\left(p_t,\Omega_t,\tilde{\theta}_t\right)E\left[\varepsilon-\varepsilon_t^*|\varepsilon>\varepsilon_t^*\right]\right]}=\frac{\phi_t}{1-\phi_t}. \tag{22}
$$

This pins down the fraction of buyers who go to submarket f, ϕ_t. V_t^b can then be rewritten as:

$$
V_t^b=b+\beta E_t V_{t+1}^b+\frac{1}{\phi_t\theta_t}E_{\Omega_t}\left[d^f\left(p_t,\Omega_t,\tilde{\theta}_t\right)E\left[\varepsilon-\varepsilon_t^*|\varepsilon>\varepsilon_t^*\right]\right]. \tag{23}
$$

Sellers have rational expectations but set their list price before $\eta_{h,t}$ is realized and without knowing the valuation of the particular buyer who visits their house. The demand curve they face is $d\left(p_t,\Omega_t,\tilde{\theta}_t\right)$, so the seller's value function is:

$$
V_t^s=s+\beta E_t V_{t+1}^s+\max_{p_t}\left\{d\left(p_t,\Omega_t,\tilde{\theta}_t\right)\left[p_t-s-\beta E_t V_{t+1}^s\right]\right\}. \tag{24}
$$

I solve for the seller's optimal price in the next subsection.

I parameterize the model by assuming distributions for $F(\cdot)$, the distribution of idiosyncratic match quality, and $G(\cdot)$, the noise in the signal. I assume that $F(\varepsilon_m)$ is a uniform distribution on $[\underline{\varepsilon},\bar{\varepsilon}]$ with a mass point of mass χ at $\bar{\varepsilon}$. In Figure 3, the demand curve for below average priced homes is very flat., which implies a very low density of $F(\cdot)$ at the margin. If there were not a mass point at the top of $F(\cdot)$, the low density would imply a very large upper tail conditional expectation $E\left[\varepsilon-\varepsilon^*|\varepsilon>\varepsilon^*\right]$, which in turn implies a very high value of future search to buyers. Adding a mass point allows me to control the value of future search and obtain a realistic buyer

Figure 5: The Concave Demand Curve in the Model

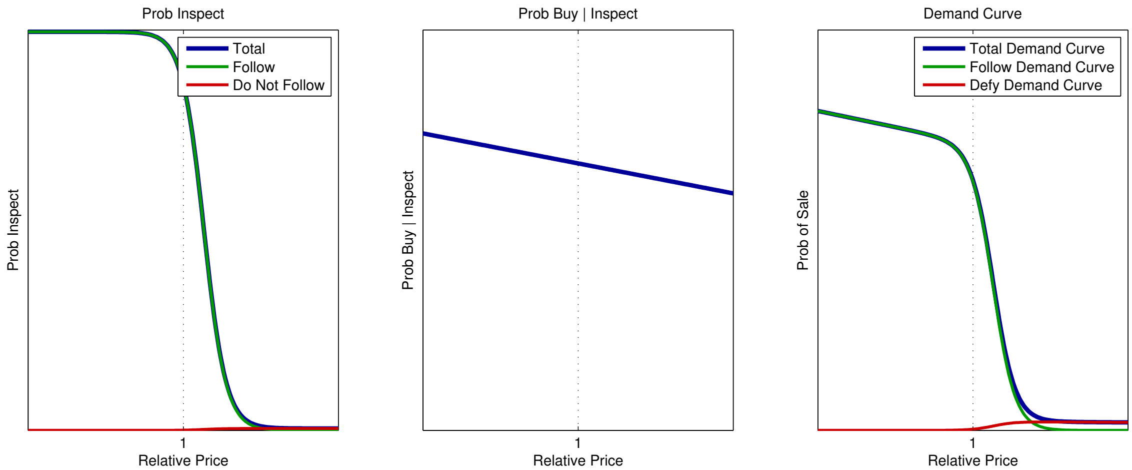

Notes: The figures are generated using calibration described in Section 5. All probabilities are calculated assuming all other sellers are setting the steady state price and considering the effect of a unilateral deviation.

search cost. In practice, this means that there are many buyers who like the property, many who do not, and a few in between. For $G(\eta_{h,t})$, I choose a logistic with mean zero and variance $\sigma^2 \frac{\pi^2}{3}$. These two distributions are both flexible and tractable, and as I show below they provide a close fit to the microdata. I also assume that the matching function is Cobb-Douglas $q^m(\theta) = \xi^m \theta^\gamma$, as is standard in the search literature, with $\xi^f > \xi^d$. While these assumptions matter somewhat for the precise quantitative predictions of the model, they are not necessary for the intuitions it illustrates.

This setup leads to a locally concave demand curve with considerable curvature in the neighborhood of the average price. There are many fewer buyers in the d submarket than the f submarket in equilibrium. At below average prices, the house receives a good signal with near certainty and the demand curve is dominated by the trade-off between idiosyncratic match quality ε_m and price, so demand is less elastic. At prices just above the average, the demand curve is dominated by the fact that perturbing the price affects whether the house gets a good signal or a bad signal, in which case it ends up in a market with few buyers who match with less efficiency. To illustrate this, Figure 5 shows the shapes of the probability of inspection $q\left(\theta_t^f\right)\left(1 - G\left(p_t - E_{\Omega_t}[p_t] - \mu\right)\right) + q\left(\theta_t^d\right) G\left(p_t - E_{\Omega_t}[p_t] - \mu\right)$, the probability of purchase conditional on inspection $1 - F\left(\varepsilon^*\left(p_t\right)\right)$, and the overall demand curve faced by sellers $d\left(p_t, \Omega_t, \theta_t\right)$, equal to the product of the first two panels. (Note that the axes are swapped from the traditional Marshallian supply and demand diagram in order to be consistent with the empirical analysis in Section 3.)

4.3 Optimal Price Setting

Sellers do not internalize that their choice of p_t affects the average price, which they treat as given. I focus on a symmetric equilibrium, and although the seller's problem is not globally concave I focus on an interior optimum and later check that the interior optimum is the global optimum by simulation. Seller optimization thus implies:

Lemma 2. *The seller's optimal list price at the interior optimum is:*

$$
\begin{aligned}
p_t &= s + \beta E_t V_{t+1}^s + E_t \left[\frac{-d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\frac{\partial d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{\partial p_t}} \right] \\
&= s + \beta E_t V_{t+1}^s + \frac{1}{\frac{f(\varepsilon_t^*)}{1-F(\varepsilon_t^*)} + \frac{g\left(p_t - E_{\Omega_t}[p_t] - \mu\right)}{1-G\left(p_t - E_{\Omega_t}[p_t] - \mu\right)} - \frac{g\left(p_t - E_{\Omega_t}[p_t] - \mu\right)}{1-G\left(p_t - E_{\Omega_t}[p_t] - \mu\right)} \frac{1}{G\left(p_t - E_{\Omega_t}[p_t] - \mu\right)} \frac{d^d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}},
\end{aligned}
\tag{25}
$$

where d^d is defined by (14) and d is defined by (15). With a logistic distribution for $G(\cdot)$, this simplifies to:

$$
p_t = s + \beta E_t V_{t+1}^s + \frac{1}{\frac{f(\varepsilon_t^*)}{1-F(\varepsilon_t^*)} + \frac{1}{\sigma} G\left(p_t - E_{\Omega_t}[p_t] - \mu\right) - \frac{1}{\sigma} \frac{d^d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}{d\left(p_t, \Omega_t, \tilde{\theta}_t\right)}}.
$$

Proof. See Appendix D. □

Sellers have monopoly power due to costly search. The optimal pricing problem they solve is the same as that of a monopolist facing the demand curve d except that the marginal cost is replaced by the seller's outside option of searching again next period. The optimal pricing strategy is a markup over the outside option $s + \beta V_{t+1}^s$. In equation (25) it is written as an additive markup equal to the reciprocal of the semi-elasticity of demand, $\frac{-d(p_t, \Omega_t, \theta_t)}{\frac{\partial d(p_t, \Omega_t, \theta_t)}{\partial p_t}}$. The semi-elasticity, in turn, is equal to the sum of the hazard rates of the idiosyncratic preference distribution $F(\cdot)$ and the distribution of signal noise $G(\cdot)$ adjusted for the share of sales that occur in the d submarket.

This creates a local strategic complementarity in price setting because the optimal price depends on relative price $p_t - E_{\Omega_t}[p_t]$ through the hazard rate of the signal $G(\cdot)$. In particular, the elasticity of demand rises as relative price increases, causing the optimal additive markup to fall. The markup thus pushes sellers to set prices close to those of others.

However, in a rational expectations equilibrium in which all sellers can set their price flexibly, all sellers choose the same list price and $p_t = E_{\Omega_t}[p_t]$, so there is no relative price to affect the markup. A shock to home values thus causes list price to jump proportionally to the seller's outside option, and there is no momentum. In the terminology of Ball and Romer (1990), concave demand is a real rigidity that only amplifies nominal rigidities. Consequently, I introduce a small amount of insensitivity of prices to generate some initial momentum.

4.4 Source of Insensitivity: A Small Fraction of Rule-of-Thumb Sellers

There are several sources of initial price insensitivity that likely exist in housing markets. These include price adjustment costs and infrequent adjustment more generally, information frictions, and backward looking sellers. In this paper, I focus on a small fraction of rule-of-thumb sellers (a previous working paper draft considered staggered pricing, which was amplified substantially but too small of an initial friction to generate the momentum in the data). I leave the analysis of other sources of insensitivity to future work.

Since Case and Shiller (1987), sellers with backward-looking expectations have been thought to play an important role in housing markets. Previous models assume that all agents have backward-looking beliefs (*e.g.*, Berkovec and Goodman, 1996), but some observers have found the notion that the majority of sellers are non-rational unpalatable given the financial importance of housing transactions for many households. Some fraction of sellers, however, may not find it worthwhile to scrutinize current market conditions due to information costs. I introduce a small number of rule-of-thumb sellers, as in Campbell and Mankiw (1989), and assess quantitatively what fraction of sellers is needed to be non-rational to explain the momentum in data, similar to Gali and Gertler (1999).

I assume that at all times a fraction $1-\alpha$ of sellers are of type R (rational) and set their list price p_t^R rationally according to Lemma 2 and (25) but a fraction α of sellers are of type E (extrapolator) and use a backward-looking rule of thumb to set their list price p_t^E. Specifically, they follow an AR(1) rule:

$$p_t^E = \frac{p_{t-2}+p_{t-3}+p_{t-4}}{3} + \psi\left(\frac{p_{t-2}+p_{t-3}+p_{t-4}}{3} - \frac{p_{t-5}+p_{t-6}+p_{t-7}}{3}\right), \tag{26}$$

where p_t is the transaction-weighted average price at time t:

$$p_t = \frac{\alpha d_t^E p_t^E + (1-\alpha)\, d_t^R p_t^R}{\alpha d_t^E + (1-\alpha)\, d_t^R}. \tag{27}$$

I use three-month lag to match the lag with which house price indices are released.[23] Such an AR(1) rule is a common assumption in models with backward-looking expectations and in the New Keynesian literature and is frequently motivated by limited knowledge, information costs, and extrapolative biases (*e.g.*, Hong and Stein, 1999; Fuster et al. 2010). I microfound such a rule using a model of limited information and near rationality in Appendix D.

For calibration, I asume that ψ that is attenuated relative to what one would find if one ran a quarterly AR(1) in the model environment. This is consistent with Case et al. (2012), who survey home buyers for four metropolitan areas from 2003 to 2011 and show that the average predicted amount of price appreciation at a one-year horizon is approximately 43 percent of the actual amount of appreciation. An attenuated AR(1) coefficient is also consistent with psychological theories in which agents overweight and “anchor” on recent observable prices (see Barberis et al., 1998).

I make two additional assumptions for tractability and parsimony that are not crucial for the results. First, I assume that regardless of whether rational or backward-looking sellers sell faster, inflows adjust so that α of the active listings are houses owned by backward-looking sellers at all

[23] I use three-month averages to correspond to how price indices like the closely watched Case-Shiller index are constructed and to smooth out saw-tooth patterns that emerge with non-averaged multi-period lags. A shorter AR(1) lag would require more backward-looking sellers to match the data.

times. Second, I assume that entry occurs according to the threshold rules (9) and (10) using rational value functions. I also assume that all sellers of each type are alike, and I continue to assume a symmetric equilibrium. To conserve notation, S_t refers to the total number of sellers, but V_t^s is now the value function for the rational sellers. V_t^s remains as in equation (24), and the value function of a buyer remains (23), but now there are two prices p_t^E and p_t^R in the market.

4.5 Equilibrium

I add a stochastic shock process to the model to examine its dynamic implications. The propagation mechanism for momentum does not qualitatively depend on any particular shock. However, the positive correlation between price and volume in the data implies that demand-side shocks dominate. Although the particular type of demand shock introduced to the model is not important for the results, I use a shock to the flow utility of being a renter u that changes the relative value of homeownership for potential entrants. This takes a cue from Wheaton and Lee (2009), who show that changes in the frequency of transitions between renting and owning due to credit conditions are a precipitating shock for housing cycles. An example of such a shock would be a change in credit standards for new homeowners. I implement the shock by assuming that $u_t = u + x_t$, where x_t is an AR(1) process understood by the forward-looking agents:

$$x_t = \rho x_{t-1} + \eta_t \text{ with } \eta_t \sim N\left(0, \sigma_\eta^2\right) \text{ and iid.} \tag{28}$$

An equilibrium with a fraction α of backward-looking sellers is then defined as:

Definition 1. Equilibrium with a fraction α of backward-looking sellers is a set of prices p_t^i, demands $d\left(p_t^i, \Omega_t, \tilde{\theta}_t\right)$, and purchase cutoffs $\varepsilon_t^{*,i}$ for each type of seller $i \in \{E, R\}$, a transaction-weighted average price p_t, rational seller, buyer, homeowner, and renter value functions V_t^s, V_t^b, V_t^h, and V_t^r, a probability that buyers follow the signal ϕ_t, entry cutoffs c_t^* and k_t^*, stocks of each type of agent B_t, S_t, H_t, and R_t, and a shock to the flow utility of renting x_t satisfying:

1. Optimal pricing for rational sellers (25) and the pricing rule (26) for backward-looking sellers, which depends on lagged transaction-weighted average prices (27);
2. Optimal purchasing decisions by buyers: $\varepsilon_t^{*,i} = p_t^i + b + \beta V_{t+1}^b - V_t^h$;
3. The demand curve for each type of seller $i \in \{E, R\}$ in the f submarket (13), the d submarket, (14), and the aggregate (15), all of which result from buyer search behavior;
4. Optimal entry decisions by homeowners and renters who receive shocks (9) and (10);
5. The value functions for buyers (23), rational sellers (24), homeowners (11), and renters (12);
6. The laws of motion for buyers (16) and sellers (17) and the closed system conditions for homes (18) and people (19) that implicitly define the laws of motion for homeowners and renters;

7. Buyers are indifferent across markets (22);

8. All agents have rational expectations that $u_t = u + x_t$ where the shock x_t evolves according to the AR(1) process (28).

The model cannot be solved analytically, so I simulate it numerically using a log-cubic approximation pruning higher order terms as in Kim et al. (2008) around a steady state described in Appendix D in which $x_t = 0 \;\forall\; t$ implemented in Dynare (Adjemian et al., 2013).[24] The Appendix shows that the impulse responses are similar in an exactly-solved model with a permanent and unexpected shock.

5 How Much Can Concave Demand Amplify Momentum?

To quantitatively assess the degree to which concave demand curves amplify house price momentum, this section calibrates the model to the empirical findings presented in Section 3 and a number of aggregate moments.

5.1 Calibration and Estimation

In order to simulate the model, 26 parameters listed in Table 7 must be set. This section describes the calibration procedure and targets, with details deferred to Appendix E.

Three components of the calibration control the shape of the demand curve and thus have a first-order impact on momentum: the local density of the idiosyncratic quality distribution $F(\cdot)$ controls the elasticity of demand for low-priced homes that are certain to be visited; σ, the logistic variance parameter of the signal, controls how much the elasticity of demand changes as relative price increases; and μ, the threshold for being overpriced, controls where on the curve the average price lies. The other parameters affect momentum mainly through equilibrium feedbacks and largely have a second order effect on momentum. Consequently, the first step of the calibration estimates the parameters that control the local density of $F(\cdot)$, μ, and σ to match the instrumental variable micro estimates presented in Section 3. The second step calibrate the rest of the model to match steady state and time series aggregate moments.

For the first step, I approximate the probability of sale as a function of a few key parameters, the relative list price I observe in the data, and a fixed effect that absorbs the aggregate market conditions so that the model can be directly compared to my empirical specification. This allows me to approximate the model with the heterogeneity in the data out of steady state for the purposes of calibration and then conduct dynamic simulations with the heterogeneity suppressed to maintain a tractable state space. Specifically, Appendix D shows that the probability of sale at the time the

[24] Because of the mass point in the $F(\cdot)$ distribution, the model is not smooth. However, a perturbation approach is appropriate because the mass point at $\bar{\varepsilon}$ is virtually never reached (less than 0.01 percent of the time in simulations).

list price is posted can be approximated as:

$$d\left(p_t - E_{\Omega_t}\left[p_t\right]\right) \approx \kappa_t \left(1 - F\left(\varepsilon^*_{mean} + p_t - E_{\Omega_t}\left[p_t\right]\right)\right) \times \left[\begin{array}{c} \left(\frac{\phi_{mean}}{E_{\Omega_t}\left[1-G\left(p_t - E_{\Omega_t}[p_t]-\mu\right)\right]}\right)^{\gamma}\left[1 - G\left(p_t - E_{\Omega_t}\left[p_t\right] - \mu\right)\right] \\ +\frac{\xi^d}{\xi^f}\left(\frac{(1-\phi_{mean})}{E_{\Omega_t}\left[G\left(p_t - E_{\Omega_t}[p_t]-\mu\right)\right]}\right)^{\gamma} G\left(p_t - E_{\Omega_t}\left[p_t\right] - \mu\right) \end{array}\right], \tag{29}$$

where κ_t is a fixed effect that summarizes the state of the market at time t, and ε^*_{mean} and ϕ_{mean} are the mean values of these variables over the cycle. Appendix D also explains why the approximation error is small.

Given data on $p_t - E_{\Omega_t}\left[p_t\right]$, a parameterization for $F(\cdot)$, σ, μ, $\frac{\xi^d}{\xi^f}$, ε^*_{mean}, and κ_t, I can then solve for ϕ_{mean} by approximating it by its steady state value which can be found using a steady-state version of (22). (29) can then be used to simulate a demand curve that can be compared to the IV binned scatter plot.

This approach requires assuming values for several parameters. I assume a value for $\varepsilon^*_{mean} =$ \$100k, which I show in robustness tests is a normalization that has no impact on the economics of the model. I also assume a value of $\xi^d/\xi^f = 1/2$. This is a parameter that has no analog in the data, and I show in robustness checks that it is of minimal importance for my quantitative results. I choose a value of 1/2 to limit the degree of non-convexity in the model.[25] Recalling that $F(\cdot)$ is a uniform distribution on $[\underline{\varepsilon}, \bar{\varepsilon}]$ with a point of mass χ at $\bar{\varepsilon}$, I introduce two additional moments that along with the assumed ε^*_{mean} and density pin down the three parameters of $F(\cdot)$. These are the average fraction of home inspections that lead to a purchase $1 - F\left(\varepsilon^*_{mean}\right)$ and the average mean excess function $E\left[\varepsilon - \varepsilon^*_{mean} | \varepsilon > \varepsilon^*_{mean}\right]$. I set the average fraction of home inspections that lead to a purchase to 1/10 to match buyer surveys from National Association of Realtors surveys analyzed by Genesove and Han (2012). I choose a target for the mean excess function to match a target for the buyer flow cost of $b = -\$10,000$. There is little evidence on costs incurred by buyers, so I set the search cost equal to that for sellers described below. I show in robustness checks that this calibration target does not affect my results on momentum, which are unchanged with a search cost of \$5,000.

Given these assumed parameters, I use equation (29) to calibrate the local density of $F(\cdot)$, μ, and σ to 25 ordered pairs (p_b, d_b) corresponding to the log relative markup plus the mean log price in the market and probability of sale within 13 weeks for each of 25 bins b of the distribution of the relative markup. The data is identical to Figure 3, except I drop 2.5 percent of the data from each end instead of 1 percent to minimize the importance of outliers. I solve for κ_t to match the average probability of sale, and use the steady state version of (22) and (29) to simulate $d(p_b)$ in the model for each p_b. Because the zero point corresponding to the average price is not precisely

[25] As discussed above, the seller's problem is non-convex. If ξ^d is high enough relative to ξ^f, sellers will have an incentive to deviate and set a price of $\bar{\varepsilon}$ when prices are rising and buyers visit rarely. Appendix D shows numerically that with $\xi^d/\xi^f = 1/2$ this happens less than .001 percent of the time and is thus not an issue. Empirically, I rule out a substantially higher ξ^d because the stronger incentive to deviate would generate asymmetries in house prices not present in the data.

Figure 6: Model Fit Relative to Instrumental Variable Estimates

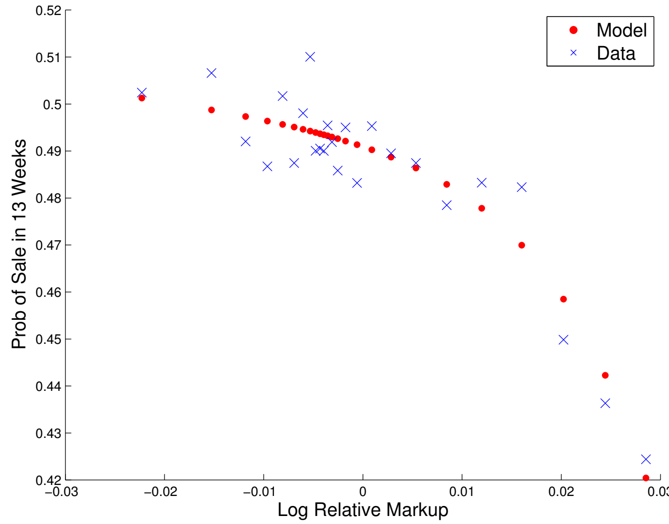

Notes: The Xs are the binned scatter plot from the IV specification with 2.5% of the data from each end excluded to reduce the effects of outliers. The dots are the simulated probabilities of sale at each price calculated using (29) and approximating ϕ_t by its steady state value using (22) as described in the text.

estimated and depends on the deadline used for a listing to count as a sale, I choose the average price so that the elasticity of demand implies a monthly seller search cost of approximately $10,000 based on evidence from Genesove and Mayer (1997) and Levitt and Syverson (2008) described in Appendix E.[26] I conduct robustness checks on this figure below. Conditional on the average price, the best fit $(density, \sigma, \mu)$ is chosen to minimize the mean squared error $\Sigma_b \left(d_b - d^{3\ month}\left(p_b\right)\right)/25$ where $d^{3\ month}(\cdot)$ is a simulated 3-month sale probability based on (29).

Figure 6 shows the IV binned scatter plot in Xs and the model's predicted $d\left(p_b\right)$ for the $(density, \sigma, \mu)$ that minimize the distance between the model and the data in circles. The fit suggests that the demand curve in the calibrated model captures the curvature in the data well.

The second step in the calibration is to match a number of aggregate steady state and stochastic moments given $\bar{\varepsilon}$, $\underline{\varepsilon}$, χ, σ, and μ and the other calibration targets from the first step.

I set the remaining parameters to match steady state moments listed in the first three panels of Table 6 as detailed in Appendix E. These targets are either from other papers or are long-run averages for the U.S. housing market, such as the homeownership rate, the average amount of time between moves for buyers and renters, and the average time on the market for buyers and sellers. A few parameters are based on limited data or assumed, and the results are not sensitive to the chosen values as shown in Section 5.4. I also match three time series moments as indicated by the bottom panel of Table 6. The monthly persistence of the shock is set to match the persistence of local income shocks as in Glaeser et al. (2014). The final parameters are set to match the

[26]The seller search cost is likely large because of the nuisances and uncertainties involved and the need to move quickly. Another factor, highlighted by Anenberg and Bayer (2014), is the high cost of simultaneously holding two homes, which pushes households to sell quickly before buying.

Table 6: Calibration Targets

Steady State Parameter or Moment	Value	Source / Justification
	Parameters	
γ (Matching Function Elasticity)	.8	Genesove and Han (2012)
L (Prob. Stay in MSA)	.7	Anenberg and Bayer (2013)
	Aggregate Targets	
Annual Discount Rate	5%	
Time on Market for Sellers	4 Months	Approx average parameter value in literature
Time on Market for Buyers	4 Months	$\approx$ Time to sell (Genesove and Han, 2012)
Homeownership Rate	65%	Long run average, 1970s-1990s
Time in House For Owner Occupants	9 Years	American Housing Survey, 1997-2005
Time Between Moves for Renters	29 Months	American Housing Survey, 1997-2005
c^* (Cost Marginal H Pays to Avoid Move)	\$37.5k	Moving cost 5% of price (Haurin & Gill, 2002)
k^* (Cost Marginal R Pays to Avoid Buying)	-\$20k	Tax benefit of owning 29 months (Poterba & Sinai, 2008)
Prob Purchase | Inspect	0.1	Mean buyer visits 10 homes (Genesove and Han, 2012)
	Assumed Values	
Time Between Shocks for Homeowners	29 Months	Same as renter
Steady State Price	\$732k	Average transaction price in IV sample
h (Flow Utility of Homeowner)	\$6.5k	2/3 of house value from expected flow util
ξ^d/ξ^f	0.5	Limited Incentive to "Fish"
ε^* in steady state	\$100k	Normalization
b (Flow Utility of Buyer)	$\approx$-\$10k	Same as seller
s (Flow Utility of Seller)	$\approx$-\$10k	Genesove and Mayer (1997), Levitt and Syverson (2008)
	Time Series Moments	
SD of Annual Log Price Changes	.065	CoreLogic national HPI adjusted for CPI, 1976-2013
ρ (Monthly Persistence of AR1 Shock)	.990	Persistence of income shocks (Glaeser et al., 2014)
Price Elasticity of Seller Entry	.878	CoreLogic, Census, and NAR, 1976-2013

standard deviation of annual log price changes and the elasticity of seller entry with respect to price in stochastic simulations.[27]

Finally, I set the AR(1) coefficient ψ to 0.4 following evidence from Case et al. (2012). Using surveys of home buyers they show that regressing realized annual house price appreciation on households' *ex-ante* beliefs yields a regression coefficient of 2.34. I use this survey evidence to calibrate the beliefs of the backward-looking sellers by dividing the approximate regression coefficient one would obtain in quarterly simulated data (approximately 0.94) by their coefficient.

I adjust α and recalibrate the model until the impulse response to the renter flow utility shock matches the matches the 36 months of positively autocorrelated price changes in the AR(5) impulse

[27] Because the stock of buyers is not observed, I cannot similarly calibrate for the buyer entry elasticity. Consequently, I assume the density of buyer entry costs is the same as the density of seller entry costs. Seller entry tends to track volume, so buyer entry cannot have a substantially different density.

Table 7: Calibrated Parameter Values for Rule of Thumb Model

Parameter	Interpretation	Value	Parameter	Interpretation	Value
β	Monthly Discount Factor	0.996	V^0	NPV of Leaving MSA	\$3,921k
γ	Matching Fn Elasticity	0.8	h	Flow Util of H	\$6.5k
ξ^f	Matching Fn Efficiency, Follow	2.570	u	Flow Util of R	\$2.084k
ξ^d	Matching Fn Efficiency, Defy	1.285	b	Flow Util of B (search cost)	-\$9.960k
λ^h	Monthly Prob H Moving Shock	0.035	s	Flow Util of S (search cost)	\$-10.157k
λ^r	Monthly Prob R Moving Shock	0.035	$\bar{\varepsilon}$	Upper Bound of F Dist	\$215.7k
$\bar{c}$	Upper Bound, H Entry Cost Dist	\$458k	$\underline{\varepsilon}$	Lower Bound of F Dist	-\$7,915k
$\underline{c}$	Lower Bound, H Entry Cost Dist	-\$1,108k	χ	Weight of Mass Point at $\bar{\varepsilon}$	0.087
$\bar{k}$	Upper Bound, R Entry Cost Dist	\$407k	σ	Variance Param of Signal Noise	7.083
$\underline{k}$	Lower Bound, R Entry Cost Dist	-\$1,159k	μ	Threshold for Signal	13.307
N	Population	1.484	σ_η	SD of Innovations to AR(1) shock	0.270
L	Prob Leave MSA	0.7	ρ	Persistence of AR(1) shock	0.990
ψ	AR(1) Param in Rule of Thumb	0.4	α	Fraction Backward Looking	0.375

Notes: The calibration is monthly.

response estimated on the CoreLogic national house price index in Section 2. Table 7 summarizes the calibrated parameter values.

5.2 Amplification of Momentum in the Calibrated Model

The model reaches a 36-month impulse response as in the data when 37.5 percent of sellers are backward looking.[28] To assess the degree of amplification of momentum in the calibrated model, Figure 7 compares the model to the estimated impulse response from the data, to a model without backward-looking sellers, and to a model with backward-looking sellers but without concave demand. To do so, it presents the impulse response to the model shock to the flow utility of renters. The impulse response is computed as the average difference between two sets of simulations that use the same sequence of random shocks except for one period in which an additional standard deviation shock is added. I contrast the model impulse responses with the impulse response to a one standard deviation price shock to the quarterly CoreLogic national house price index estimated from an AR(5), as in Section 2.

The solid line shows the model impulse response with 37.5 percent backward-looking sellers, while the dotted line shows the estimated AR(5) impulse response from the CoreLogic data with the 95 percent confidence interval shown as thin grey lines.[29] The two impulse responses are similar, although the model impulse response grows less at the beginning and is slightly more S-shaped than the AR(5) impulse response. This is the case because backward-looking sellers are insensitive to the shock for several months and so the growth rate of prices takes a few months to accelerate.

[28] With additional initial sources of price insensitivity it is likely that the 37.5 percent figure could be reduced even further. Intuitively, concave demand creates an incentive to price close to others that interacts with *any* source of heterogenous price insensitivity to create additional momentum.

[29] Appendix F shows the impulse responses to shocks that cause price declines are symmetric, as in the data.

Figure 7: Price Impulse Response Functions: Model and Data

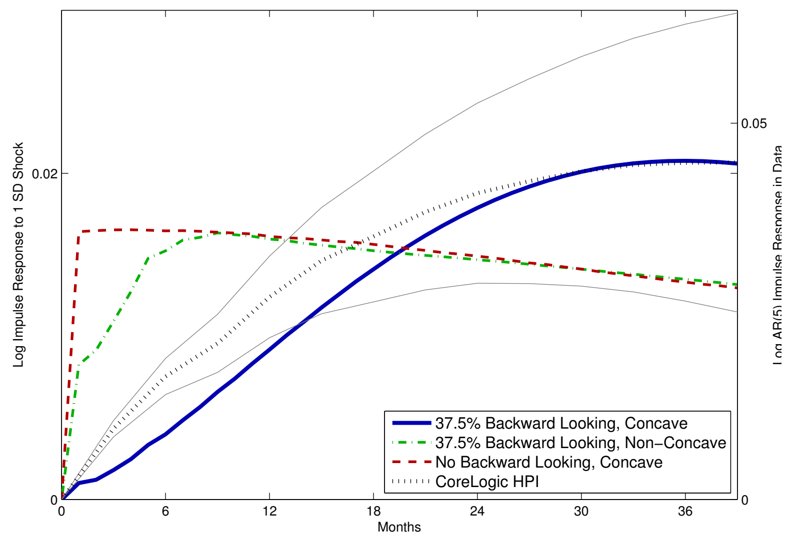

Notes: The figure shows the impulse responses to a one standard deviation shock to the flow utility of renting in the backward-looking model with and without concavity as well as the fully-flexible model. The calibration removing concave demand maintains the average elasticity of demand and is detailed in Appendix D. Also shown in the figure in the dotted black line and with grey 95% confidence intervals and on the right axis is the impulse response to a one standard deviation price shock estimated from a quarterly AR(5) for the seasonally and CPI adjusted CoreLogic national house price index for 1976-2013, as in Figure 1. Simulated impulse responses are calculated by differencing two simulations of the model from periods 100 to 150, both of which use identical random shocks except in period 101 in which a one standard deviation negative draw is added to the random sequence, and then computing the average difference over 100 simulations.

The dashed line shows the model without backward-looking sellers but with concave demand. In this case, the optimal price jumps nearly all the way to its peak level and then gradually mean reverts. There is almost no momentum because there is no initial stickiness for the strategic complementarity to amplify.

Finally, the dash-dotted line shows a model without concave demand calibrated so that the elasticity of demand for the average house is the same as in the model with concave demand, as detailed in Appendix D. Prices jump over half of the way to their peak value upon the impact of the shock and continue to rise for five months, at which point they begin to mean revert. One can generate a 36 month impulse response without concave demand, but this requires 73 percent of sellers to be backward looking.

Far fewer backward-looking sellers are needed to match the data with concave demand because the strategic complementarity creates a two-way feedback. When a shock occurs, the backward-looking sellers are not aware of it for several months, and the rational sellers only slightly increase their prices so that they do not dramatically reduce their chances of attracting a buyer. When the backward-looking sellers do observe increasing prices, they observe a much small increase than in the non-concave case and gradually adjust their price according to their AR(1) rule, reinforcing the incentives of the rational sellers not to raise their prices too quickly.

Table 8: Robustness to Calibration Targets

Calibration	α for 36-Month IRF	Calibration	α for 36-Month IRF
Baseline	0.375	$h = \$7.5\text{k}$	0.375
$c^* = \$0$, $k^* = \$0$	0.375	$\xi^d/\xi^f = 3/4$	0.385
$\varepsilon^*_{mean} = \$50\text{k}$	0.375	$b = -\$5\text{k}$	0.375
$\varepsilon^*_{mean} = \$150\text{k}$	0.375	$s = -\$5\text{k}$	0.410
$h = \$5.5\text{k}$	0.375	$s = -\$5\text{k}$ and $b = -\$5\text{k}$	0.410

Notes: Each line alters a single calibration target from the baseline, fully recalibrates the model, and reports the α necessary to generate a 36-month impulse response.

5.3 Loss From Failure to Optimize

In equilibrium, the loss the backward-looking sellers experience due to their failure to optimize is typically small. To calculate this loss, I define a value of being a backward-looking seller analogously to (24) as:

$$V_t^{s,E} = s + \beta V_{t+1}^{s,E} + d\left(p_t^E, \Omega_t, \tilde{\theta}_t\right)\left(p_t^E - s - \beta V_{t+1}^{s,E}\right). \tag{30}$$

The median loss from being an extrapolative seller $V_t^s - V_t^{s,E}$ is \$2,674 or about 0.37 percent of the average sales price, while the mean loss is \$5,449 or 0.74 percent of the average sales price. This implies a median loss from posting an extrapolative price per month of about \$670. While the loss rises in down markets when backward-looking sellers have high prices and take much longer to sell, the on-average small loss illustrates that the explanation for momentum I propose is not one of mass non-rationality by a few—which Glaeser and Nathanson (2015) have criticized—but rather one of minor non-rationality by a few.

5.4 Robustness

Several of the targets used in the calibration use assumed values or have limited empirical support. Table 8 shows that the degree of momentum generated by the model is not sensitive to these calibration targets, which are largely normalizations. In particular, it shows that the assumed calibration targets for c^* in steady state, k^* in steady state, ε^*_{mean}, h, and b have absolutely no effect on the fraction of backward-looking sellers needed to generate a 36-month impulse response. ξ^d/ξ^f does have a small effect, as if this is increased from 0.5 to 0.75 one percent more backward-looking sellers are needed to generate a 36-month impulse response.

The most important parameter is the target for s. If this falls from -\$10k to -\$5k, 41 percent of sellers need to be backward-looking to generate a 36-month IRF. This is because the markup in steady state falls with a lower s and thus the markup has less ability to fluctuate counter-cyclically to generate momentum.

Figure 8: Impulse Response Functions of Non-Price Variables

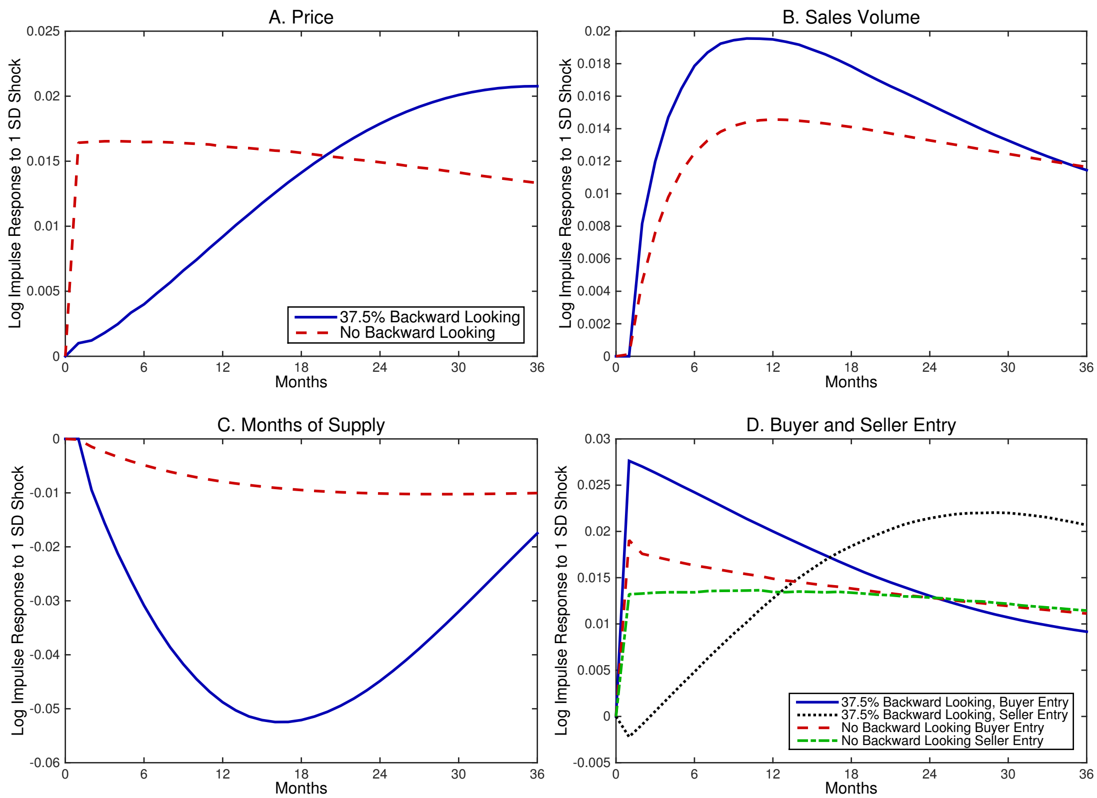

Notes: Each panel plots the indicated impulse response to a one standard deviation shock for the no backward-looking and 37.5 percent backward-looking models with identical calibrations except for α. Simulated impulse responses are calculated by differencing two simulations of the model from periods 100 to 150, both of which use identical random shocks except in period 101 in which a one standard deviation negative draw is added to the random sequence, and then computing the average difference over 100 simulations.

6 Can Momentum Explain the Cyclicality of Price and Inventory?

This section argues that momentum helps explain the excess volatility of inventory and the housing Phillips curve introduced in Section 2. These features arise because some buyers and sellers re-time their sales and purchases in light of predictable price changes, creating volatility in the inventory of homes for sale that is correlated with price changes.

6.1 Intuition From Impulse Responses

To illustrate the intuition behind these results, Figure 8 shows the impulse responses of price, sales volume, months of supply, and buyer and seller entry in a model without backward-looking sellers (dashed) and in a model with 37.5 percent backward-looking sellers (solid). Recall that the shock reduces the value of being a renter and increases the incentives to enter the market to buy.

Without backward-looking sellers, price jumps immediately and gradually returns towards the stochastic steady state, so there is not a strong incentive to buy or sell today relative to tomorrow. Buyer entry and seller entry, shown in panel D in dashed and dash-dotted lines, respectively, both jump on impact due to the change in the relative value of homeownership and the elevated house price. Buyer entry is slightly higher for 18 months as the ratio of buyers to sellers slowly rises until it settles on a stable transition path to the stochastic steady state. The slow adjustment of market tightness, in turn, causes a gradual increase in volume and decrease in months of supply.

By contrast, the momentum generated with a small fraction of backward-looking sellers makes price changes predictable. This creates an incentive for potential buyers on the margin of entering to enter today rather than waiting and for sellers on the margin of entering to wait to do so until prices rise.[30] The entry responses are shown as the solid (buyer entry) and dotted (seller entry) lines in panel D. The gap between these lines is substantial, particularly relative to the no-backward-looking case. The excess buyer entry causes inventory to fall substantially for 18 months as shown in panel C, which is why inventory becomes so much more volatile. Because inventory adjusts rapidly while prices gradually adjust, this creates a negative relationship between price changes and inventory levels. Indeed, inventory begins to mean revert when buyer entry falls below seller entry, which happens almost exactly as the speed of price adjustment stops rising and starts falling.

6.2 The Excess Volatility of Inventory

In housing search models without momentum, inventory is too smooth. The top half of Table 9 shows the standard deviation of annual changes of log price, log volume, and log months of supply in the data and three versions of the model. In the model without backward-looking sellers, months of supply is about a sixth as volatile and sales volume less a third as volatile as the data.

This is not unique to my particular model. In a broad class of housing search models, combining the steady-state value of being a seller with the steady-state price and differentiating yields a condition of the form:

$$\frac{dp}{d\Pr\left[Sell\right]} = \frac{Seller\ Surplus}{1-\beta}.$$

This steady state response illustrates that if the seller surplus is not miniscule, price is very sensitive to the probability of sale, which is mechanically related to inventory and sales volume.[31] The relative volatility of volume and inventory are also low due to the gradual dynamic adjustment of market tightness as shown in the impulse responses.

[30] In the model, this operates through the entry cutoffs c_t^* and k_t^*, which are defined by differences of value functions in equations (9) and (10). For instance, the cutoff cost for a renter to enter $k_t^* = V_t^r - V_t^b$. Because the value function of being a renter V_t^r accounts for the likelihood of getting a shock and entering as a buyer in the future, when prices are expected to rise V_t^r falls relative to V_t^b, k_t^* falls, and the mass of buyer entrants, which is proportional to $1 - K\left(k_t^*\right)$, rises.

[31] Diaz and Jerez (2013), for example, explain the relative volatilities in a model without momentum by using a calibration in which the seller surplus is 0.5% of the purchase price. Consequently, they argue that price is too insensitive and volume and time on the market are too sensitive to shocks and introduce a model with amplified price volatility. My calibration, which uses a seller surplus that is approximately 7.0% of the steady state price, implies that price is too volatile without additional frictions. Head et al. (2014) make a similar point that momentum reduces price volatility.

Table 9: Quantitative Performance of Calibrated Models

	Data	No Backward Looking	37.5% Backward Looking
SD Annual $\Delta \log$ (Real Price)	0.065	0.066	$0.065^{\diamond}$
SD Annual $\Delta \log$ (Sales)	0.143	0.055	0.076
SD Annual $\Delta \log$ (Inventory)	0.207	0.034	0.250
Regression Coefficient of	-.140	0.146	-0.231
$\log$ (Inventory) on $\Delta_{yr} \log$ (Real Price)	(.015)	(.026)	(.007)
Regression R^2	.543	0.039	0.787

Notes: $\diamond$ indicates the model is calibrated to match the data. All are statistics calculated as means of 200 random simulations of 500 years. The standard deviations of annual log changes in the model are calculated by collapsing simulated data to the quarterly level, taking logs, and reporting the quarterly standard deviations of annual differences. Inventory is measured as months of supply. The regression of log price changes on log inventory levels is as in equation (31), with inventory measured as months of supply at the midpoint of the year differenced to calculate the log change in price. The model without backward looking sellers uses the same calibration as the the 37.5% backward looking calibration except for α.

The low volatility of inventory is directly analogous to labor search models. Shimer (2005) shows that unless the employer surplus is tiny, labor search models have difficulty accounting for the volatility of unemployment because most of the response to a shock is absorbed by the wage. Here, the unemployment rate is analogous to inventory and the wage rate is analogous to price.

Like sticky wages in labor search models, momentum makes house prices adjust more slowly and slightly reduces price volatility. Quantities adjust slightly more and inventory adjusts substantially more, as shown in the impulse responses and Table 9, which shows the standard deviation of annual changes for log price, log sales, and log inventory averaged over 200 500-year simulations. Inventory is a bit too volatile in the calibrated model, although it is of the same order of magnitude as the data in contrast to the model without backward-looking sellers. The volatility of sales volume, on the other hand, falls short of the data, which suggests that other factors, such as lock in due to equity (Stein, 1995), may play a role in amplifying volume volatility. Appendix B shows that the model's strongest prediction about relative volatilities—that momentum and inventory volatility are positively correlated—is borne out in a cross-section of cities.

6.3 Housing Phillips Curve

In the data, price changes are strongly negatively correlated with inventory levels, creating a "housing Phillips curve." Without backward-looking sellers, price changes are negatively correlated with inventory changes, albeit weakly because the inventory response is delayed due to gradual entry and search frictions. With persistent but mean reverting shocks, this generates a positive correlation between price changes and inventory levels because when inventory is high, prices are low and tend to rise towards the stochastic steady state. This can be seen in the bottom half of Table 9, which shows a regression coefficient β_1 in:

$$\Delta_{t,t-4} \log(p) = \beta_0 + \beta_1 \log(MS_{t-2}) + \varepsilon, \tag{31}$$

estimated on simulated quarterly data. For a model without pricing frictions, the regression coefficient is significantly positive, although with a small R-squared.

With momentum, inventory rapidly adjusts and then mean reverts while price appreciation grows and then gradually weakens. This creates a strong negative correlation between price changes and inventory levels. Table 9 shows that with 37.5 percent backward-looking sellers, a robust negative relationship emerges. In fact, the relationship is slightly stronger than in the data with a larger regression coefficient and an R-squared of nearly 0.8. Appendix B shows that the housing Phillips curve is stronger in terms of the magnitude of β_1 and explanatory power in cities with more momentum, which is consistent with the model.

7 Conclusion

The degree and persistence of autocorrelation in house price changes is one of the housing market's most distinctive features and greatest puzzles. This paper introduces a mechanism that amplifies small frictions that have been discussed in the literature into substantial momentum. Search frictions and concave demand in relative price together imply that increasing one's list price above the market average is costly, while lowering one's list price below the market average has little benefit. This strategic complementarity induces sellers to set their list prices close to the market average. Consequently, modest price insensitivity to changes in fundamentals can lead to prolonged periods of autocorrelated price changes as sellers slowly adjust their list price to remain close to the mean.

I provide evidence for concave demand in micro data and introduce an equilibrium search model with concave demand that is calibrated to match the amount of concavity in the micro data. The amount of amplification provided by the strategic complementarity is significant: if just 37.5 percent of sellers use a backward-looking rule of thumb, the impulse response to a shock lasts for three years. Without concave demand, 73 percent of sellers would have to be backward-looking to obtain the same result. Importantly, in equilibrium the backward-looking sellers' failure to optimize is on average not very costly. The explanation for momentum proposed here is thus one of small non-rationality by the few.

Concave demand also amplifies other pricing frictions beside backward-looking sellers. In particular, it interacts with any friction that creates heterogeneity in the speed of price adjustment because the incentive to price close to the average makes sellers who would change their price quickly instead respond sluggishly. Assessing whether a model with concave demand and other frictions, such as learning about fundamentals or gradually-spreading belief disagreement, can quantitatively explain momentum without appealing to non-rationality is an important path for future research.[32]

[32] Incomplete information and learning may be particularly potent. Strategic complementarities in such a framework can cause very gradual price adjustment even if first-order learning occurs rapidly because the motive to price close

Beyond the housing market, this paper shows how a central idea in macroeconomics—that strategic complementarities can greatly amplify modest frictions—can be applied in new contexts. These contexts can, in turn, serve as empirical laboratories to study macroeconomic phenomena for which micro evidence has proven elusive. In particular, many models with real rigidities (Ball and Romer, 1990) use a concave demand curve. This paper provides new evidence that a concave demand curve in relative price is not merely a theoretical construct and can have a significant effect on market dynamics.

to others makes higher order beliefs matter. Learning about higher order beliefs is more gradual—which in turn makes price adjustment more gradual—because agents must learn not only about fundamentals but also about what everyone else has learned as in a Keynesian beauty contest.

References

Adjemian, S., H. Bastani, M. Juillard, K. Frederic, F. Mihoubi, G. Perendia, J. Pfeifer, M. Ratto, and S. Villemot (2011). Dynare: Reference Manual, Version 4.

Albrecht, J., P. A. Gautier, and S. Vroman (2014). Directed Search in the Housing Market.

Anenberg, E. (2014). Information Frictions and Housing Market Dynamics.

Anenberg, E. and P. Bayer (2014). Endogenous Sources of Volatility in Housing Markets: The Joint Buyer-Seller Problem.

Ball, L. and D. Romer (1990). Real Rigidities and the Non-Neutrality of Money. *The Review of Economic Studies 57*(2), 183–203.

Barberis, N., A. Shleifer, and R. Vishny (1998). A Model of Investor Sentiment. *Journal of Financial Economics 49*, 307–343.

Berkovec, J. A. and J. L. Goodman (1996). Turnover as a Measure of Demand for Existing Homes. *Real Estate Economics 24*(4), 421–440.

Burnside, C., S. Rebelo, and M. Eichenbaum (2015). Understanding Booms and Busts in Housing Markets.

Campbell, J. Y. and N. G. Mankiw (1989). Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence. In O. J. Blanchard and S. Fischer (Eds.), *NBER Macroeconomics Annual 1989*, Volume 4, pp. 185–246. Cambridge, MA: MIT Press.

Caplin, A. and J. Leahy (2011). Trading Frictions and House Price Dynamics. *Journal of Money, Credit and Banking 43*(7), 283–303.

Capozza, D. R., P. H. Hendershott, and C. Mack (2004). An Anatomy of Price Dynamics in Illiquid Markets: Analysis and Evidence from Local Housing Markets. *Real Estate Economics 32*(1), 1–32.

Carrillo, P. E. (2012). An Empirical Stationary Equilibrium Search Model of the Housing Market. *International Economic Review 53*(1), 203–234.

Case, K. E. and R. J. Shiller (1987). Prices of Single-Family Homes Since 1970: New Indices for Four Cities. *New England Economic Review* (September), 45–56.

Case, K. E. and R. J. Shiller (1989). The Efficiency of the Market for Single-Family Homes. *American Economic Reiwew 79*(1), 125–137.

Case, K. E., R. J. Shiller, and A. K. Thompson (2012). What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets. *Brookings Papers on Economic Activity* (Fall), 265–298.

Chen, Y. and R. W. Rosenthal (1996). Asking Prices as Commitment Devices. *International Economic Review 37*(1), 129–155.

Cho, M. (1996). House Price Dynamics: A Survey of Theoretical and Empirical Issues. *Journal of Housing Research 7*(2), 145–172.

Cutler, D. M., J. M. Poterba, and L. H. Summers (1991). Speculative Dynamics. *Review of Economic Studies 58*(3), 529–546.

Diaz, A. and B. Jerez (2013). House Prices, Sales, and Time on the Market: A Search-Theoretic Framework. *International Economic Review 54*(3), 837–872.

Fehr, E. and J.-R. Tyran (2005). Individual Irrationality and Aggregate Outcomes. *Journal of Economic Perspectives 19*(4), 43–66.

Frazzini, A. (2006). The Disposition Effect and Underreaction to News. *The Journal of Finance 61*(4), 2017–2046.

Fuhrer, J. C. (2011). Inflation Persistence. In B. Friedman and M. Woodford (Eds.), *Handbook of Monetary Economics*, Volume 3, Chapter 9, pp. 423–486. Elsevier.

Fuster, A., D. Laibson, and B. Mendel (2010). Natural Expectations and Macroeconomic Fluctuations. *Journal of Economic Perspectives 24*(4), 67–84.

Gali, J. and M. Gertler (1999). Inflation Dynamics: A Structural Econometric Analysis. *Journal of Monetary Economics 44*, 195–222.

Genesove, D. and L. Han (2012). Search and Matching in the Housing Market. *Journal of Urban Economics 72*(1), 31–45.

Genesove, D. and C. Mayer (1997). Equity and Time to Sale in the Real Estate Market. *American Economic Reivew 87*(3), 255–269.

Genesove, D. and C. Mayer (2001). Loss Aversion and Seller Behavior: Evidence From the Housing Market. *Quarterly Journal of Economics 116*(4), 1233–1260.

Glaeser, E. L., J. Gyourko, E. Morales, and C. G. Nathanson (2014). Housing Dynamics: An Urban Approach. *Journal of Urban Economics 81*, 45–56.

Glaeser, E. L. and C. G. Nathanson (2015). An Extrapolative Model of House Price Dynamics.

Gopinath, G. and O. Itskhoki (2010). In Search of Real Rigidities. In D. Acemoglu and M. Woodford (Eds.), *NBER Macroeconomics Annual*, Volume 25, pp. 261–310. Chicago: University of Chicago Press.

Guren, A. and T. McQuade (2015). How Do Foreclosures Exacerbate Housing Downturns?

Haltiwanger, J. and M. Waldman (1989). Limited Rationality and Strategic Complements: the Implications for Macroeconomics. *The Quarterly Journal of Economics 104*(3), 463–483.

Han, L. and W. C. Strange (2012). What is the Role of the Asking Price for a House.

Haurin, D. R. and H. L. Gill (2002). The Impact of Transaction Costs and the Expected Length of Stay on Homeownership. *Journal of Urban Economics 51*, 563–584.

Haurin, D. R., J. L. Haurin, T. Nadauld, and A. Sanders (2010). List Prices, Sale Prices and Marketing Time: An Application to U.S. Housing Markets. *Real Estate Economics 38*(4), 659–685.

Hausman, J. A., H. Ichimura, W. K. Newey, and J. L. Powell (1991). Identification and Estimation of Polynomial Errors-In-Variables Models. *Journal of Econometrics 50*(3), 273–295.

Head, A., H. Lloyd-Ellis, and H. Sun (2014). Search, Liquidity, and the Dynamics of House Prices and Construction. *American Economic Review 104*(4), 1172–1210.

Hong, H. and J. C. Stein (1999). A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. *The Journal of Finance 54*(6), 2143–2184.

Kim, J., S. Kim, E. Schaumburg, and C. A. Sims (2008). Calculating and Using Second-Order Accurate Solutions of Discrete Time Dynamic Equilibrium Models. *Journal of Economic Dynamics and Control 32*(11), 3397–3414.

Kimball, M. (1995). The Quantitative Analytics of the Basic Neomonetarist Model. *Journal of Money, Credit and Banking 27*(4), 1241–1277.

Klenow, P. J. and J. L. Willis (2006). Real Rigidities and Nominal Price Changes.

Krainer, J. (2001). A Theory of Liquidity in Residential Real Estate Markets. *Journal of Urban Economics 49*(1), 32–53.

Levitt, S. D. and C. Syverson (2008). Market Distortions When Agents Are Better Informed: The Value of Information in Real Estate Transactions. *The Review of Economics and Statistics 90*(4), 599–611.

Meese, R. A. and N. Wallace (1997). The Construction of Residential Housing Price Indices: A Comparison of Repeat-Sales, Hedonic-Regression, and Hybrid Approaches. *Journal of Real Estate Finance and Economics 14*, 51–73.

Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen (2012, May). Time Series Momentum. *Journal of Financial Economics 104*(2), 228–250.

Nakamura, E. and D. Zerom (2010). Accounting for Incomplete Pass-Through. *Review of Economic Studies 77*, 1192–1230.

Ngai, L. R. and S. Tenreyro (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.

Peach, R. (1983). *An Econometric Model of Investment in Single Family Homes.* Ph. D. thesis, University of Maryland.

Piazzesi, M. and M. Schneider (2009). Momentum Traders in the Housing Market: Survey Evidence and a Search Model. *American Economic Review 99*(2), 406–411.

Poterba, J. and T. Sinai (2008). Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income. *American Economic Review 98*(2), 84–89.

Shimer, R. (2005). The Cyclical Behavior of Equilibrium Unemployment and Vacancies. *American Economic Review 95*(1), 25–49.

Smets, F. and R. Wouters (2007). Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach. *American Economic Review 97*(3), 586–606.

Stein, J. C. (1995). Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects. *Quarterly Journal of Economics 110*(2), 379–406.

Stiglitz, J. (1979). Equilibrium in Product Markets with Imperfect Information. *American Economic Review 69*(2), 339–345.

Taylor, C. R. (1999). Time-on-the-Market as a Sign of Quality. *Review of Economic Studies 66*(3), 555–578.

Taylor, J. B. (1980). Aggregate Dynamics and Staggered Contracts. *Journal of Political Economy 88*(1), 1–23.

Titman, S., K. Wang, and J. Yang (2014). The Dynamics of Housing Prices.

Wheaton, W. C. (1990). Vacancy, Search, and Prices in a Housing Market Matching Model. *Journal of Political Economy 98*(6), 1270–1292.

Wheaton, W. C. and N. J. Lee (2009). The Co-Movement of Housing Sales and Housing Prices: Empirics and Theory.

Yavas, A. and S. Yang (1995). The Strategic Role of Listing Price in Marketing Real Estate: Theory and Evidence. *Real Estate Economics 23*(3), 347–368.