BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2013-18

Booms and Busts in House Prices Explained by Constraints in Housing Supply

by Narayan Bulusu, Jefferson Duarte and Carles Vergara-Alert

Bank of Canada Working Paper 2013-18

June 2013

Booms and Busts in House Prices Explained by Constraints in Housing Supply

by

Narayan Bulusu,[1] Jefferson Duarte[2] and Carles Vergara-Alert[3]

[1]Funds Management and Banking Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
nbulusu@bankofcanada.ca

[2]Jesse H. Jones Graduate School of Business
Rice University
jefferson.duarte@rice.edu

[3]IESE Business School
University of Navarra
CVergara@iese.edu

ISSN 1701-9397

Acknowledgements

We thank Alan Crane, Morris Davis, Eleonora Granziera, Brian Peterson, Randall Morck, Nancy Wallace and participants of sessions at the 48th Annual AREUEA Meetings, the 19th AEFIN Finance Forum and Universitat Pompeu Fabra for their helpful comments. Carles Vergara-Alert acknowledges financial support of the Public-Private Sector Research Center at IESE and the Ministry of Economy of Spain (ref: ECO2010-10102-E, ECO2009-13169). All remaining errors are our own.

Abstract

We study the importance of supply constraints in explaining the heterogeneity in house price cycles across geographies in the United States. Comparing the equilibrium house price generated with and without supply constraints in a representative-agent model under irreversibility of housing investment, we derive a relationship between housing returns and changes in supply constraints and determinants of housing demand. Our empirical analysis shows that supply constraints play an important role in Metropolitan Statistical Areas (MSAs) with boom-and-bust behavior. We estimate that, in 19 of the largest MSAs in the United States, supply constraints contributed 25% to the dramatic rise in house prices from 2000 to 2006, and 17% to their sharp fall from 2006 to 2010.

JEL classification: R310
Bank classification: Asset pricing; Economic models

Résumé

Dans cette étude, les auteurs examinent la mesure dans laquelle les contraintes s'exerçant sur l'offre expliquent l'hétérogénéité des cycles des prix des maisons entre différentes régions aux États-Unis. Ils comparent le prix d'équilibre des maisons généré en intégrant ou non des contraintes d'offre dans un modèle comportant un agent représentatif dont l'investissement dans le logement est irréversible, et font ressortir une relation entre le rendement de l'investissement résidentiel, d'une part, et l'évolution des contraintes de l'offre de logements et des déterminants de la demande, d'autre part. Une analyse empirique leur permet de constater que les contraintes pesant sur l'offre jouent un rôle important dans les zones statistiques métropolitaines (Metropolitan Statistical Areas ou MSA) des États-Unis où l'on a observé des cycles de flambée et d'effondrement des prix. Les auteurs estiment que, dans 19 des plus grandes MSA, les contraintes de l'offre ont contribué dans une proportion de 25 % à la hausse spectaculaire des prix des maisons enregistrée de 2000 à 2006, et dans une proportion de 17 % à leur chute brutale de 2006 à 2010.

Classification JEL : R310
Classification de la Banque : Évaluation des actifs; Modèles économiques

1 Introduction

Considerable effort has been expended on explaining the drivers of aggregate U.S. house prices, and whether or not there was a bubble[1] in this market. On the one hand, there is literature that attributes the boom-and-bust pattern to investors suffering from money illusion (Brunnermeier and Julliard 2008), irrational optimism (Glaeser et al. 2008), possessing heterogeneous beliefs about house price appreciation (Burnside et al. 2011), or placing incorrect belief in the efficient markets hypothesis (Peterson 2012). On the other hand, proponents of the rational-investors framework argue that high uncertainty in firms' profitability (Pástor and Veronesi 2006), perceived changes in relative trend productivity (Kahn 2008) or changing credit conditions (Chu 2012; Garriga et al. 2012; Favilukis et al. 2012) could generate the aggregate prices observed in the United States.

Even though these mechanisms are all plausible for the aggregate, they face considerable challenges in explaining observed differences across Metropolitan Statistical Areas (MSAs).[2] For example, why do investors in San Francisco suffer from money illusion while those in Atlanta do not? Alternatively, is it reasonable to assume that credit conditions change in the appropriate manner over the cycle in Miami, but not in Cleveland?

We show that constraints on housing supply – which are local in nature – are a natural way of generating heterogeneity in house price cycles across MSAs, while staying within the classical rational-investors framework.[3] More concretely, we attempt to quantify the contribution of the constraints on housing supply in a given MSA and year to the observed price in the same MSA and year. In doing so, we suggest a systematic method to control for variations in non-housing consumption, housing stock and housing supply capacity that can be used in empirical studies that aim to differentiate among competing hypotheses for the causes of booms and busts in house prices. Our focus on explaining house price cycles differs from that of the literature using supply constraints, which shows either that the cross-section of house prices varies positively with measures of geographical (Saiz 2010) and regulatory (Glaeser et al. 2006) restrictions on land development, or that supply constraints are positively related to the first and second moments of the returns to housing (Davis and Heathcote 2005; Paciorek 2012; Murphy 2012).

We formally incorporate the effects of supply constraints on house prices by comparing the house prices generated when exogenously imposed supply constraints do, and do not,

[1]Since we do not contribute to the debate on bubbles, we use the Brunnermeier (2008) characterization that bubbles are typically associated with dramatic asset price increases, followed by a collapse. We therefore use the terms 'boom and bust' and 'bubble' interchangeably in the text.

[2]Sinai (2012) documents that not only is the amplitude and timing of house price cycles different across MSAs, but also that there is greater cross-sectional variation in house price rises (falls) during booms (busts).

[3]Housing supply limitations could arise due to restrictions placed by geographical features (steeply sloping land, presence of water bodies, etc.), regulations on land use or availability of labor for construction, all of which differ across MSAs.

bind. This gives us the specification relating housing returns in a given geographical area to the effect of changes to supply constraints and their interaction with changes to determinants of housing demand. We model an infinitely lived rational representative agent with utility for housing services provided by the stock of housing H_t and non-housing consumption c_t, as in Piazzesi et al. (2007). At each period, from the stock of non-housing wealth K_t, the agent consumes c_t and transfers ψ_t to augment the stock of housing. This investment is irreversible ($\psi_t \geq 0$), and is subject to an exogenously determined constraint ($\psi_t \leq \Psi^{\max}$).[4] The growth of the stocks of non-housing wealth K_t and housing H_t have exogenously given drift, and suffer independent shocks. We assume that there are substantial barriers to moving between cities and abstract from the issue of migration between MSAs. Cross-location variation in house price patterns is driven, therefore, by differences in the evolution of supply characteristics across MSAs.

The mechanism by which the model generates booms and busts in house prices is simple: an increase in demand for housing that can be fully supplied by new construction would see modest changes in prices; the presence of a binding constraint results in a larger movement in prices. A fall in demand that takes it back below the threshold that can be fully met by additional supply would result in a return to the 'normal regime,' but would show up as a large fall in prices.

Using data from 19 MSAs in the United States to test these relationships, we find that a significant part of the boom and bust in prices observed between 2000 and 2011 can be explained by supply constraints.[5] For instance, aggregate returns to housing in the 19 MSAs were 85% (84%) during the period 2000–06 (2000–07). This growth is not fully captured by empirical specifications that do not incorporate constraints on housing supply, which generate at most 36% (42%) returns to housing during this period. Including the effect of supply constraints, on the other hand, generates 61% (69%) returns to housing. More importantly, only by considering supply constraints are we able to replicate the observed pattern of house prices: periods of high growth, followed by sharp falls. Prices explained by empirical specifications that do not control for changes in supply conditions exhibit a tendency to increase almost monotonically.

At the individual MSA level, house prices in MSAs that did not see large booms and busts during the sample period are reasonably well explained by models that do not incorporate housing supply constraints. An empirical specification without supply constraints explains

[4]Note that a capacity constraint is different from a time-to-build restriction. In a model with time-to-build (e.g., Kydland and Prescott 1988), agents know the amount of time necessary to build a desired amount of housing. In a model with capacity constraints, agents know how much housing can be built in a certain time, but they do not know ex ante the amount of time necessary to build the desired amount of housing. As a result, a model with capacity constraints can generate price increases that are not easily generated in a traditional model with time-to-build.

[5]The choice of our sample period is constrained by issues of data availability, which we discuss in detail in later sections of the paper.

at least 63% of the price rise between 2000 and 2006; adding supply constraints does not improve explanatory power for such geographies.[6] On the other hand, between 70% and 81% of the price rise observed in MSAs which experienced significant house price boom and bust (e.g., Los Angeles, Miami, Phoenix, San Diego, San Francisco and Seattle) is explained when a measure of supply constraints is taken into account. Without this key measure, only 24% to 54% of the observed price rise is explained.[7] Overall, including supply constraints explains between 7% and 44% more of the observed price rise in this period, depending on the MSA; the greatest addition to explanatory power is in the MSAs encompassing the cities of Miami, Phoenix, San Diego, San Francisco and Seattle.

Key to the empirical results are measures we construct that proxy for housing supply constraints. Consistent with the evidence in Albouy and Ehrlich (2012) highlighting the importance of other inputs, especially labor, in the construction process, our measures incorporate proxies for both land and labor supply.[8] The inclusion of both these factors of production turns out to be quite important in some MSAs. For example, Phoenix has neither significant geographical nor regulatory constraints to new housing; however, its house prices rose by 109% during 2000–06. We find that using a measure of supply constraints that includes a proxy for constraints on the availability of labor helps explain an increase of 81%; without including housing supply constraints, the empirical analysis explains only a 42% rise. Our empirical analysis also shows that the presence of supply constraints is unlikely to be the driver behind the pattern of house prices observed in Detroit and Portland: prices are rather high in the period 2000–06 even after controlling for this effect.

This paper is related to two strands of literature on housing: one that incorporates supply issues in understanding house prices, and one that participates in the debate as to whether there was a 'housing bubble' between the years 2000 and 2010.

Following the influential review by DiPasquale (1999), there has been increasing interest in the factors that affect housing supply, and the effects of the peculiarities of supply on house prices. A long list of studies – including Mayer and Somerville (2000), who studied the empirical relationship between new construction and price and cost changes, and Grenadier (1996), who modeled the strategic interaction considerations in an individual's land development decision – have attempted to understand the determinants of housing supply. The drivers of supply considered include demographics, time-to-build restrictions, dynamics of

[6]The MSAs containing the cities of Cleveland, Dallas, Detroit, Houston and Minneapolis-St. Paul did not show bubble-like behavior; these are also the cities that have relatively low scores for the Saiz (2010) and Gyourko et al. (2008) measures of geographical and regulatory constraints on land, respectively.

[7]These are also the MSAs with relatively high scores of the aforementioned measures of residential land supply constraints.

[8]The magnitude of the contribution of various sources to the cost of housing is an empirical issue that is the subject of debate; see, for instance, Wheaton and Simonton (2007). We assume that access to capital does not impose significant constraints on the production of housing, given the relatively greater financial capacity of firms, and therefore do not include it as a component of our measure.

the construction industry, and geographical and regulatory restrictions on the development of land for residential use (see Haughwout et al. 2012 and the citations therein). Unlike in Paciorek (2012) and Murphy (2012), we do not explicitly model the supply. Instead, following the approach taken in Chu (2012), we attempt to understand the effects on prices of (endogenously generated) demand in the presence of supply constraints, which are exogenous to the model.

Our paper is also related to the literature on housing bubbles. In particular, we contribute to the literature that attempts to understand whether assumptions on investor behavior are necessary to generate the price patterns observed in the decade 2000–10. We do not discount the contribution of the alternative explanations put forward in the literature for the paths of house prices; however, our results show that it is not necessary to rely on non-standard investor behavior or heterogeneity of beliefs to explain at least part of the rises and falls of house prices. Moreover, because our model relies on supply imperfections that vary across regions, it can generate cross-regional differences in price appreciations and hence is a good starting point for understanding whether real estate prices can be explained within a rational-agent framework.

The remainder of the paper is organized as follows. In section 2, we set up the model, study its equilibrium and explain the technique used to obtain its numerical solution. We also discuss the results of the baseline model and their economic intuition. In section 3, we derive a reduced-form specification from the model and implement this specification to the main 19 MSAs in the United States using data on consumption, housing assets, non-housing assets and house prices at the MSA level. Finally, section 5 concludes the paper.

2 The model

In this section, we set up and solve a general-equilibrium model to study the dynamics of house prices, consumption and housing investment. Our work builds on the representative-agent models of Dunn and Singleton (1986), Yogo (2006), Piazzesi et al. (2007), and Gomes et al. (2009). However, while these models focus only on the role of the separation between durable (i.e., housing) and non-durable (i.e., non-housing) goods, our model highlights the importance of irreversibility and capacity constraints on the supply side of the housing markets as an economic mechanism that is able to generate 'bubble-like' house price dynamics. The model delivers most of our key empirical findings in a very simple rational-expectations setting.

2.1 Model set-up

Consider an economy with an infinitely lived representative agent. There are two goods in this economy: a non-housing consumption good (the numeraire good) and a housing good.

The agent consumes housing services (shelter) provided by the stock of housing H_t, and non-housing consumption, c_t, which represents the consumption of all non-durables and services except housing services.[9] Preferences over consumption take the following standard form:

$$U(c_t, H_t) = \frac{\left[\left(c_t^{\frac{\epsilon-1}{\epsilon}} + \omega H_t^{\frac{\epsilon-1}{\epsilon}}\right)^{\frac{\epsilon}{\epsilon-1}}\right]^{(1-\gamma)}}{1-\gamma}. \tag{1}$$

The parameter ω is the utility multiplier that converts units of housing services into units of the non-housing good. Let ϵ denote the intratemporal elasticity of substitution between housing services and the non-housing good and $1/\gamma$ denote the intertemporal elasticity of substitution.[10] The agent maximizes the net present value of future utility flows given by (1). There are two productive sectors in the economy: one that produces non-housing goods (including services) and another that produces houses. Each sector consists of a representative firm that takes input and output prices as given. Let us assume that the representative agent fully owns the two representative firms in the economy. Consequently, we can set up the general-equilibrium model as the central planner's problem.[11] Let us also assume that the only input for production is capital (e.g., no labor). Sector K represents the gross of the economy (e.g., all the economy but the housing industry) and sector H represents the housing industry. Sector K presents a perfectly reversible technology that has constant returns to scale. Let K_t denote the stock of non-housing capital good. Between times t and $t + dt$, a uniform flow per unit time interval of ψ_t units of the non-housing good are invested in the construction of houses; that is, an amount $\psi_t \cdot dt$ of the capital in sector K is transferred to sector H at time t. Under this set-up, the central planner chooses optimal flow of consumption of the numeraire good per unit time interval, c_t with $c_t \geq 0$, and housing investment ψ_t in order to maximize the household's objective function. These choices are affected by the non-negativity constraint of capital in the non-housing sector, $K_t \geq 0$, where K_t evolves according to the process

$$dK_t = K_t(\alpha_K dt + \sigma_K dW_t^K) - (c_t + \psi_t)dt, \tag{2}$$

where α_K and σ_K are the drift and diffusion, respectively, of the returns per unit capital of, and dW_t^K is a standard Brownian motion that drives shocks to, sector K. We assume that

[9]The stock of housing services is measured in units of area; that is, H_t could be the number of square feet of housing of a standard quality of construction. The agent derives utility from living in a house of a particular size H_t.

[10]For low values of ϵ, agents are not willing to substitute between the consumption of non-housing goods and housing services in the same period. For high values of γ, agents are not willing to substitute aggregate consumption of non-housing goods and housing services over time.

[11]The classic models in Lucas Jr and Prescott (1971) and Prescott and Mehra (1980) also make this assumption and provide some related technical details.

investment in the production of housing is irreversible.[12] The flow of housing investment is irreversible; in other words, it has the restriction $\psi_t \geq 0$. While this feature may appear an extreme version of the relative illiquidity of housing investment, it is not unreasonable in a representative-agent economy, because in the aggregate, houses cannot be converted back into non-housing consumption goods. The stock of housing held by the representative agent evolves according to

$$dH_t = H_t(\alpha_H dt + \sigma_H dW_t^H) + \frac{1}{P_t}\psi_t dt, \tag{3}$$

where α_H and σ_H are the drift and diffusion, respectively, of the rate of change of the housing stock. The drift α_H is interpreted as including depreciation (recall that the housing stock is defined in terms of a particular standard of construction quality), effects in housing size due to expected demographic changes such as population growth, etc. σ_H is the factor that amplifies the unexpected changes in the environment (these exogenous shocks are represented by dW_t^H); this could be from population changes due to unexpected changes in economic conditions, or other factors that affect the demand for housing. We impose limits to the supply of housing by setting the upper bound of ψ_t to be an exogenously determined quantity, $\Psi^{\max}$.

Let P_t denote the price of the housing good in units of the numeraire (non-housing) good; for example, how many units of numeraire we need to buy one square foot of housing. We allow the two Brownian motions dW_t^H and dW_t^K to have an instantaneous correlation coefficient of ρ, which allows transmission of shocks from the productivity of sector K to sector H, and vice-versa. For instance, if it is hypothesized that a large positive shock to sector K's productivity might encourage a large inward migration, which might reduce the per capita area of housing, ρ is assumed to be negative.

By appropriately choosing the controls c_t and ψ_t, the representative agent seeks to maximize the lifetime discounted expected utility:

$$J(K_t, H_t; c_t^*, \psi_t^*) = \sup_{\substack{c_s \geq 0; 0 \leq \psi_s \leq \Psi^{\max} \\ K_s \geq 0}} \int_t^\infty e^{-\rho(s-t)} U(c_s, H_s) dt. \tag{4}$$

2.2 Equilibrium

There exists no known analytical solution for the optimal value function J^* and controls c^* and ψ^* under the general formulation in section 2.1. We therefore solve the model numerically. To do this, however, we first need to solve for the equilibrium price P_t by finding the first-order conditions from equation (4). Since this is an infinite-horizon problem, the solution is independent of time, and we will ignore the time subscript in what follows. The Hamilton-

[12] Houses once built cannot be converted back into the numeraire good. We make this assumption to keep the model simple, and it is quite reasonable in our representative-agent economy. In any case, relaxing the assumption to make housing investment partially reversible will raise the optimal investment level in housing. This raise increases the equilibrium price of housing, making our results even stronger.

Jacobi-Bellman equation of the solution of the optimal control problem can be written as

$$J(K,H;c^*,\psi^*) = \sup_{\substack{c\geq 0;0\leq\psi\leq\Psi^{\max}\\ K_s\geq 0}} \left[U(c^*,H)dt + \mathcal{L}J^*\right],$$

where the infinitesimal generator $\mathcal{L}J^*$ is given by

$$\begin{aligned}\mathcal{L}J^* &= [(K\alpha_K - c^* - \psi^*)J^*_K + (H\alpha_H + (P^*_t)^{-1}\psi^*)J^*_H \\ &\quad + \frac{1}{2}K^2\sigma_K^2 J^*_{KK} + KH\rho\sigma_K\sigma_H J^*_{KH} + \frac{1}{2}H^2\sigma_H^2 J^*_{HH}]dt,\end{aligned}$$

and where the value function J^* with subscripts refers to partial derivatives of J^* with respect to the variable in the subscript.[13] Consequently, the first-order conditions of this optimal control problem are given by the following two expressions:

$$J^*_K = U^*_c(c^*,H), \tag{5}$$

$$P^*_t = \frac{J^*_H}{J^*_K}. \tag{6}$$

Equation (5) is the envelope condition and establishes the equality between marginal consumption and marginal investment in the numeraire good at equilibrium. Equation (6) provides the equilibrium house price in units of the numeraire good, P^*_t, as the ratio of marginal investment in housing to marginal investment in the numeraire good. The algorithm we use to obtain a numerical solution of the investor's problem is standard, and we therefore relegate details to the appendix.

2.3 Numerical results for the baseline model

This section reports the numerical results for a baseline model that we use to determine the functional form of the relationship between house prices and supply constraints. We choose plausible parameter values, summarized in Table 1.[14] Housing investment decisions are not taken very frequently; we assume that the time between successive decisions is three months. β is chosen to set the annual subjective discount factor as 0.98 annually. The drift and volatility of the return on the numeraire asset are assumed to be equal to those of the return of the U.S. stock market. Two factors contribute to our assumption that housing stock falls by 5% annually: population growth and our measure of housing stock is in units of constant quality. Annual volatility of the per capita housing stock is calibrated equal to its sample counterpart, 12%. We set values of γ and ϵ consistent with the ones used in Hall (1988), Ogaki and Reinhart (1998) and Piazzesi et al. (2007).

[13] For example, J^*_K is the partial derivative of J^* with respect to K, and J^*_{KK} is the second derivative of J^* with respect to K.

[14] The qualitative results of the analysis performed below are not dependent on this set of parameters.

We solve for house price P, investment in housing ψ and non-housing consumption c as functions of the stock of non-housing wealth K and housing H. The choice of the boundaries of the two state variables, and the grid size, are influenced by two factors. First, keeping the grid size constant, computational complexity grows as the range of K and H increases. Second, choosing disproportionately high values of K (H) relative to H (K) results in a high (low) value of ψ throughout the state variable space, which hinders getting a complete picture of the shape of the control variable. With these considerations in mind, we choose $K, H \in [1, 2]$, and use a grid size of 0.2. Finally, after observing that the maximum housing investment ψ in the unconstrained case is about 0.37, we choose $\Psi^{\max} = 0.2$ to ensure that the supply constraint is binding for at least some values of the state variables.

It bears emphasizing that we do not calibrate the values of K, H or Ψ, since we use the numerical results with the sole purpose of determining the functional form of house prices, particularly since we recognize the stylized nature of our model. We therefore do not draw conclusions about the magnitude of the impact of supply constraints from the numerical solutions, and conduct a rigorous empirical study instead.

Figure 1 shows the numerical solution to the optimization problem. Panel A plots the equilibrium house price P^*, Panel B shows the optimal investment in housing ψ^*, and Panel C shows the optimal consumption of non-housing goods c^*, all as functions of housing stock (H) and the stock of the numeraire (K). Within each panel, the figure on the left-hand side is the solution to the problem where the supply constraints do not bind; we set the value of $\Psi^{\max}$ to a large number in the optimization algorithm to get these results. Figures on the right-hand side show the optimal solution under binding supply constraints.

We begin by discussing the case in which supply constraints are not binding. As is reasonable, house prices are high when the agent has a lot of non-housing wealth but has little housing, and low when the agent has a lot of housing but little non-housing wealth.[15] Given that the elasticity of substitution between housing and non-housing consumption is low, for a given housing stock, increasing wealth induces greater housing purchases, which increases house prices. For a given level of wealth, increasing the stock of housing decreases the incentive to purchase more housing, which reduces its price; note that the investment in housing increases with K and decreases with H in the left-hand figure of Panel B. The opposite of this is shown in the graph of non-housing consumption in Panel C, which increases with K for a given level of H, and increases with H for a given K.

To investigate the impact of a binding supply constraint, we reduce $\Psi^{\max}$ to the value of 0.20 and rerun the algorithm maintaining all other parameters as above. Next, we compare the three graphs on the left-hand side in each panel of Figure 1 to the corresponding graphs on

[15]The terms 'a lot' and 'too little' are used with reference to the optimal mix of goods, which depends on the preference parameters governing the agent's utility function, and the parameters governing the stochastic processes of K and H.

its right-hand side. Panel B shows that the constraint is binding: the investment in housing is capped at the 0.20 level, while the optimal in the absence of this limit is to invest a greater amount in housing. The investment cap shows up as a large spike in price in Panel A in the exact region where the constraint binds. We see little impact of the constraint on non-housing consumption; this is reasonable given the low substitutability between the two goods.

One of our main arguments is that the large spikes and sharp declines in house prices observed in some MSAs of the United States may be working through the channel depicted in Panel A of Figure 1. Cities with constraints on housing supply, such as Miami, have significant geographical limitations on new construction and could conceivably move from an environment where demand for new housing can be fully satisfied by new construction to periods in which demand rises above the capacity constraint.[16] This shift in housing demand will result in a large spike in prices, or a 'boom.' The 'bust' in prices occurs when the demand goes from being higher than supply to one that can be met by supply.

Our model provides two novel empirical implications that have not so far been, to the best of our knowledge, explored in the literature. Firstly, the model predicts that real estate prices should be more sensitive to increases in non-housing consumption when the capacity constraint binds. When demand for housing is above capacity limits of construction, the rise in house prices is steep. This is because the low elasticity of substitution between non-housing consumption and housing investment means that increasing non-housing consumption is not a substitute for the inability to increase housing investment. Consequently, the derivative of real estate prices with respect to non-housing consumption is higher when housing capacity constraint binds. Secondly, the model predicts that real estate prices should be less sensitive to increases in housing stock when the supply capacity constraint binds. Increases in housing stock decrease real estate prices and when supply is constrained an increase in housing stock leads to a smaller decrease in house prices due to the fact that it is not possible to build as quickly as desired.

The model suggests a novel way to understand real estate bubbles, which is based on the model-implied relationship between house prices, consumption, housing stock and housing supply constraints. Specifically, the model implies that house price booms and busts happen together with expansions and contractions in non-housing consumption when housing supply constraints bind. On the other hand, when housing supply constraints do not bind, expansions and contractions in non-housing consumption are accompanied by modest variations in house prices. As a result, an empiricist trying to differentiate among explanations for booms and busts in house prices would need variations of housing supply constraints across time and

[16] In our model, spikes in demand for housing arise from significant imbalances in the composition of K and H driven either by a large positive shock to K, or a large negative shock to H. A similar mechanism drives a fall in housing investment. Outside our model environment, however, there could be other reasons (e.g., changes in credit supply, or technological changes such as the invention of the automobile) that influence the dynamics of housing demand.

locations as well as variations in non-housing consumption. For instance, an empiricist may wish to test whether governmental incentives toward home ownership could have caused the boom and bust in house prices that were observed between 2000 and 2011. Such incentives would alter consumers' preference away from non-housing consumption and toward housing consumption and potentially increase house prices. If this were indeed the case, the empirical application suggested by our model would reveal that house price increases observed between 2000 and 2006 would not be completely explained by changes in non-housing consumption and constraints in housing supply.

3 The empirical specification

Our model shows that house prices P can be expressed as a function of non-housing wealth K, housing stock, H, and whether the housing supply constraint binds. Taken literally, the model offers a non-linear relationship between the slope of the price function and changes in K and H for each MSA (see Panel A of Figure 1). We use three implications from the model to capture its central insights without leaning too heavily on the overstylized framework.

First, even if non-housing wealth and the stock of housing remain constant over time in a specific MSA, housing returns could be produced by switching between the left and right graphs in Panel A of Figure 1. Thus, if the measure of housing supply constraints, D, varies within an MSA over time, ΔD is a factor that drives housing returns (we use Δ to denote change over time). Second, even if D remains constant over time, changes in H or K affect housing returns. Third, from a comparison of both figures of Panel A of Figure 1, the slope (with respect to changes to both K and H) is different in the model in which supply constraints are not binding from the model in which supply constraints bind. The simplest functional form that represents all these relationships is

$$\Delta p_{i,t} = u_i + \alpha_{i,0}\Delta D_t + \alpha_{i,1}\Delta k_{i,t} + \alpha_{i,2}D_t\Delta k_{i,t} + \alpha_{i,3}\Delta h_{i,t} + \alpha_{i,4}D_t\Delta h_{i,t} + \xi_{i,t},$$

where $\Delta p_{i,t} = log(P_t/P_{t-1})$; $\Delta k_{i,t} = log(K_t/K_{t-1})$; $\Delta h_{i,t} = log(H_t/H_{t-1})$; and u_i is the term that captures all the unobserved factors that affect prices in MSA i and $\xi_{i,t}$ is the error term that captures all other effects. We assume that errors are uncorrelated across time $\xi_{i,t} \overset{i.i.d}{\sim} \mathcal{N}(0, \sigma_{\xi_i})$ and also across MSAs.[17] Notice that we must convert house prices into

[17]It is equally possible to hypothesize that the relationship between Δp and Δk or Δh is not linear, but also includes higher-order terms, in which case the relationship can be written as (only including second-order terms)

$$\begin{aligned}\Delta p_{i,t} \;=\;& u_i + \alpha_{i,0}\Delta D_t + \alpha_{i,1}\Delta k_{i,t} + \gamma_{i,1}\Delta k_{i,t}^2 + \alpha_{i,2}D_t\Delta k_{i,t} + \gamma_{i,2}D_t\Delta k_{i,t}^2 + \alpha_{i,3}\Delta h_{i,t} + \gamma_{i,3}\Delta h_{i,t}^2 + \\ & \alpha_{i,4}D_t\Delta h_{i,t} + \gamma_{i,4}D_t\Delta h_{i,t}^2 + \alpha_{i,5}\Delta(k_{i,t}h_{i,t}) + \gamma_{i,5}D_t\Delta(k_{i,t}h_{i,t}) + \xi_{i,t}. \qquad (7)\end{aligned}$$

We use the parsimonious specification, since it is easier to interpret its results, has fewer parameters that need to be estimated and avoids the possibility of overfitting the model to the data.

returns to be able to carry out empirical tests. This transformation is necessary because P and K are non-stationary in the data, while their returns are stationary.

Given that data for non-housing wealth, K, are not available at the MSA level, we propose an alternative specification that uses data on consumption expenditures, which are reported at the MSA level. Panel C of Figure 1 shows that optimal consumption is a smooth, linear function of non-housing wealth (i.e., $\Delta c = \theta_i \Delta k$). This allows us to write the following testable implication:

$$\Delta p_{i,t} = u_i + \beta_{i,0}\Delta D_t + \beta_{i,1}\Delta c_{i,t} + \beta_{i,2} D_t \Delta c_{i,t} + \beta_{i,3}\Delta h_{i,t} + \beta_{i,4} D_t \Delta h_{i,t} + \lambda_t, \tag{8}$$

where $\lambda_t \overset{i.i.d}{\sim} \mathcal{N}(0, \sigma_\lambda)$ is the error term. In our empirical analysis, we focus on testing the specification on equation (8), and use the reliable data on non-housing consumption[18] at the MSA level.[19]

3.1 Data and sources

To test the model-implied relationship between the returns on housing (Δp), the growth of non-housing consumption (Δc), the growth of housing stock (Δh) and the change in the supply constraint (ΔD), our choice of MSAs is constrained by the availability of data on all these variables; our sample has data from 1987 to 2011 for 19 MSAs that encompass the cities of Atlanta, Baltimore, Boston, Chicago, Cleveland, Dallas, Detroit, Houston, Los Angeles, Miami, Minneapolis-St. Paul, New York, Philadelphia, Phoenix, Portland, San Diego, San Francisco, Seattle and Washington, DC.

The Federal Housing Finance Agency (FHFA) house price index is available from 1977 at the MSA level, from which we calculate housing returns. Data on annual per capita consumption and consumption of housing services (shelter) at the MSA level are provided from 1987 to 2011 by the Bureau of Labor Statistics, which we use to estimate annual per capita non-housing services consumption expenditure.

Data on the stock of the housing asset are obtained from the American Housing Survey conducted – in different years at different locations – by the U.S. Census Bureau (for the full

[18] In this paper, we use 'consumption' interchangeably with 'non-housing consumption'; there should be no confusion since there is no other type of consumption in the model, even if utility is derived from investment in the stock of housing.

[19] In a more general context, these relationships could also be thought of as being generated by approximations of the Taylor expansions of the function f^U (f^C) that maps $(\Delta c, \Delta h)$ into Δp, for the case without (with) supply constraints. Ignoring terms above the first order, we get

$$\Delta p = \pi_0^U + \pi_1^U \Delta c + \pi_2^U \Delta h \tag{9}$$
$$\Delta p = \pi_0^C + \pi_1^C \Delta c + \pi_2^C \Delta h, \tag{10}$$

for the case when supply constraints do not bind (superscript U) and when they do (superscript C), which can be summarized in a single equation by equation (8).

list, see the link provided in footnote 22). The per capita square footage of housing stock is obtained by multiplying the number of housing units by the median square footage of a housing unit, and dividing by the population of the MSA, given by the U.S. Census Bureau. Intersurvey estimates are obtained using linear interpolation; those for the years between the last survey and 2011 are estimated by using the average growth rate of housing size of all MSAs over the period 1987 to the last survey period.

The first set of columns in Table 2 report the mean and standard deviation of log returns on housing, growth of log per capita consumption and of the log per capita square footage of housing by MSA for the sample period 1987 to 2011. The MSA that includes Portland, OR has seen average annual returns of 5.3% per annum over the past 25 years – the highest in the sample – and San Francisco and Seattle are not far behind, having experienced average housing returns of about 4.6% per annum. At the other end of the spectrum, the MSAs including Atlanta and Cleveland have seen the lowest average annual housing returns, with 2.1% and 2.5%, respectively. The 'boom-and-bust' behavior of house prices in Los Angeles, Miami and Phoenix is seen clearly in the rather large time-series variation in housing returns of the MSAs containing these three cities: over 11% per annum. On the other hand, a standard deviation in housing returns of just above 3% was observed for the sample period in Cleveland, Dallas and Houston. There is much less cross-sectional variation in the average growth in non-housing consumption or its volatility, all MSAs experiencing an average annual growth of between 2% and 3%, and a standard deviation of between 3% and 7% for most MSAs. Finally, per capita square footage of housing fell by 3% in Seattle and Boston; Chicago, Detroit, Los Angeles and New York also experienced modest declines of less than 1% annually. The rest of the MSAs saw the size of housing per capita grow less than 1%, except Dallas, San Francisco and Washington, DC, where it grew by marginally over 1% per annum.

3.2 Proxies for the supply constraint measure

For simplicity, our model assumed that housing supply constraints were either binding, in which case $D_t = 1$, or not, in which case $D_t = 0$. Our empirical analysis takes a more nuanced view of the constraints on housing supply. Our aim is to construct an index of how difficult it is to add to the stock of housing in each MSA. Among the factors of production required for construction, as a first approximation, we postulate that availability of land and labor are the most crucial factors determining the supply of housing. Prior literature has convincingly argued that two major factors that affect availability of land for construction are the regulations that govern residential land use and the feasibility of constructing given the geographical features; Gyourko et al. (2008) have developed an index for the former, and Saiz (2010) has constructed a measure of the latter. As a proxy for the ease of availability of labor, we use the percentage of employees in the private sector employed in the residential construction sector as a fraction of the maximum value it attains in the sample; we call this measure $DMPEC$.[20] To capture the intuition that the constraints on housing supply depend jointly on the supply of land and labor, we create two measures: (i) $CAPSAIZ$, which is obtained by multiplying $DMPEC$ and the Saiz (2010) index of geographical constraints to housing supply ($SAIZ$); and (ii) $CAPWRI$, obtained by multiplying $DMPEC$ and the Wharton Residential Land Use Regulation Index ($WRLURI$) developed in Saks (2008).[21]

The $SAIZ$ and $WRLURI$ measures are available in the original papers that developed them. These measures are estimated using data around the time the papers were written; however, lacking a better solution, and given our relatively short sample period of 25 years, it appears reasonable to assume that the measures of geographical and regulatory constraints are constant in the sample period. Data on the time series of overall employment in the private sector, and data in the residential construction sector, are obtained from the U.S. Bureau of Economic Analysis.[22]

The last columns of Table 2 report the summary statistics of the various housing supply constraint measures used in the text. $SAIZ$ is the measure of the percentage of the area in an MSA that is unavailable for residential or commercial real estate development within the boundaries of the MSA; we see that MSAs containing Miami and Seattle have the greatest geographical constraints, while those containing Atlanta, Houston and Dallas have the least restrictions. Regulatory burdens for residential development are the tightest in Boston, Philadelphia and San Francisco, and the lightest in Dallas, Houston and Cleveland. The percentage of employees in the private sector working in construction as a fraction of its maximum value is, on average, between 0.7 and 0.9 for all the MSAs under construction, with small variance across time. Finally, the composite measures that we use as indices of housing supply constraints, $CAPWRI$ and $CAPSAIZ$, all have high sample means compared to their variance in time.

4 Empirical results

In this section, we show that incorporating supply constraints can explain a large proportion of the observed pattern of house prices, especially in MSAs that experienced sharp rises and

[20] We rerun the analyses below with other variants of this measure: (i) the percentage of population employed in the residential construction sector as a fraction of its maximum value ($DMPPC$); (ii) to avoid possible look-ahead biases, we also calculate the ratios $DMPEC$ and $DMPPC$ using the running maximum in the denominator. The main results discussed are robust to using these alternative measures.

[21] The higher the value of the index, the more onerous the regulation on new residential developments; this index is normalized and can take negative values.

[22] Links to data sources mentioned above follow. FHFA housing indices: http://www.fhfa.gov/Default.aspx?Page=87. Consumption: http://www.bls.gov/cex/csxmsa.htm#top. Housing data: number of housing units (Table 1-1), median square footage of housing unit (Table 1-3); http://www.census.gov/housing/ahs/data/metro.html. Population: http://www.census.gov/popest/data/historical/1990s/metro.html. Employment by industry: Table CA04 at http://bea.gov/iTable/iTable.cfm?ReqID=70&step=1&isuri=1&acrdn=5.

subsequent rapid falls in house prices. Our model abstracts from the effects of migration or other cross-geographical effects.[23] Accordingly, we begin with testing equation (8) separately for each MSA; we do this by running a separate OLS regression for each MSA.

To establish the base case that allows comparison with models incorporating supply constraints, we begin by estimating equation (8), but without the supply constraint proxies, and summarize the results in Panel A of Table 3. The rather poor fit is apparent from a glance at the adjusted R^2 values: all except two MSAs have values below 37% (the exceptions being Dallas, with 59%, and Los Angeles, with 53%); Atlanta, Baltimore and Miami even have negative R^2s. In the vast majority of MSAs, house price rises with consumption growth. A one standard deviation increase in consumption growth is associated with a house price rise ranging from 7% in Los Angeles and Phoenix, 5% in San Diego, to 1% in San Francisco and Boston, and less than 0.5% in Dallas, Miami and Houston. Consumption growth is associated with a decline in house prices in Atlanta, Baltimore and Washington, DC; however, these estimates are not different from zero at the 5% significance level. Growth in housing stock is negatively associated with house price growth in 11 of the 19 MSAs we consider, but the confidence interval of the coefficients does not include zero in only four of them. A one standard deviation in housing stock changes house prices between -4% in San Francisco and 5% in Phoenix and Washington, DC. It is interesting to note that in cities that do not have large boom-and-bust type behavior, such as Atlanta, Baltimore, Cleveland, Dallas, Minneapolis-St. Paul and Philadelphia, house prices grow by a constant 3%-4% annually, while other MSAs do not have a significant constant term.

Comparing these results to those of the simplest model that includes the supply constraint $CAPSAIZ$ shown in Panel B of Table 3, we find a significant improvement, to begin with, in the model fit. The adjusted R^2 values now lie between 1% for Washington, DC, and 68% for Dallas. A one standard deviation increase in the constraints on housing supply (as measured by $CAPSAIZ$) has an extremely large – and statistically significant – impact on house prices: they rise from between 1% in Houston and 2% in Cleveland and Washington, DC, to 9% in Miami and 10% in Phoenix. In general, the magnitude of this impact tends to be higher the higher the $SAIZ$ or $WRLURI$ index of the MSA. The constant term is also positive and significant for all MSAs, and indicates a growth of between 3% and 6% annually.

Finally, in Panel C of Table 3, we report the results of the model specified in equation (8), using $D = CAPSAIZ$. The full model has significantly improved fit for all MSAs, except New York: the range of R^2 now lies between 12% and 79%, with 9 MSAs at over 50%. The improvement in fit over the base case is particularly pronounced in cities that experienced

[23] The literature following the spatial equilibrium tradition, such as Glaeser et al. (2006) and Glaeser et al. (2011), explicitly models the migration between geographies, and shows that this explains some of the cross-sectional differences in house prices observed across MSAs in the United States. Our model does not preclude this effect from occurring; it simply shows that, even without accounting for this, supply constraints are of first-order importance in explaining the differences in time-series dynamics across MSAs.

'boom-and-bust' type behavior. While the coefficients of many terms have large standard errors, which lead to their being statistically indistinguishable from zero, the coefficients of the term ΔD retain their significantly positive effect; a one standard deviation increase in the supply constraint measure increases house prices by between 1% in Houston and Cleveland to 9% in Miami and Phoenix. The constant term loses its significance in many MSAs, and remains relevant mostly in the MSAs such as Atlanta, Baltimore, Cleveland, Dallas and Philadelphia, which do not exhibit a bubble-like pattern.

The overwhelming evidence from the regressions at the individual MSA level is therefore that supply constraints are sufficient in explaining the observed 'boom-and-bust' pattern in many MSAs, and explain a large fraction of this pattern in the rest. Supply constraint considerations are relatively less important in cities that do not experience significant constraints to new residential development.

The importance of including supply constraints to explain observed prices is shown in Figure 2, which plots the observed house price index, and the price indices fitted by two models – including, and excluding, measures of supply constraints – along with their 95% confidence intervals, where the indices are aggregated over the 19 MSAs.[24] The full model that includes supply constraints follows the same 'boom-and-bust' pattern, even if it does not fully match the observed price indices. In contrast, the model that does not consider the effects of supply constraints rises almost monotonically during the entire sample, and does not exhibit the 'boom-and-bust' pattern. Additionally, the model without supply constraints fits the observed indices quite well between 1987 and 2000, even though it fails in the latter half of the sample. This is significant: the mere presence of constraints to supply does not imply that they play an important role – recall the similarity between the graphs in Panel A of Figure 1 in the region excepting that where H is low and K is high; it is important for the supply constraint to bind for it to have an effect on prices.

That these results hold not only at the aggregate, but also at the individual MSA level, is seen by comparing Figures 3 and 4, where the former plots the indices of the observed and model-explained house prices for the full model without supply constraints, and the latter for the full model with $CAPSAIZ$ as the measure of supply constraints. Figure 3 reiterates that the model which does not take into account constraints on housing supply shows a nearly monotonic pattern that does not explain the spikes and steep declines in the index observed in the large majority of the MSAs. However, not accounting for supply constraints does not seem to matter much for cities that do not have a 'boom-and-bust' pattern, such as Atlanta, Cleveland, Dallas, Houston and Philadelphia. These are precisely the cities that do not have large supply constraints, at least as measured by the $SAIZ$ and $WRLURI$, which is in line

[24]The aggregate indices are calculated as follows. First, a weighted average of the house prices – whether using the original or the model-generated price series – for each MSA is calculated, where the weights for an MSA are the median size of housing as a fraction of the sum of median housing sizes across the 19 MSAs. The value of the aggregate in 1987 is set to 100, and the value of the index is calculated for the years 1988–2011.

with our argument. Contrasting these results with those plotted in Figure 4, it is immediately clear that introducing supply constraints into the regression helps explain nearly all of the bubble-like behavior in Los Angeles, New York, Phoenix, San Diego, San Francisco, Seattle and Washington, and explains a significant proportion of the observed price index in Miami, all cities that have a high $SAIZ$ (or $WRLURI$) score.

4.1 Regressions with MSA fixed effects

Many of the coefficient estimates in the individual MSA level regressions have large standard errors, and we suspect that this was largely due to small-sample-size issues. For instance, the lack of a longer time series means that we estimate seven parameters for the full model including supply constraints using a sample of 24 observations.

In what follows, we make the rather stringent assumption – not implied by the model, it must be emphasized, which treats each MSA as an unconnected island – that the coefficients that determine the response of housing returns to all independent variables are equal across MSAs. In the language of equation (8), we assume that $\beta_{i,0} \equiv \beta_0$, $\beta_{i,1} \equiv \beta_1$, $\beta_{i,2} \equiv \beta_2$, $\beta_{i,3} \equiv \beta_3$ and $\beta_{i,4} \equiv \beta_4$. However, aware of the fact that there might be heterogeneity unaccounted for by the independent variables we consider in the study, we conduct a panel regression with MSA fixed effects:

$$\Delta p_{i,t} = u_i + \beta_0 \Delta D_{i,t} + \beta_1 \Delta c_{i,t} + \beta_2 D_{i,t} \Delta c_{i,t} + \beta_3 \Delta h_{i,t} + \beta_4 D_{i,t} \Delta h_{i,t} + \lambda_t. \tag{11}$$

Given the results of the regressions for individual MSAs – which showed that the parameter estimates were different, sometimes even in sign, across MSAs – we would expect that the restrictions implied by the fixed-effects regressions would lead the models to perform significantly worse overall. The comparison of the panel data and individual regression analysis informs us about the importance of including coefficients that vary across MSAs. The larger sample size helps to obtain better parameter estimates, which allows us to get a more accurate picture of the importance of including the terms that interact the supply constraint proxy with the other independent variables in the regressions.

As before, the base for comparison is the analysis that does not include proxies for supply constraints, which is shown in the columns labeled [1]-[3] under the 'No supply constraint' panel of Table 4. As expected, consumption growth is positively related with housing returns, either when used as the only explanatory variable, or when used in conjunction with growth in housing stock. While the coefficient is common across MSAs, the consumption growth process is different across MSAs; this makes a one standard deviation increase in consumption translate to house price increases between 1% in New York and 3% in San Diego. Consumption growth averages around 2% for most MSAs; this translates to an average return of about 0.80% for housing. Growth in housing stock is associated with a fall in housing returns, both when used individually, or together with consumption growth, as the explanatory variable. Ceteris paribus, growth in housing stock reduces the need to purchase housing, thus reducing the returns to this asset; this effect turns out to be of the order of 29 basis points for a 1% change in housing stock. The average annual growth of housing stock is of the order of about 0.5%, which implies that the annual impact of the growth in housing stock is to reduce housing returns by about 15 basis points. The per capita housing size is not very volatile; this makes the impact on house prices from this channel small – less than 0.5% for a one standard deviation in housing stock in 10 of the 19 MSAs under study, and the maximum impact is slightly over 2% in Seattle. The (common) intercept in these fixed-effects regressions is highly significant, and is of the order of 2% per annum. The intercept for Dallas is lower by 1.77% – the returns to housing are about 0.23% in the year when neither the non-housing consumption nor the stock of housing changes in Dallas; it is lower by about 1.24% in Detroit and 1.15% in Atlanta. Among cities that have a higher intercept of housing returns than 2% are Portland (1.30% higher), Miami (1.19% higher) and San Francisco (1.10% higher). The models have a rather poor fit, with overall R^2 values of between 1% and 7%.

The set of models labeled [1]-[5] in the second panel of Table 4 use $CAPSAIZ$ as the proxy for the constraints on housing supply. The large increase in model fit is obvious: R^2 values are over 32% for any specification that incorporates supply constraints as an explanatory variable. The most important variable in improving the model is the change in the supply constraints (ΔD): its coefficient is positive and highly significant in every case. A one standard deviation change in supply constraints changes prices by 10% in Miami, 8% in San Diego, 7% in Seattle and 5% in Los Angeles – these large variations occurring precisely in those MSAs that have high $SAIZ$ or $WRLURI$ scores. MSAs that have less-significant supply constraints are less impacted by a 1σ change in $CAPSAIZ$ (Atlanta -0.4%; and Dallas, Houston, Philadelphia, Washington, DC -1%). Consumption (housing) growth has an impact whose magnitude is much smaller: for the full model ([5]), a one standard deviation changes house prices by between 0.5% and 1.7% (0.2% and 3.3%). Given the extremely wide variation across MSAs in the coefficient of the $CAPSAIZ \times (\Delta c)$ ($CAPSAIZ \times (\Delta h)$) term in the results of the individual regressions (i.e., between -10 and +66 (-226 and +151) in Table 3), we expect that the restriction of equal coefficients across MSAs imposed by the panel regression would prove especially costly for these variables. As expected, these coefficients are not large, and neither are they significantly different from zero; their impact is rarely above 0.5%.

The panel regressions confirm the importance of including constraints in supply to explain observed price patterns, and also confirm the importance of allowing the impact of the intercept terms to vary across regions: the slope of house price growth with respect to consumption (housing) growth varies with supply constraints, but with different magnitudes across MSAs. The fit at the aggregate level obtained by alternative models estimated using MSA level fixed effects is shown in Figure 5, which plots the aggregate price index, and the indices fitted by the full model without supply constraints, and the full model that uses $CAPSAIZ$ as the proxy for supply constraints, along with the 95% confidence intervals of the latter two series. This plot shares the three principal features of Figure 2: (i) the monotonic rise in the index that does not account for supply constraints, (ii) the 'boom-and-bust' behavior exhibited by the model that accounts for supply constraints, and (iii) the rather good fit of the model without supply constraints in the first part of the sample. The cost of imposing the restrictions on the parameters is shown in the less-steep rise and fall of the index that incorporates supply constraints.[25]

The graphs in Figure 6 – which plot observed price indices and those fitted by the model implied by the MSA fixed-effects regression (equation (11)) using $CAPSAIZ$ as the measure of supply constraints – show that the impact of constraints on the coefficients is felt largely in the MSAs that experienced a remarkable boom and bust in house prices. The model with MSA fixed effects fails to explain the price dynamics observed in Phoenix and Washington (while the separate OLS models fit them very well); the fit is also worse in Baltimore, Los Angeles, San Francisco and Minneapolis-St. Paul. The two alternative analyses share the failure to explain prices in Detroit and Portland.

Despite the relatively poor performance when compared to separate regressions for each MSA, the results using MSA fixed effects also highlight two issues. First, the importance of accounting for constraints on housing supply in explaining the booms and busts in house prices. Second, the fact that not accounting for supply constraints does not matter in MSAs that do not suffer significant risk that demand is likely to move between the "barrier" represented by the constraint on housing supply. This is the essence of our central message: any meaningful discussion about housing bubbles or boom and bust in house prices should begin with a systematic consideration of the empirical impact of constraints on housing supply.

We also provide a simple method implied by a model with supply constraints to empirically implement this systematic consideration. While our method explains a substantial proportion of the observed price indices in many MSAs, the observed housing returns in Detroit and Portland are far higher than those generated by our model. Table 5 compares the fit of the booms and busts in house prices provided by the model with supply constraints to the fit of the model without supply constraints. The inability of the model incorporating supply constraints does not automatically qualify the resulting price pattern as a bubble, however: we view our analysis as a starting point for a discussion about the causes of the swings that remain unexplained by our framework and methods.

[25]The flip side of imposing these restrictions – improved parameter estimation – is seen in the far tighter 95% confidence bands of the estimated indices.

4.2 Using alternative proxies for constraints in housing supply

We check that our results hold using a variety of other proxies for constraints in housing supply, involving interactions between variables of labor availability for the construction sector, and the geographical or regulatory measures of land restrictions to residential housing construction.[26] We briefly discuss the results using $CAPWRI$, which show that while the overall direction of results remains unaltered, some measures of supply restrictions may be better suited for some regions than others.[27]

The similarity of results using this alternative measure is shown in Table 4 under the heading $D = CAPWRI$. While the fit is slightly lower than for the corresponding model using $CAPSAIZ$, the single most important variable is the change in the supply constraint. The magnitudes of the coefficients are different because $WRLURI$ and $SAIZ$ have different magnitudes, but the t-statistics are comparably strong. Consumption growth has a positive impact on housing returns, and the coefficient remains significant even after the introduction of the interaction term $CAPWRI \times \Delta c$. Growth in housing stock continues to exercise a negative impact on housing returns. The interaction term between $CAPWRI$ and growth in housing stock is also negative, and the magnitude and significance of the constant term are very similar.

Comparing the observed price indices with those generated by the model using $CAPWRI$ (not reported), we find that the most striking difference is that using $CAPWRI$ results in a larger rise and fall in the model-fitted price index in Phoenix. Since the only difference is the use of the $WRLURI$ – as against $SAIZ$ – the intuitive idea that supply restrictions in different regions arise due to different causes is strongly suggested. While Phoenix has no significant geographical restrictions (it is among the five least geographically constrained MSAs in our sample), it has the sixth-strongest regulations of land use for residential building; it could be argued that this makes the $CAPWRI$ measure more appropriate for Phoenix. In the same vein, using $CAPWRI$ does not help explain the price spike and decline in MSAs such as Chicago, Los Angeles, Miami and San Diego as well as using $CAPSAIZ$ does. What is common to these MSAs is that they are coastal cities with significant geographical constraints, but have lower ranking in the restrictions due to regulations in the sample cross-section.

These results suggest that the construction of the measure reflecting constraints on housing supply needs to take into account the properties of the particular area in question, or alternatively, that a composite measure that takes into account both geographical and regulatory constraints may be more appropriate.

[26] Results are not reported, but are available on request.

[27] $CAPSAIZ$ interacts the index of geographical constraint with an index of labor constraints; $CAPWRI$ interacts an index of land use regulations with the index of labor constraints. As shown in Frech III and Lafferty (1984), it is quite conceivable that regulatory constraints may impact house prices even in the absence of geographical constraints; thus, $CAPSAIZ$ and $CAPWRI$ may have different impact in different MSAs.

5 Conclusion

We analyze the importance of including a key friction in housing markets in a representative-agent two-goods economy. We present evidence that the presence of constraints to housing supply could generate 'boom-and-bust' behavior in house prices: when demand rises above the level allowed by supply constraints, house prices rise sharply, and when the demand falls back below this key level, prices plummet. There exists strong evidence of the importance of constraints in explaining a large proportion of the observed price patterns in the data. Our empirical specification also provides a simple method that allows researchers to control for supply constraints when analyzing real estate bubbles.

Table 1: **Baseline parameters.** This table lists the parameters used to obtain the baseline numerical solution.

Parameter description	Notation	Value
Drift of the process for the growth of liquid capital	α_K	0.030
Drift of the process for the growth of housing stock	α_H	-0.013
Volatility of the process for the growth of liquid capital	σ_K	0.115
Volatility of the process for the growth of liquid stock	σ_H	0.060
Instantaneous correlation between shocks to liquid and housing stocks	ρ	-0.900
*Inter*temporal elasticity of substitution	γ	2.900
*Intra*temporal elasticity of substitution	ϵ	1.250
Parameter	ω	0.800
Subjective discount rate	β	0.100
Time interval	"dt"	0.250

Table 2: **Summary statistics.** This table provides the sample mean and standard deviation of the time series of the log returns of the median house price per square foot Δp, log growth of per capita non-housing consumption Δc, log growth of per capita housing stock Δh, and the various measures used to construct proxies for constraints of housing supply – the Saiz (2010) measure of geographical constraints to housing construction ($SAIZ$), the Gyourko et al. (2008) Wharton Residential Land Use Regulation Index ($WRLURI$), the percentage of employees in the private sector employed in the construction industry as a fraction of its maximum value in the sample ($DMPEC$) – and the measures we use as proxies in the main text, $CAPWRI = DMPEC \times WRI$, and $CAPWRI = DMPEC \times SAIZ$. Note that the standard deviations of $SAIZ$ and $WRLURI$ are omitted, since these indices are not time-varying. The sample consists of data from 1988 to 2011 at the Metropolitan Statistical Area (MSA) level. The details on the data sources are discussed in the main text.

MSA	Δp (\$/sq ft)		Δc (\$)		Δh (sq ft)		$SAIZ$	$WRLURI$	$DMPEC$		$CAPSAIZ$		$CAPWRI$	
	μ	σ	μ	σ	μ	σ	μ	μ	μ	σ	μ	σ	μ	σ
Atlanta	0.021	0.042	0.022	0.047	0.007	0.024	0.041	0.040	0.870	0.103	0.035	0.004	0.035	0.004
Baltimore	0.040	0.067	0.032	0.068	0.004	0.009	0.219	0.790	0.876	0.075	0.191	0.016	0.692	0.059
Boston	0.033	0.065	0.039	0.040	-0.030	0.031	0.339	1.540	0.823	0.133	0.279	0.045	1.268	0.204
Chicago	0.034	0.054	0.027	0.042	-0.002	0.005	0.400	0.060	0.886	0.086	0.355	0.034	0.053	0.005
Cleveland	0.025	0.032	0.021	0.055	0.000	0.008	0.405	-0.160	0.908	0.068	0.368	0.028	-0.145	0.011
Dallas	0.017	0.035	0.026	0.035	0.013	0.029	0.092	-0.350	0.810	0.086	0.074	0.008	-0.284	0.030
Detroit	0.023	0.069	0.028	0.036	-0.004	0.009	0.245	0.120	0.846	0.107	0.207	0.026	0.102	0.013
Houston	0.029	0.032	0.026	0.044	0.006	0.006	0.084	-0.190	0.910	0.053	0.076	0.004	-0.173	0.010
Los Angeles	0.041	0.116	0.025	0.041	-0.007	0.030	0.525	0.510	0.805	0.105	0.422	0.055	0.411	0.053
Miami	0.039	0.112	0.005	0.068	0.002	0.034	0.766	0.370	0.802	0.105	0.615	0.080	0.297	0.039
Minneapolis-St.Paul	0.033	0.054	0.030	0.046	0.005	0.004	0.192	0.340	0.816	0.119	0.157	0.023	0.277	0.041
New York	0.037	0.067	0.027	0.025	-0.004	0.025	0.404	0.630	0.870	0.092	0.352	0.037	0.548	0.058
Philadelphia	0.039	0.058	0.026	0.054	0.001	0.007	0.102	1.030	0.836	0.077	0.085	0.008	0.861	0.079
Phoenix	0.022	0.111	0.024	0.043	0.006	0.007	0.140	0.700	0.753	0.142	0.105	0.020	0.527	0.100
Portland	0.053	0.065	0.036	0.034	0.006	0.006	0.375	0.290	0.859	0.078	0.322	0.029	0.249	0.023
San Diego	0.040	0.102	0.022	0.083	0.005	0.004	0.634	0.480	0.795	0.135	0.504	0.086	0.382	0.065
San Francisco	0.047	0.091	0.030	0.039	0.015	0.036	0.436	1.010	0.826	0.124	0.360	0.054	0.834	0.125
Seattle	0.046	0.068	0.034	0.044	-0.032	0.069	0.731	0.900	0.796	0.084	0.583	0.061	0.717	0.075
Washington, DC	0.044	0.085	0.030	0.039	0.010	0.007	0.140	0.330	0.700	0.083	0.098	0.012	0.231	0.028

Table 3: **Independent regressions for each MSA.** This table summarizes the results of the independent regressions, for 19 MSA in the United States, of housing log returns on alternative specifications involving changes in the measure of supply constraints ΔD, non-housing consumption log growth Δc and growth of housing (size) stock Δh. Panel A shows results of the model that do not use measures of supply constraints. Panels B and C show results where the supply constraint measure is $CAPSAIZ$, the product of $(DMPEC)$ – the percentage of employees in the private sector working in the construction industry as a fraction of its sample maximum within the MSA – and the Saiz (2010) index of geographical constraints ($SAIZ$). Each column reports results for an MSA; quantities below each regression coefficient in parentheses represent the corresponding t-statistics. The sample consists of data from 1988 to 2011.

	Atlanta	Baltimore	Boston	Chicago	Cleveland	Dallas	Detroit	Houston	Los Ang	Miami	Minn.-St.Paul	New York	Philadelphia	Phoenix	Portland	San Diego	San Fran	Seattle	Wash, DC
Panel A: Model without supply constraints																			
Δc	-0.21	-0.01	0.30	**0.66**	-0.03	0.11	0.67	0.02	**1.81**	0.02	**0.57**	0.89	0.36	**1.64**	**0.80**	**0.61**	0.29	**0.78**	-0.51
	(-1.09)	(-0.05)	(0.96)	(2.82)	(-0.27)	(0.71)	(1.75)	(0.16)	(4.34)	(0.06)	(2.31)	(1.53)	(1.67)	(3.23)	(2.16)	(2.82)	(0.64)	(2.72)	(-1.33)
Δh	0.01	2.02	**-1.02**	-2.37	-1.28	**-1.01**	1.18	**2.82**	**-1.20**	0.68	-3.88	0.18	2.99	6.57	**5.06**	7.07	**-1.16**	-0.06	**6.51**
	(0.03)	(1.33)	(-2.52)	(-1.23)	(-1.61)	(-5.36)	(0.76)	(2.63)	(-2.10)	(0.91)	(-1.43)	(0.31)	(1.78)	(0.38)	(2.03)	(1.56)	(-2.41)	(-0.33)	(3.18)
$Const$	**0.03**	**0.03**	-0.01	0.01	**0.03**	**0.03**	0.01	0.01	-0.01	0.04	**0.04**	0.01	**0.03**	-0.04	-0.01	-0.01	**0.06**	0.02	-0.01
	(2.62)	(2.09)	(-0.40)	(1.05)	(3.68)	(4.75)	(0.48)	(1.23)	(-0.69)	(1.42)	(2.20)	(0.70)	(2.15)	(-0.62)	(-0.58)	(-0.25)	(2.47)	(1.04)	(-0.35)
$AdjR^2$	-0.04	-0.01	0.17	0.23	0.03	0.59	0.07	0.19	0.53	-0.06	0.14	0.02	0.11	0.34	0.37	0.32	0.15	0.20	0.29
Panel B: Model with only supply constraints																			
ΔD	**14.75**	**4.23**	**2.22**	**2.32**	**1.01**	**5.98**	**4.85**	**3.96**	**2.17**	**1.81**	**4.99**	**1.79**	**6.15**	**10.27**	**2.65**	**1.75**	**2.52**	**1.54**	2.86
	(4.45)	(2.95)	(4.62)	(4.23)	(2.66)	(6.98)	(4.56)	(2.02)	(2.64)	(5.23)	(6.19)	(2.36)	(2.25)	(6.19)	(4.55)	(4.52)	(3.18)	(6.84)	(1.08)
$Const$	**0.03**	**0.05**	**0.04**	**0.04**	**0.02**	**0.02**	**0.03**	**0.03**	**0.05**	**0.06**	**0.04**	**0.04**	**0.05**	**0.03**	**0.05**	**0.05**	**0.05**	**0.05**	**0.05**
	(4.32)	(3.98)	(4.06)	(4.97)	(4.34)	(5.29)	(2.79)	(4.76)	(2.52)	(3.59)	(5.97)	(3.32)	(4.02)	(2.19)	(4.76)	(3.34)	(3.13)	(6.12)	(2.75)
$AdjR^2$	0.45	0.25	0.47	0.42	0.21	0.68	0.46	0.12	0.21	0.57	0.62	0.17	0.15	0.65	0.50	0.46	0.28	0.67	0.01
Panel C: Full model																			
ΔD	**16.18**	**4.75**	**2.09**	**1.99**	0.55	**4.29**	**4.65**	**3.86**	1.26	**1.83**	**4.70**	1.60	**6.00**	**9.39**	**1.92**	**1.52**	**2.28**	**1.35**	**5.87**
	(4.94)	(3.28)	(3.65)	(3.31)	(1.25)	(3.67)	(3.85)	(2.17)	(1.92)	(4.62)	(5.15)	(1.90)	(2.31)	(4.32)	(2.88)	(2.99)	(3.29)	(4.10)	(2.64)
Δc	-1.90	-2.54	-2.13	-1.30	3.71	**-2.90**	1.37	**-5.06**	**-8.54**	1.30	-2.18	-4.25	-2.83	-1.25	-3.55	0.08	1.69	-1.55	-0.93
	(-1.73)	(-1.22)	(-1.40)	(-0.32)	(1.52)	(-3.29)	(0.53)	(-2.15)	(-2.63)	(0.87)	(-1.39)	(-0.85)	(-1.22)	(-0.58)	(-0.72)	(0.05)	(0.70)	(-0.80)	(-0.22)
$D \times \Delta c$	51.28	12.41	6.88	4.63	-9.94	**38.53**	-4.78	**66.00**	**22.46**	-3.09	15.13	14.03	36.36	14.18	11.94	0.54	-4.52	2.89	4.62
	(1.71)	(1.21)	(1.43)	(0.44)	(-1.53)	(3.28)	(-0.41)	(2.15)	(3.15)	(-1.23)	(1.44)	(0.99)	(1.34)	(0.70)	(0.80)	(0.17)	(-0.70)	(0.85)	(0.11)
Δh	**7.05**	30.56	1.43	-0.24	-32.99	2.47	7.34	2.24	-2.33	-6.60	-12.94	8.56	-5.15	0.51	10.10	-33.72	**-7.09**	3.20	-10.59
	(2.07)	(1.66)	(0.80)	(0.00)	(-1.59)	(1.49)	(0.37)	(0.17)	(-0.27)	(-0.77)	(-0.75)	(0.62)	(-0.29)	(0.01)	(0.39)	(-1.53)	(-2.24)	(0.82)	(-0.50)
$D \times \Delta h$	**-226.84**	-143.28	-5.63	-2.01	87.49	-37.92	-31.48	3.76	2.61	10.39	101.26	-22.33	81.24	96.99	-11.26	58.08	17.26	-5.98	151.29
	(-2.18)	(-1.55)	(-0.83)	(-0.01)	(1.56)	(-1.68)	(-0.31)	(0.02)	(0.13)	(0.80)	(0.89)	(-0.61)	(0.40)	(0.34)	(-0.13)	(1.56)	(1.92)	(-0.82)	(0.86)
$Const$	**0.03**	**0.04**	0.04	**0.03**	**0.02**	**0.03**	0.02	0.01	0.00	**0.06**	0.02	0.02	**0.03**	-0.02	0.00	0.06	**0.05**	**0.04**	0.02
	(4.21)	(3.26)	(1.79)	(1.97)	(3.62)	(6.44)	(1.51)	(1.59)	(-0.07)	(2.61)	(1.76)	(0.86)	(2.23)	(-0.40)	(-0.18)	(1.90)	(2.72)	(3.22)	(0.43)
$AdjR^2$	0.56	0.30	0.51	0.45	0.21	0.79	0.42	0.41	0.66	0.58	0.66	0.12	0.26	0.67	0.57	0.55	0.46	0.62	0.41

Table 4: **Panel regressions with MSA fixed effects.** This table summarizes the results of the regression of housing log returns on alternative specifications involving changes in the measure of supply constraints ΔD, non-housing consumption log growth Δc and growth of housing (size) stock Δh. The first three columns show results of models that do not use measures of supply constraints. The subsequent five columns report results where the supply constraint measure is $CAPSAIZ$, the product of $DMPEC$ – the percentage of employees in the private sector employed in the construction industry as a fraction of its maximum within the MSA in the sample – and the Saiz (2010) index of geographical constraints ($SAIZ$). The last five columns report the results where the supply constraint measure is $CAPWRI$, the product of $DMPEC$ and the Gyourko et al. (2008) Wharton Residential Land Use Regulation Index ($WRLURI$). Each column reports results with a different choice of the independent variables; quantities in the second row in parentheses below the point estimates are the corresponding t-statistics. The overall model R^2 is also reported. The sample consists of data from 19 MSAs in the United States, from 1988 to 2011.

	No supply constraint			$D = CAPSAIZ$					$D = CAPWRI$				
	[1]	[2]	[3]	[1]	[2]	[3]	[4]	[5]	[1]	[2]	[3]	[4]	[5]
ΔD				**2.06**	**1.96**	**1.95**	**1.97**	**1.94**	**1.07**	**1.00**	**0.99**	**0.99**	**1.03**
				(14.62)	(13.57)	(13.45)	(13.23)	(12.97)	(11.38)	(10.68)	(10.49)	(10.55)	(10.58)
Δc	**0.39**		**0.39**		**0.17**	**0.18**	0.20	**0.22**		**0.26**	**0.27**	**0.18**	**0.21**
	(5.36)		(5.44)		(2.77)	(2.84)	(1.77)	(1.97)		(4.00)	(4.05)	(1.98)	(2.24)
$D \times \Delta c$							-0.09	-0.13				0.20	0.15
							(-0.29)	(-0.42)				(1.21)	(0.88)
Δh		**-0.29**	-0.31			-0.20		-0.48			-0.17		**-0.58**
		(-2.03)	(-2.20)			(-1.69)		(-1.61)			(-1.38)		(-2.48)
$D \times \Delta h$								0.67					0.74
								(1.02)					(2.06)
Const	**0.02**	**0.03**	**0.02**	**0.04**	**0.04**	**0.04**	**0.04**	**0.04**	**0.04**	**0.03**	**0.03**	**0.03**	**0.03**
	(6.23)	(10.06)	(6.22)	(13.76)	(10.41)	(10.37)	(10.37)	(10.38)	(12.35)	(8.82)	(8.77)	(8.74)	(8.97)
R^2	0.06	0.01	0.07	0.32	0.33	0.33	0.33	0.33	0.22	0.25	0.25	0.26	0.27

Table 5: **Comparison of model performance in capturing price rise in 2000-2006.** The table reports the house price price growth in 19 MSAs between the years 2000 and 2006, along with the house price growth explained by alternative models in the same period. Columns 2-3 contain growth rates implied by separate regressions for each MSA, while columns 4-5 contain those explained by panel regressions with MSA fixed effects. Returns are calculated using indices constructed from the fitted house price returns in each year in each MSA. As a measure of supply constraints, columns 2 and 4 use $CAPSAIZ$, the product of $DMPEC$, the percentage of employees in the private sector working in the construction industry as a fraction of the maximum value of this quantity in the MSA, and $SAIZ$, the Saiz (2010) measure of geographical constraints. Models in columns 3 and 5 do not use constraints on housing supply. Model estimation is conducted over the sample of 19 MSAs for the years 1988 to 2011.

	Original	Individual regressions		MSA fixed effect	
	series	$CAPSAIZ$	No constr	$CAPSAIZ$	No constr
Atlanta	30%	25%	13%	17%	16%
Baltimore	109%	73%	43%	31%	21%
Boston	63%	59%	38%	38%	31%
Chicago	56%	40%	29%	34%	26%
Cleveland	18%	18%	12%	12%	14%
Dallas	23%	20%	17%	14%	11%
Detroit	13%	4%	15%	10%	16%
Houston	32%	28%	25%	19%	20%
Los Angeles	152%	132%	83%	74%	39%
Miami	163%	105%	39%	99%	21%
Minneapolis-St.Paul	55%	44%	34%	27%	24%
New York	92%	40%	26%	34%	22%
Philadelphia	81%	30%	22%	30%	27%
Phoenix	109%	81%	42%	27%	20%
Portland	74%	40%	24%	46%	35%
San Diego	112%	90%	41%	74%	33%
San Francisco	65%	50%	21%	43%	32%
Seattle	73%	51%	22%	51%	20%
Washington, DC	124%	67%	37%	35%	27%

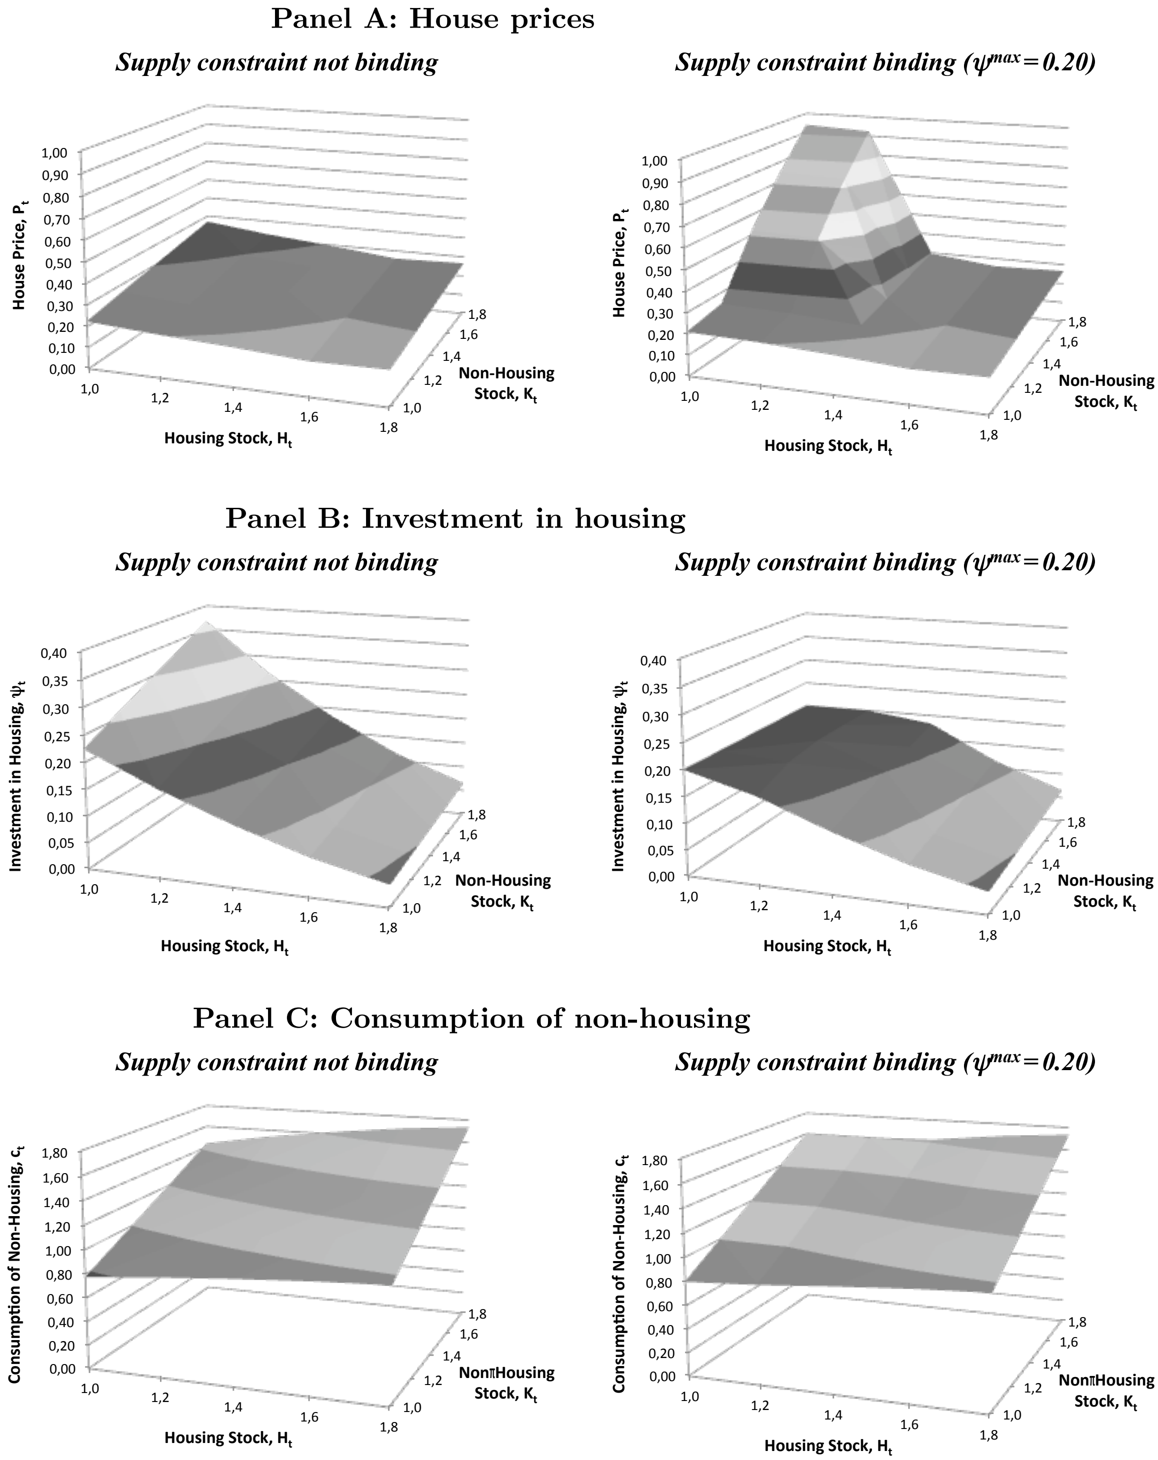

Figure 1: **Baseline numerical results.** House prices (Panel A), investment in housing (Panel B), and consumption of the non-housing good (Panel C) obtained from the representative agent's optimization problem. Each panel shows the results when no supply constraints are binding (left figure) and the results when the supply constraints are binding (right figure.)

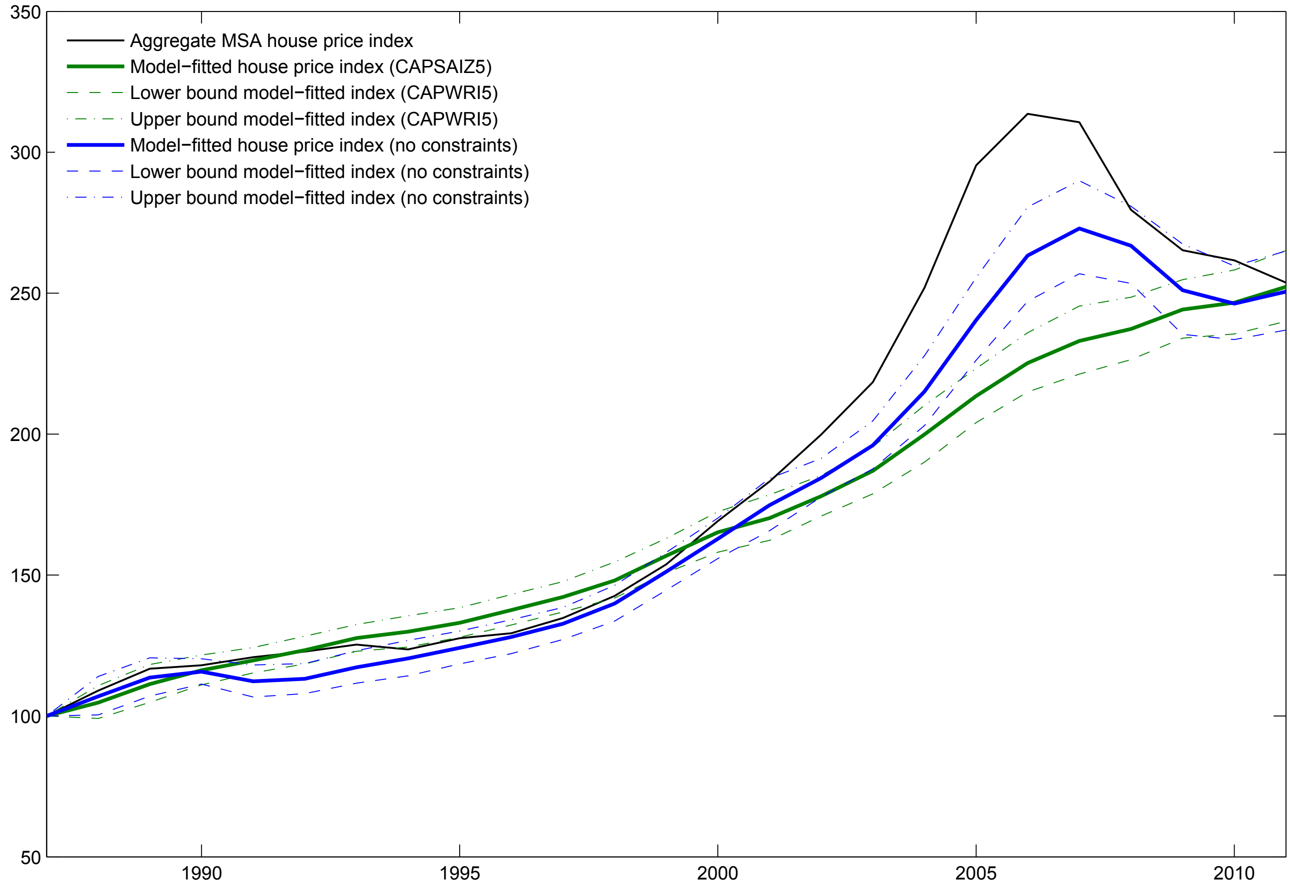

Figure 2: **Aggregate observed index and house price index explained by separate regressions for each MSA.** The figure plots the time series of the house price index for the United States obtained by aggregating the price series for the 19 MSAs in our sample, the aggregate house price series obtained by fitting the regression $\Delta p_i = u_i + \beta_{i,1}\Delta c + \beta_{i,3}\Delta h$ for each MSA separately, and that obtained through aggregation after fitting the values obtained using separate MSA-level regressions $\Delta p_i = u_i + \beta_{i,0}\Delta D + \beta_{i,1}\Delta c + \beta_{i,2}D\Delta c + \beta_{i,3}\Delta h + \beta_{i,4}D\Delta h$, where the proxy of housing supply constraints $D = CAPSAIZ$ is the product of the percentage of employees in the private sector employed in the construction industry as a fraction of its sample maximum in the MSA ($DMPEC$) and the Saiz (2010) index of geographical constraints ($SAIZ$). 95% confidence intervals for the two indices obtained using model-fitted returns are plotted using dashed lines of the corresponding color. To obtain the U.S.-level indices, individual MSA prices are weighted by the ratio of the housing size in the MSA to the total housing size across MSAs, and aggregated. The resulting prices are then reconverted into indices, where the value of each series in 1987 is set to 100. The estimation of the parameters is based on data from 1988 to 2011.

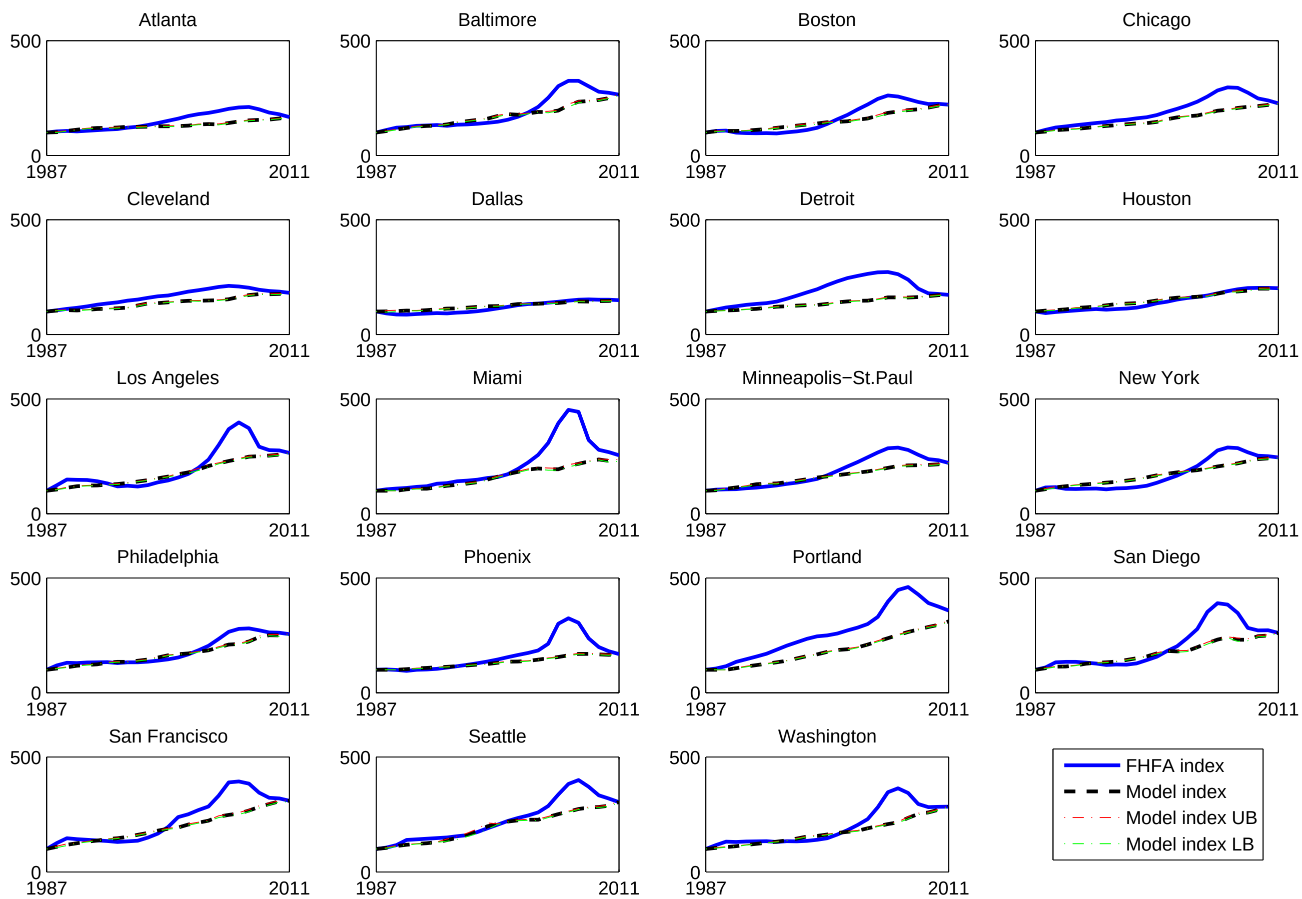

Figure 3: **MSA-level comparison of observed index and house price index fitted by separate regressions for each MSA without supply constraints.** The figure plots the observed house price index for each of 19 MSAs, and that implied by house price returns fitted by $\Delta p_i = u_i + \beta_{i,1}\Delta c + \beta_{i,3}\Delta h$, where the coefficients are estimated by separate regressions for each MSA, along with its 95% confidence interval bands (plotted using dashed lines). The value in 1987 of each series is set to 100. The sample contains data for the period 1988 to 2011; the graph title is shorthand for the MSA to which it belongs. The solid line plots the FHFA index, and the dashed line that explained by the model.

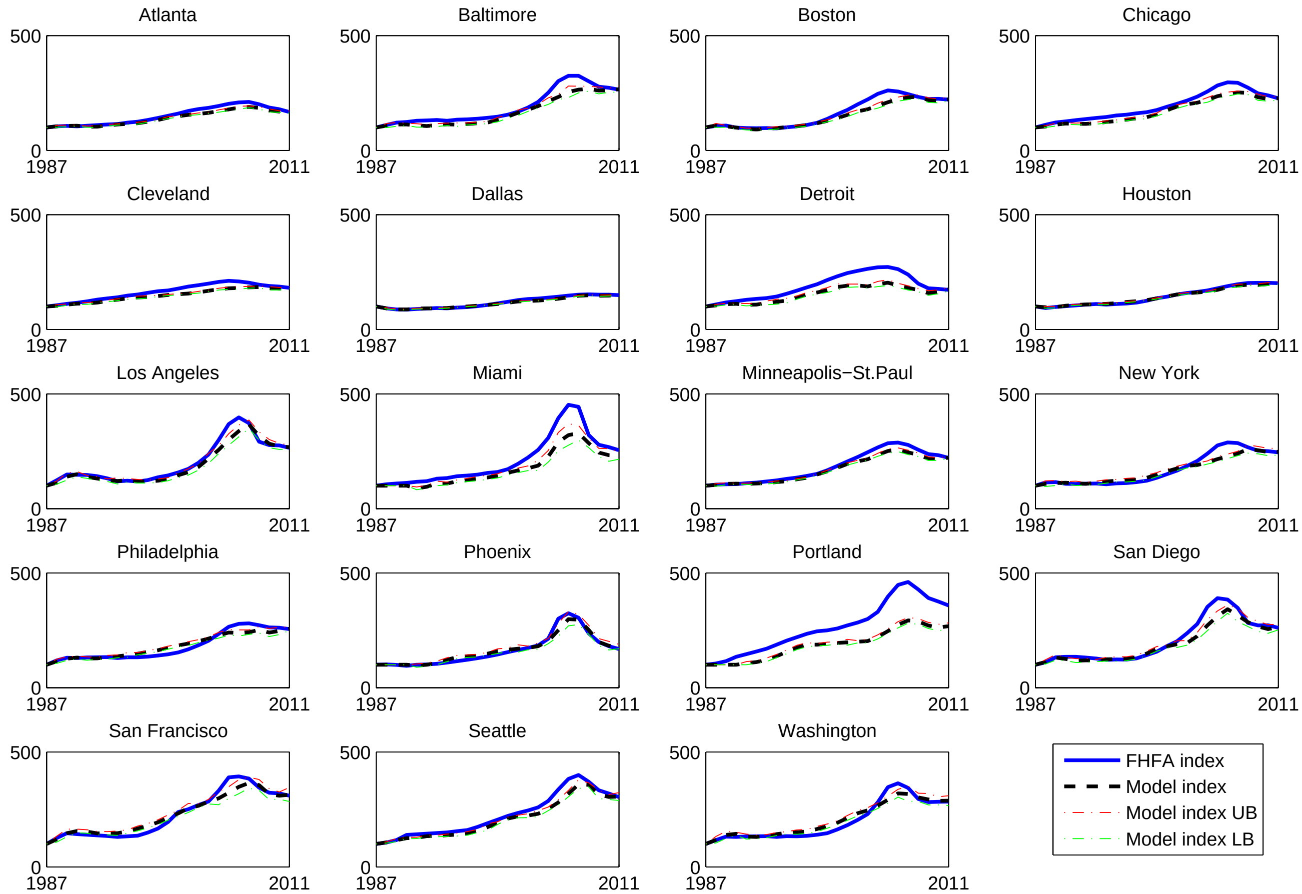

Figure 4: **MSA-level comparison of observed index and house price index fitted by separate regressions for each MSA using supply constraints.** The figure plots the time series of the FHFA house price index (normalized to 100 in 1987) and the house price index (1987 value set to 100) implied by taking the fitted house price returns by $\Delta p_i = \alpha_i + \beta_{0,i}\Delta D + \beta_{i,1}\Delta c + \beta_{i,2}D\Delta c + \beta_{i,3}\Delta h + \beta_{i,4}D\Delta h$ for each of 19 MSAs, where the coefficients are estimated by using separate regressions for each MSA. The measure of housing supply constraints is $CAPSAIZ$, the product of the percentage of employees in the private sector employed in the construction industry as a fraction of its sample maximum ($DMPEC$) and the Saiz (2010) index of geographical constraints ($SAIZ$). The 95% confidence interval of the model-fitted price is plotted in dashed lines. Estimation of the parameters is based on data from 1988 to 2011. The title of each graph is shorthand for the MSA to which it belongs; the solid line plots the FHFA index, and the dashed line that explained by the model.

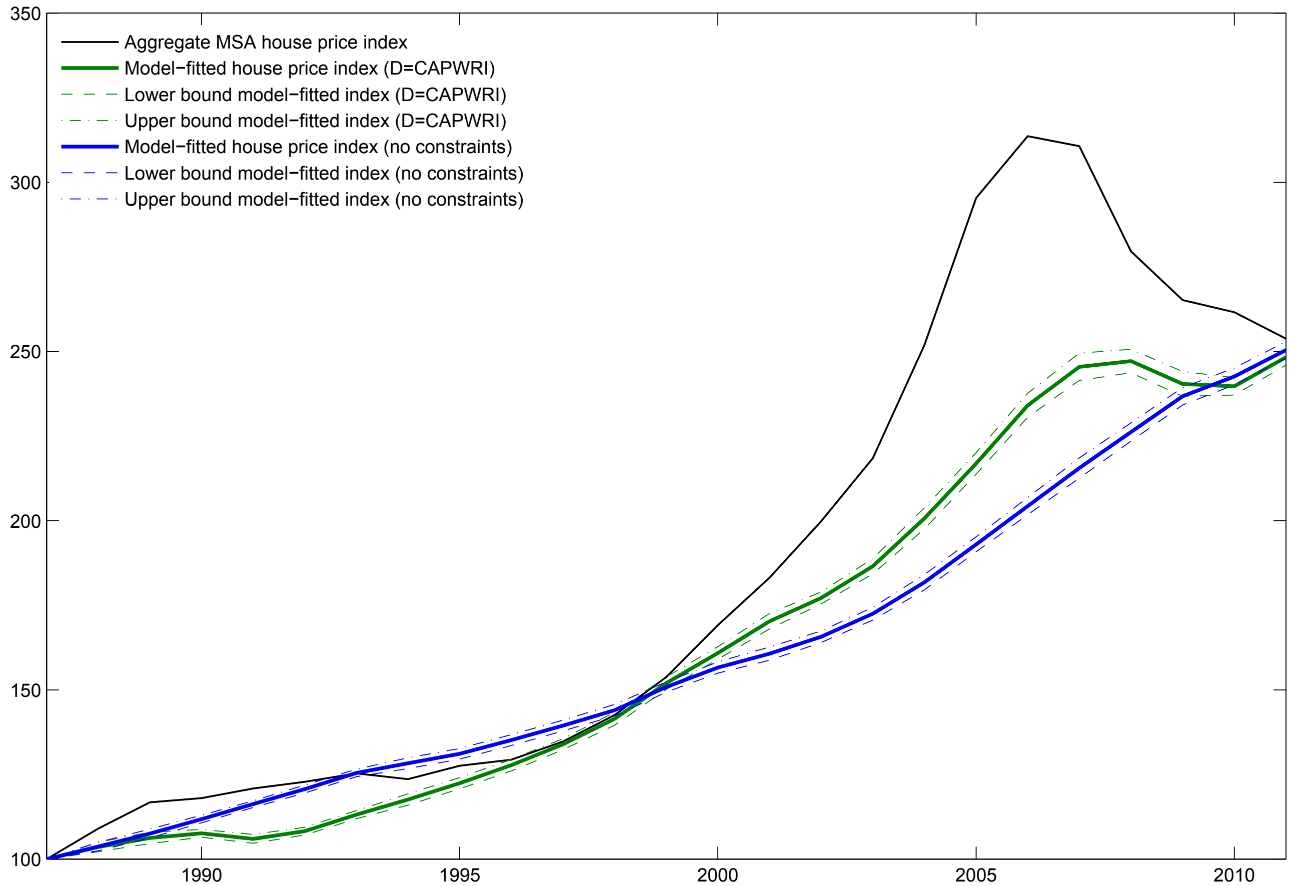

Figure 5: **Aggregate observed index and house price index fitted using panel regression.** The figure plots the time series of the house price index for the United States obtained by aggregating the price series for the 19 MSAs in our sample, the aggregate house price series implied by taking the house price returns fitted by a model without supply constraints $\Delta p_i = u_i + \beta_1 \Delta c + \beta_3 \Delta h$, and those explained by a model using $CAPSAIZ$, the product of the percentage of employees in the private sector employed in the construction industry as a fraction of its sample maximum ($DMPEC$) and the Saiz (2010) index of geographical constraints ($SAIZ$), as the supply constraint measure in $\Delta p_i = u_i + \beta_0 \Delta D + \beta_1 \Delta c + \beta_2 D \Delta c + \beta_3 \Delta h + \beta_4 D \Delta h$, where the coefficients are estimated by using a fixed-effects regression. 95% confidence intervals for the two indices obtained using model-fitted returns are plotted using dashed lines of the corresponding color. The fitted returns are used to find the time series of model-implied house prices for each MSA, which are then aggregated by taking a weighted sum across MSAs in each year, where the weight for each MSA is the ratio of the housing size in the MSA to the total housing size across MSAs. The resulting prices are then reconverted into indices, where the value of 1987 is set to 100. The estimation of the parameters is based on a sample of 19 MSAs listed in the text from 1988 to 2011.

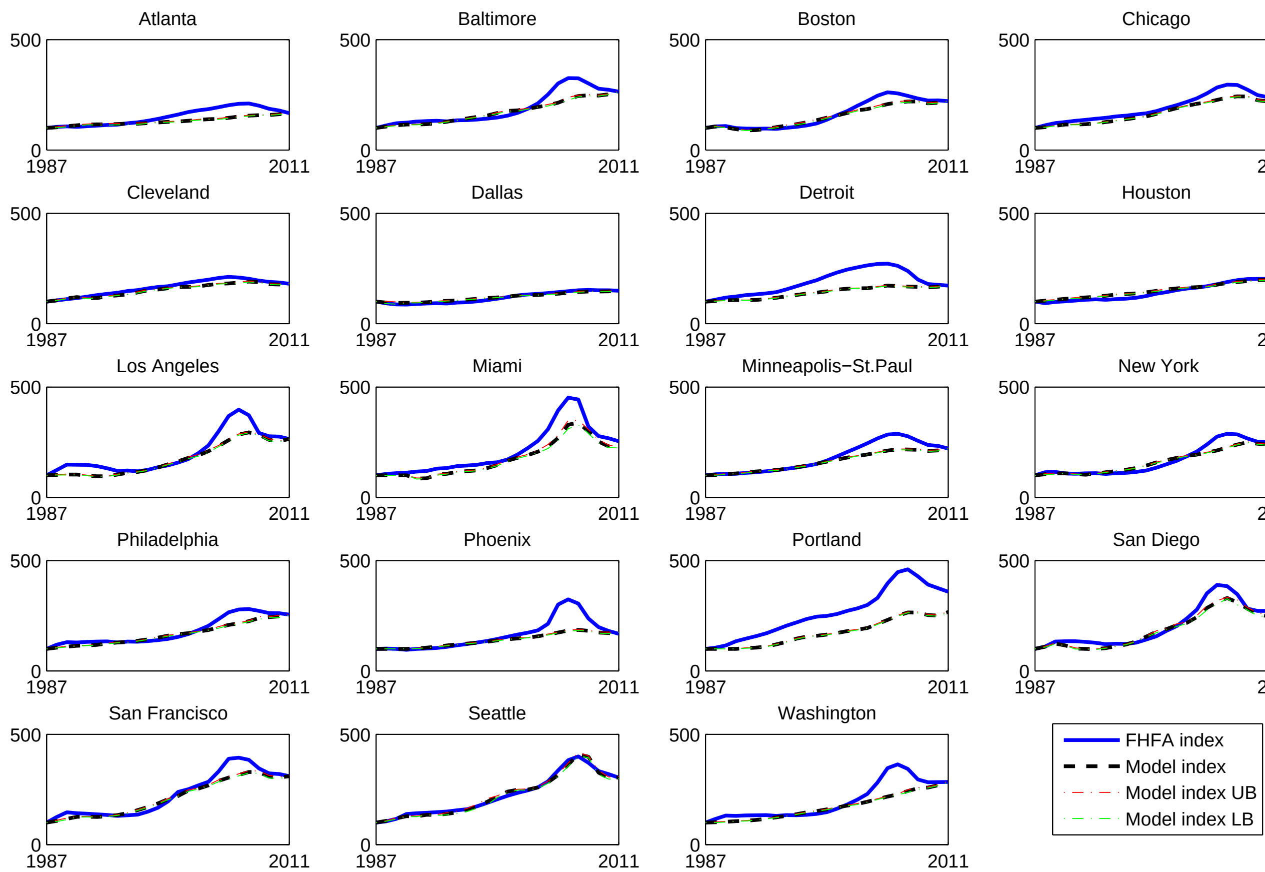

Figure 6: **MSA-level comparison of observed index and house price index fitted by MSA fixed-effect regressions using supply constraint measures.** The figure plots the time series of the FHFA house price index (normalized to 100 in 1987) and the house price index (1987 value set to 100) implied by taking the fitted house price returns by $\Delta p_i = u_i + \beta_0 \Delta D + \beta_1 \Delta c + \beta_2 D \Delta c + \beta_3 \Delta h + \beta_4 D \Delta h$ for 19 MSAs, where the coefficients are estimated by using a fixed-effects regression. The measure of housing supply constraints is $CAPSAIZ$, the product of the percentage of employees in the private sector employed in the construction industry as a fraction of its sample maximum ($DMPEC$) and the Saiz (2010) index of geographical constraints ($SAIZ$). The 95% confidence interval of the model-fitted price is plotted in dashed lines. Estimation of the parameters is based on data from 1988 to 2011. The title of each graph is shorthand for the MSA to which it belongs; the solid line plots the FHFA index, and the dashed line that explained by the model.

A Appendix

A.1 Discretization of the problem

To search for a numerical solution, we discretize both the state variable and the time dimension. The value function then becomes

$$\begin{aligned} J(K_t, H_t; c_t, \psi_t) &= \max_{\substack{c_s \geq 0; 0 \leq \psi_s \leq \Psi^{\max} \\ K_s \geq 0}} E[U(c_t, H_t)\phi + e^{-\beta \times \phi} J(K_{t+\phi}, H_{t+\phi}; c_{t+\phi}, \psi_{t+\phi})] \\ &= \max_{\substack{c_s \geq 0; 0 \leq \psi_s \leq \Psi^{\max} \\ K_s \geq 0}} [U(c_t, H_t)\phi + \alpha \mathbf{\Pi} J(K_{t+\phi}, H_{t+\phi}; c_{t+\phi}, \psi_{t+\phi})], \end{aligned}$$

where ϕ is the distance between two successive time periods in the discretized time space, $\alpha = e^{-\beta \times \phi}$ and $\mathbf{\Pi}$ is the matrix of transition probabilities between all possible states in the discretized state space. $\mathbf{\Pi}$ is a function of the chosen controls c_t, ψ_t and the house price P_t. We describe, in appendix A.2, the procedure to obtain the transition matrix given c_t, ψ_t and P_t. The problem now is to find the optimal value function $J^*(K_t, H_t)$ and the optimal control $[c^*(K_t, H_t),\ \psi^*(K_t, H_t)]$. Since this is an infinite-horizon discounted cost problem, the value function, the (stationary) transition matrix and the control variables are all independent of t. Thus, we can write the optimal value function

$$J(K, H; c^*, \psi^*) = U(c^*, H)\phi + \alpha \mathbf{\Pi}^* J(K, H; c^*, \psi^*), \tag{12}$$

where $\mathbf{\Pi}^*$ is the transition matrix under the optimal control. With the transition matrix in hand, we can write $[I - \alpha \mathbf{\Pi}^*] J^* = U^* \phi$ and thus

$$J^* = [I - \alpha \mathbf{\Pi}^*]^{-1} U^* \phi. \tag{13}$$

In appendix A.3, we describe the policy iteration algorithm to find successive approximations of the value function and the control policy. The procedure involves the following steps: *(i)* start with a guess of the optimal control policy; *(ii)* calculate the transition matrix and find an intermediate value function for the given controls using equation (13); *(iii)* find updated controls satisfying the constraints $c \geq 0$, $0 \leq \psi \leq \Psi^{\max}$ and $c + \psi \leq K$ that maximize the right-hand side of equation (12), where, although the value function is calculated using the previous controls, the utility function and the matrix of transition probabilities are both functions of the unknown controls; *(iv)* use the new controls as inputs for steps *(ii)* and *(iii)*; repeat the process until the solution converges.[28] The output of this algorithm is the optimal control vector (c^*, ψ^*), the optimal value function J^* and the house price P^* as a function of the stock of the numeraire and housing assets, K and H, respectively. Note that since the

[28] See Bertsekas (1995) and references therein for proofs of convergence of this algorithm and other technical details.

optimization is an infinite-horizon problem, we obtain a steady-state solution that gives us time-independent optimal controls that are only a function of the stock of the numeraire and housing assets. Given a defined initial range of endowments, we use these optimal controls and the optimal value function for each combination of a discrete number of initial endowments to obtain the endogenously generated house price using equation (6).

A.2 Transition probability matrix

Each node i is defined by its coordinates (K^i, H^i), the value of the first and second state variables, respectively. Given the controls (c^i, ψ^i) and the house price P^i at node i, the discretized equations of motion are

$$\begin{aligned} K^j &= K^i + (K^i\alpha_K - c^i - \psi^i)\phi + K^i\sigma_K\Delta W^K \\ H^j &= H^i + \left(H^i\alpha_H + \frac{1}{P^i}\psi^i\right)\phi + H^i\sigma_H\Delta W^H, \end{aligned}$$

where

$$\begin{bmatrix} \Delta W^K \\ \Delta W^H \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \phi & \rho\phi \\ \rho\phi & \phi \end{bmatrix}\right).$$

Except for some special cases, which we discuss below, given the width of the grid 2δ along two dimensions denoted by K and H, we can approximate the transition probability from node i to j by

$$\pi_{ij} = \Pr\left(\begin{bmatrix} \xi^l_{ij} \\ \eta^l_{ij} \end{bmatrix} < \begin{bmatrix} \Delta W^K \\ \Delta W^H \end{bmatrix} < \begin{bmatrix} \xi^h_{ij} \\ \eta^h_{ij} \end{bmatrix}\right), \tag{14}$$

which can be evaluated since we know the distribution of ΔW^K and ΔW^H, where

$$\begin{bmatrix} \xi^l_{ij} \\ \xi^h_{ij} \\ \eta^l_{ij} \\ \eta^h_{ij} \end{bmatrix} = \begin{bmatrix} \frac{(K^j - K^i - \delta) - (K^i\alpha_K - c^i - \psi^i)\phi}{K^i\sigma_K} \\ \frac{(K^j - K^i + \delta) - (K^i\alpha_K - c^i - \psi^i)\phi}{K^i\sigma_K} \\ \frac{(H^j - H^i - \delta) - (H^i\alpha_H + (1/P^i)\psi^i)\phi}{H^i\sigma_H} \\ \frac{(H^j - H^i + \delta) - (H^i\alpha_H + (1/P^i)\psi^i)\phi}{H^i\sigma_H} \end{bmatrix}. \tag{15}$$

In our analysis, we restrict the lower boundary of both K and H to be strictly greater than zero. This is done entirely for convenience: we would need to add special cases in calculating the transition probability matrix if either, or both, of them were zero. This restriction, however, does not result in any loss of generality; results when the lower bounds are set to zero share all the properties described in the text, and are available on request.

The special cases where the above expression for the transition probability is inadequate arise when the node j lies on either the lower or upper bound of either coordinate of the grid. If we use the above expression at all nodes, we shall be restricting the range of the normal distributions ΔW^K and ΔW^H for each i to lie between $[\xi^l_{i,LB} - \delta, \xi^h_{i,UB} + \delta]$ and $[\eta^l_{i,LB} - \delta, \eta^h_{i,UB} + \delta]$, respectively, where the subscripts LB and UB indicate, respectively, the lower and upper bounds of the corresponding coordinate. Since the grid is finite whereas the normal distribution can take values between $-\infty$ and ∞, we must make an adjustment to ensure that $\sum_{j=1}^{N} \pi_{ij} = 1$.

There are two possible ways to do this. In the first method, we keep the range of the distributions of ΔW^K and ΔW^H restricted as mentioned above, and calculate the transition probability from node i to j as

$$\pi^*_{ij} = \frac{\pi_{ij}}{\sum_{j=1}^{N} \pi_{ij}}.$$

This, however, makes the distribution non-normal. To avoid this distortion of the probability measure, we 'force' the range of the distribution to lie between $-\infty$ and ∞ using the following adjustments: if the L-coordinate of the j node is the lower (upper) bound, we set $\xi^l_{ij} = -\infty$ ($\xi^h_{ij} = \infty$); if the H-coordinate of the j node is the lower (upper) bound, we set $\eta^l_{ij} = -\infty$ ($\eta^h_{ij} = \infty$).

A.3 Policy iteration algorithm

Step 1 Let the state space have N nodes, where nodes are formed by the intersection of the discretized state space. Guess an initial control $[c^i_0, \psi^i_0]$ and an initial relative price $[P^i_0]$ for each node $i \in \{1, 2, ..., N\}$.

Step 2 Find the intermediate value function $J(K, H; c_{k-1}, \psi_{k-1})$ using the formula

$$\begin{bmatrix} J(K^1, H^1; c^1_{k-1}, \psi^1_{k-1}) \\ J(K^2, H^2; c^2_{k-1}, \psi^2_{k-1}) \\ \vdots \\ J(K^N, H^N; c^N_{k-1}, \psi^N_{k-1}) \end{bmatrix} = [I_{N\times N} - \alpha \mathbf{\Pi}]^{-1} \begin{bmatrix} U(c^1_{k-1}, \psi^1_{k-1}) \\ U(c^2_{k-1}, \psi^2_{k-1}) \\ \vdots \\ U(c^N_{k-1}, \psi^N_{k-1}) \end{bmatrix}.$$

The procedure to calculate the transition matrix is described in the next subsection; for now, we just note that it depends on the initial and final nodes, the control applied and the relative price at the initial node.

Step 3 Using the value function obtained from Step 2, find the control $[c^i_k, \psi^i_k]$ for each node $i \in \{1, 2, ..., N\}$ – imposing the constraints $c^i_k \geq 0$ and $0 \leq \psi^i_k \leq \Psi^{\max}$ in the numerical procedure – as

$$\underset{c_k^i \geq 0, 0 \leq \psi_k^i \leq \Psi^{\max}}{\arg\max} [U(K^i, H^i; c_k^i, \psi_k^i) + \alpha \sum_{j=1}^{N} \pi_{ij}(c_k^i, \psi_k^i) \times J(K^j, H^j; c_{k-1}^j, \psi_{k-1}^j)],$$

where $\pi_{ij}(c_k^i, \psi_k^i)$ is the transition probability of reaching node j beginning at node i when the controls applied are (c_k^i, ψ_k^i).

Step 4 Using the controls obtained in Step 3, find the new intermediate value function $J_k = J(K, H; c_k, \psi_k)$ using

$$\begin{bmatrix} J(K^1, H^1; c_k^1, \psi_k^1) \\ J(K^2, H^2; c_k^2, \psi_k^2) \\ \vdots \\ J(K^N, H^N; c_k^N, \psi_k^N) \end{bmatrix} = [I_{N \times N} - \alpha \mathbf{\Pi}]^{-1} \begin{bmatrix} U(c_k^1, \psi_k^1) \\ U(c_k^2, \psi_k^2) \\ \vdots \\ U(c_k^N, \psi_k^N) \end{bmatrix},$$

where the new transition matrix is calculated under the controls (c_k^i, ψ_k^i) obtained in Step 3 and using the new approximation of P^i at each node. From equations (5) and (6), we have that $P = \frac{J_H}{J_K} = \frac{J_H}{U_c(c,H)}$; we approximate J_H at node i by $\frac{J(i')-J(i)}{(i'-i)}$ where node i' is one whose ordinate is the same as that of node i but whose abscissa is that of node i plus the step size of the grid. Finding the denominator is straightforward; substituting this we are able to find the approximate relative price:

$$P^i = \frac{\partial J_{k-1}^i / \partial H^i}{\left((c^i)^{\frac{\epsilon-1}{\epsilon}} + \omega (H^i)^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{1-\epsilon\gamma}{\epsilon-1}} (c^i)^{-\frac{1}{\epsilon}}}.$$

This iteration is repeated until the maximum difference between the controls obtained in consecutive steps is less than a tolerance limit.[29]

A.4 Model-fitted price paths with their confidence intervals

We begin by equating the model price P with the observed price in 1987 (‘t’) for the corresponding MSA. The house price in 1988 for MSA i is given by

$$P_{i,t+1} = P_{i,t} e^{u_i + \beta' X_{i,t+1}}, \tag{16}$$

[29] The convergence of the algorithm under this approximation is to be formally proved. However, it appears to be logical since with each consecutive iteration the value function and control policy are closer to the optimal, which should imply that P^i, which is derived using the improved value function and control policy, should converge to its optimal value, too. Convergence of the relative prices should, in turn, result in the convergence of the algorithm.

where u_i is the fixed effect for MSA i, and $X_{i,t}$ are the independent variables (including a constant) in MSA i at $t+1$. These parameters are estimated; as we show in what follows immediately, their estimation error is accounted for in our analyses. We assume that the fitted value of the returns at $t+1$ is distributed as $\mathcal{N}(\mu_{i,t+1}, \sigma_{i,t+1})$, where $\mu_{i,t+1} = \hat{u}_i + \hat{\beta}' X_{i,t+1}$ with the 'hat' notation used to represent point estimates, and $\sigma_{i,t+1} = \sqrt{Var(\beta' X_{i,t+1})}$. Now,

$$E\left[P_{t+1}|P_t\right] = P_t e^{\mu_{i,t+1} + \frac{\sigma^2_{i,t+1}}{2}}. \tag{17}$$

The upper bound of this point estimate is obtained by $P^{UB}_{t+1} = P_t Z_{1-\alpha/2}$ and the lower bound by $P^{LB}_{t+1} = P_t Z_{\alpha/2}$, where $Z_\alpha = P(z : F(Y \leq z) = \alpha)$, where Y is distributed as a lognormal with mean $\mu_{i,t+1} + \sigma^2_{i,t+1}/2$ and variance $e^{\sigma^2_{i,t+1} - 1}(2\mu_{i,t+1} + \sigma^2_{i,t+1})$. All subsequent price paths and confidence intervals are calculated conditional on the values just prior to that being estimated.

References

D. Albouy and G. Ehrlich. Metropolitan Land Values and Housing Productivity. 2012.

D.P. Bertsekas. Dynamic Programming and Optimal Control, vol. 2a. *Athena Scientific*, 1995.

M.K. Brunnermeier. The New Palgrave Dictionary of Economics. *Chap. Bubbles*, 2008.

M.K. Brunnermeier and C. Julliard. Money Illusion and Housing Frenzies. *Review of Financial Studies*, 21(1):135–180, 2008.

C. Burnside, M. Eichenbaum, and S. Rebelo. Understanding Booms and Busts in Housing Markets. Technical report, National Bureau of Economic Research, 2011.

Y. Chu. Credit Constraints, Inelastic Supply, and the Housing Boom. *Working Paper*, 2012.

M.A. Davis and J. Heathcote. Housing and the Business Cycle. *International Economic Review*, 46(3):751–784, 2005.

D. DiPasquale. Why don't we know more about housing supply? *Journal of Real Estate Finance and Economics*, 18(1):9–23, 1999.

K.B. Dunn and K.J. Singleton. Modeling the Term Structure of Interest Rates Under Non-Separable Utility and Durability of Goods. *Journal of Financial Economics*, 17(1):27–55, 1986.

J. Favilukis, S.C. Ludvigson, and S. Van Nieuwerburgh. The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium. 2012.

H. Frech III and R.N. Lafferty. The Effect of the California Coastal Commission on Housing Prices. *Journal of Urban Economics*, 16(1):105–123, 1984.

C. Garriga, R.E. Manuelli, and A. Peralta-Alva. A Model of Price Swings in the Housing Market. *Federal Reserve Bank of St. Louis Working Paper Series*, 2012.

E.L. Glaeser, J. Gyourko, and R.E. Saks. Urban Growth and Housing Supply. *Journal of Economic Geography*, 6(1):71–89, 2006.

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

E.L. Glaeser, J. Gyourko, E. Morales, and C.G. Nathanson. Housing Dynamics. 2011.

J.F. Gomes, L. Kogan, and M. Yogo. Durability of Output and Expected Stock Returns. *Journal of Political Economy*, 117(5):941–986, 2009.

S.R. Grenadier. The Strategic Exercise of Options: Development Cascades and Overbuilding in Real Estate Markets. *Journal of Finance*, 51(5):1653–1679, 1996.

J. Gyourko, A. Saiz, and A. Summers. A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index. *Urban Studies*, 45(3):693, 2008.

R.E. Hall. Intertemporal Substitution in Consumption. *Journal of Political Economy*, 96(2): 339–357, 1988.

A. Haughwout, R.W. Peach, J. Sporn, and J. Tracy. The Supply Side of the Housing Boom and Bust of the 2000s. In *Housing and the Financial Crisis*. University of Chicago Press, 2012.

J.A. Kahn. What Drives Housing Prices? *Federal Reserve Bank of New York Staff Reports*, 2008.

F.E. Kydland and E.C. Prescott. Time to Build and Aggregate Fluctuations. *Econometrica*, 50(6):1345–1370, 1988.

R.E. Lucas Jr and E.C. Prescott. Investment Under Uncertainty. *Econometrica*, 39(5):659–681, 1971.

C.J. Mayer and C.T. Somerville. Land Use Regulation and New Construction. *Regional Science and Urban Economics*, 30(6):639–662, 2000.

A. Murphy. A Dynamic Model of Housing Supply. 2012.

M. Ogaki and C.M. Reinhart. Measuring Intertemporal Substitution: The Role of Durable Goods. *Journal of Political Economy*, 106(5):1078–1098, 1998.

A.D. Paciorek. Supply constraints and housing market dynamics. *Finance and Economics Discussion Series*, 2012.

L. Pástor and P. Veronesi. Was There a Nasdaq Bubble In the Late 1990s? *Journal of Financial Economics*, 81(1):61–100, 2006.

B.M. Peterson. Fooled by Search: Housing Prices, Turnover and Bubbles. *Bank of Canada Working Paper 2012-3*, 2012.

M. Piazzesi, M. Schneider, and S. Tuzel. Housing, Consumption and Asset Pricing. *Journal of Financial Economics*, 83(3):531–569, 2007.

E.C. Prescott and R. Mehra. Recursive Competitive Equilibrium: The Case of Homogeneous Households. *Econometrica*, 48(6):1365–1379, 1980.

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

R.E. Saks. Job Creation and Housing Construction: Constraints on Metropolitan Area Employment Growth. *Journal of Urban Economics*, 64(1):178–195, 2008.

T.M. Sinai. House Price Moments in Boom-Bust Cycles. 2012.

W.C. Wheaton and W.E. Simonton. The Secular and Cyclic Behavior of True Construction Costs. *Journal of Real Estate Research*, 29(1):1–26, 2007.

M. Yogo. A Consumption-Based Explanation of Expected Stock Returns. *Journal of Finance*, 61(2):539–580, 2006.