

Speculators and Middlemen: The Strategy and Performance of Investors in the Housing Market

Patrick Bayer

Duke University and NBER

Christopher Geissler

ISO New England

Kyle Mangum

Federal Reserve Bank of Philadelphia

James W. Roberts

Duke University and NBER

Using data from the Los Angeles area from 1988 to 2012, we study the behavior and sources of returns of individual investors in the housing market. We document the existence of two distinct investor types. The first act as middlemen, purchasing substantially below and reselling above market prices throughout the cycle, improving liquidity and the existing capital stock in the process. The second act as speculators, who primarily enter during the boom, buying and selling at essentially market prices. Neither type anticipated the housing bust. We document similar behavior by speculators and middlemen in 96 other U.S. metro areas. (*JEL* D40, D84, R30)

Received July 21, 2017; editorial decision October 26, 2019 by Editor Stijn Van Nieuwerburgh. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

We thank Ed Glaeser, Matt Kahn, Chris Mayer, Fernando Ferreira, Donghoon Lee, Andy Haughwout, Wilbert Van der Klauww, Steve Ross, Amine Ouazad, Todd Sinai, and Joe Gyourko and the seminar participants at Harvard University, University of California at Berkeley and the Haas School of Business, London School of Economics, Mannheim, Toulouse, Zurich, University of Georgia, INSEAD, CREST, Northwestern University and Kellogg Business School, Duke-ERID Housing Market Dynamics Conference, HULM conference, and the UCLA Housing Market Conference for their useful feedback on the paper. We are grateful for Duke University's financial support. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors are our own. Supplementary data can be found on *The Review of Financial Studies* web site. Send correspondence to Patrick Bayer, Duke University, 213 Social Sciences Building, 419 Chapel Drive, Durham, NC 27708-0097; telephone: 919-660-1800. E-mail: patrick.bayer@duke.edu.

The Review of Financial Studies 33 (2020) 5212–5247

© The Author(s) 2020. Published by Oxford University Press on behalf of The Society for Financial Studies. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com.

doi:10.1093/rfs/hhaa042

Advance Access publication April 6, 2020

The recent housing boom and bust in the United States affected millions of households and hindered the country's emergence from the Great Recession.¹ Housing market investors accounted for a large fraction of transactions during the boom and recent evidence suggests they causally fomented the steep runup in prices.² How did these investors behave and how did they affect the actual functioning of the housing market itself? Modern investment theory provides a variety of models that admit a wide range of investor strategies and influences on a market. However, despite the richness of investor behavior and strategies in the theory literature, little evidence exists describing how investors *actually* behave in many market settings. How well-informed real investors are, whether they behave in a manner consistent with naïve or rational decision-making, and whether the activity of various types of investors is sufficient to be quantitatively important for equilibrium price dynamics remain key, and open, empirical questions.

In this paper we use transaction-level data from the Los Angeles metro area³ between 1988 and 2012 to characterize the behavior, performance, and strategies of investors in the housing market. In so doing we marshal three key broad forms of evidence about the role of investors in this market.

First, we document that a large share of the properties that were purchased in the Los Angeles market late in the boom were bought by those already holding another property in the area, presumably as an investment property. These findings are in line with those of Haughwout et al. (2011), Bhutta (2015), and DeFusco, Nathanson, and Zwick (2017). The sheer magnitude of the activity of novice investors suggests they play a nonnegligible role in the market and supports at a basic level their use in theory.

Second, we identify different investor strategies in the housing market. For this portion of our analysis, we focus primarily on the behavior of a set of individuals that we observe reselling investment properties after short holding periods, using the colloquial name “flippers.” Relative to most financial products, tracking investment behavior in a durable good like housing is complicated by the fact that owners may invest in the good before reselling. In that event, a portion of any price appreciation is likely due to such improvements, not just any particular ability of the investor to buy cheap, or sell dear. To address the challenge this creates for understanding the sources of returns, we introduce a novel research design using properties that sell repeatedly during the study period to decompose the observed price growth during the flipper's holding period into four components: (1) the discount relative to market price at the time of purchase, (2) the premium relative to

¹ See, for example, Mayer, Pence, and Sherlund (2009), Mian and Sufi (2014), and Rognlie, Shleifer, and Simsek (2018).

² See, for example, Chinco and Mayer (2016) and Gao, Sockin, and Xiong (2017).

³ Although Los Angeles is the primary focus of the paper, we provide evidence from nearly 100 other metro areas that the patterns we document in Los Angeles generalize across a wide range of other U.S. cities.

market at the time of sale, (3) the market return during the holding period, and (4) physical improvements made to the property by flippers. Our research design distinguishes any costly improvements that a flipper may have made (which are not directly observed in the data) by measuring the extent to which any above-market appreciation that a flipper earns at sale persists through a subsequent sale of the same property.

This analysis illuminates two distinct types of investors in the housing market. The first type acts as an intermediary or *middleman* who purchases properties at prices well-below market value and resells them quickly at or above market prices. The steep discount that they receive by locating “good deals” and buying from “motivated” sellers accounts for the majority of their returns. Market timing, on the other hand, is not an important source of their returns; in fact, they operate more intensely during periods when prices are stagnant or declining and systematically target submarkets that are appreciating more slowly than the rest of the metro area.

Their behavior contrasts sharply with that of the novice investor or *speculator* that entered the market in droves during the boom. Relative to market prices, these investors do not buy at much of a discount or sell at much of a premium, suggesting that most are not inordinately skilled real estate professionals. Instead these novice investors earn most of their return through the market appreciation over the period that they hold the property. Speculators operate primarily during boom times and purchase homes in submarkets of the Los Angeles area that experience both an above average rate of appreciation in the short term (1–2 years) and a sharp decline in the intermediate term (3–5 years). After controlling for transaction costs, we find that realized returns for middlemen are nearly twice that of speculators.⁴

To further support and provide external validity for our findings about the Los Angeles housing market, we extend our framework to 96 other U.S. metro areas.⁵ We show that speculators and middlemen operate in other markets as well. Our results suggest that the investor strategies and behaviors we document in Los Angeles are representative of broader patterns in the housing market.

Third, having established distinct investor types that follow different strategies, we ask how these investor types interact with the broader housing market. Here, we focus on three channels. First, we examine whether the novice investors that purchased homes late in the boom were able to anticipate the market peak, or instead seem to exhibit trend chasing behavior. Remarkably, they continued to purchase homes at near-record rates right up to the peak, and there is no change in the rate at which they sold their existing holdings, despite a clear financial incentive to do so. Second, we show that middlemen

⁴ As we note elsewhere, because we cannot observe rental income or any particular investor’s specific tax rate, we are somewhat limited in our ability to accurately measure the precise levels of returns.

⁵ The selection of the other metro areas for analysis is mainly determined by data availability. See the discussion in Section 4.

improve market liquidity as they are more likely than a typical homeowner to buy a property not currently in use. In contrast, speculators are no more likely than a typical homeowner to buy such properties. Third, we show more direct evidence that middlemen improve the quality of the housing stock by investing in the properties they own, as measured by the frequency with which they are issued permits for construction and home improvements. Speculators, on the other hand, take out permits no more frequently than does a normal homeowner.

Taken as a whole, the various forms of evidence that we present provides a comprehensive picture of the activity of investors in the housing boom that is consistent with the behavior captured in many theoretical models of financial markets. We find no indication that the speculators that poured into the market late in the boom had access to superior information, which is consistent with the notion that many of these investors had very limited experience and may have simply been swept up in the exuberance of the boom. This is of course also consistent with famous accounts of speculative activity by Charles Mackay⁶ and Kindelberger (1978), who in his “anatomy of a typical crisis” notes that bubbles are frequently characterized by “More and more firms and households that previously had been aloof from these speculative ventures” beginning to participate in the market. Some of these speculators, to be sure, earned substantial returns via successful market timing. But the evidence that speculators were not particularly more well-informed casts doubt on the notion that they improved efficiency by transmitting any valuable information to the market. Middlemen, on the other hand, appear to operate throughout housing cycles and likely improve the market’s liquidity and capital stock. More generally, that these flippers were party to such a large number of transactions during the final few years of the housing boom not only justifies their use as a theoretical device but also suggests that their activity may be of first-order importance from a policy perspective, a topic we turn to in the paper’s conclusion.

1. Related Literature

Our paper is connected to and contributes to a number of literatures. Housing markets are a classic example of a thin market for high-valued durable goods in which sellers with heterogenous holding costs (because of, say, differing relocation or borrowing cost needs)⁷ and buyers with heterogeneous values for any given home (because of, say, its location or horizontally differentiated characteristics) must engage in costly search in order to transact. Flippers

⁶ In his well-known description of the boom and bust in the 1637 Dutch tulip market Mackay commented that at its peak, “Nobles, citizens, farmers, mechanics, seamen, footmen, maid-servants, even chimney-sweeps and old clotheswomen, dabbled in tulips” (see Mackay (1841), p. 94).

⁷ Springer (1996) finds that distressed sellers deal more quickly and sell for less than other sellers. Glower, Haurin, and Hendershott (1998) find that when a seller takes a new job, she sells faster than average, indicating a higher holding cost.

who purchase a property with plans to put the house back on the market are acclimated to such a market. We provide evidence that flippers participated in a sizable share of housing transactions during the recent boom and bust. In this way, our results complement those in Haughwout et al. (2011) and Bhutta (2015). Though the absolute magnitude of the overall housing ownership by investors may not be large, as Piazzesi and Schneider (2009) argue, even a small number of optimistic investors that participate in housing transactions can meaningfully affect housing prices. Indeed, recent work has argued that real estate investors causally increased housing prices during the boom (Chinco and Mayer 2016; Gao, Sockin, and Xiong 2017).

How did these investors behave? Modern finance theory admits a wide range of possible answers to this question and our paper explores this topic in detail. In general, the economic function of flippers that buy properties from especially motivated sellers, hold them for a short period, and then sell them to a buyer that places a sufficiently high value on the property is that of a middleman.⁸ Beyond the need to set prices and clear markets or guarantee quality, a crucial role for middlemen is to provide liquidity and immediacy to transactions (see, e.g., Demsetz 1968; Garman 1976; Amihud and Mendelson 1986; Copeland and Galai 1983; Glosten and Milgrom 1985; Kyle 1985). In durable goods markets they can also invest in the asset during the holding period. For middlemen, opportunities to buy may occur under any market conditions, provided they are able to identify especially motivated sellers (those with higher holding costs than their own), and we provide evidence that many flippers operate as middlemen throughout the housing cycle and perform these central functions.

Flippers can also act as speculators in the housing market. If flippers have access to better information than the broad set of agents participating in a market, which is plausible given the decentralized nature of the housing market and the infrequency with which a typical household moves, they may be able to exploit arbitrage opportunities. In the classic theory of efficient markets, speculators, acting on the basis of their superior information, serve to align prices more closely with market fundamentals, generally improving the efficiency of the market (Fama 1965). But modern finance theory admits a wider range of strategies for speculators and a more ambiguous understanding of their impact on welfare and efficiency.^{9,10} In the presence of noise traders that behave according to heuristics like chasing trends, prices can temporarily deviate from fundamentals and the rational speculator may not wish to simply take a short position. In fact, it can be optimal to pursue a much wider range of strategies. If, for example, noise traders engage in positive feedback trading, that is, have a tendency to extrapolate or to chase the trend, it can be optimal

⁸ Spulber (1996) provides a nice discussion of the roles of middlemen.

⁹ See Shleifer and Summers (1990), Barberis and Thaler (2003), and Shiller (2003) for summaries of this literature.

¹⁰ Differing levels of risk aversion, which may vary over time, may also lead some people to buy an asset. We thank an anonymous referee for encouraging us to point this out.

for rational speculators to jump on the bandwagon (DeLong et al. 1990). In this case, rational speculators take advantage of the noise traders by strategically selling before the noise traders realize the bubble is about to burst. In this way, the welfare consequences of the existence of speculators need not be positive. In general, speculators will require expected market appreciation to be sufficiently high to justify their purchases and, therefore, will be active in only those times and places where conditions are right. Here, we show dynamic patterns of a subset of investors that align with these predictions.¹¹

Finally, our paper is connected to a rich strand of research investigating specific episodes throughout history to understand the actions of investors in a variety of financial markets. For example, we show that investors were not able to anticipate the housing market's peak. Our findings complement those of Cheng, Raina, and Xiong (2014), who demonstrate that agents working in securitization finance, and thus who were likely to be as informed about the housing market's fundamentals as anyone, and which is a market that has been linked to the boom and bust in housing, were not able to predict the market's crash. In contrast to these results, Temin and Voth (2004) study a sophisticated investor who successfully profited from "riding" the South Sea bubble and Brunnermeier and Nagel (2004) find that hedge funds were able to reduce their exposure to tech stocks before the dot-com bubble burst. Other studies of bubble-like episodes include Garber (1989), who studies the Dutch tulip bubble, Greenwood and Nagel (2009) and Griffin et al. (2011), who, along with others, study investor behavior during the dot-com bubble, and Rajan and Ramcharan (2015), who study the rise and fall of U.S. farmland prices in the 1920s. In general we view the investigation of investor behavior in the housing market as particularly important for a variety of reasons, not the least of which is that housing is the asset which typically comprises the greatest share of a household's wealth.¹² Other recent work exploring the quantitative importance of housing market speculators includes Mian and Sufi (2018).

2. Data

The primary data set that we have assembled for our analysis is based on a large database of residential housing transactions compiled by Dataquick Information Services, a national real estate data company. Dataquick acquires data from public sources like local tax assessor offices, and they have provided us with the complete census of housing transactions in the five largest counties in the Los Angeles metropolitan area (Los Angeles, Orange, Riverside, San

¹¹ In a recent paper, DeFusco, Nathanson, and Zwick (2017) also find evidence that inexperienced speculators disproportionately enter the market in boom times, and the authors also provide a theoretical explanation for why this is the case.

¹² Bostic, Gabriel, and Painter (2009) document that, by 2004, housing had grown to compose more than 50% of a typical household's wealth.

Bernardino, and Ventura), between 1988 and the third quarter of 2012. For each transaction, the data contain the names of the buyer and seller, the transaction price and date, the address, and numerous characteristics including, for example, square footage, year built, number of bathrooms and bedrooms, lot size, and whether the house has a pool. While we are able to observe the date, price, and names of the buyer and seller for every transaction in the data, a drawback of the data is that Dataquick only maintains a current assessor file and overwrites historical information on house characteristics. This means we observe housing characteristics as they were in 2012, and consequently we cannot see how they may have evolved over time. The assessor data does, however, include a field with the year of permitted improvements to the property, if any, but it does not detail the change in property characteristics. This data limitation will partially motivate our research design to control for unobserved investment in houses that is explained below, although a research design to address the possibility of unobserved improvements to properties would be necessary even if Dataquick kept track of housing attributes on a continuous basis, as many home improvements (e.g., a renovated kitchen or bathroom) would not generally affect the more basic attributes of the home (e.g., lot size, square footage, room partitions) collected by the tax assessor.

From the original census of transactions, we drop observations if a property was subdivided or split into several smaller properties and resold, the transaction was flagged as not being arm's-length, the price of the house was \$1 or less,¹³ the house sold more than once in a single day, the property was especially unusual in that the price or square footage was in the top or bottom 1% of the sample, or there is a potential inconsistency in the data, such as the transaction year being earlier than the year the house was built (which may be land transactions). By observing the full transaction history for each property, we are able to see whether the buyer resold the home (versus holding to the end of the sample or transferring the property in a non-arm's-length transaction or foreclosure) and, if so, how long it was held in the interim.

Table 11 in Appendix 1 provides summary statistics for our primary data set based on a full sample of over 4.75 million transactions between 1988 and 2012. Figure 1 shows the dynamics of prices and transaction volume for the Los Angeles metropolitan area over the study period. The price index is computed with our data using a standard repeat sales method that we describe in Section 2.4. Following a rapid increase in prices in the late 1980s, the early 1990s were a "cold" market period for Los Angeles, with prices declining by roughly 30% between 1992 and 1997 and transaction volume averaging only a little more than 30,000 houses per quarter during this period. Starting in the late 1990s and continuing until early 2006, the Los Angeles housing market experienced a major boom, with house prices more than tripling and transaction volume nearly

¹³ Such a low price indicates that the seller did not put the house on the open market and instead transferred ownership in a non-arm's-length transaction.




Figure 1
Price and quarterly transaction volume

The figure reports quarterly transaction volume and a log price index for arm's-length sales in the greater Los Angeles metropolitan area. The price index is derived from a repeat sale regression described in Section 2.4, with normalization to one in 1988:Q1.

doubling. Just 2 years later almost all the appreciation in house prices from the previous decade had evaporated and transaction volume had fallen to record low levels (less than 20,000 houses per quarter). In the analysis below, we will reference the three key market periods evident in Figure 1: the “cold” market period (1992–1998), the “hot” or boom market period (1999–2006Q2), and the “post-peak” period (2006Q3–2009). While we generally will exclude the periods 1988–1991 and 2010–2012 from our analyses of investment purchases, it is useful to have both pre- and post-study period data because measurement of many of our definitions and outcomes rely on the entire history of the buyer and seller and the properties themselves. Our main analyses, in particular, employ a repeat sales framework, so it is useful to observe more transactions of each property.

2.1 Flippers

A basic measurement challenge for anyone wishing to study the behavior of investors in the housing market in these data is identifying such agents in the first place.¹⁴ Figure 2 reports the time series for four measures of investor behavior in the Los Angeles market between 1992 and 2012 derived from our transaction data set. The first of these, labeled “Second Homes,” is constructed in the spirit of Haughwout et al. (2011) by identifying individuals that own two

¹⁴ Appendix 2 discusses other approaches in the literature for identifying housing market investors.

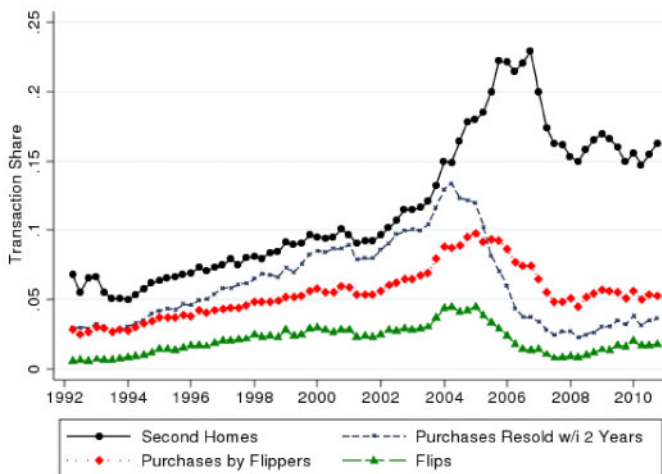


Figure 2
Quarterly residential real estate investment activity

The figure plots the quarterly share of investment activity under multiple (nonexclusive) definitions of real estate investors using transactions data for the greater Los Angeles metropolitan area. See the main text for more details.

homes at the same time. In particular, we categorize a property as a second home if the buyer’s name matches that of an individual that we also observe to be simultaneously holding another property in our data set.¹⁵ The use of buyer names is a potential source of error that we address in detail below. An advantage of our procedure is that it relies on only the deeds history file itself, and not external sources like property tax address or credit histories. A fundamental problem with this definition, of course, is that for an individual to be observed as a homeowner at all, they need to have purchased a home since the beginning of our data in 1988. Thus, our measure of “Second Homes” is likely to understate the amount of actual second home purchases, especially near the beginning of the sample period (note the figure’s series begins in 1992). For this reason, one should not overinterpret the trend in the measure. However, even subject to these various limitations, our measure of second home purchases closely tracks that of Haughwout et al. (2011), rising to a peak of nearly 24% of the market in 2006.

Many housing market investors ought not to be considered “flippers,” because motivations aside from speculation, arbitrage, or property improvement can drive investment in housing. One concrete example is the aim to rent the property. We are particularly interested in individuals buying with the intention of reselling the property after a relatively short holding period. To identify flippers, therefore, we look for evidence that an individual is generally engaged

¹⁵ This count ignores properties that overlap for 6 months or less to allow for a typical search time during household moves.

in such a strategy. In this spirit, we generate a second time series shown in Figure 2, “Purchases Resold Within Two Years”, which simply reports the fraction of all homes (regardless of the buyer) purchased in a given quarter that are subsequently resold within 2 years time. Nearly 15% of all homes purchased near the peak of the boom in 2003–2005 were resold within 2 years, a rate that is more than triple the corresponding rate for the cold market of 1992–1994 or after the 2006 price peak, periods when home prices were declining. While certainly a portion of the buyers that resell homes within 2 years of purchase are owner-occupants rather than investors, this time series provides an initial proxy for flipper-like behavior in the market throughout the cycle and demonstrates that it is economically significant.

For much of our analysis, we focus not on flipped homes *per se*, but on a set of individuals and firms that we identify as flippers. We designate a buyer to be a flipper from one of two (not mutually exclusive) conditions from the investment proxy series just presented. One is an overlapping property tenure definition, that is, the “second home” method where the buyer is observed to already hold another property in the Los Angeles market, and the additional property is sold within Y years. For most of our analysis we use $Y=2$, but we will examine how outcomes vary with respect to hold time. Another definition is to observe the buyer engaged in the buying and selling at least X different properties while holding them for less than Y years, as in the “Purchases Resold Within Y ” definition. For this definition, we usually set $X=2$ and again $Y=2$, though a major focus of our analysis is how flipper behavior, strategy, and returns vary with X .¹⁶ The difference from the first definition is that the second does not require contemporaneous overlap in property holdings, but the individual must conduct a “flip” multiple times.¹⁷

Limiting our definition of flippers in these ways provides a conservative measure of flipper activity, as we miss any individuals who engage in this activity only once during the sample period without an overlap with an existing property, or those who tend to hold properties for slightly longer periods of time, or any out-of-town investors who purchase only one investment. We do these to make sure that we avoid (as much as possible) counting normal owner-occupants as flippers. To the extent that we categorize some investors as normal owner-occupants, we would expect such contamination to cause attenuation in our measures of outcomes. Each of these measures relies on the name matches available in the transactions register and we detail this procedure in Appendix 3 and perform a number of robustness checks to possible complications associated with the procedure.

¹⁶ The latter condition would catch investors whose primary residence is “out of town” if they conduct sufficiently many flips within Los Angeles. For more on out-of-town investors, see Favilukis and Van Nieuwerburgh (2018) or Chinco and Mayer (2016).

¹⁷ Table 12 in Appendix 1 provides the purchase frequency distribution for individuals we have classified as normal owner-occupants, investors (not flippers), flippers, and institutional buyers with nonpersonal names (such as businesses, various government entities, and places of worship).

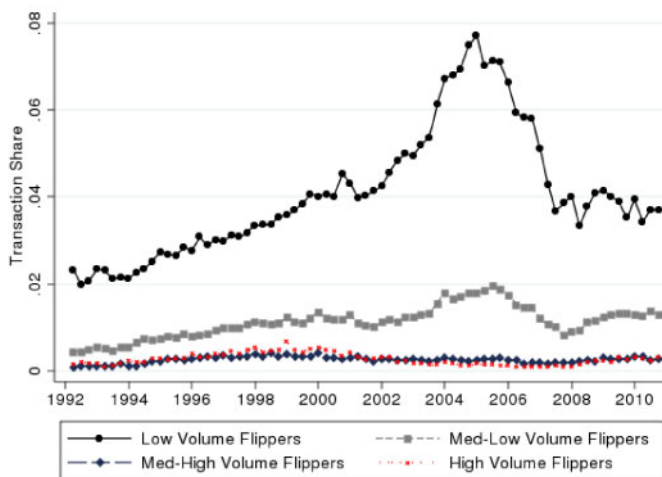


Figure 3
Flipper purchases over time by flipper type
 The figure plots the quarterly share of flipper purchases by flipper volume type. Low volume: 1-2 flips in available data; medium-low volume: 3-5 flips; medium-high Volume: 6-9 flips; high volume: 10 or more flips.

Figure 2 shows the flipper data series. “Purchases by Flippers” reports the fraction of housing transactions in each quarter that were made by individuals that we define as flippers. Note that this measure includes all homes purchased by flippers regardless of how quickly these homes are resold. This time series generally tracks housing market conditions, peaking at nearly 10% of all purchases from 2004 to 2006, a rate 3 times higher than the rate of flipper activity in the early 1990s. The last series, “Flips,” represents properties held less than 2 years and purchased by individuals designated as flippers. As the intersection of the last two series, the pattern over time shows similar cyclical behavior to all purchases by flippers, though the peak comes sooner in about 2005. It is plausible that some 2006 purchases were intended to be flips, only to see prices fall precipitously in the coming 1 to 2 years and then not resold.¹⁸

Overall, the four broad metrics of aggregate investor or flipper activity shown in Figure 2 depict a consistent pattern of procyclical behavior, with the share of purchases by these agents reaching a maximum at the peak of the housing boom, at levels that are roughly 3 times the level of activity observed during the market trough in the early 1990s. Below, we will distinguish between investor types whose strategies and participation over the housing cycle differ from one another.

¹⁸ We examine sale propensity over the cycle in Section 5.2.

2.2 Purchase activity by flipper type

In the analysis that follows, we document considerable heterogeneity in flipper behavior, strategy, and outcomes that is strongly associated with the individual flipper's volume of investments. Figure 3 shows the share of transactions in a given quarter by flippers in four volume categories. We define low-volume flippers as those flipping one or two properties, medium-low as three to five, medium-high as six to nine, and high volume as 10 or more. For the purposes of this definition, we count a purchase as a flipped home if it was resold within 2 years, and we categorize flippers on the basis of their activity over the full data period. The reported data series in Figure 3 includes all purchases (flips and others), and their sum produces the total count of flipper purchases shown in Figure 2.¹⁹

Figure 3 clearly illustrates that the dynamics of flipper activity depends on the individual's total observed volume. The purchase activity by higher-volume flippers is relatively constant over the study period, actually peaking in the colder market period of the mid-1990s. This pattern of activity is consistent with a view that these flippers tend to operate as middlemen, looking for opportunities to buy from motivated sellers with higher holding costs than their own, opportunities that are just as (or perhaps more) likely to arise in cold versus hot market conditions.

In contrast, the purchase activity by lower-volume flippers, especially the lowest category, is highly procyclical, rising from a very small percentage of the overall market in the early- to mid-1990s to almost 8% of the market in 2004–2006. This pattern of activity is consistent with the view that many novice investors were drawn into the market during the boom period. While this measure of activity is not enough to establish the motives of these flippers, the timing of their purchases is certainly consistent with a view that they are seeking to make a quick speculative gain on the basis of market appreciation.

It is worth noting that our categorization of flippers is not perfect. In particular, our volume measure is based on activity over the full study period. Thus, many of the flippers that we categorize as low volume may, in fact, ultimately become more experienced if they continue to flip homes after our study period ends, a concern mitigated by using data after the market peak. Moreover, survival in the flipping business is likely to be nonrandom, with more profitable flippers being more likely to survive long enough in the business to reach the higher-volume categories. In our analysis below, we will explicitly address these and other issues that arise due to our definition of flipper type.

A final, important aspect of flipper purchase activity to describe at the outset of our analysis is the heterogeneity in the attributes of homes purchased by flippers of each type. To this end, Table 13 in Appendix 1 summarizes some basic characteristics of the homes purchased by flippers

¹⁹ Series for flips-only show similar patterns.

of each type. As the table makes clear, flippers, especially high-volume flippers, generally purchase properties that are older and smaller than the remainder of homes that sell at least twice during our study period. These differences remain in the lower panel, which uses a regression with flipper type dummies and controls for ZIP-code-level fixed effects. Thus, flippers are buying smaller and older properties than others in the properties' surrounding neighborhoods. The research design that we present below for estimating the sources of flipper returns is motivated in large part by the possibility that flippers may systematically purchase older homes or "fixer uppers" that can benefit from substantial renovations or improvements before being resold. We also take the additional steps described below to ensure that we compare the sources of returns for flippers across comparable houses.

2.3 Flipper holding times

Before turning to our analysis of the sources of flipper returns, we present a final descriptive characterization of the heterogeneous behavior of flippers of each type. Table 1 reports the fraction of homes sold by the investors we have classified as flippers of each type within 1–4 years of the purchase. The table reveals significant heterogeneity in holding periods by flipper type. Higher-volume flippers, in particular those in the highest volume category, are much more likely to resell homes after very short holding periods, selling close to 80% of all of the homes they purchase within the first year and more than 90% within 4 years.²⁰ This suggests that these flippers purchase properties with the intent to put them immediately back on the market and is consistent with the notion that they serve the economic function of middlemen, seeking to buy cheaply from motivated sellers and resell quickly. By contrast, flippers in the lowest volume category sell only 22% of their purchases within a year of purchase, about half within 2 years, and only 62% by the 4-year mark.²¹ The overall pattern of behavior is consistent with a strategy of buying properties with the intention of capturing market appreciation, a strategy which, of course, requires a reasonable holding period.

2.4 Measuring the sources of flipper returns: Research design

Having documented the pattern of purchase activity by flippers, we turn next to an analysis of the sources of their returns. We will not be able to calculate the actual profit that a flipper earns on a particular investment because we do not observe several components affecting an individual investment holding, including whether the property is rented to a tenant during a holding period,

²⁰ Table 14 in Appendix 1 shows an expanded version of the table with subpanels devoted to hot and cold periods. This behavior is slightly more pronounced in the cold market period than the hot market.

²¹ Low-volume types not only have fewer properties held less than 4 years (which is somewhat mechanical, by our definitions), but fewer held for 1 year or less, though the flip definition extends to 2 years holding time.

Table 1
Property holding times by flipper type

Flipper type (no. of flips)	Own-occ. 0	Low 1-2	Med-low 3-5	Med-high 5-9	High 10+
Combined periods: 1992-2009					
Purchases	1,762,943	111,714	31,253	7,522	8,403
Frac. sold within:					
1 yr	0.036	0.224	0.398	0.614	0.786
2 yrs	0.098	0.452	0.588	0.748	0.875
3 yrs	0.189	0.529	0.654	0.800	0.906
4 yrs	0.265	0.587	0.698	0.827	0.919

The table reports the fraction of the homes purchased by flippers of different types sold with 1, 2, 3, and 4 years, respectively, in the greater Los Angeles area.

the particular transaction costs that a flipper might pay while buying and selling a house, or the borrowing costs that a flipper faces when procuring a mortgage in order to purchase a property.

Instead, we will focus only on the components of the returns that are associated directly with the purchase, holding, and sale of the property, and how these might correspond to the roles of speculator or middleman. In particular, we seek to identify (1) the discount that flippers receive (relative to the expected sales price in the market in the corresponding period at the time of purchase), (2) the market return that they earn over the period that they hold the property, and (3) the premium that they receive at the time of sale (again relative to the expected sales price in the market at the time). By measuring these sources of flipper returns, we seek to categorize flippers on the basis of their motivation and strategy to identify whether they appear to be operating as middlemen or speculators.

An important complicating factor is that flippers may systematically make physical improvements to the properties that they purchase, improvements which are unobserved in our data set for the reasons mentioned in Section 2. The concern is that a naive analysis of the sources of flipper returns from buying, holding, and selling a property might wind up counting money that flippers invested in improving a property as part of their return. This would be especially worrisome for our purposes if flippers of different types modified the properties in different ways.

To address this problem, we develop a research design that aims to uncover the sources of flipper returns from buying, holding, and selling a property in the (potential) presence of unobserved investment. The design is a modified version of the classic repeat sales method (Case and Shiller 1987):

$$\log(p_{it}) = \alpha_t + \gamma_i + \beta_{1k} b_{kit} + \beta_{2k} s_{kit} + \beta_{3k} a_{kit} + \varepsilon_{it}. \quad (1)$$

In Equation (1), α_t are time period (e.g., quarter) fixed effects and γ_i are house-level fixed effects. Exponentiating the estimated time fixed effects gives the price index for each quarter, which can be normalized to 1 in any quarter. The first departure from the standard repeat sales method comes from including b_{kit} ,

a dummy for if the buyer is a flipper of type $k = \{low, medium-low, medium-high, high\}$, and s_{kit} , a dummy for if a flipper of type k is the seller. These estimated coefficients related to flipper activity will provide estimates of the discount that flippers receive when buying (should $\hat{\beta}_{1k} < 0$) and the premium they command when selling (should $\hat{\beta}_{2k} > 0$), provided that house quality is constant over time. If, however, flippers purchase houses and then invest heavily to improve them before putting them back on the market, these parameter estimates will be biased. In particular, we would expect β_{1k} to be negative because the true house quality in this period would be less than the estimated quality. Similarly, β_{2k} would likely be positive because the true quality in this period would be greater than the quality estimated. The researcher may, therefore, infer that flippers are buying at a discount and selling at a premium when they are simply investing more than the average homeowner. This is why we further depart from the standard repeat sales method by including a_{kit} , which is equal to one if, in any previous period, we see a flipper of type k purchase house i , and thus allows us to control for any improvements made by the flipper that extend beyond average homeowner investment because β_{3k} captures the change in house quality between when the flipper purchased and sold the home.²²

In the standard repeat sales framework, a house only helps to identify the time series of market appreciation when it sells at least twice; otherwise it can only identify its corresponding house fixed effect. However, to identify the coefficients corresponding to the sources of flipper returns and investment, β_{1k} , β_{2k} and β_{3k} in Equation (1), homes must sell at least 4 times, with at least one nonflipper to nonflipper transaction before and after a flipper buys and sells the house in order to identify its own fixed effect and the flipper coefficients. We term this “ABCD” structure: at A both transacting parties are nonflippers; at B the house is sold to a flipper by the nonflipper; at C the flipper sells the house to a nonflipper; and at D it is sold to a nonflipper by the nonflipper. The observation before the flipper buys is used to identify the original house quality and the observation after the flipper sells is used to identify the new house quality, while concurrent nonflipper to nonflipper transactions on other properties identify the market trend. Appendix 4 further illustrates the identification argument.

Note that because our estimates of the sources of flipper returns will be based on properties that sell between normal owner occupants and fit this ABCD structure, then by construction, the period of time that the previous owner held a property before selling to a flipper is limited (as the sale at point A must be within the sample period). This excludes houses that may have been neglected over a long period of time by an owner from contributing to our estimates of the

²² Using information in the tax assessor data, we also include controls for whether the transaction is observed before or after a permitted improvement to the property. However, because there may be heterogeneity in the value of investment not captured by a permit-issuance dummy and some investments do not require permits, we include the a_{kit} controls.

sources of flipper returns.²³ While flippers, especially those seeking to make significant physical improvements, may target such homes, they generally will not be the ones that identify the sources of returns given our research design.

In the analysis that follows, we report results for two slight adjustments to the specification shown in Equation (1). First, we include a series of dummy variables for how many times we have seen a given property previously transacted in the study period. In general, sellers make some home improvements at the time of a sale so that a house will show well. Thus, we include these additional sales number dummy variables to avoid systematically overstating the performance of homes that meet the ABCD structure simply because they sell at least 4 times during the study period. Second, as shown in Table 13, flippers (especially experienced flippers) tend to purchase homes that are slightly older and smaller than the average homes that are sold in the market. Therefore, to ensure that we are comparing apples to apples, we report results for a second specification of Equation (1) that interacts the three key flipper variables with demeaned measures of housing attributes, reporting the flipper coefficients at the mean attributes of the homes sold in the study period. This ensures that all comparisons of sources of returns are done for the same type of property, even though flippers of different types may purchase properties that are somewhat heterogeneous.

3. Sources of Flipper Returns: Baseline Results

3.1 Average flipper returns

We now provide estimates of the sources of flippers' returns using the above research design. Table 2 presents our initial results for all flipper types. The table abbreviates the main coefficients of interest (Appendix Table 15 reports expanded results with control variables). The analysis includes all the data from 1988 to 2012 to account for as many repeat sales as possible, but flipper activity is based on purchases made in the years between 1992 and 2009. We exclude the ends of the sample from flip designation because we need some buyer history to designate investments and then must follow the properties for at least 2 (or more) years to identify it as a flip.

Column 1 presents estimates from a simplified version of specification (1) which omits a_{kit} . The controls for improvements made by the flipper in the simple specification are the permit issuance dummy (which is significant and positive), and sale number dummies. The coefficient estimates indicate that a purchase designated as a flip implies an 11% discount, and the sale of the flip comes at a 7% premium relative to the expected market value of the property

²³ In fact, a comparison of the housing attributes of homes that meet the ABCD structure reveals less heterogeneity in the properties that flippers purchase versus the average homes that sell in the market as a whole. For example, the average age of the homes purchased by high-volume flippers decreases from 44 to 41 years when the sample is limited to just homes that meet the ABCD structure.

Table 2
Augmented repeat sales regression results: All flipper types combined

	1	2	3	4	5	Cold 1992-'98	Hot 1999-'06Q2	Post peak 2006Q3-'09
Flip qtr of purch.	1992-'09	1992-'09	1992-'09	1992-'09	1992-'09	1992-'98	1999-'06Q2	2006Q3-'09
Flip purchase (β_1)	-0.1152 (0.0013)	-0.1200 (0.0014)	-0.1028 (0.0013)	-0.1118 (0.0015)	-0.0933 (0.0021)	-0.2063 (0.0036)	-0.0714 (0.0017)	-0.1859 (0.0054)
Flip sale (β_2)	0.0762 (0.0011)	0.0594 (0.0012)	0.0173 (0.0011)	0.0684 (0.0014)	0.0784 (0.0019)	0.0509 (0.0026)	0.0540 (0.0016)	0.0447 (0.0070)
Flip investment (β_3)				0.0207 (0.0016)	0.0187 (0.0017)	0.0333 (0.0032)	0.0021 (0.0018)	0.0700 (0.0058)
Interact property char.		Yes	Yes	Yes	Yes		Yes	Yes
Institutional buyer controls			Yes	Yes	Yes		Yes	Yes
Institutional seller controls			Yes	Yes	Yes		Yes	Yes
Limit ABCD sample					Yes			
N/T	3,785,362	3,785,362	3,785,362	3,785,362	3,685,345		3,785,362	
N	1,300,031	1,300,031	1,300,031	1,300,031	1,286,874		1,300,031	
Flip N	80,904	80,904	80,904	80,904	30,306	21,133	57,406	2,365

The table gives estimates of Equation (1) for all flippers regardless of type. Standard errors are in parentheses. See the main text for a discussion of the specifications. Appendix Table 15 reports expanded results with control variables.

at the time of transaction.²⁴ Column 2 adds indicator controls for whether the buyer was an institution (a trust, for-profit entity, or government/nonprofit entity) to ensure that the measured market rate is a normal owner-occupant to owner-occupant transaction, and Column 3 adds these controls for the seller. These transactions are associated with below-market discounts on both sides. The flipper purchase discount is little affected by their inclusion, though the sale premium is reduced somewhat. We proceed with a baseline of using the entity purchase controls, but leaving out the entity sales designations, because as we show below, purchasing distressed properties is one component of flipper returns, especially for high-volume types.

The results from the simple approach indicate that flippers transact at purchase discounts and sales premiums like the depiction in appendix Figure 41. We proceed with a more rigorous analysis of these apparent excess returns. Column 4 introduces our preferred method, the ABCD research design, Equation (1), with a post-flip dummy to capture the effect of physical improvements and the interaction of house characteristics (age and size) with flipper dummies to report results at mean house characteristics. The coefficient on *Flip Investment* indicates that flippers are not investing much more than 2% of a house's value beyond the normal transaction's "sprucing up" for sale. Purchase discounts are reduced slightly but qualitatively similar to Column 3. Thus, unobserved investment is not the primary driver of the purchase discount or sales premium estimates. As investment does not make up a sizable portion of flipper returns, in the subsequent tables we continue to control for it but suppress the estimates of coefficients for the sake of exposition. Column 5 is the most stringent specification of the ABCD design in that it drops any flips that fail to satisfy the ABCD framework, which selects the sample more towards the center of the data period.²⁵ This reduces the flip purchase coefficient slightly, but it remains statistically and economically significant.

Finally, specification 6 (spread over three columns) splits the flip purchase and sale dummies to properties acquired in the cold (1992–1998), hot (1999–2006Q2), and post-peak (2006Q3–2009) eras. The results are significantly heterogeneous over time, especially in the purchase discounts that are substantially larger in the periods of declining prices. This pattern is consistent with two notions of the role of flippers, collectively. In the cold market periods, flippers make their return by operating as middlemen, buying low and selling at a premium, relative to the average sales price in the market at the time. In hot

²⁴ Subsequent sales are associated with slight appreciation consistent with "sprucing up" of a property in staging it for market, though this levels off by about the third sale. The purchase and sale of a property under the "Second Home" investment property definition (i.e., held concurrently with other properties, but not flipped within 2 years) is associated with statistically significant but economically small purchase discount and sales premiums, which is consistent with a number of other objectives of real estate investment (e.g., held for rental income, for occupation by family members, or as vacation properties) rather than capital returns. See Table 15.

²⁵ The post-flip indicator acts as an unobserved investment control that helps to identify the flipper purchase and sales coefficients even when some flipped properties are not observed to satisfy either the A or D transaction.

Table 3
Regression coefficients by flipper type

Flip qtr of purchase	1	2	3	4	5		
	1992-2009	1992-2009	1992-2009	1992-2009	Cold 1992-1998	Hot 1999-2006Q2	Post peak 2006Q3-2009
Flip purchase, by volume							
Low	-0.0686 (0.0016)	-0.0590 (0.0017)	-0.0526 (0.0017)	-0.0362 (0.0024)	-0.1414 (0.0050)	-0.0347 (0.0019)	-0.1592 (0.0062)
Med-low	-0.1548 (0.0032)	-0.1466 (0.0035)	-0.1237 (0.0034)	-0.1193 (0.0045)	-0.2665 (0.0072)	-0.1027 (0.0038)	-0.2531 (0.0101)
Med-high	-0.2593 (0.0068)	-0.2625 (0.0071)	-0.2204 (0.0069)	-0.2733 (0.0103)	-0.3636 (0.0108)	-0.2135 (0.0080)	-0.2728 (0.0154)
High	-0.3184 (0.0066)	-0.3115 (0.0067)	-0.2393 (0.0067)	-0.3135 (0.0083)	-0.3466 (0.0096)	-0.2880 (0.0073)	-0.3923 (0.0157)
Flip sale, by volume							
Low	0.0577 (0.0014)	0.0682 (0.0017)	0.0330 (0.0016)	0.0783 (0.0024)	0.0540 (0.0045)	0.0571 (0.0018)	0.1060 (0.0070)
Med-low	0.0639 (0.0025)	0.0730 (0.0029)	0.0287 (0.0029)	0.0789 (0.0041)	0.0651 (0.0061)	0.0575 (0.0033)	0.1078 (0.0126)
Med-high	0.0640 (0.0056)	0.0614 (0.0061)	0.0024 (0.0060)	0.0612 (0.0096)	0.0607 (0.0094)	0.0339 (0.0065)	0.1038 (0.0341)
High	0.0500 (0.0046)	0.0577 (0.0049)	-0.0093 (0.0048)	0.0719 (0.0059)	0.0874 (0.0070)	0.0153 (0.0057)	0.0672 (0.0215)
Other controls:							
Post permit	Yes	Yes	Yes	Yes		Yes	
Transaction count	Yes	Yes	Yes	Yes		Yes	
Entity buyers	Yes	Yes	Yes	Yes		Yes	
Entity sellers			Yes				
Distress indicators			Yes				
Interact property char.	Yes	Yes	Yes	Yes		Yes	
Flipper investment		Yes	Yes	Yes		Yes	
Limit ABCD sample				Yes			
<i>NT</i>	3,785,362	3,785,362	3,785,362	3,685,345		3,785,362	
<i>N</i>	1,300,031	1,300,031	1,300,031	1,286,874		1,300,031	

The table gives estimates of Equation (1) when the coefficients associated with flipper activity are separately estimated by type. Standard errors are in parentheses. Interacting property characteristics indicates that the mean house characteristics for the sample are subtracted from individual house characteristics and these values are interacted with the flipper dummies. Other controls, including estimates of the investment coefficients (which also vary by flipper type), and the sale number dummy variables are suppressed for expositional sake.

market conditions, in which prices were increasing rapidly and sales volume was much greater, flippers have the potential to make returns by purchasing houses at times and in locations where expected market appreciation is high. Thus, as the parameter estimates in Table 2 show, during the hot period flippers on average do not capture a particularly low price when buying, a discount of now only 7% instead of 20 during this period.

3.2 Returns by Flipper Type

Next, we investigate the differential sources of returns across flipper volume types, using the same four categories defined above in Section 2.2. Table 3 presents parameter estimates for a set of specifications that allow the coefficients related to flipper discount, premium and investment to vary by flipper type. The sale order, post-permit, post-flip, and investor and institutional transaction

dummy variables, and interactions of property characteristics are included in the specifications reported in Table 3, but the parameter estimates (which are similar to those reported in Table 2) are not reported for ease of exposition. The post-flip controls and property characteristics are also interacted with flipper type to account for possible differences between types in the average investment level and property attribute selection, respectively.

Our main result, seen in the basic specification in Column 1 and our preferred augmented method in Column 2, is a striking heterogeneity in the sources of returns across flipper types. While all flippers buy below expected market price, higher-volume flippers buy at a much deeper discount relative to expected market prices. For the whole sample period, the highest volume flippers receive a discount at purchase of approximately 31% compared to 6% for the lowest volume types. Across specifications, discounts are monotonically increasing in volume type. Sales premiums are comparatively smaller and more similar across types.

Column 3 adds indicators for entity sellers. This preserves the ordering of discount size across type but diminishes the magnitude for the highest volume types. This indicates that, to some extent, the highest volume types earn discounts by buying from nonoccupant sellers, who presumably have greater holding costs. In Section 5.3 we examine the implications this propensity has for market liquidity. We leave out these effects for our headline calculations but bear them in mind for interpretation.

Column 4 limits the sample to properties satisfying both ends of the ABCD restriction, showing the striking heterogeneity is robust to the most stringent specification controlling for physical investment. Finally, specification 5, spanning three columns, separately estimates the flipper type purchase and sale dummies by time period. Discounts are larger for all flipper types during cold market periods, though the heterogeneity between types is still evident. The higher-volume flippers exhibit the smallest differences in discounts between cold and hot market periods. Steep discounts at the time of purchase are consistent with these flippers operating as middlemen, finding properties to buy cheaply and operating during any market conditions. Low-volume flippers, on the other hand, generally do not buy at much of a discount, especially in hot market conditions. Like their cyclical purchase activity shown in Figure 3, this is consistent with the idea that they are generally seeking to profit from market appreciation as speculators rather than arbitraging as middlemen.²⁶

Using the results from the estimates of specification 2 in Table 3, we can report the source of a flipper's return for each flipper type: decomposing the fraction that stems from buying cheaply, selling high, and simply earning the

²⁶ The reduction in discounts during the hot period is a function of a compression of the dispersion in prices during periods of greater appreciation, but for low-volume flippers, of their buying closer to the center of the distribution during the hot period. The behavior is consistent with a strategy of earning more return from market appreciation, or at least, being less selective about finding "good deals" when the overall market is appreciating. Higher-volume flippers consistently buy at the lowest percentiles of the distribution at each place in the cycle.

Table 4
Sources of return by flipper type

Volume type	1 Flips <i>N</i>	2 Buyer discount	3 Seller premium	4 Market growth	5 Transaction cost	6 Nominal rate of return	7 Quarters held	8 Annualized mkt growth	9 Annualized rate of return
Low	49,660	-0.059	0.068	0.130	-0.104	0.153	4.074	0.127	0.150
Med-low	17,892	-0.147	0.073	0.095	-0.112	0.202	3.151	0.120	0.256
Med-high	5,473	-0.263	0.061	0.058	-0.123	0.259	2.388	0.097	0.434
High	7,253	-0.312	0.058	0.037	-0.126	0.281	1.932	0.077	0.581

The table shows the sources of returns on flips by flipper type. The discounts, premiums, and market growth are calculated from specification 2 of Table 3, and quarters held is simply the mean number of quarters held. The gross rate of return is generated by totaling the sources of return, and the annualized return comes from dividing the gross return by the mean years held. See the text for more details.

market return during the holding period. Table 4 provides these results. We include estimates of flipper rates of return based on time held, market growth, and the residual transaction prices. Column 2 gives purchase discounts, Column 3 the sale premiums, and Column 4 an estimate of market appreciation. We also include in Column 5 an estimate of transaction costs typically associated with the buying and selling of real estate, including broker fees, deed recording and local tax costs, and loan procurement (weighted by the probability that a loan is used to finance the transaction). The table consolidates all of these into one column for brevity, and we explain our computation of these costs in Appendix 5.²⁷ Higher-volume types exhibit slightly larger transaction costs mostly because broker fees at time of sale represent a larger fraction of the heavily discounted purchase price. Column 6 reports the sum of gross returns.²⁸ The returns for high-volume types are nearly twice that of the lowest types, and essentially all of their return owes to purchase discounts. Furthermore, Column 7 shows a large disparity in the time held (recalling Table 1). High-volume flipper quickly resell their houses, whereas low-volume flippers hold them almost twice as long. Besides indicating a difference in strategy, this has implications for magnitude of annualized rates of return, which are that much larger for high-volume types, and the market appreciation component occupies a smaller share even at annualized rates.

Although it is interesting that higher-volume types earn higher returns on average (even in gross annualized rates), for the purpose of understanding their role in the functioning of the housing market, we find it most important that the source of their returns is primarily the steep discounts they earn at purchase. That is, the decomposition in Table 4 highlights how flippers of different types

²⁷ It is also important to keep in mind that investors may face capital gains taxes after selling a property, but reliable computation of such taxes is difficult for a variety of reasons. One consideration, in particular, is the possibility that capital gains taxes can be deferred if a “like-kind” asset is purchased within one-half a year (the 1031 exception). We have found in our data that, unsurprisingly, higher-volume flippers are more likely to purchase again in this window, at a rate of 65% compared to 18 percent among low-volume flippers, which, all else equal, would further separate their effective returns.

²⁸ We say “gross returns,” distinct from net, to emphasize that we do not have measures of holding costs, costs of capital, or tax liability.

and roles find their returns from different sources, and provides strong evidence that some flippers act as speculators while others operate as middlemen. Low-volume flippers do not buy at an especially low price and, as a result, their gross rate of return is primarily driven by overall market growth: 84% of their return stems from market growth. High-volume flippers, on the other hand, earn most of their return by buying at prices below average market prices and quickly reselling so that only a small fraction (13%) of their return stems from overall market growth, whereas purchasing cheaply generates 54% of their return.

Taken together, the evidence on purchase activity, holding times, and sources of returns paints a very consistent picture: high-volume flippers generally act as middlemen and low-volume flippers as speculators. We now establish (a) the robustness of these results to a number of potential concerns and (b) the external validity of the findings by demonstrating that the patterns documented so far for Los Angeles are also evident in many other U.S. cities.

3.3 Robustness

In this subsection we summarize the analyses we have carried out to establish the robustness of the results presented above to a number of the assumptions that underlie our analysis. To help with exposition we present the full details of the analyses in Appendix 6.

We begin by focusing on the definition of a flip and a flipper. In Appendix 6.1 we show that our results are robust to other sensible definitions. For example, we show that the results are consistent whether flips are identified by buyers having at least two short tenured-sales or as second homes. We also show that the results are robust to other filters meant to address our misclassifying two different people as the same investor if they share a common name. We demonstrate similar robustness if we restrict our identification of a flipper based on the geography over which we allow them to operate (i.e., how far apart two transactions are allowed to be for the same flipper), or the maximum length of time between transactions. In summary, while our identification of flippers may be imperfect, the primary result of the heterogeneity of returns by volume type holds up to alternative definitions. Indeed, it would be remarkable if such a stark pattern were somehow the result of using buyer names to identify investors. Improper aggregation should result in owner-occupiers being designated flippers, and low-volume flippers being designated as higher volume, and so on, which should homogenize the types and attenuate the estimates, the opposite of our findings.

In the results above, to focus on properties turned over quickly rather than retained for rental income, we examined the sources of returns for houses that were resold in less than 2 years. Of course, the timing of the decision to resell the property is an endogenous choice made by the investor, likely influenced by the appreciation of the property and the cost of capital. By limiting the sample to only those homes that were resold in the first 2 years, we may be inadvertently

focusing on a selected sample of homes that performed well in terms of market appreciation. Appendix 6.2 shows that our result that high- and low-volume flipper types pursue different strategies is not sensitive to alternative restrictions on how fast a home must be resold to count as a flip.

As we discussed in detail above, a broad challenge in examining the sources of flipper returns is the possibility that flippers invest significant amounts of money to improve properties, investment that is unobserved by the researcher. While several aspects of our baseline analysis have been designed to minimize this concern, we further pursue this possibility in Appendix 6.3. We show that if we further limit the time between sales at points A and B, and C and D, so as to minimize the possibility of unobserved investment, our results remain unchanged.²⁹

For our baseline results, we estimated a single housing price index for the Los Angeles metropolitan market and used that to measure the rate of market appreciation while flippers held on to properties, as reported in Table 4. An issue with using a single aggregate price index for our analysis is that flippers might be able to identify and target submarkets or neighborhoods that appreciate faster than the metropolitan area as a whole. In one sense, this is not a problem but a matter of interpretation, as the submarket targeting might be the means by which volume flippers act as middlemen. But as a matter of measurement, it might lead to an understatement of market appreciation and an overstatement of the excess returns component of flipper returns. In Appendix 6.4 we show that our results are robust to alternative approaches for estimating “neighborhood-level” price indices.

Finally, we show in Appendix 6.5 that our results are not based on identifying heterogeneous flipper types based on volume or experience. Rather, we show that these types are persistent throughout a flipper’s tenure. For example, we show that even on their first two houses flipped, high-volume flippers received substantial discounts at purchase. Across flipper types, purchase discounts and sales premiums are remarkably consistent throughout their career profile.

4. External validity: Flippers in other U.S. metro areas

We focus our analysis on Los Angeles, a large, diverse, and volatile housing market, for which there is available a long transaction record with which to conduct our analysis. To demonstrate the external validity of these results, in this section we show the patterns of flipper behavior that we documented for Los Angeles are evident in many other U.S. metro areas.

²⁹ Nothing in our analysis implies that flippers do not purchase fixer uppers that could be physically improved in a profitable way. Rather, our research design ensures that such properties do not contribute to the identification of the sources of flipper returns that stem from buying cheaply and selling at a premium.

We construct data on flips and flippers for as many other large metro areas possible during the time period we study. Dataquick collected transaction register and county tax assessor data from public records in most large counties in the United States, although available time periods and reporting practices vary somewhat from county to county. Notably, the state of Texas suppresses transaction prices, so we exclude Texan cities from our analysis. Dataquick was expanding collection areas throughout the 1990s, and many metro areas do not come online until late in the decade. Our identification of flippers and flipper types relies on observing an individual's activity over a sufficiently long time period, so we limit our analysis to cities with continuous coverage as of 1998. We focus on larger metro areas, which we define as those with at least 100,000 households (in the 2010 census) and an average of 500 transactions per quarter over its coverage period. For years after 1998, our sample totals 96 different U.S. metro areas besides Los Angeles. See Appendix Table 74 for a list.

For the sake of exposition, in this analysis and for the rest of the paper we group flippers into two categories based on experience. High-volume flippers, or *middlemen*, are those who engage in three or more flips over our sample (in the language above, these are medium-low, medium-high, and high) and low-volume flippers, or *speculators*, are those who flip two or fewer times during our sample. Given the findings above, this categorization divides flippers according to two distinct investment strategies.

Figure 4 shows the summary of flipper activity in our national sample of metro areas excluding Los Angeles. It displays the purchases by flipper type as a share of total transaction volume, pooling across all metro areas with available data in the quarter.³⁰ The figure shows that both types of flippers are operating nationally in ways largely similar to the Los Angeles market. Speculators sharply increase their purchasing during the housing boom, while middlemen operate at similar volumes throughout the cycle. Although investors' shares of transactions elsewhere are, on average, a bit lower than in Los Angeles, their presence is still sizable, and as we will show, the patterns of returns that we documented in Los Angeles will be apparent nationally.

As the level of activity by each flipper type varies across cities, we explore the connections between their participation and features of the local housing in Appendix 7. The results from that analysis show that speculators operate at higher rates in markets with greater turnover in the housing market. Middlemen show a much weaker association to transaction rates. We also show that speculator activity is higher in markets with greater price appreciation, while middlemen exhibits a weaker correlation.³¹ These findings indicate that the characterization of speculators and middlemen from the Los Angeles market is a good depiction of their behavior more generally. Speculators operate cyclically

³⁰ A separate series shows the number of metro areas in the sample at each quarter. By our criteria, this number reaches its maximum of 96 in the year 1998.

³¹ This finding is in line with that of Adelino, Schoar, and Severino (2017).




Figure 4
Speculator and middlemen buying activity nationally

Calculations exclude the greater Los Angeles market and cities without sufficient transactions data as described in main text.

in the national time series as well as being more prevalent in the markets with more pronounced cyclical activity, while middlemen operate counter- or acyclically in the time series and in the cross section.

One of the key differences we highlighted between speculators and middlemen in Los Angeles was the source of their returns; namely that middlemen purchase properties at a large discount relative to the market and in comparison to speculators. We now show that this pattern holds outside of Los Angeles. We reestimate Equation (1) for each metro area.³² This generates a large number of buying and selling coefficients by flipper type, so we summarize the results by presenting the distribution of point estimates in Figure 5. The left panel of Figure 5 shows that flippers of each type buy properties at substantial discounts,³³ although the distribution of middlemen coefficients is clearly to the left of the speculators'. The right panel shows that flippers earn small premiums at the time of sale, and the average sale premiums earned by middlemen is slightly greater than that earned by speculators. For reference we label the point estimates for the Los Angeles market. Table 72 in Appendix 7 presents estimates from a selection of notable markets. It is apparent that the pattern of purchase discounts documented above for Los Angeles is common across several major cities, with Los Angeles itself near the center of the national distribution. In Appendix 7 we relate the coefficient estimates for each city with

³² including Los Angeles, which we do, is not a redundant exercise, because the categorization of middlemen is coarser than that of Table 3.

³³ The results include a few cities with point estimates near or above zero, but none of the above-zero estimates are statistically significant at the 5% level.




Figure 5

Distribution of speculator and middlemen coefficients from repeat sale regressions across MSAs

The figure plots kernel density estimates of the distribution of point estimates for speculator (S) and middlemen (M) purchase (left plot) and sale (right plot) coefficients for the metro areas that meet our selection criteria described in the main text. The specific point estimates for Los Angeles are denoted by text.

local housing market characteristics. The analysis there shows that speculators seem to be receiving smaller discounts in hotter markets, whereas middlemen again appear to be operating roughly acyclically.

5. Speculators, Middlemen, and the Housing Market

Up to now our analysis has focused on documenting stark differences in the strategies employed by speculators and middlemen. In this section we take a broader look at the way they interact with and potentially affect the housing market in general.

5.1 Local price dynamics and flipper activity

We begin this section by exploring the relationship between speculator activity and housing price dynamics in the regions of the market in which they are active. The aim of this portion of our analysis is to study what happens to housing prices in the subregions of the Los Angeles metropolitan area that experienced greater levels of speculator activity. To do so, we regress annual neighborhood-level price appreciation where, as described in Appendix 6.4, neighborhoods are defined by 3-digit ZIP codes and price appreciation at the neighborhood level is computed using neighborhood-level price indices, on lagged quarterly purchases by speculators and middlemen, as well as lagged annual neighborhood-level price appreciation. The regressions also include quarter fixed effects. Table 5 provides the results for the differing lags.^{34, 35}

It is well known (e.g., Case and Shiller 1989) that lagged price appreciation predicts future price appreciation in housing markets. This can be seen in the

³⁴ We obtain similar results if lagged flipper holdings are used instead of purchases.

³⁵ We note that our analysis in this section is related to that in Chinco and Mayer (2016), although they are more focused on establishing a causal link between speculative purchases and price growth.

Table 5
Flipper activity and local price dynamics

	1	2	3	4	5
Price appreciation	0.197 (0.0203)	-0.122 (0.0209)	-0.186 (0.0216)	-0.182 (0.0247)	-0.134 (0.0333)
Purchases by:					
Speculators	1.736 (0.2161)	1.928 (0.2239)	0.321 (0.2315)	-1.784 (0.2339)	-1.950 (0.2444)
Middlemen	-0.459 (0.1608)	-0.465 (0.1673)	-0.741 (0.1737)	-0.397 (0.1742)	0.283 (0.1805)
Obs	2,366	2,262	2,158	2,054	1,950
R ²	.8621	.8613	.8628	.8681	.8641
Lags	1	2	3	4	5

The table shows regressions of neighborhood yearly price appreciation on quarterly speculator and middlemen purchase activity and neighborhood-level yearly price appreciation. Each specification lags the covariates a different amount, which is given in the last row of the table.

top row of Table 5, which establishes positive short-term persistence and long-term mean reversion in house price appreciation at the neighborhood level.³⁶ However, purchase activity of speculators is also associated with future price appreciation. The first three columns show that a greater number of purchases by inexperienced, speculative flippers are correlated with above average rates of price appreciation over the following 2 years, moderating in the third, while the last two columns show that these short-term gains are followed by below average returns over the following 3 years.

Middlemen, on the other hand, operate in areas where prices are rising more slowly than the rest of the metropolitan area, and so their activity is not predictive of future prices. Because middlemen primarily earn their returns by finding “good” deals when purchasing and selling for high prices relative to the market, it is not surprising to find them operating in hot and cold neighborhoods, just as they operate in hot and cold portions of the housing cycle.

5.2 How well informed are speculators?

The evidence presented so far establishes that a substantial number of amateur speculators entered the market in the recent housing boom and that their entry into a submarket was associated with the formation and bursting of market-wide and local price bubbles over the next several years. In this subsection, we ask whether the fact their purchases are associated with future price appreciation in the short- to medium-term reflects their ability to predict and take advantage of short-term price spikes due to superior information about local market fundamentals. If not, and speculators do not transmit valuable information to the market or act with any special knowledge about market fundamentals, then their participation in and of itself may not improve market efficiency, and their activity may have, in and of itself, further fueled the mispricing in the housing

³⁶ To our knowledge, the fact that, even at the neighborhood level, lagged appreciation strongly predicts future price appreciation is undocumented elsewhere in the literature.

sector.³⁷ Other ways that they may have improved the market's functioning include providing liquidity and improving the capital stock through investment, and we also will investigate these possibilities below in Sections 5.3 and 5.4.

While some might take the inexperience of the speculators in the data as *prima facie* evidence that they are not especially well informed, we offer a more formal analysis of their behavior in this subsection.³⁸ In particular, we present evidence from the timing of purchases and sales as the Los Angeles market neared and went over its peak in 2006 that suggests that flippers, and speculators in particular, failed to anticipate the market's peak in any way.

We think it is especially important³⁹ to evaluate speculators vis-a-vis middlemen on this front because they were so reliant on market timing. That speculators could not anticipate the peak makes it difficult to rationalize the timing and location of their activity as being driven by superior information and instead makes it more likely that the causation could run in the opposite direction.⁴⁰ Of course, price appreciation and speculator participation may have been jointly influenced by other factors, such as a loosening in overall lending standards, and, in this paper, we do not provide evidence to the contrary. For now, our goal is to explore whether speculators were particularly well informed about market fundamentals.

For the following analysis we construct a sample of all homes purchased by each type of flipper, regardless of how long the property was eventually held, and for comparison, we include a measure of purchasing activity for owner-occupants. To display all on the same scale, we normalize the quarterly purchase volume by the average quarterly purchases in 1998, the last year of the cold market, so that a value of one means at the 1998 average. Figure 6a plots the resultant indices for each type of buyer. The figure includes a shaded region denoting the period of greatest 2-year appreciation as the "hottest" period. A flip held for 8 quarters would yield the greatest level of market appreciation if purchased at the start of this band and sold at its conclusion, and thus provides a rough guide of the beginning of the end of ideal market timing for a flipper.

The figure shows the substantial increase in speculator purchasing as prices rose. Their level of activity in 2005 was 2.5 times that of 1998. That middlemen and owner occupiers showed little increase in activity indicates that much of

³⁷ Our analysis in this subsection is similar in spirit to that of Griffin et al. (2011), who study the activity of institutional and individual investors around the peak of the dot-com bubble.

³⁸ Other researchers (e.g., Greenwood and Nagel 2009) have found that inexperienced traders engage in the type of trend chasing behavior exhibited here. A multitude of evidence of this phenomenon hails from lab and retail investor survey settings. See, for example, Smith, Suchanek, and Williams (1988), Haruvy, Lahav, and Noussair (2007), or Vissing-Jorgensen (2003).

³⁹ See also DeFusco, Nathanson, and Zwick (2017).

⁴⁰ Our conclusion that speculators may have helped fuel the bubble is consistent with and complements the evidence presented in Chinco and Mayer (2016), who develop a research design that makes use of data from multiple MSAs to identify the impact of long-distance investors on metropolitan price bubbles.

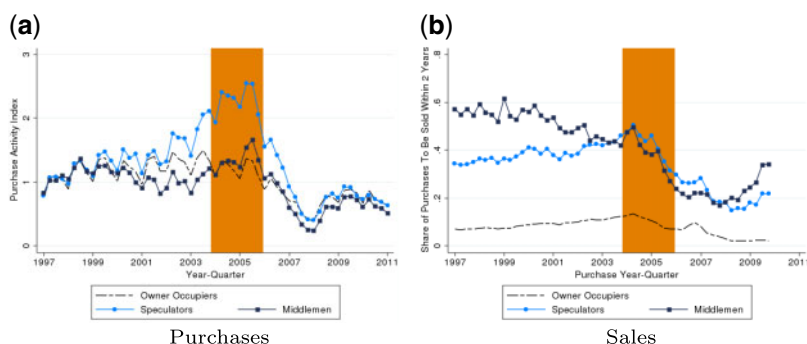


Figure 6
Purchase and sale dynamics by flipper type

The figure displays (a) an index of purchasing activity for each flipper type and for owner occupiers, and (b) the percentage of homes bought in any quarter that were sold within 2 years from that quarter for each flipper type and for owner occupiers. The index is normalized to one at the average of the cold market period of the mid- to late 1990s. The shaded region denotes the 8-quarter period of highest 2-year market appreciation.

the extra volume occurring in the runup was due to the speculator types and other investors (see also Figure 2).⁴¹

There is nothing necessarily irrational about speculators buying more often as prices rose; indeed, some surely made decent returns. The issue we are concerned with is whether these speculators had superior information or whether they were trend chasing. The fact that purchases increased well into the hot period even after crossing the point of optimal market timing suggests they were doing more of the latter. Moreover, if speculators had anticipated the downturn in housing, they should have curtailed their purchasing activity well in advance of the market’s peak. The figure shows that only very late did speculators begin to reduce their buying activity ahead of the market’s decline. Given the time involved in both buying and selling properties, had they anticipated the market slowdown, they certainly should have curtailed their buying activity much sooner.

While Figure 6a suggests that speculators did not anticipate the market’s peak, a confounding factor is that many of them might not have been putting much of their own money at risk. If banks were making mortgage loans without requiring these investors to contribute much of a downpayment, they may simply have been gambling with the bank’s money and, consequently, had little to lose should it turn out that they bought too late in the boom. To investigate this possibility, we examined the average combined LTV for all mortgages held by speculators on the properties that they purchased in the boom. See Figure 81 in Appendix 8. While the average LTV did rise for speculators in the boom (as

⁴¹ The slight blip in middlemen purchasing late in the hot period is entirely due to the lowest volume category (3–5 flips) within the group (see Figure 3). It seems possible that several more speculator types purchased an additional property and crossed the threshold of our definitions.

it also did for traditional home buyers), it remained around 90% at the peak, and was at a rate comparable to that of traditional homebuyers.⁴² Thus, speculators were still putting a significant amount of their own money at risk and certainly had a financial incentive to stop making new purchases ahead of the peak.

The dynamics of flipper selling activity provides further evidence that they did not anticipate the market peak. Figure 6b shows the fraction of homes purchased by speculators, middlemen, and owner occupiers in each quarter that were sold within 2 years from that quarter. Not only did prices decline rapidly starting in 2006, but transaction volume had slowed considerably by late 2005, making it difficult to unload any remaining inventory.

The figure indicates that, of the homes that were purchased prior to the second quarter of 2004, speculators typically sold 40% to 50% within 2 years of purchase, and middlemen 50 to 60. However, starting with homes purchased in the second quarter of 2004, 2 years prior to the price peak, this percentage declined, tipping drastically as transaction volumes contracted. If speculators had anticipated the market peak, we would have expected them to have sold a much greater percentage of the homes that they purchased in that quarter of 2004, for example, within 2 years in order to avoid holding inventory as prices fell precipitously. Instead, they were left with their holdings.

The pattern is not confined to speculators, however, as middlemen and owner occupiers were similarly left with purchases made later in the cycle, indicating that none of these buyers unloaded inventory in anticipation of a market turn. Note however that Figure 6b conditions on purchase. That is, if speculators demonstrated less-than-average prescience, it was in their elevated purchases through the market peak, not in their propensity to sell holdings already in inventory.

5.3 Flippers and liquidity

Even if speculators were uninformed market participants, they may have improved the housing market's functioning by increasing overall liquidity. In this subsection we investigate this possibility. Specifically, we look at the parties from whom the flippers are buying, and in particular, whether they are more likely than an average homeowner to purchase a property not in use (held by an institution like a bank or government entity or one that has been foreclosed on). Table 6 reports the probability that the flipper's purchase transaction indicated a seller we could identify via name text as a financial company (primarily banks), a government sponsored enterprise (Fannie Mae or Freddie Mac), or a government entity like Housing and Urban Development (HUD) or state, county, or municipal governments. Finally, we also report purchases

⁴² Figure 81 also shows some interesting heterogeneity in how flippers use financing. Middlemen buying in cold markets are much more likely to pay cash than speculators, but conditional on purchasing during the hottest market period, they used loans at similar rates. Throughout, conditional on using a loan, they exhibit slightly higher rates of LTV than speculators or owner occupants.

Table 6
Probability of purchasing distressed properties

Period	Buyer type:	Owner- occupier	Other investors	Flippers	
	Seller type			Speculators	Middlemen
Purchases, 1992-2009					
	<i>N</i>	2,748,701	262,692	155,766	94,341
	Bank/financial Co.	6.25	7.02	7.36	12.35
	Government entity	2.54	3.04	4.51	9.84
	GSE	2.26	2.21	2.46	4.03
	Foreclosed property	11.99	15.07	16.03	26.64
Cold: 1992-1998, 2006Q3-2009					
	<i>N</i>	1,401,335	132,708	64,283	43,944
	Bank/financial Co.	10.53	12.15	14.27	20.00
	Government entity	3.10	3.51	5.87	10.95
	GSE	3.88	3.84	5.07	7.32
	Foreclosed property	17.71	22.69	25.94	36.79
Hot: 1999-2006Q2					
	<i>N</i>	1,347,366	129,984	91,483	50,397
	Bank/financial Co.	1.80	1.78	2.50	5.67
	Government entity	1.95	2.56	3.56	8.87
	GSE	0.59	0.55	0.64	1.17
	Foreclosed property	6.04	7.30	9.07	17.78

The table reports the mean probabilities of the observed groups, by column, of conducting the action denoted in each row. All figures, except sample counts, are expressed as percentages (%). “GSE” refers to a government-sponsored enterprise, namely, Fannie Mae or Freddie Mac.

of properties that we identify as being previously foreclosed up, meaning we observe an institution taking ownership of the deed from a previous arm’s-length buyer, even if the institution is not observed as the seller of record (the foreclosure designation may overlap with institutional seller designations).⁴³

Speculators make purchases of institutionally held properties slightly more often than owner occupiers and other investors. Middlemen, however, are about twice as likely to take part in these distressed transactions. The levels of these purchases rise during cold market periods, comprising a substantial fraction of middlemen’s purchases. This fact is consistent with the role of middlemen taking properties off the market from high-holding-cost parties and providing liquidity, especially during downturns.

5.4 Flippers and investment

Another way that speculators could improve the housing market is if they improve the overall capital stock by investing in the properties they flip. To investigate this possibility, we use the assessor’s data to measure the rate of permit issuance during the holding period of the flippers. This is an imperfect measure of capital improvements. Sometimes permits are required for trivial matters (to the market), and, likely more importantly, many market-valued

⁴³ We use this definition to catch properties that move hands multiple times before we observe a market rate transaction. The propensities are slightly smaller than when using Dataquick/Corelogic’s internal “distress sale” indicators, although qualitatively, the patterns by type of buyer hold with either definition.

Table 7
Probability of improving properties and buying new or renovated properties

Period	Buyer type Purchase/action Type	Owner- occupier	Other investors	Flippers	
				Speculators	Middlemen
Purchases, 1992-2009					
	<i>N</i>	2,748,701	2,748,701	155,766	94,341
	Take permit	0.06	0.08	0.13	0.23
	Buy recent permit	2.83	2.33	1.96	1.29
	Buy from real estate developer	3.54	2.36	2.12	1.61
	Buy new construction	9.44	6.13	5.28	3.45
Cold: 1992-1998, 2006Q3-2009					
	<i>N</i>	1,401,335	132,708	64,283	43,944
	Take permit	0.05	0.08	0.10	0.26
	Buy recent permit	3.38	2.72	2.31	1.47
	Buy from real estate developer	2.96	2.11	1.80	1.43
	Buy new construction	9.59	6.07	5.13	3.22
Hot: 1999-2006Q2					
	<i>N</i>	1,347,366	129,984	91,483	50,397
	Take permit	0.07	0.08	0.14	0.20
	Buy recent permit	2.25	1.93	1.72	1.13
	Buy from real estate developer	4.14	2.63	2.35	1.77
	Buy new construction	9.28	6.18	5.38	3.64

The table reports the mean probabilities of the observed groups, by column, of conducting the action denoted in each row. All figures, except sample counts, are expressed as percentages (%).

cosmetic improvements do not require the issuance of a permit. To this point, the flipper coefficients in the augmented repeat sales regressions (buy, sell, and post-flip) help to capture things valued by the market, whether or not permitted. Moreover, from these regressions, the value of a property rises substantially after issuance of a permit, on average about 13%. So despite its limitations, permit issuance seems a relevant feature to examine.

Table 7 reports the probability that the buyer is issued a permit during his tenure of holding, a simple mean of the indicator variable and not adjusted for holding length (i.e., not a hazard rate). This is a relatively rare occurrence, with less than one-half of a percent of holdings involving a permit issuance.⁴⁴ Speculators are about twice as likely than owner-occupiers or other investors to obtain a permit, but middlemen are roughly twice as likely as speculators and 4 times as likely as owner occupiers.^{45, 46} Moreover, the second and third panels show that permit rates for middlemen are even higher during cold market periods—in contrast to speculators—suggesting that capital improvements are a relatively larger part of their investment strategy.

The other rows of the table show that speculators are less likely than owner occupiers or other investors to buy a recently permitted (i.e., recently renovated)

⁴⁴ This is itself likely an undercounting of the true prevalence of permit issuance, because the assessor data is observed once, so we only observe the most recent permit. For example, if an owner was issued a permit in 1995, and another owner of the property issued another permit in 2005, the data report only the 2005 permit.

⁴⁵ Flipper rates are even higher as hazard rates, because flippers hold the properties for less time.

⁴⁶ We note but do not report that that among middlemen, the rate of permit issuance rises with flipper volume.

property, and less likely to buy new construction (as measured by sale and year-built fields fewer than 2 years apart), or properties sold by real estate development companies (e.g., K Hovnanian). Middlemen are even less likely to buy new or renovated structures. Hence, consistent with Table 13, flippers, but especially middlemen, are buying and selling older properties in greater need of rehabilitation.

6. Conclusion

Making use of a large transactions database and a novel research design, this paper provides the first comprehensive study of intermediaries (middlemen and speculators) in the housing market: identifying their activity, their sources of the returns and apparent strategies, the quality of their information with and the extent to which their activity affects liquidity and the overall stock of housing.

A main contribution of our analysis is a detailed characterization of the role that speculators played in the recent housing bubble. In particular, we document substantial entry by amateur speculators at the height of the boom, entry that was strongly associated with sharp short-term increases and intermediate-term declines in local housing prices. Their collective inexperience and complete inability to anticipate the market peak by either curtailing their purchases or selling their inventory suggests that these speculators were not acting with superior information, but were instead simply betting that the boom would continue for a while longer.

Our analysis goes beyond a retelling of the role of speculation in the recent boom and bust. By examining the behavior of intermediaries over nearly a quarter century, we also identify real estate investors clearly operating in the role of middlemen. In this capacity, we document investors buying homes at prices well below market value and quickly reselling after short holding periods at, or above, market value. These middlemen operate in all market conditions and, if anything, are more active during busts, when the liquidity they provide to the market is likely to be especially valuable.

In this way, our analysis suggests that the impact of real estate investors on the market—especially whether their activity increases or decreases social welfare—likely depends critically on the market cycle. In busts, middlemen almost certainly help to stabilize the market, effectively putting a floor on price declines and providing liquidity for many homeowners who may be desperate to sell quickly following economic shocks. In booms, rampant speculation may have a de-stabilizing role that significantly amplifies price volatility. Even well-informed speculators may have deleterious effects on the market if they use their informational advantage to fuel and ride the bubble as in DeLong et al. (1990), although we find no such evidence of an informational advantage. Of course, an open question is the extent to which our results extend to other settings where the goods are not durable like housing. We view this as a fruitful avenue for future research.

The possible contribution of real estate investors to a housing bubble has led some policy makers to propose restrictions on their activity.⁴⁷ The appropriateness of such restrictions depends in part on the information content of the speculation, and we provide evidence in this paper that speculators did not, in fact, appear to be well informed of market fundamentals. Then again, such restrictions may inadvertently hamper the many welfare-enhancing roles that intermediaries play in markets that are subject to important search and informational frictions. In particular, we also show that middlemen improve the market's liquidity and capital stock, and policies that impinge on their behavior may be welfare reducing. A crucial element to evaluating welfare effects of flippers is the holding costs of sellers from whom they buy, a topic left for future research.

While regulating investor behavior as a function of the market cycle may be difficult, other policy tools may be available in the future. For example, by tracking investor activity as in our paper and Haughwout et al. (2011), it may be possible to generate real-time measures of speculative activity based on data from housing transactions and mortgages. Such measures might provide an indication that a speculative boom was underway in a given market, providing the just the kind of warning that was not widely available during the 2000s.

References

- Adelino, M., A. Schoar, and F. Severino. 2017. Loan originations and defaults in the mortgage crisis: The role of the middle class. *Review of Financial Studies* 29:1635–70.
- Amihud, Y. and H. Mendelson. 1986. Asset pricing and the bid-ask spread. *Journal of Financial Economics* 17:223–49.
- Barberis, N. and R. Thaler. 2003. A survey of behavioral finance. In *Handbook of the economics of finance*, eds. G. Constantinides, M. Harris, and R. Stulz, vol. 1. Amsterdam, the Netherlands: Elsevier.
- Bayer, P., K. Mangum, and J. W. Roberts. 2019. Speculative fever: Investor contagion in the housing bubble. Working Paper, Duke University.
- Bhutta, N. 2015. The ins and outs of mortgage debt during the housing boom and bust. *Journal of Monetary Economics* 76:284–98.
- Bostic, R., S. Gabriel, and G. Painter. 2009. Housing wealth, financial wealth, and consumption: New evidence from micro data. *Regional Science and Urban Economics* 39:79–89.
- Brunnermeier, M. K. and S. Nagel. 2004. Hedge funds and the technology bubble. *Journal of Finance* 59:2013–40.
- Case, K. E. and R. J. Shiller. 1987. Prices of single-family homes since 1970: New indexes for four cities. *New England Economic Review* Sept/Oct:45–56.
- . 1989. The efficiency of the market for single-family homes. *American Economic Review* 79:125–37.

⁴⁷ For example, a 2006 HUD regulation (Federal Register, volume 71, p. 33,138) prevented FHA financing for houses sold within 90 days of purchase. Partly in response to the weak housing market, HUD waived this restriction in 2010 (Federal Register, volume 75, p. 28,633). More broadly, antispeculative policy prescriptions, such as transaction taxes, have been suggested in other speculative markets. See, for example, Tobin (1974), Tobin (1978), Eichengreen, Tobin, and Wyplosz (1995), or Summers and Summers (1988).

- Cheng, I.-H., S. Raina, and W. Xiong. 2014. Wall street and the housing bubble. *American Economic Review* 104:2797–829.
- Chinco, A. and C. Mayer. 2016. Misinformed speculators and mispricing in the housing market. *The Review of Financial Studies* 29:486–522.
- Copeland, T. E. and D. Galai. 1983. Information effects on the bid-ask spread. *Journal of Finance* 38:1457–69.
- DeFusco, A. A., C. G. Nathanson, and E. Zwick. 2017. Speculative dynamics of prices and volume. Working Paper, University of Chicago.
- DeLong, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann. 1990. Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance* 45:379–95.
- Demsetz, H. 1968. The cost of transacting. *Quarterly Journal of Economics* 82:33–53.
- Eichengreen, B., J. Tobin, and C. Wyplosz. 1995. Two cases for sand in the wheels of international finance. *Economic Journal* 105:162–72.
- Fama, E. F. 1965. The behavior of stock-market prices. *Journal of Business* 38:34–105.
- Favilukis, J. Y. and S. Van Nieuwerburgh. 2018. Out-of-town home buyers and city welfare. Working Paper, University of British Columbia.
- Gao, Z., M. Sockin, and W. Xiong. 2017. Economic consequences of housing speculation. Working Paper, Chinese University of Hong Kong.
- Garber, P. M. 1989. Tulipmania. *Journal of Political Economy* 97:535–60.
- Garman, M. B. 1976. Market microstructure. *Journal of Financial Economics* 3:257–75.
- Glosten, L. R. and P. R. Milgrom. 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14:71–100.
- Glomer, M., D. R. Haurin, and P. H. Hendershott. 1998. Selling time and selling price: The impact of seller motivation. *Real Estate Economics* 26:719–40.
- Greenwood, R. and S. Nagel. 2009. Inexperienced investors and bubbles. *Journal of Financial Economics* 93:239–58.
- Griffin, J. M., J. H. Harris, T. Shu, and S. Topaloglu. 2011. Who drove and burst the tech bubble? *Journal of Finance* 66:1251–90.
- Griffin, J. M. and G. Maturana. 2016. Who facilitated misreporting in securitized loans? *Review of Financial Studies* 29:384–419.
- Haruvy, E., Y. Lahav, and C. N. Noussair. 2007. Traders' expectations in asset markets: Experimental evidence. *American Economic Review* 97:1901–20.
- Haughwout, A., D. Lee, J. Tracy, and W. van der Klaauw. 2011. Real estate investors, the leverage cycle, and the housing market crisis. Working Paper, Federal Reserve Bank of New York.
- Kindelberger, C. P. 1978. *Manias, panics and crashes: A history of financial crises*. New York: John Wiley & Sons.
- Kyle, A. S. 1985. Continuous auctions and insider trading. *Econometrica* 53:1315–35.
- Mackay, C. 1841. *Extraordinary popular delusions and the madness of crowds*, 1932 edition. London: Richard Bentley.
- Mayer, C., K. Pence, and S. M. Sherlund. 2009. The rise in mortgage defaults. *Journal of Economic Perspectives* 23:27–50.
- Mian, A. and A. Sufi. 2014. What explains the 2007-2009 drop in employment? *Econometrica* 82:2197–223.

- . 2017. Fraudulent income overstatement on mortgage applications during the credit expansion of 2002 to 2005. *Review of Financial Studies* 30:1832–64.
- . 2018. Credit supply and housing speculation. Working Paper, Princeton University.
- Piazzesi, M. and M. Schneider. 2009. Momentum traders in the housing market: Survey evidence and a search model. *American Economic Review* 99:406–11.
- Piskorski, T., A. Seru, and J. Witkin. 2015. Asset quality misrepresentation by financial intermediaries: Evidence from the rmbs market. *Journal of Finance* 70:2635–78.
- Rajan, R. and R. Ramcharan. 2015. The anatomy of a credit crisis: The boom and bust in farm land prices in the united states in the 1920s. *American Economic Review* 105:1439–77.
- Rognlie, M., A. Shleifer, and A. Simsek. 2018. Investment hangover and the great recession. *American Economic Journal: Macroeconomics* 10:113–53.
- Shiller, R. J. 2003. From efficient markets theory to behavioral finance. *Journal of Economic Perspectives* 17:83–104.
- Shleifer, A. and L. H. Summers. 1990. The noise trader approach to finance. *Journal of Economic Perspectives* 4:19–33.
- Smith, V. L., G. L. Suchanek, and A. W. Williams. 1988. Bubbles, crashes, and endogenous expectations in experimental spot markets. *Econometrica* 56:1119–51.
- Springer, T. M. 1996. Single-family housing transactions: Seller motivations, price, and marketing time. *Journal of Real Estate Finance and Economics* 13:237–54.
- Spulber, D. F. 1996. Market-making by price setting firms. *Review of Economic Studies* 63:559–80.
- Summers, L. H. and V. P. Summers. 1988. When financial markets work too well: A cautious case for a securities transactions tax. *Journal of Financial Services Research* 3:261–86.
- Temin, P. and H.-J. Voth. 2004. Riding the south sea bubble. *American Economic Review* 94:1654–68.
- Tobin, J. 1974. *The new economics one decade older*. Princeton, NJ: Princeton University Press.
- . 1978. A proposal for international monetary reform. *Eastern Economic Journal* 4:153–9.
- Vissing-Jorgensen, A. 2003. Perspectives on behavioral finance: Does irrationality disappear with wealth? Evidence from expectations and actions. In *NBER Macroeconomics Annual*, eds. M. Gertler and K. Rogoff. Cambridge: MIT Press.