Arrested Development: A Theory of Supply-Side Speculation in the Housing Market*

Charles G. Nathanson
Harvard University
cnathans@fas.harvard.edu

Eric Zwick
Harvard University
zwick@fas.harvard.edu

February 4, 2012

Abstract

How were there large house price bubbles in cities with historically elastic housing supply? High raw land prices capitalizing optimistic beliefs about future housing demand curtailed supply in these cities. In cities with excess land relative to the current population, optimistic land speculators are the marginal buyers of real estate, making these cities more prone to housing bubbles than fully developed cities. In the latter, the marginal buyers are homeowners, who derive flow benefits from holding land in addition to prospective capital gains and so need not be especially optimistic. This theory matches the joint cross section of house and land prices during the recent U.S. housing bubble. Home builders, who were in the position to arbitrage high home prices by selling more houses, acted like land speculators by taking large, unhedged positions many years in advance of plans to build and sell. Less developed neighborhoods within fully built cities also show larger boom-bust cycles.

Keywords: bubbles, durable goods, heterogeneous beliefs, housing supply.

Preliminary and incomplete

*We thank Chris Allen, Nick Barberis, Jediphi Cabal, John Campbell, Raj Chetty, Fritz Foley, Drew Fudenberg, Robin Greenwood, Sam Hanson, Erik Hurst, Josh Lerner, Chris Mayer, Nicolas Retsinas, David Scharfstein, Arthur Segel, Andrei Shleifer, Alp Simsek, Amir Sufi, Adi Sunderam, and especially Ed Glaeser, David Laibson, and Jeremy Stein for their comments and support. We also thank Harry Lourimore, Joe Restrepo, Hubble Smith, Jon Wardlaw, Anna Wharton, and CoStar for enlightening conversations and data. Nathanson thanks the NSF Graduate Research Fellowship Program for financial support.

In the recent housing boom there were large house price bubbles[1] in cities without barriers to construction. Many of these cities had experienced decades of high population growth without an increase in house prices. For instance, Figure 1 shows that during the 1980-2000 period Las Vegas experienced enormous construction rates relative to other cities, while real house prices stayed nearly constant. Then, prices doubled between 2000 and 2006, while construction remained high. Far from being an isolated case, Las Vegas is an example of a peculiar stylized fact of the recent housing bubble: many of the *largest* bubbles occurred in cities with elastic housing supply. Figure 2 shows that many of the largest bubbles occurred in metro areas with the largest construction booms. How is a housing bubble possible in a city where builders can build an unlimited quantity of housing at a constant unit cost? Why would bubbles in such places ever be larger than bubbles in cities like Boston where the housing stock is roughly fixed?

Figure 1: Las Vegas House Prices and Construction

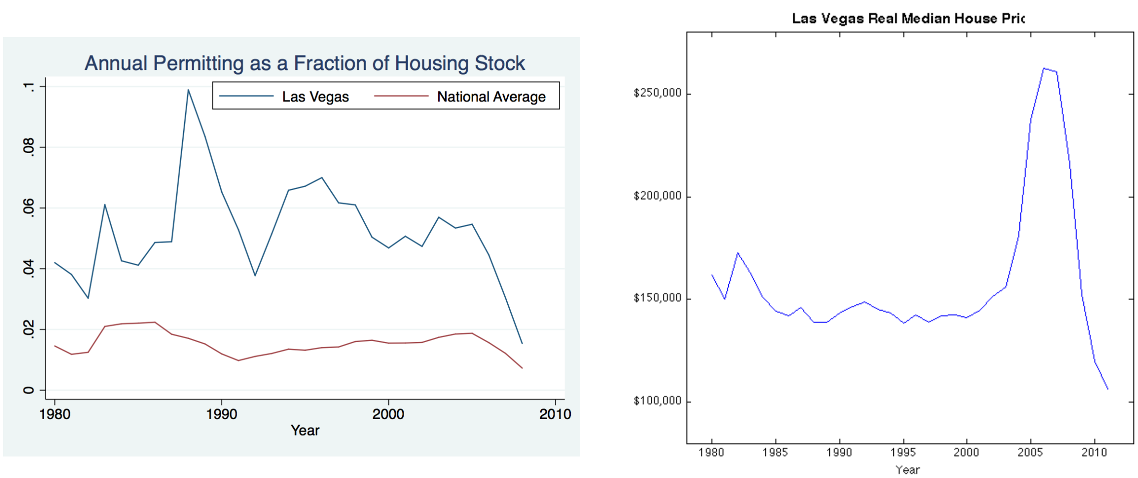

Unlike other housing inputs, land available for home building in a given city is in fixed supply. Were land in truly unlimited supply, competition would drives its price down to its outside use value, regardless of housing market conditions. In reality, regulations often restrict the amount of land that builders can develop, even in geographically "flat" cities with ample land. For instance, Phoenix is surrounded by Indian reservations and national forests, and the federal government refuses to sell land outside a 16-mile radius around Las Vegas. In both these cities, fringe land was trading at over $200,000 an acre in 2006, far

[1] The term "bubble" connotes a degree of irrationality in the market that we do not attempt to document nor wish to suggest. In fact, our theory features a competitive market with rational maximization by sophisticated actors, who happen to disagree about an uncertain, unrealized future. However, the term is so pervasive that adopting a new terminology would be more confusing.

Figure 2: Building Booms and Price Bubbles.

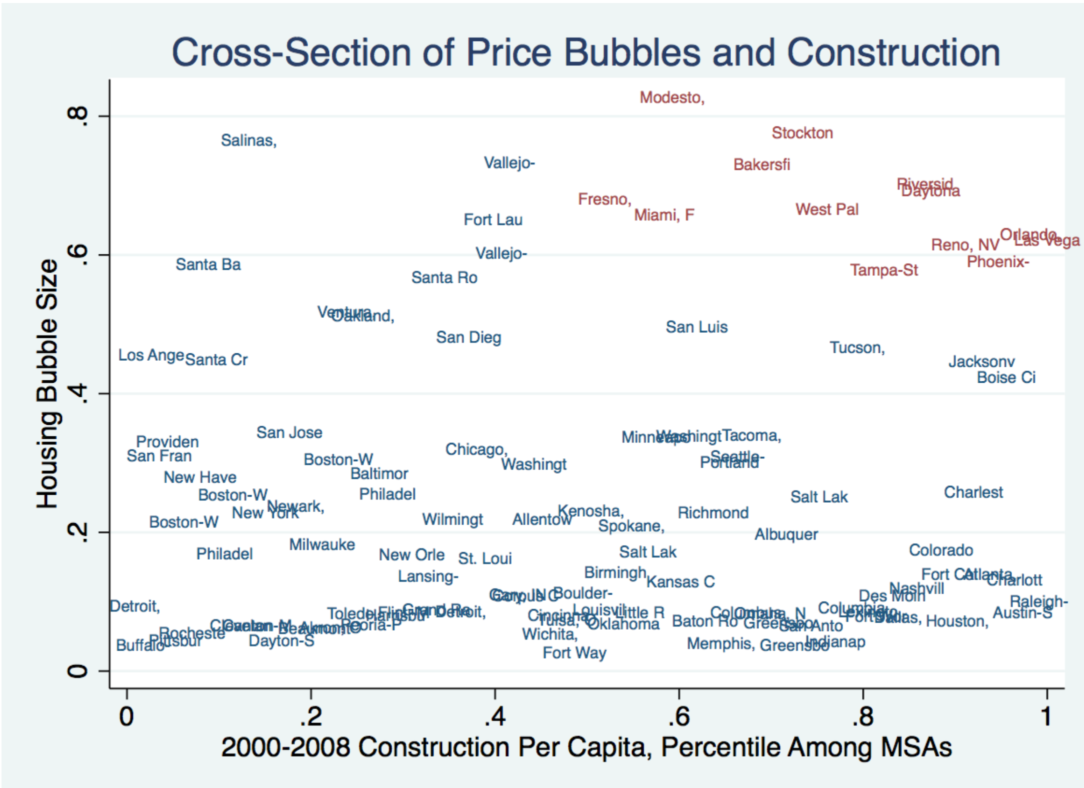

Notes: For each MSA, *Housing Bubble Size* is the minimum of the log price increase from 2000 until the local peak and the log price decrease from the local peak until 2011.

surpassing any possible outside use value for that land.

In this paper we develop a model in which homeowners and home builders/developers compete for a fixed quantity of land. For a homeowner, the value of land is the flow utility benefit he realizes from living in the city, plus the land's future resale value. Only the future resale value matters to a builder. The real estate development decision is dynamic: builders have the option to build on land immediately or to hold it for future investment. The value of land therefore equals expected discounted future house prices less construction costs. A builder who is optimistic about future house prices will have a high reservation price for using his land to build a house today. Therefore, builder optimism about future house prices can push up prices today.

Both homeowners and builders can disagree on the future growth of housing demand. We define the *bubble size* to be the ratio of the market-clearing price for housing to the average price implied by the demand forecasts of the homeowners and builders. We categorize cities based on the amount of land in the city beyond current homeowner demand, or *excess land*. Our main result is that the bubble size is *increasing* in excess land over a relevant range, in contrast to the standard model. Partially developed cities like Las Vegas and Phoenix are more prone to destabilizing optimism than fully developed cities like Boston and New York.

In markets with heterogeneous actors, the marginal buyers set the market price, while others may be far from an indifference point with respect to price. Whether the marginal buyers are representative of the population as a whole depends upon the joint distribution of beliefs, capital, and storage technology, as well as on asset float. With short-sale constraints, the marginal buyers may be far from representative.[2] Prices display more *optimism sensitivity* when a larger fraction of marginal buyers are optimists. When excess land is zero, homeowners are the exclusive landowners, because they have higher flow utility than the builders. With some excess land, builders enter. The first builders to enter are the most optimistic ones, so optimism sensitivity increases relative to the no land case.

The value of land, like any scarce resource, equals the utility flow from when supply is exhausted, discounted to the present. With no excess land, buyers can disagree on the growth path of this utility stream. With some excess land, buyers can disagree on *both the growth rate of the eventual stream and the discount to apply.* With more scope for disagreement, prices implied by optimistic beliefs are further from the average price consistent with everyone's beliefs. We refer to this ratio as *optimism significance.* Optimism significance goes to zero as excess land and hence the time until exhaustion continue to grow. We stress this limiting result as the primary reason that bubble size goes to zero for the least developed cities.

Both optimism sensitivity and optimism significance rise and fall with excess land, implying that bubbles are likely to be larger in partially developed cities. We confirm the predictions of our model in the cross-section of metropolitan areas, using MSA-level price indices for raw land developed by Nichols et al. (2010). Land price returns were large during the 1996-2008 period, averaging an annual 9 percent real return. Land price returns exceeded house price returns, consistent with the idea of land price increases as a driving force behind the housing bubble. Among metro areas that experienced housing bubbles, the ones with more land experienced larger bubbles. Metros with lower initial housing prices in 2000 and thus more land experienced larger bubbles. This relationship holds nationally, within California, as well as within fully developed cities. In the latter setting, neighborhoods with a smaller frequency of owner-occupied units had much larger booms. In addition, the more established neighborhoods in terms of owner-occupancy at the end of the boom experienced smaller crashes.

We argue that builder optimism about future housing demand was an important supply constraint in the less developed cities. When builder-speculators dominate the margin, their beliefs determine house prices. Housing demand shocks cause increased construction, but

[2]See, e.g., Harrison and Kreps (1978), Fostel and Geanakoplos (2008), Simsek (2011).

outside of their effect on builder beliefs, housing demand shocks do not raise prices. Several recent papers have focused on homeowner demand to explain the housing bubble.[3] These demand shocks may have raised prices through influencing the demand forecasts of builders, but not through a direct effect on prices. As we document, builders assembled substantial unhedged land portfolios. According to financial statements, the land holdings of the eight largest home builders tripled between 2001 and 2005, and increased fifty percent controlling for increasing construction needs. The seventy percent decline in market equity values for these firms during the subsequent period is mostly due to land losses.

1 Related Literature

Our work contrasts with Glaeser et al. (2008), who show that, on average, less elastic cities had larger housing bubbles. In their model, home builders build as much as is profitable in response to a housing demand shock, given increasing construction costs that result from regulation or technological constraints. Our theory is consistent with the general cross-sectional relationship between elasticity and bubble size that they find, but we succeed in explaining the major outliers—elastic cities that experienced large bubbles. We model regulation as limiting a stock variable, the long-run quantity of buildable land, whereas they model regulation as limiting a flow variable, the quantity of new construction each period. By focusing on the regulation of a stock, we capture the relationship between demand forecasts and prices. While we do not doubt the importance of construction inelasticities in explaining the cross-section of housing prices in the United States, flow constraints would not predict rising land prices, which we show explain much of the housing bubble. Our work also departs from Glaeser et al. (2008) in highlighting the belief heterogeneity of home builders. We argue *against* the image of home builders as arbitrageurs who make enormous profits during the bubble on the spread between housing prices and construction costs.

Our paper fits into a behavioral finance literature in which investors with heterogeneous beliefs compete for an asset. Because they cannot short this asset, optimists price the asset and push its value above some average expectation of the market.[4] Other papers in this literature give further reasons why short-sales constraints can lead prices to diverge from fundamentals. In Harrison and Kreps (1978) and Scheinkman and Xiong (2003), the option to resell the asset to a greater optimist pushes up the asset's price. Simsek (2011) argues that optimists' use of credit allows them to push up asset prices. Scharfstein and Stein

[3]Case and Shiller (2010) document extraordinary homeowner beliefs (as reported in surveys) about future house price appreciation, and Himmelberg et al. (2005) and Mian and Sufi (2009) provide theoretical and empirical evidence that increased credit for homeowners led house prices to rise.

[4]Miller (1977), Hong and Stein (2003).

(1990) and Chevalier and Ellison (1999) argue that career concerns of managers prevent them from deviating from the consensus view of their competitors. All of these mechanisms would strengthen our results.

2 Optimism Sensitivity

We model the housing market in a general way that applies to any financial market in a durable consumption good. In our model, a consumption good is produced from an input that is in fixed supply. The production process is completely elastic. At a constant per-period cost, a consumer can turn the input into the consumption good. The input is completely durable. In the housing market, the input is land that can be used for construction in a metro area. The consumption good is a house.

In our model, the input is a traded financial asset. There are two types of agents who compete for this asset: "end-users" and "speculators." Both end-users and speculators care about the future resale value of the asset. But end-users, unlike speculators, receive flow utility from holding the asset. This flow utility equals the flow utility an end-user receives from the consumption good, minus the cost of assembling the consumption good from the input asset. In the case of housing, end-users are residents living in a metro area. The flow utility they receive includes access to a city's labor market and to the city's amenities. Speculators are land developers and home builders that hold land portfolios as investments.

The two key properties of flow utility in our model are that it (1) is heterogeneous, and (2) exhibits diminishing marginal returns over the quantity of the asset. Heterogeneity guarantees that differences in flow utility, and not just differences in beliefs about future prices, determine which end-users hold the asset. When flow utility is very heterogeneous, it is the sole factor that determines which end-users end up with the asset. The marginal end-users need not be especially optimistic. The diminishing marginal returns feature guarantees that asset float can increase to the point where end-users no longer desire the asset and optimistic speculators become the marginal investors. As float increases, more optimistic speculators enter and the asset's optimism sensitivity increases.

In reality, the end-user and speculator classes overlap somewhat. Homeowners in many cities acted as speculators by flipping houses as well as by entering the land market directly. These homeowners acted as end-users and as speculators. In our model, what matters are the number of houses demanded by end-users, and the amount of speculative capital held by speculators. A homeowner who also bought a second house to speculate on real estate would count toward both populations. We adopt linear utility so that the utility is separable over a first unit of the asset that delivers flow utility and a second unit that does not.

Our model contains two periods, 0 and 1. In period 0, end-users and speculators compete for a fixed quantity L of the asset. The market-clearing price in period 0 is p_0, and an investor's expectation of the period 1 price is $\tilde{p_1}^i$. For now, these expectations are exogenous, but we will endogenize them in Section 3.

2.1 End-users

End-users receive idiosyncratic flow utility v^i in period 0. We model diminishing marginal returns bluntly by assuming that an end-user receives flow utility from holding only one unit of the asset. The total number of end-users at time 0 is N, and the distribution of flow utility is given by a continuous distribution with a probability density function f_v and a cumulative distribution function F_v. In this section, this distribution takes the form

$$v^i \sim U[0, v^H]$$

for some $v^H > 0$. End-users are risk-neutral, fully forward-looking, and do not discount the future. We use risk-neutrality to stress that diminishing marginal flow utility, and not finite risk capacity, limits end-user demand for the asset. For the reasons discussed above, we rule out the option for an end-user to buy more than one unit of the asset. Thus, his decision reduces to the binary choice of whether or not to buy one unit of the asset. He buys if

$$v^i + \tilde{p_1}^i > p_0.$$

2.2 Speculators

Speculators are firms that raise capital from investors who want exposure to the asset. They differ from end-users in two ways. First, these firms do not derive any flow utility from the asset; they hold it only for capital gain. Second, speculators buy multiple units of the asset. There are no diminishing returns in the capital gain dimension. The cost of raising capital limits speculators. Each speculator is simply limited to k units of capital. We let $m = k/p_0$ denote the maximum quantity of the asset a speculator can purchase with his limited capital. The quantity m is endogenous to the market outcome, but we can compare speculators and end-users more easily when we write the speculator constraint in terms of m instead of k. Like end-users, each speculator holds an idiosyncratic belief $\tilde{p_1}^j$ about the period 1 price of the asset. There are M/m firms, so that the maximum quantity of the asset speculators can hold is M.

A speculator chooses an asset position q^j to maximize expected profits,

$$\Pi \equiv \left(p_0 - \tilde{p_1}^j\right) q^j,$$

subject to the capital constraint $q^j \leq m$ and the no-shorting constraint $q^j \geq 0$. The resulting demand schedule is

$$q^j = \begin{cases} m & p_0 \leq \tilde{p_1}^j \\ 0 & p_0 > \tilde{p_1}^j. \end{cases}$$

2.3 Beliefs

End-users and speculators hold heterogeneous beliefs about p_1. We allow the distribution of beliefs within each group to be distinct so that we can discuss the relevance of each group's beliefs for the market outcome. The cumulative distribution functions of these beliefs are F_p^U and the F_p^S for the end-users and speculators, respectively. Both distributions have compact support. The support for F_p^U is $[\tilde{p_1}^{L,U}, \tilde{p_1}^{H,U}]$ and the support for F_p^S is $[\tilde{p_1}^{L,S}, \tilde{p_1}^{H,S}]$. In the spirit of Harrison and Kreps (1978) and Morris (1996), agents "agree to disagree." They refuse to update their beliefs based on the beliefs of other market participants, even though the information structure is common knowledge. These heterogeneous beliefs arise as a result of some unprecedented information shock of the sort described by Morris (1996).

Flow utility is important enough so that the investors who value the asset the most are end-users. This assumption guarantees that only end-user purchase the asset when the float is sufficiently small. The inequality

$$v^H > \tilde{p_1}^{H,S} - \tilde{p_1}^{H,U} \tag{1}$$

holds, which guarantees that the highest value investors are always end-users.

2.4 Market Equilibrium

Total end-user demand is given by the number of end-users whose beliefs and flow utility satisfy $v^i + \tilde{p_1}^i > p_0$:

$$D^U(p_0) = N \int_{\tilde{p_1}^{L,U}}^{\tilde{p_1}^{H,U}} \left(1 - F_v(p_0 - \tilde{p_1})\right) f_p^U(\tilde{p_1}) d\tilde{p_1}.$$

Recall that speculator j buys m units of the asset if $\tilde{p_1}^j > p_0$, and 0 units otherwise. Therefore, total speculator demand equals the fraction of speculators for which $\tilde{p_1} > p_0$, times the total quantity M of the asset speculators can buy:

$$D^S(p_0) = M \int_{p_0}^{\tilde{p_1}^{H,S}} f_p^S(\tilde{p_1}) d\tilde{p_1}.$$

The market clears when p_0 satisfies $D^U(p_0) + D^S(p_0) = L$.

When L is sufficiently low, only end-users will hold the asset. End-users receive flow utility that speculators do not, and inequality (1) guarantees that this flow utility is large enough so that the highest value investors are end-users. When $p_0 > \tilde{p_1}^{H,S}$, speculator demand will be 0, but end-user demand will be positive.

Lemma 1. *There exists an asset float L^* such that speculator demand is 0 if $L < L^*$ and positive if $L > L^*$. This cutoff asset float is given by $L^* \equiv D^U(\tilde{p_1}^{H,S}) > 0$.*

The market is fundamentally different depending on whether the float L exceeds the cutoff L^*. When $L < L^*$, end-users are the sole holders of the asset. The asset is allocated according to which end-users have the highest sum of flow utility and belief about p_1. When flow utility is much more heterogeneous than beliefs, flow utility becomes the sole determinant of which end-users hold the asset. As we will show, prices reflect the average end-user belief in this case, not the belief of some top portion of the belief distribution. When $L > L^*$, speculators and end-users hold the asset. Both optimistic speculators and pessimistic end-users are the marginal investors. When the marginal speculators outnumber the marginal end-users, prices will look very optimistic. We formalize this intuition by examining the market equilibrium in each of these two cases.

2.4.1 Small Asset Float: $L < L^*$

Flow utility becomes more heterogeneous when v^H gets large. The following result states that when flow utility heterogeneity is large relative to belief heteorgeneity, the price reflects sorting among the flow utility dimension and not along the belief dimension.

Proposition 1. *Assume that the spread in flow utility is much larger than the spread in beliefs: $v^H \gg \tilde{p_1}^{H,U} - \tilde{p_1}^{L,U}$. Then the current asset price reflects the average end-user belief:*

$$p_0 = F_v^{-1}\left(1 - \frac{L}{N}\right) + \mu_p^U,$$

where μ_p^U is the average value of $\tilde{p_1}$ among end-users.

Proposition 1 shows that information is perfectly aggregated when flow utility heterogeneity is the dominant heterogeneity in the market. Even though the asset permits no short-selling and agents differ in their beliefs, prices depend only on the average belief.

The intuition for this result is as follows. The result must hold when beliefs are homogeneous. In this case, there are L units of the asset to be allocated among $N > L$ end-users. The marginal end-user has flow utility equal to $F_v^{-1}(1 - L/N)$, and his total value for the asset is this flow utility plus the common belief μ_p^U. Now consider the change in end-user demand when small, mean-preserving heterogeneity in $\tilde{p}_1$ is added. End-users whose beliefs fall above μ_p^U will increase their demand, meaning that optimistic end-users with lower flow utility will begin purchasing the asset. Similarly, the demand of end-users whose beliefs fall below μ_p^U will reduce their demand. Because flow utility is uniform, the change in demand for a given belief is linear in the difference between the belief and the average belief. Therefore, all the demand changes cancel, and the market remains at equilibrium when belief heterogeneity is introduced.

This argument fails when belief heterogeneity is so large that pessimist demand hits zero. In this case, only the end-users with the highest flow utility *and* the highest beliefs hold the asset. When flow utility heterogeneity is much larger than belief heterogeneity, the asset float must be extremely small for the asset to be allocated in this manner. To see this, note that the price point at which some pessimists sit out the market is $p_0 = v^H + \tilde{p}_1^{L,U}$. By Proposition 1, the asset float at this price is

$$L' = \frac{\mu_p^U - \tilde{p}_1^{L,U}}{v^H} N.$$

When flow utility heterogeneity is much larger than belief heterogeneity, the asset float L' will be nearly 0, because $v^H \gg \tilde{p}_1^{H,U} - \tilde{p}_1^{L,U} \geq \mu_p^U - \tilde{p}_1^{L,U}$. The higher the ratio of flow utility heterogeneity to belief heterogeneity, the smaller asset float must be before the asset is rationed according to beliefs. Even when the asset float is very small, prices still reflect the average end-user belief if flow utility is sufficiently heterogeneous.

2.4.2 Large Asset Float: $L > L^*$

When asset float is large, both speculators and end-users hold the asset. Their demands jointly determine prices via the equation $D^S(p_0) + D^U(p_0) = L$. *When marginal speculators are much more frequent than marginal end-users, prices are sensitive to speculator optimism.* To understand this result, suppose that there are no end-users on the margin, so that end-user demand is fixed at D^U. Then speculators compete to own the remaining $L - D^U$ units of the asset. Speculators differ only on their beliefs about p_1. Therefore, the most optimistic speculators end up owning the asset. Speculators can hold a maximum of M units of the asset, so the price is given by

$$p_0 = \left(F_p^S\right)^{-1}\left(1 - \frac{L - D^U}{M}\right) \tag{2}$$

where F_p^S is the cumulative distribution function of speculator beliefs. The asset price simply equals the marginal speculator's belief about p_1. This pricing equation sharply differs from the pricing equation when the asset float is large given by Proposition 1. When asset float is small, the price reflects the *flow utility* of the marginal end-user, plus the average end-user belief. When asset float is large, the marginal investors no longer receive any flow utility because they are speculators. The asset is instead allocated according to whichever speculators are the most optimistic. Price reflects the belief of the marginal speculators.

When no end-users are on the margin, the supply of the asset to end-users appears elastic, in that increases in end-user demand are matched one-for-one with increases in the quantity of the asset held by end-users. Indeed, if no end-users are marginal, then new end-users who arrive after a demand shock will all want to purchase the asset at the current price. Speculators will sell to these new end-users. In the example of housing, such an event is a construction boom. When home builder speculators sell their land to new end-users, construction takes place. A construction boom is consistent with land prices that reflect a high degree of optimism.

Supply also appears elastic when the speculator capital M is large. When speculator capital is large, the speculators who hold the asset do not vary in their beliefs. Indeed, the most optimistic speculators hold all of the asset, so that the heterogeneity of beliefs among speculator asset holders is low. Therefore, speculators are willing to sell a large quantity of the asset at the market price; supply appears elastic. When end-user demand increases, the quantity of the asset held by the end-users increases, but the asset price does not. The most extreme example of this phenomenon occurs when M is so large that the single most optimistic speculator firm is the only speculator firm holding the asset, and is willing to sell an unlimited quantity of land at its reservation price.

We now formally analyze the case in which some end-users are on the margin. The first result measures the responsiveness of price to changes in speculator optimism.

Proposition 2. *Suppose the beliefs of all speculators who are at least as optimistic as the marginal speculator increase. If speculator capital is large, or if few end-users are on the flow utility margin, then this belief shock fully passes through to prices. Specifically, let x_0 be the percentile in the belief distribution of the marginal speculator: $x_0 = 1 - \left(L - D^U(p_0)\right)/M$. Suppose the speculator belief distribution F_p^S is replaced by F_p^{S*} in which the beliefs of speculators at least as optimistic as the marginal speculator increase by β: $\left(F_p^{S*}\right)^{-1}(x) = \beta \mathbf{1}_{x \geq x_0} +$*

$\left(F_p^S\right)^{-1}(x)$. *Then*

$$\frac{\partial p_0}{\partial \beta} = \frac{M}{M + \lambda N}$$

where

$$\lambda = \frac{f_v}{f_p^S(p_0)} F_p^U(p_0).$$

If speculator capital M is much larger than the number of end-users N, of if λ is near 0, then the pass-through of the belief shock $\partial p_0/\partial \beta = 1$.

The intuition of Proposition 2 involves two steps. First, raising the beliefs of the speculators above x_0 in the belief distribution has the same effect on price as raising the beliefs of *all* speculators. Indeed, only the speculators above x_0 influence prices. Therefore, when speculators above x_0 in the belief distribution increase their optimism by β, the entire speculator demand curve shifts out by β, and the new equilibrium satisfies $D^S(p_0 - \beta) + D^U(p_0) = L$. The comparative static in the proposition follows from this equation:

$$\frac{\partial p_0}{\partial \beta} = \frac{\partial D^S/\partial p_0}{\partial D^S/\partial p_0 + \partial D^U/\partial p_0} = \frac{M f_p^S(p_0)}{M f_p^S(p_0) + N F_p^U(p_0) f_v} = \frac{M}{M + \lambda N}.$$

The greater speculator price sensitivity is relative to end-user sensitivity, the higher this comparative static is. When speculators are relatively more price sensitive, movements in their willingness to pay pass through more fully to prices. Speculators are more price sensitive under two circumstances. The first circumstance is when the aggregate size M of speculators in the market is much larger than the number of end-users N. Price sensitivity scales with M and N, so when $M \gg N$, speculators are much more price sensitive as a population than end-users. The second circumstance under which speculators are more price sensitive is when the share of speculators on the margin is larger than the share of end-users on the margin. The parameter λ captures this relationship. The more heterogeneous flow utility is relative to the heterogeneity in speculator beliefs, the lower is f_v/f_p^S and the closer λ is to zero. In the extreme case where f_v approaches 0, there are no homeowners on the margin, and speculators belief shocks pass through fully to prices.

Speculator capital increases the sensitivity of prices to the beliefs of speculators at least as optimistic as the marginal speculator. In addition, speculator capital increases the optimism of the marginal speculator. When speculator capital is high, the most optimistic speculators can control more of the asset. The next proposition makes this point.

Proposition 3. *As speculator capital increases, the marginal speculator becomes more optimistic. As speculator capital goes to infinity, the marginal speculator becomes the most optimistic speculator in the population. Specifically, if $x_0 = 1 - \left(L - D^U(p_0)\right)/M$, then*

$\partial x_0/\partial M > 0$ *and* $\lim_{M\to\infty} x_0 = 1$.

Proposition 2 shows that when speculators dominate the margin, movements in their beliefs pass through fully to prices. The next proposition shows that under these conditions, movements in end-user beliefs have no direct effect on prices. When speculators are very price sensitive, movements in end-user demand change quantities, not prices. Suppose the beliefs of all end-users shift up by β, while the beliefs of speculators stay unchanged. The new market equilibrium is given by $D^S(p_0) + D^U(p_0 - \beta) = L$. Differentiating this equation with respect to β gives the next proposition:

Proposition 4. *Suppose end-user beliefs increase. If speculator capital is large, or if few end-users are on the flow utility margin, then this end-user belief shock has no effect on prices, but does increase the quantity of the asset held by end-users. Specifically, suppose all end-user beliefs increase by β: F_p^U is replaced by $F_p^{U^*}$ satisfying $\left(F_p^{U^*}\right)^{-1}(x) = \beta + \left(F_p^U\right)^{-1}(x)$. Then*

$$\frac{\partial p_0}{\partial \beta} = \frac{\lambda N}{M + \lambda N}$$

and

$$\frac{\partial D^U}{\partial \beta} = -\frac{M}{M + \lambda N}\frac{dD^U}{dp}.$$

If speculator capital M is much larger than the number of end-users N, of if λ is near 0, then the price pass-through of the belief shock $\partial p_0/\partial \beta = 0$, and end-users increase their holdings of the asset.

Speculators dominate the margin when $M \gg N$ or when λ is near 0. In either case, a shock to end-user beliefs leaves prices unchanged and increases the quantity of the asset they hold. Indeed, when $M/(M + \lambda N) = 1$, an end-user belief shock is like a price cut for end-users that does not appear in posted prices. Proposition 4 shows the conditions under which homeowner beliefs lacked a direct price effect during the housing bubble. If speculator capital was large or homeowner flow utility heterogeneity was large, then optimistic homeowner demand simply caused additional construction, not higher prices. Under these conditions, homeowner beliefs were important only insofar as they affected builder beliefs. The homeowner beliefs themselves had no direct effect on prices.

When speculators dominate the margin, a shock to the number of end-users also passes through fully to new quantities. We capture this shock as a percentage increase in N, the number of end-users who receive positive flow utility from the asset.

Proposition 5. *Suppose the number of end-users N increases. If speculator capital is large, or if few end-users are on the margin, then this demand shock passes through one-for-one to* *the quantity of the asset held by end-users:*

$$\frac{\partial \log D^U}{\partial \log N} = \frac{M}{M + \lambda N}. \tag{3}$$

If speculator capital is large, then this end-user demand shock has no effect on prices:

$$\frac{\partial p_0}{\partial \log N} = \frac{D^U}{M + \lambda N}.$$

The intuition of Proposition 5 is as follows. When λ is near 0, very few end-users are on the margin. Therefore, when additional end-users demand the asset, they outbid the speculators who are the marginal investors. Hence, $\partial \log D^U / \partial \log N = 1$ when $\lambda = 0$. Prices rise because speculators give up some of the asset, so that the marginal speculator is now more optimistic. When M is large, this price effect is small, because all of the speculators who hold the asset look like the marginal speculator. They are all very optimistic. Therefore, $\partial p_0 / \partial \log N = 0$ only when M is large.

Table 1 summarizes the market effects of large speculator capital M and relatively inelastic end-user demand represented by λ. Large speculator capital makes the asset price fully reflect very optimistic speculator beliefs. When speculator capital is large, the market appears elastic and end-user demand has no effect on prices. The fewer end-users are on the margin, the stronger this elasticity effect is.

Table 1: Summary of Optimism Sensitivity Theory

	Large M	Small λ
Prices		
Marginal speculator very optimistic	X	
Full pass-through of speculator belief shock	X	X
No direct effect of end-user belief shock	X	X
No pass-through of fundamental end-user demand shock	X	
Quantities		
Full pass-through of end-user demand shock	X	X

3 Optimism Significance

In the previous section, we derived conditions under which the asset's price reflects only the beliefs of the most optimistic investors. In this section, we explore when this bias is relevant for the price of the consumption good, which equals the asset price plus an assembly cost. We define *optimism significance* as the ratio of the good's price implied by optimistic beliefs to the average of the prices implied by people's beliefs. Formally, if p is the good's price and people's beliefs are drawn from a distribution with a cumulative distribution function of F and a probability density function of f, then

$$\text{os}(x) \equiv \frac{F^{-1}(x)}{\int_{-\infty}^{\infty} pf(p)dp}. \tag{4}$$

Our main result in this section is that as the asset's float goes to infinity, optimism significance goes to 1 for any value of x. The good's price equals the assembly cost, and bias in the asset's price becomes irrelevant. We also prove an additional result: under certain conditions, optimism significance can be *increasing* in asset float over some range.

To analyze optimism significance, we build a dynamic exhaustible resource model of the asset's price. The price depends on the float of the asset, the current end-user demand, and a forecast of future end-user demand. The model contains overlapping generations of end-users and speculators who always sell to investors in the next period. End-user demand grows over time. Therefore, there exists a future time at which end-users will hold all of the asset. At that point, the asset must be rationed among end-users. The number of end-users receiving flow utility in excess of the assembly cost exceeds the float of the asset. Therefore, the asset begins to pay dividends equal to the flow utility of the marginal investor each period. The faster this marginal flow utility is expected to grow, the higher the asset's price is today. When some speculators still hold the asset, the asset's price equals the eventual dividend stream, times a discount factor accounting for the waiting time until the asset pays these dividends.

We solve for the good's price in this model when the future path of demand is completely certain. We then explore how sensitive the good's price is to changes in this known future demand path. This sensitivity goes to zero as the asset float goes to infinity. We also show that over some range, this sensitivity increases with the asset float. By solving the model under certainty, we abstract from complications such as strategic trading and overbuilding and instead focus on the effect of the demand forecast on the good's price.

The previous section explored the dynamics of the asset price p. Here we are also interested in the final good price. We distinguish these prices with superscripts: p^A denotes the asset price and p^G denotes the good price.

3.1 End-users

An end-user who holds the asset pays a per-period assembly cost c to consume the good. He receives utility equal to $u^i = v^i + c$ from consuming the good for one period. In the previous section, v^i denoted the flow utility an end-user received from holding the asset. We maintain that notation by introducing u^i, the flow utility the end-user receives from consuming the good. In period t, there are N_t end-users for whom $u^i > c$. The distribution of flow utility among these end-users has a cumulative distribution function F_t^u. The future path of N_t and F_t^u affect today's asset and good prices. Both N_t and F_t^u are always growing and never decrease. End-users are risk-neutral and forward looking, but unlike the previous section's model, end-users now discount the future at a constant factor $\delta < 1$. Without this discount factor, asset prices would be infinite in the presence of demand growth. An end-user can buy at most one unit of the asset. He buys if

$$v^i + \delta p_{t+1}^A > p_t^A.$$

The price of the good equals the asset price plus the cost of assembling the good forever:

$$p_t^G = p_t^A + \frac{c}{1-\delta}. \quad (5)$$

We can use this equation to restate the end-user decision in terms of the good cost. An end-user buys the asset if

$$u^i + \delta p_{t+1}^G > p_t^G.$$

The utility from the good must exceed the difference between today's good price and tomorrow's discounted good price.

3.2 Speculators

Speculators are firms that hold the asset but receive no flow utility from it. These firms hold the asset only for capital gain. Like end-users, they discount the future with a factor δ. This discount factor comes from the existence of a risk-free bond in elastic supply whose gross per-period interest rate is $1/\delta$. Each speculator holds k units of capital. We let $m_t = k/p_t^A$ denote the maximum quantity of the asset a speculator can purchase with his limited capital. Because there is no uncertainty, each speculator simply chooses an asset holding q to maximize

$$\Pi \equiv (p_t^A - \delta p_{t+1}^A)q_t$$

subject to the capital constraint $q_t \leq m_t$ and the no-shorting constraint $q_t \geq 0$. The resulting demand schedule is

$$q_t = \begin{cases} m_t & p_t^A < \delta p_{t+1}^A \\ \in [0, m_t] & p_t^A = \delta p_{t+1}^A \\ 0 & p_t^A > \delta p_{t+1}^A. \end{cases}$$

The number of speculator firms is M_t/m_t, so that the maximum quantity of the asset speculator firms can hold at time t is M_t. We assume that at any time t, $M_t > L$, the total float of the asset. This inequality guarantees that $p_t^A < \delta p_{t+1}^A$ can never hold in equilibrium. Were this relationship between p_t^A and p_{t+1}^A to hold, total speculator demand M_t would exceed the float of the asset L and the market would not clear. In equilibrium, speculators hold the asset if and only if $p_t^A = \delta p_{t+1}^A$. The asset's price path is certain, so speculators hold the asset if and only if it yields the risk-free rate.

3.3 Market Equilibrium

When $N_t > L$, the number of end-users who receive positive flow utility from holding the asset exceeds the float of the asset. Therefore, only end-users hold the asset. At this point, the price path of land must obey

$$L = \left(1 - F_t^u(p_t^G - \delta p_{t+1}^G)\right) N_t,$$

so that

$$p_t^G = \delta p_{t+1}^G + (F_t^u)^{-1} \left(1 - \frac{L}{N_t}\right). \tag{6}$$

Define

$$u_t^* \equiv (F_t^u)^{-1} \left(1 - \frac{L}{N_t}\right).$$

The quantity u_t^* represents the dividend that the good pays at time t. It equals the flow utility of the good to the marginal end-user. Note that $u_t^* > c$ when $N_t > L$: the marginal end-user receives flow utility in excess of the assembly cost because the good is being rationed. Equations (5) and (6) therefore confirm that no speculator will hold the asset when $N_t > L$:

$$p_t^A - \delta p_{t+1}^A = p_t^G - \delta p_{t+1}^G - c = u_t^* - c > 0.$$

Using (6), we can write the good's price as

$$p_t^G = \sum_{i=0}^{\infty} \delta^i u_{t+i}^*.$$

The price of the good equals the discounted sum of the future dividend stream. Suppose we know that u_t^* will grow with a growth factor $g > 1$. This growth represents either a rising number of end-users N_t who will want the good in the future, or rising flow utilities of end-users in the future, captured by an outward shift of F_t^u. Then

$$p_t^G = D_g u_t^*$$

where $D_g \equiv (1 - \delta g)^{-1}$. The quantity D_g captures beliefs about future demand.

When $N_t < L$, some speculators hold the asset. Because N_t is growing, they know that at some future point, $N_t = L$. Suppose this equality will hold in τ periods. At the point this equality holds, $u_{t+\tau}^* = c$ because there are just enough end-users with flow utility for the good exceeding c to purchase the entire float of the asset. Speculators must be indifferent between selling the asset at time t or holding it until $t + \tau$. Therefore $p_t^A = \delta^\tau p_{t+\tau}^A$, which implies that

$$p_t^G - \frac{c}{1-\delta} = \delta^\tau \left(p_{t+\tau}^G - \frac{c}{1-\delta} \right) = \delta^\tau \left(D_g c - \frac{c}{1-\delta} \right),$$

or more simply that

$$p_t^G = \left((1 - \delta^\tau) D_0 + \delta^\tau D_g \right) c \tag{7}$$

where $D_0 = (1 - \delta)^{-1}$. The good price is a weighted average of the price that would obtain if price equaled construction cost forever, $D_0 c$, and the price that obtains when land runs out, $D_g c$. The weight depends on the amount of time τ until land runs out.

3.4 Beliefs

We now examine the effect of beliefs about future demand on the good's price. Two factors in the pricing equations capture these beliefs: D_g and δ^τ. D_g represents beliefs about the growth rate of marginal flow utility once end-users hold all the asset. In terms of housing, D_g corresponds to a price-rent ratio. It is the ratio of the market price for housing to the imputed rent of the marginal homebuyer. This ratio is higher the more quickly the number of people who want to live in the city is expected to grow, or the more quickly the willingness to pay to live in the city is expected to grow. The second factor representing beliefs about demand growth is δ^τ. The factor δ^τ represents beliefs about the length of time until end-users purchase the entire float of the asset. In terms of housing, δ^τ represents beliefs about population growth in cities with excess land. The greater is expected population growth, the higher δ^τ will be.

An investor's belief about τ will change depending on L. For instance, suppose an investor believes that N_t will grow with a growth rate g_N forever. Then $\tau = \log_{g_N} (L/N_t)$, and τ will rise with L. Note that in this example, $\lim_{L\to\infty} \tau = \infty$ for any value of g_N. As L increases, the belief of a given investor about τ continues to increase as well. The key assumption we make is that this limiting relationship holds for whatever beliefs investors have about τ:

Assumption 1. *A given investor's belief about τ has the property that* $\lim_{L\to\infty} \delta^\tau = 0$.

When this assumption holds, the optimism reflected in the asset price becomes irrelevant for the good price as L gets large. Indeed, equation (7) shows that as $\delta^\tau \to 0$, the good's price goes to $D_0 c$, the cost of assembling the good forever. For any level of optimism reflected in prices, good prices always approach the long-run assembly cost as L gets sufficiently large. The following proposition formally states this conclusion.

Proposition 6. *As the float of the asset gets large, beliefs about future end-user demand become irrelevant for pricing the consumption good. For any distribution of beliefs,* $\lim_{L\to\infty} p_t^G = D_0 c$, *the long-run assembly cost of the good.*

Proposition 6 shows that the results of Section 2 are irrelevant for cities with enough land. Cities where land prices are low cannot experience a large house price bubble because land prices, which capitalize optimistic beliefs, are a small portion of house prices.

We now show that under certain conditions, optimism significance *rises* with L over a range. This result demonstrates that the limiting result of Proposition 6 does not set in immediately as the asset float increases. When $L < N_t$, investors disagree only about D_g. When L grows larger past N_t, investors now disagree about D_g and δ^τ. The added dimension of disagreement can make optimistic beliefs deviate even more from the average beliefs. We show with a simple calibration that this effect was relevant in the housing bubble.

Let k denote δ^τ. Suppose beliefs about k and D_g are perfectly correlated: investors who believe k is high believe D_g is high as well. Both factors depend on the growth in the number of end-users who receive positive flow utility from the asset, so they are naturally correlated. Let F_D and F_k be the cumulative distribution functions for the beliefs about D_g and k. For $x \in [0,1]$, we define $D(x) = F_D^{-1}(x)$ and $k(x) = F_k^{-1}(x)$. Using the definition of optimism significance in Equation (4) and the pricing equation (7), we calculate optimism significance as

$$\text{os}(x) = \frac{(1-k(x))D_0 + k(x)D(x)}{(1-\mu_k)D_0 + \mu_k\mu_D + \sigma_k\sigma_D}.$$

We think of x as near 1 so that this fraction reflects the ratio of optimistic beliefs to the average belief. When $L < N_t$, $k(x) = 1$ for all x: end-users already are the exclusive holders of the asset. In this case, the optimism significance ratio equals $D(x)/\mu_D$. We know from Proposition 6 that as L increases to ∞, this ratio goes to 1. But as L increases past N_t, can this ratio first increase before decreasing? It increases if $\partial \text{os}(x)/\partial L > 0$ when $L = N_t$. Proposition 7 states the condition under which this derivative is positive.

Proposition 7. *There exist a condition under which optimism significance increases with the asset float L in the range where the asset float exceeds the number of end-users N_t. This condition is*

$$\frac{\partial k(x)/\partial L}{\partial \mu_k/\partial L} < \frac{D(x)}{\mu_D}\frac{\mu_D - D_0}{D_g - D_0} + \frac{\partial \sigma_k/\partial L}{\partial \mu_k/\partial L}\frac{\sigma_D}{\mu_D(D_g - D_0)}.$$

Optimism significance initially increases with L if $k(x)$ decreases much less with L than μ_k decreases with L. If optimists assign a much higher value to δ^τ than the average belief when L is near N_t, then optimism significance will be increasing in L. A simple example calibrated off of Las Vegas illustrates this effect. Suppose $\delta = 0.93$ so that $D_0 = 14$. Optimists speculated that Las Vegas would run out of land in 8 years, whereas if Las Vegas continued to grow at historical rates, land exhaustion would not occur for 30 years. These forecasts correspond to $k = 0.55$ and $k = 0.1$. Suppose 10 percent of the investors are optimists who believe that $k = 0.55$ and $D_g = 40$, and 90 percent are pessimists who believe that $k = 0.1$ and $D_g = 25$. With this belief structure, optimism significance at the 90 percent level would equal $\text{os}(0.9) = D(x)/\mu_D = 1.5$ if Las Vegas had had no excess land. But with the added uncertainty about δ^τ arising from the extra land in Las Vegas, optimism significance rises to $\text{os}(0.9) = 1.7$.

Proposition 7 documents a mechanism that attenuates the negative effect increasing L has on optimism significance. Because investors disagree about δ^τ, a higher asset float initially has a muted effect on optimism significance as L exceeds N_t; Proposition 7 shows that the sign of this relationship can even reverse. Increasing the asset float does not render optimism sensitivity irrelevant immediately.

4 Empirical Results

4.1 Data Sources

In MSA-level analysis, we use the house price indices from the Federal Housing Finance Agency (FHFA). Our zip-code level house price data come from Zillow.com, a housing market research firm. Zillow estimates a market value for every house in their sample using a proprietary method. The zip code indices are the medians of these estimates. The correlation of price changes in Zillow's MSA-level data and the FHFA data is 0.94.

Our land price data come from Nichols et al. (2010), who use land transaction data from the CoStar Group, a commercial real estate information firm. CoStar maintains a database of more than one million commercial and residential transactions throughout the United States. Nichols et al. (2010) use the 170,000 land transactions in CoStar to construct MSA-level land price indices for 23 MSAs from 1995 through the first half of 2009. They run hedonic regressions of log price on time fixed effects and parcel characteristics, including detailed within-city location dummies. The time coefficients from these regressions are the index for each MSA.

We compute owner-occupancy rates at the zip code level from the 2000 and 2010 decennial censuses. Finally, we draw data on home builder land holdings and balance sheet losses from their 10-K reports.

4.2 Across-MSA Cross Section

The basic prediction of our model is that optimistic beliefs capitalized in land prices pushed up house prices during the housing bubble. In our model, land price increases explain all of the increase in housing prices, while construction costs stay constant. Land prices must rise more than one-for-one with housing to be consistent with this story. Table 2 shows that across all metro areas for which we have data, land prices did rise more than house prices during the 1996-2008 period. On average, land prices rose three times as much as housing prices.

Our theory predicts that among MSAs that experienced bubbles, the ones with *lower* land prices in 2000 would experience the larger bubbles, controlling for increasing fundamental demand from homebuyers. We control for increasing fundamental demand by defining bubble size to be the minimum of the log house price increase and decrease during the 2000-2011 period. This metric makes sense as long as increases in fundamental demand are permanent. We use the median house price reported by the census to proxy for land prices, as we lack direct data on land price levels. House price levels are an imperfect proxy for land prices because construction costs and house sizes can vary across metro areas. Nevertheless, Figure 3 show that among metros that did experience bubbles, metro areas that had higher 2000 house price experienced smaller bubbles. This relationship is the key prediction of our optimism sensitivity model. The top panel of Figure 3 shows that this relationship holds strongly among MSAs in California. The bottom panel shows that the relationship holds among MSAs even when California cities are excluded.

4.3 Within-City Evidence

We now test our theory on the cross-section of neighborhoods within a city. Suppose different neighborhoods in a city are imperfect substitutes for each other. Then a city resident's choice

Table 2: Annualized Land and House Price Returns (%), 2000-2006

	Land	Housing
Las Vegas	23.5	7.7
Chicago	18.8	4.5
Los Angeles	18.1	10.9
Sacramento	18	9.6
San Francisco	17.8	8.6
Phoenix	16.9	8.3
Seattle	16.4	6
Miami	16.3	10.4
New York	16.2	8.2
Washington D.C	16.1	9.3
Tampa	15.7	8.7
Orlando	14.3	8.4
Boston	14	7.3
San Diego	14	10.4
Philadelphia	13.8	5.9
Houston	13.6	2.6
Portland	13.4	5.4
Baltimore	12.5	7.5
Dallas	9.8	1.5
Atlanta	9.4	2.4
Denver	9.3	3.4
Detroit	4.4	1.4
Mean	**14.7**	**6.8**

Notes: Land price returns are calculated using the price index series from Nichols et al. (2010) based on CoStar data. House price returns are calculated from the OFHEO index.

of which neighborhood to live in is like a homeowner's choice of which metro area to live in. The "consumption good" in a neighborhood setting is owner-occupied housing, and the "input asset" is any other building, such as rental housing or parking garages. Our distinction between owner-occupied units as the consumption good and rental units as the input good reflects the imperfect substitutability of these units. As highlighted by Glaeser and Gyourko (2007), owner-occupied and rental units feature different characteristics, and characteristics of residents who choose each type of unit differ significantly as well.

The first prediction of our model is that neighborhoods with a smaller share of owner-occupied housing would experience a larger housing bubble. Neighborhoods with a high owner-occupancy share would experience little price appreciation, because housing is allo-

Figure 3: Housing Bubble Size and 2000 House Prices

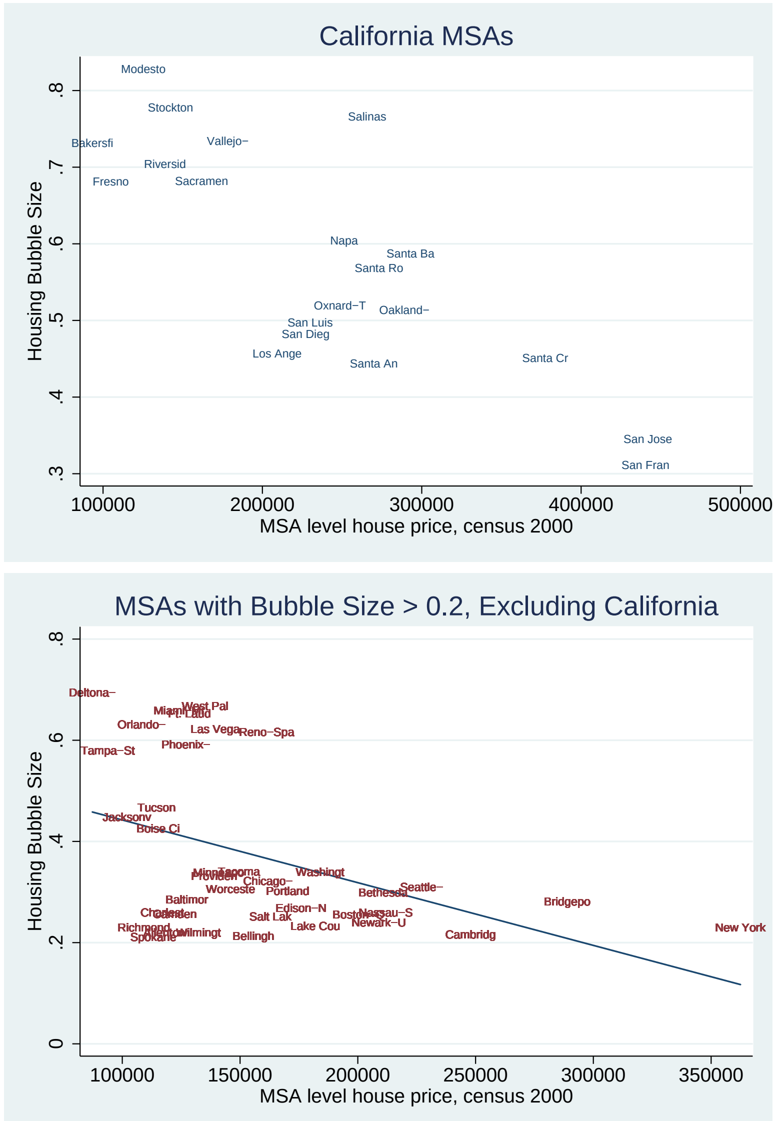

Notes: For each MSA, *Housing Bubble Size* is the minimum of the log price increase from 2000 until the local peak and the log price decrease from the local peak until 2011.

cated according to flow utility and not beliefs. In neighborhoods with a smaller owner-occupied share, developers who hold rental units, parking garages, and other non-owner-occupied structures are the marginal holders of real estate. The price of an owner-occupied unit reflects the beliefs of these developers, who supply the market with new owner-occupied units.

Figure 4 displays a negative relationship between the price increase during the bubble and the share of owner-occupied housing in a zip code for Boston and Chicago. On the

x-axis is the share of the non-vacant housing stock in 2000 that is owner-occupied, given by the Census. The y-axis displays real house price appreciation from 1996 to the market peak.

Figure 4: Owner-Occupancy and Price Increase: Boston and Chicago

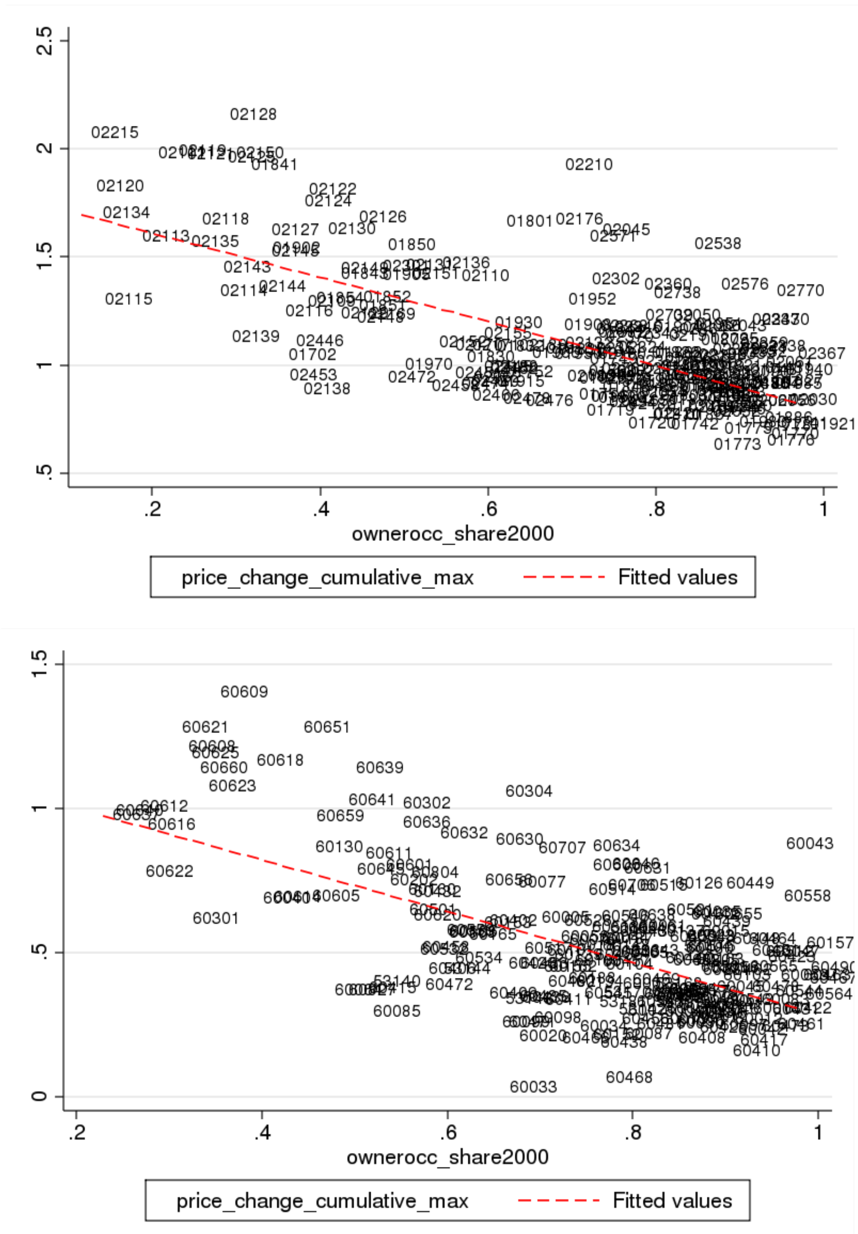

Notes: Using Zillow data, we compute the cumulative log real appreciation from 1996 to the zip code peak (y-axis). We plot this against the fraction of the occupied housing stock classified as owner-occupied in the 2000 Census (x-axis).

We document this relationship in all metro areas for which we have census data on owner occupancy and zip code house price data. For each zip code, we regress real house price appreciation from 1996 until the market peak on the owner-occupancy share in 2000. All regressions use MSA-level fixed effects. Table 3 shows the results of these regressions. Column 1 presents the baseline regression, while column 2 presents a weighted least squares regression with the zip code housing stock as weights. We expect the estimated coefficient to underestimate the strength of our theory's mechanism, because we include many sprawling metro areas where neighborhoods are in fact perfect substitutes. In column 3, we restrict the sample to cities with low new building from 1996 to 2006 relative to the 1996 stock. The estimates in all three specifications are large and significant. For instance, the third specification's estimate of -0.29 shows that a zip code where only half the 2000 housing stock was owner-occupied had a 15 percentage point larger price increase than a zip code where all of the 2000 housing stock was owner-occupied.

Table 3: Price Increase and Owner-Occupancy in 2000

	Cumulative Price Increase, 1996 to Zip Peak					
2000 Owner-Occupancy	**-0.18**	**-0.25**	**-0.29**	**-0.11**	**-0.19**	**-0.22**
	(0.03)	(0.03)	(0.05)	(0.03)	(0.04)	(0.04)
2000 Housing Price				**-0.09**	**-0.09**	**-0.10**
				(0.01)	(0.02)	(0.02)
Constant	1.26	1.38	1.47	1.81	1.86	1.93
	(0.05)	(0.04)	(0.08)	(0.16)	(0.17)	(0.24)
MSA Fixed Effects	X	X	X	X	X	X
Weighted by Housing Stock		X			X	
Low Construction			X			X
Observations	7,975	7,975	3,037	7,975	7,975	3,037
R^2	0.73	0.79	0.73	0.74	0.80	0.75

Notes: This table contains cross-sectional regressions at the zip code level. House price data come from Zillow. The dependent variable is the cumulative log price increase from 1996 until the zip code's respective peak. The owner-occupancy variable measures the fraction of the occupied housing stock classified as owner-occupied in the 2000 Census. Standard errors in parentheses are clustered at the MSA-level. Bolded coefficients are significant at the 1 percent level.

An alternative explanation for these results is based on the argument in Guerrieri et al. (2010). Their model of endogenous gentrification implies that low price areas experience greater price appreciation than high price areas as rich people move into the neighborhoods previously occupied by poor people. Columns 4 through 6 of Table 3 replicate our zip code regressions with an additional control for the initial price level in the zip code. Our theories overlap in this context, so to an extent, we may be over-controlling by including both proxies of relative underdevelopment. Nevertheless, both variables enter as significant and important correlates of the price increase across cities. Again, this effect is concentrated in the MSAs with lower supply elasticities.

We now test for a more subtle prediction of our theory. Imagine two neighborhoods that have the same owner-occupancy share in 2000 and experience identical price increases during the bubble. Suppose the owner-occupancy share of one neighborhood increases during the bubble, while the owner-occupancy share of the other neighborhood remains unchanged. We expect the neighborhood that added owner-occupants to have a smaller crash. That neighborhood has more owner-occupants on the margin, and prices there reflect the flow utility of these marginal owner-occupants more than they reflect the beliefs of developers, who have exited the market by selling to owner-occupants. If a crash occurs when optimists realize that their beliefs were wrong, then the crash would be smaller in the neighborhood that added owner-occupants, because prices are less reflective of beliefs there.

We test this hypothesis in Table 4. We run fixed effect regressions of the price fall from the peak for each zip code on the 1996-to-max price increase, the initial owner-occupied share, and the difference in the owner-occupied share. As in Table 3, our first specification weights all zip codes equally, our second specification weights zip codes according to their 2000 housing stock, and our third specification restricts to MSAs with low amounts of construction during the 2000 decade. The coefficient on the difference in owner-occupied share is negative and highly significant in all three specifications. For instance, the coefficient of -0.48 implies that if a neighborhood changed from a 25 percent owner-occupied share to a 75 percent owner-occupied share, its crash would be 25 percentage points smaller.

4.4 Public Home Builders as Speculators

Our model predicts that optimistic speculators purchased large quantities of land during the housing bubble, driving up land prices. We document this behavior among one class of speculators: public home builders. Home builders expanded their land holdings in excess of their immediate construction needs during the bubble. In addition, most of the losses suffered by home builders after the bubble came from land write-downs, consistent with the position of home builders as optimists.

Covering the years between 2001 and 2010, we collect data from the annual reports of the eight largest public builders as of 2005: Centex, Pulte, Lennar, D.R. Horton, K.B. Homes, Toll Brothers, Hovnanian, and Southern Pacific. The focus on a few important firms allows us to collect a number of details, including homes closed, land controlled through ownership and purchase options, deposits paid for optioned land and the remaining purchase price, as well as the value of writedowns taken on land inventories as the boom reversed. These eight

Table 4: Price Crash and Owner-Occupancy Change

	Price Fall from Peak		
Log Price Increase	**0.29**	**0.24**	**0.25**
	(0.05)	(0.06)	(0.10)
Δ Owner-Occupancy	**-0.50**	**-0.68**	**-0.48**
	(0.08)	(0.11)	(0.14)
2000 Owner-Occupancy	-0.04	-0.05	-0.00
	(0.03)	(0.04)	(0.04)
Constant	0.30	0.34	0.25
	(0.03)	(0.04)	(0.06)
MSA Fixed Effects	X	X	X
Weighted by Housing Stock		X	
Low Construction			X
Observations	7,975	7,975	3,037
R^2	0.74	0.75	0.53

Notes: This table contains cross-sectional regressions at the zip code level. House price data come from Zillow. The dependent variable is the cumulative log price decrease from the zip code's respective peak through 2010. The owner-occupancy variables measure the fraction of the occupied housing stock classified as owner-occupied in the 2000 Census and the change in this measure between 2010 and 2000. Standard errors in parentheses are clustered at the MSA-level. Bolded coefficients are significant at the 1 percent level.

firms delivered between 10 and 15 percent of all homes built in the U.S. during the 2001-2010 period.

Anecdotal evidence of a shift to land speculation comes from Pulte's annual reports. In 2001, they describe domestic home building as their core business and mention that this includes land acquisition and development as well as construction on this land. In 2005, Pulte replaced this language with the more direct statement, "We consider land acquisition one of our core competencies." This language appeared until 2008, when it reverted again to, "the Homebuilding operations represent our core business."

Figure 5 shows that these firms significantly increased their land holdings during the bubble, both in absolute terms and relative to what they were using for construction. The home builders held land both directly and through options granting them the exclusive right to purchase a tract of land at a certain price. The left panel shows time series aggregates of the eight firms' direct land holdings, optioned land holdings, and home building. At the peak of the cycle in 2005, these eight builders controlled approximately 1.8 million lots across the country, while delivering 250,000 homes. The right panel of Figure 5 shows that the builders' landholdings were increasing even relative to their construction needs. Most of this increase came through the quantity of land builders controlled through options. That land holdings increased even relative to current construction suggests that these builders were buying land because they thought prices would rise even more in the future.

Figure 5: Home Closings and Land Holdings by the Big Eight, 2001-2010

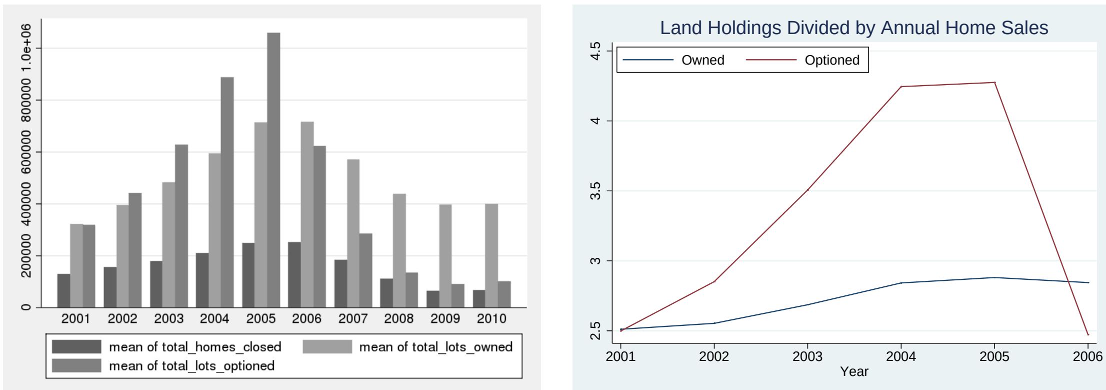

Notes: Data come from 10-K filings for Centex, Pulte, Lennar, D.R. Horton, K.B. Homes, Toll Brothers, Hovnanian, and Southern Pacific.

The use of options increased the amount of land builders could control with a fixed amount of capital. In terms of our model, the use of options increased M. We have data from five of these builders on the remaining purchase price of land controlled through options and the amount paid in deposits. In 2005, these five controlled 490,000 lots through options with a remaining purchase price of $25.9 billion. At that point, they had only paid $908 million, or 3.5 percent of the remaining obligation. D.R. Horton's 2002 annual report states:

> We also use lot option contracts, in which we purchase the right, but not the obligation, to buy building lots at predetermined prices on a takedown schedule commensurate with anticipated home closings. Lot option contracts generally are on a nonrecourse basis, thereby limiting our financial exposure to earnest money deposits given to property sellers. *This enables us to control significant lot positions with a minimal capital investment* and substantially reduces the risks associated with land ownership and development. [emphasis added]

Our model is more compelling when speculator capital, M, is very large. The ability of home builders to use options to leverage their land positions 30-to-1, in addition to their standard operating leverage, gives evidence that M was indeed large.

In spite of the market power that might arise from being such large buyers of raw land, the average price paid per lot more than doubled from $24,900 per lot in 2001 to $54,200 in 2006. Evidence from the annual reports implies that this increase was due to increased competition for land and not to differences in land quality over time. In 2006, Pulte stated, "Over the past few years, we have experienced an increase in competition for suitable land as a result of land constraints in many of our markets." In 2007, Lennar stated, "In addition to competition for homebuyers, we also compete with other home builders for desirable properties, raw materials and reliable, skilled labor." Such competition was likely an important force in driving up prices for raw land in markets that were not fully built out.

These firms faced large losses in their land portfolios. These losses are consistent with our model, in which optimists push up land prices by buying land, and then face losses when future demand growth is realized and is below their optimistic expectations. The left panel of Figure 6 shows the large losses these firms faced between 2005 and 2010. We measure firm value in three ways: book equity, market equity, or enterprise value (the market value of equity plus debt). Each of these measures falls between 65 percent and 75 percent between 2005 and 2010. The right panel of Figure 6 shows that much of these losses came from land impairments reported on the balance sheets of these firms. Land impairments account for nearly all of the book equity losses reported during this period. Cumulatively, land impairments account for 55 percent of the decline in enterprise value, 70 percent of the decline in market equity, and 120 percent of the decline in book equity. To a first approximation, these eight home builders were land speculators.

Figure 6: Big Eight Land Inventory Impairments and Franchise Value, 2001-2010

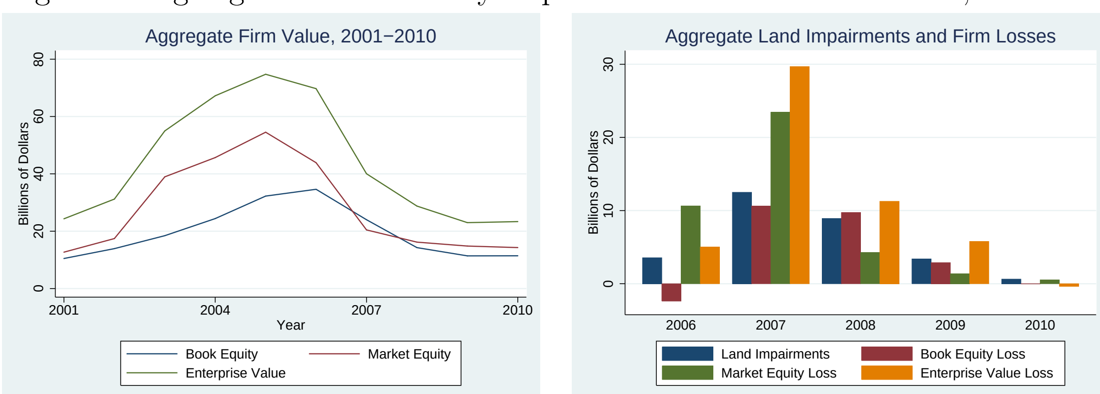

Notes: Data come from 10-K filings for Centex, Pulte, Lennar, D.R. Horton, K.B. Homes, Toll Brothers, Hovnanian, and Southern Pacific. The plot in the right panel begins in 2006 because that is the first year with reported land impairments.

4.5 The Case of Las Vegas

Informal theorizing about the causes of the Las Vegas episode has produced a slew of plausible stories. No one story can explain the whole bubble, but rapidly rising land prices explain a significant part of it. In spite of the availability of millions of acres of undeveloped land, very little of that land was developable because the federal government refused to sell it.

According to Saiz (2010), 68 percent of the land in a 50-kilometer circle around Las Vegas (1.3 million acres) is developable. Data from the Clark Country demographer suggest that the city was far from this boundary, as less than 300 thousand acres had been developed by 2006. Moreover, the demographer counts another 3 million acres of land beyond the Saiz radius as buildable but largely unbuilt.

However, most of this undeveloped land is owned not by private, profit-maximizing citizens, but by the U.S. government. And in the late 1990s, Congress passed the Southern Nevada Public Land Management Act (SNPLMA), which drew a development boundary around the city, demanded that land within the boundary be sold at public auction, and stipulated that land outside the boundary not be sold. This boundary restricted the amount of land in play to just 280 thousand acres, of which 162 thousand had already been built.

Given the boundary constraint, the pace of development and population were sufficient to generate expectations that it would only take a few decades and not centuries to deplete the available land. The demographer's office published a projection that the remaining land within the boundary would be consumed by 2035. This shift in expectations was followed by a frenzy for raw land. We can see this in the turnover statistics from the raw land market depicted in Figure 7, which exceeded the pace at which the Bureau of Land Management was releasing land and the pace at which builders were converting land into houses.

Figure 7: Vegas land transactions

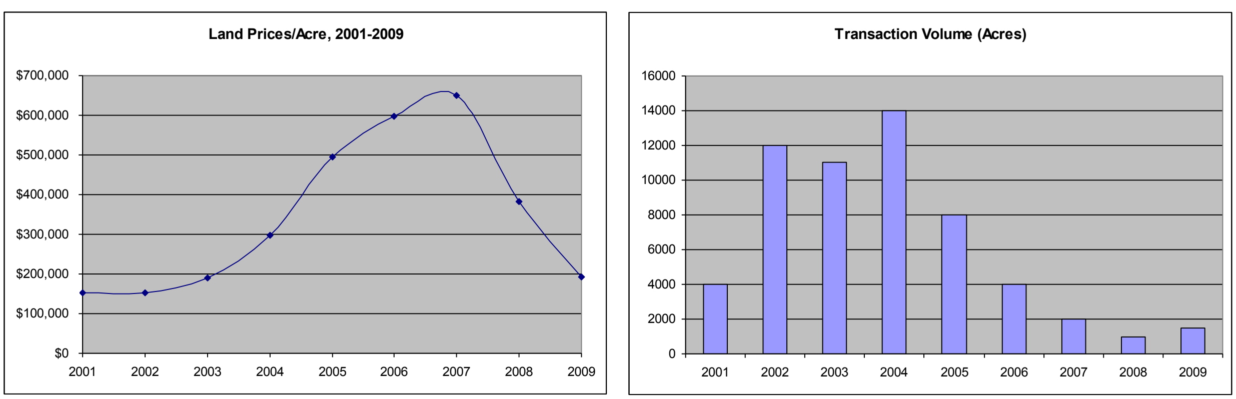

Notes: Data provided by Applied Analytics, a Las Vegas-based consulting group.

If inputs into the home building process other than land had been the binding constraints during the bubble, raw land prices would not have risen. This is why a price gap emerges between raw land and developed land in cities with tight residential zoning, where permission to build is the scarce resource.[5] In Las Vegas, much of the house price bubble can be

[5]See, e.g., Glaeser et al. (2005).

attributed to a bubble in land. There are six houses per acre in the larger developments in Las Vegas. The $500,000 price increase in land depicted in Figure 7 corresponds to an $83,000 increase in construction costs. House prices increased about $120,000 over this period.

The boundary established by SNPLMA left 120,000 acres of vacant land in 1998, comprising a 24-year supply at the going rate of land use. From this time, Las Vegas's housing stock expanded by 80 percent in just ten years. Glaeser et al. (2008) argue that the coincidence of price and supply booms in cities like Las Vegas is "a more profound signal of a bubble than any other of which we are aware." Our theory reconciles these apparent contradictions, which allows builders to disagree over future projections of when the resource constraint will bind while still continuing to build. These projections can materially alter the path of land prices even as construction grows, as long as the eventual constraint remains on the horizon.

5 Conclusion

This paper has shown that optimism captured in land prices was an important constraint in several cities during the housing bubble. Our model describes how well-capitalized speculators can bid land prices beyond fundamental value, while providing housing to homebuyers in an elastic fashion. The model's key prediction that bubble size and land availability are positively correlated in places with scarce land is borne out in both cross-MSA and within-MSA comparisons. We documented speculation among public home builders, who increased their land positions during the housing bubble and then suffered large capital losses from their land investments.

Our paper draws attention to an understudied class of actors in housing bubbles: home builders and land developers. We have argued that the beliefs of these firms are critical in determining prices, as these firms are the marginal buyers of real estate assets in the markets where they operate. Further research on the determination of these firms' beliefs, as well as on their mechanisms for raising capital, would shed additional light on the origins of housing bubbles generally.

References

Case, K. and Shiller, R. (2010). What were they thinking? home buyer behavior in hot and cold markets. *Yale School of Management mimeo.*

Chevalier, J. and Ellison, G. (1999). Career concerns of mutual fund managers. *The Quarterly Journal of Economics*, 114(2):389.

Fostel, A. and Geanakoplos, J. (2008). Leverage cycles and the anxious economy. *American Economic Review*, 98(4):1211–1244.

Glaeser, E. and Gyourko, J. (2007). Arbitrage in housing markets. *NBER Working Papers.*

Glaeser, E., Gyourko, J., and Saiz, A. (2008). Housing supply and housing bubbles. *Journal of Urban Economics*, 64(2):198–217.

Glaeser, E., Gyourko, J., and Saks, R. (2005). Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices. *Journal of Law and Economics*, 48(2):331–69.

Guerrieri, V., Hartley, D., and Hurst, E. (2010). Endogenous gentrification and housing price dynamics. *NBER Working Papers.*

Harrison, J. and Kreps, D. (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics*, 92(2):323.

Himmelberg, C., Mayer, C., and Sinai, T. (2005). Assessing high house prices: Bubbles, fundamentals and misperceptions. *Journal of Economic Perspectives*, pages 67–92.

Hong, H. and Stein, J. (2003). Differences of opinion, short-sales constraints, and market crashes. *Review of financial studies*, 16(2):487.

Mian, A. and Sufi, A. (2009). The consequences of mortgage credit expansion: Evidence from the us mortgage default crisis. *The Quarterly Journal of Economics*, 124(4):1449.

Miller, E. (1977). Risk, uncertainty, and divergence of opinion. *The Journal of Finance*, 32(4):1151–1168.

Morris, S. (1996). Speculative investor behavior and learning. *The Quarterly Journal of Economics*, 111(4):1111.

Nichols, J., Oliner, S., and Mulhall, M. (2010). Commercial and residential land prices across the United States. *Finance and Economics Discussion Series.*

Saiz, A. (2010). The Geographic Determinants of Housing Supply. *Quarterly Journal of Economics*, 125(3):1253–1296.

Scharfstein, D. and Stein, J. (1990). Herd behavior and investment. *The American Economic Review*, 80(3):465–479.

Scheinkman, J. and Xiong, W. (2003). Overconfidence and speculative bubbles. *Journal of Political Economy*, 111(6):1183–1220.

Simsek, A. (2011). Belief disagreements and collateral constraints. *Harvard mimeo.*