Supply Elasticity and the Housing Cycle of the 2000s

Thomas Davidoff
Sauder School of Business, University of British Columbia

November 21, 2011

Abstract

There is no evidence that differences in supply elasticity caused cross sectional variation among US housing markets in the severity of the 2000s housing cycle. This is true in three sets of empirical specifications: a first that assumes identical demand changes in the 2000s across markets, a second that proxies for supply elasticity and demand changes in the 2000s with estimates based on price and quantity changes in the 1980s, and a third that uses physical and regulatory constraints to proxy for supply elasticity and uses state fixed effects to capture variation in demand conditions.

JEL Codes: R21, R31 (housing demand, housing supply and markets)

1 Introduction

This paper asks whether any of the differences across US metropolitan areas in the severity of the housing price cycle of the 2000s can be attributed to differences in supply elasticity. The large average amplitude of the cycle suggests that data from the 2000s will play an important role in shaping future expectations about US housing price movements. Investors' beliefs about the mapping from underlying market fundamentals to price outcomes in the last decade may thus have significant consequences.

There are good reasons to suspect that investors in less elastically supplied markets will face greater than average price volatility. If housing supply is reversible and adjusts immediately to demand shocks, then even in the face of speculative demand pressure, price is always equal to marginal cost, and marginal cost varies less with quantity in more elastically supplied markets. This basic argument, sketched informally by Krugman (2005), can be found in the conference responses to Shiller (2003), in Himmelberg et al. (2005) and is central to the identification of the effects of mortgage credit expansion and contraction in the 2000s in Mian and Sufi (2009). Consistent with this argument, Glaeser et al. (2008) show that in the 1980s, housing cycles were far more pronounced in inelastically supplied markets.

There are also good reasons to question whether supply inelasticity caused cycle severity in the 2000s. Theoretically, Glaeser et al. (2008) observe that because supply is less elastic in the short run than the long run, prices may fluctuate widely even in elastically supplied markets; they also find that both theoretically and empirically, speculative bubbles are shorter in more elastically supplied markets. For investors and regulators, the amplitude (ratio of peak price to prices at the endpoints of a cycle) of housing cycles or volatility of prices may be as interesting as cycle frequency. Glaeser et al. (2008) conclude that the correlation between supply elasticity and cycle amplitude is ambiguously signed in their model of price dynamics.[1]

Empirically, the housing price cycle of the 2000s was characterized by relatively large price increases followed by relatively steep declines in two overlapping sets of markets that appear to have very different supply elasticities. The first set are "Coastal" markets, particularly the notoriously inelastically supplied markets on the California and Northeast Atlantic coasts; the second set, with even more severe price cycles, are markets in the "Sand States" of Arizona, California, Florida, and Nevada (see e.g. Olesiuk and Kalser (2009)). Sand state markets that are not on the California coast have seen considerable growth in supply over time, and continued to through the 2000s. In the 1980s, while markets on the California and

Northeast Atlantic coasts saw high price growth and low quantity growth, the opposite was true in Sand State markets not on the California coast.

Figures 1 and 2 display these regularities with box-and-whisker plots of the distribution of 2000s housing price cycle magnitude and quantity growth across among US metropolitan areas that are Sand, Coastal, neither, or both. Notably, the Sand State markets almost first order stochastic dominate all other markets in terms of housing cycle amplitude, so explaining their performance relative to other markets is critical to explaining cross sectional differences.

The extremely severe price cycles and high rates of supply growth in Sand States have two implications for empirical analysis. First, if relative supply elasticity played an important role in the distribution of price cycle magnitude across metropolitan areas, then there must have been important differences in demand volatility across markets. If all markets saw the same movements in demand and feature identical demand elasticities, then the non-Coastal Sand State markets' large supply growth must indicate that they are at the high end of supply elasticity. This would render their extremely large price cycles inconsistent with the hypothesized model.[2] Second, the uniformly large price cycles in the Sand States, despite considerable differences on almost all dimensions among markets within these states (e.g. coastal San Francisco is a very different place from non-coastal Bakersfield) point to the existence of state-level demand shocks that would be difficult to capture with observable characteristics alone.

Informed by the facts illustrated in Figures 1 and 2, I take three approaches to estimating the relationship between supply elasticity and the intensity of the 2000s housing price cycle conditional on some measure of demand. Within each approach, I employ three distinct measures of cycle severity, defined below. The first identification strategy, following a model sketched by Krugman (2005), assumes that all metropolitan areas experienced identical demand shocks and exhibited relatively fixed supply curves. These assumptions jointly imply that supply elasticity is monotonically increasing in quantity growth in the 2000s. While common growth pressure in the 2000s is not a reasonable assumption, the simplest estimation approach warrants brief consideration, and the expansion and contraction of mortgage credit over the 2000s was a widespread, if unequally distributed phenomenon. If we find no negative correlation between supply growth and cycle amplitude, then the data are consistent with a causal role for supply inelasticity only if inelastically supplied markets saw abnormally low volatility of demand growth over the 2000s.

The second identification strategy recognizes the persistence of impediments to housing supply growth and the likely importance of backward-looking expectations in the expansion of equity and debt investment that fueled the 2000s price boom. In this approach, I use price and quantity data from the 1980s to infer local supply elasticities and demand pressure that prevailed in the 2000s. While historical demand growth must be an imperfect measure of the intensity of demand fluctuations over the 2000s, the persistently high degree of regulation, amenity, and demand for high-skill labor in most Coastal markets and the persistent quantity growth in most Sand State markets over the last four decades suggest that the approximation may not be too far off.

Figure 3 provides visual evidence that demand growth in the 1980s is, in fact, a reasonable measure of susceptibility to a demand-driven price cycle in the 2000s. There, each metropolitan area is represented in a graph of 1980s price and quantity growth by a "bubble" that is proportional to a measures of 2000s cycle intensity. We find that demand growth in the 1980s, whether expressed through quantity or price growth, is highly correlated with the magnitude of the 2000s price cycle. For the most part, only cities that experienced both weak price and quantity growth in the 1980s escaped without a violent price cycle in the 2000s.

The third estimation approach follows a growing literature and proxies for supply elasticity with measures of physical and regulatory constraints to development provided by Rappaport and Sachs (2003), Saiz (2008), and Gyourko et al. (2006). In this setting, some kind of control for demand growth is important. As Saiz (2008) and Kolko (2008) emphasize, the steep slopes and bodies of water that generate physical constraints to supply growth are typically associated with different (higher) levels of demand and potentially differential demand volatility over the 2000s. Regulatory barriers to development are typically imposed where there is demand pressure. A simple regression of 2000s cycle amplitude on regulatory intensity would thus suffer from omitted variable bias, with the estimated coefficient on supply constraints most likely inflated upward, pursuant to Figure 3.

Rather than try to capture all sources of differential demand pressure in regressions of cycle severity on measures of supply constraints, I account only for the evidently important state-level phenomena illustrated in Figure 1. One channel through which strong state fixed effects on housing cycle amplitude likely operated is differences in the expansion and contraction of credit. Credit movements vary considerably across states, but in ways that may or may not be driven by subtle differences in regulation and policy (see, e.g. Ghent and Kudlyak (2010) and Pence (2006)). That credit conditions were linked to price volatility in the cross section and time series is demonstrated statistically by, e.g. Mian and Sufi (2009) and Levitin and Wachter (2010). Notably, the Sand States were among the only places where the eventually bankrupt Washington Mutual Bank had large market share as of 2001.[3] A direct control for metropolitan-level volatility in credit conditions would be unappealing, because changes in credit availability may reflect past or anticipated future price movements.

While the proxies for demand that I use must be imperfect measures of demand fluctuations over the 2000s, the fact that three different approaches to identification arrive at the same conclusion should provide confidence in the main result. Conditional on demand, there appears to be no negative correlation between supply elasticity and 2000s housing price cycle severity. The consistency of results across specifications leaves little room for plausible explanations of the data in which supply inelasticity was an important driver of cross-sectional variation in cycle severity. In the case of the second estimation approach, it is difficult to argue that markets such as New York and San Francisco that saw extremely high price growth and low quantity growth in the 1980s became either elastically supplied or lost relative demand growth in the 2000s (see e.g. Van Nieuwerburgh and Weill (2006) or Gyourko et al. (2004)). As for the third approach, even if, *contra* Olesiuk and Kalser (2009), fluctuations in credit markets did not drive cross sectional differences in cycle amplitude across states, the absence of a positive relationship between supply constraints and cycle severity within states is hard to square with a causal role for supply inelasticity. If anything, residual correlation between supply constraints and demand factors would be expected to generate a positive correlation with the cycle severity measures.

2 The Empirical Relationship Between Housing Cycle Severity and Supply Elasticity

2.1 Regression Framework

Suppose that both supply and demand for owner housing in market m have constant price elasticities, that the demand elasticity γ is the same for all markets at all times, and that supply curves vary across markets but not time. Ignoring differences between the short and long runs, let:

$$\ln q_{mt}^{d} = \ln \alpha_{mt} - \gamma \ln p_{mt} \tag{1}$$

$$\ln q_{mt}^{s} = \ln c_{m} + \eta_{m} \ln p_{mt}. \tag{2}$$

In demand equation (1), α_{mt} is a market- and date-specific measure of demand. Among other factors embedded in α are interest rates, appreciation expectations, mortgage underwriting standards, population, employment, and public goods. In supply equation (2), η_m is a market-specific supply elasticity.

Combining (1) and (2) to solve for price at date t in market m, in equilibria where supply equals demand:

$$\ln p_{mt} = \frac{\ln \alpha_{mt} - \ln c_m}{\eta_m + \gamma} \tag{3}$$

In log differences:

$$\ln p_{mt'} - \ln p_{mt} = \frac{\ln \alpha_{mt'} - \ln \alpha_{mt}}{\eta_m + \gamma} \tag{4}$$

Under equation (4), price cycles should have been more severe in markets where demand α grew more in the boom between 2000 and 2007, and fell more during the bust starting in mid-2007. Holding changes in demand constant across markets, there should also have been greater volatility in markets with lower supply elasticities η.

There are several reasons why equation (4) may be a poor approximation of short- or medium-term price dynamics. Importantly, short- and long-run price elasticities are not identical.[4] We want to know if gaps between the simple model and reality could have undermined the model's implication that supply elasticity is negatively correlated with price cycle amplitude. Recognizing that demand volatility and supply elasticity may be correlated at the metropolitan level, I estimate regressions of the form:

$$b_{im} = \beta_0 + \beta_1 \text{supply elasticity}_{mj} + \beta_2 \text{demand volatility}_{mj} + \epsilon_{mij}. \tag{5}$$

In (5), b is a measure of cycle severity, with different measures indexed by i. j indexes the three approaches to supply elasticity and demand volatility measurement, and ϵ is a regression residual.

2.2 Measuring Cycle Severity

I measure the severity of the recent price cycle in three ways. The first is the difference in annualized real price growth rates between the boom and bust periods:

$$b_{1m} \equiv \frac{\ln \frac{p_{m2007q2}}{p_{m2000q1}}}{7.25} - \frac{\ln \frac{p_{m2010q3}}{p_{m2007q2}}}{3.25} \tag{6}$$

In equation (6), m denotes metropolitan areas, and p_{myqx} is the FHFA (Federal Housing Finance Agency) repeated sale price in metropolitan area m in year y, quarter x. b_1 is a measure of peak-to-endpoints amplitude ($\frac{\frac{p_{2007q2}}{p_{2000q1}}}{\frac{p_{2010q3}}{p_{2007q2}}} = \frac{p^2_{2007q2}}{p_{2010q3}p_{2001q1}}$), scaled to allow for different long-run price growth rates across markets.

I choose 2007, second quarter as the peak of the housing cycle based on the mechanism outlined by Bai and Perron (1998): I minimize over the breakpoint date X the sum across metropolitan areas m of summed squared deviations of overlapping annual (4 quarter) growth rates of real metropolitan area FHFA repeated sale price indexes between the period 1999q1 to X from the mean by metropolitan area in that period, plus the sum of squared deviations over the period X through the end of available price data in 2010q3. That is, $X = 2007q2$ minimizes the objective V:

$$V = \sum_m \left[\sum_{t=2000q1}^{X} \left[\ln \frac{p_{mt}}{p_{mt-4}} - \frac{\sum_{t=2000q1}^{X} \ln \frac{p_{mt}}{p_{mt-4}}}{X - 2000q1} \right]^2 + \sum_{t=X}^{2010q1} \left[\ln \frac{p_{mt}}{p_{mt-4}} - \frac{\sum_{t=X}^{2010q3} \ln \frac{p_{mt}}{p_{mt-4}}}{2010q3 - X} \right]^2 \right]. \quad (7)$$

Regression results are fundamentally unchanged if each market is allowed a different peak date over the 2000s because the peak real price dates are densely centered around 2007 quarter 2. The correlation between b_1 and an alternative measure that allows different peaks across markets is .99. If we assume a common shock to demand α in (4), then it is appropriate to define the cycle measure over fixed dates across metropolitan areas.

The second measure of cycle amplitude, b_2, is the standard deviation of overlapping four-quarter real price growth between 1999 and 2010:

$$b_{2m} \equiv \text{standard deviation} \left(\frac{p_{mt}}{p_{mt-4}} \right), \; t \in [2000q1, 2010q3]. \quad (8)$$

Similar results hold for b_2 measured as a coefficient of variation, but the unscaled standard deviation is presumably of greater interest to investors and regulators.

The third measure of cycle severity, b_3, is the ratio of real price in 2007q2 to 2010q3. This last measure simply captures the extent of the price bust, a natural measure of excess pricing at the peak assuming 2010q3 prices reflect a somehow correct price level. As with b_1, allowing the date of peak price to vary across markets does not meaningfully affect results. For all three measures, I use FHFA quarterly repeated price sale home price data at the CBSA (or metropolitan division) level deflated by the US Consumer Price Index for all goods. The three measures, summarized along with other variables in Table 1, are highly correlated with each other, as shown in Table 2, with pairwise correlations of .97, .91, and .96.

The FHFA index has the important feature of being the longest and widest panel of home prices available. An important weakness of the FHFA index is that it excludes data from transactions that have no mortgage; or high value, high loan to value, subprime, or other "non-conforming" mortgages. In principle, this could be salient to the analysis in this paper follows because we are interested in differences in volatility between Coastal versus non-Coastal markets, and the conforming share is generally lower in Coastal markets. Homes purchased with conforming mortgages may exhibit smaller price movements than those with other mortgages, potentially confounding regression results. Comparison with the S&P Case Shiller (CS) index, which is based on a fuller sample, suggests that this may affect the magnitude of results presented here, but likely not their general direction. For example, both the CS and FHFA indexes agree that between 2007 and 2011, the price decline was larger in non-Coastal Las Vegas and Phoenix then in Coastal Los Angeles and San Francisco. However, the magnitude of all declines is larger in the CS data, and more so for Los Angeles and particularly San Francisco.

2.3 Common demand changes

The regression analysis starts under the strong assumptions that demand shocks α in the 2000s were identical across markets. In that case, combining (1) and (2) and denoting 2000 to peak differences with Δ yields:

$$\Delta \ln q_m = \frac{\eta_m \Delta \ln \alpha}{\eta_m + \gamma} \tag{9}$$

Differentiating equation (9) yields a positive correlation between log quantity growth over the 2000s and supply elasticity η:

$$\frac{d\Delta q}{d\eta} = \frac{\gamma}{[\eta + \gamma]^2} \Delta\alpha > 0. \tag{10}$$

The first specification of regression (5) thus drops demand from the right hand side under the assumption of equality across metropolitan areas and proxies for supply elasticity with the log change in quantity from 2000 to 2009. I use 2009, rather than 2007, as the end of the period of supply growth recognizing (in contrast to (2)) that supply may adjust slowly to price does not jump downward in the face of a price decline.[5] I estimate quantity changes from US Census counts of housing units; these are in turn largely determined by residential permitting. Mean quantity growth in the 2000s was 14%, as shown in Table 1.

Assuming c_m is fixed over time, equation (2) implies that we can compute local supply elasticities as the ratio of log quantity growth to log price growth. If the price growth is taken to be between 2000q1 and 2007q2, this ratio has a mean value of .65 and a standard deviation of .71. The only negative value is for New Orleans, presumably attributable to Katrina. The next smallest elasticity is Nassau-Suffolk, at .04. The largest value is for Fort Wayne, IN, at 3.15. If log price growth is taken between 2000q1 and 2009q4, much more extreme and less sensible positive and negative values are computed.

Under the assumptions of a common demand shock and heterogeneous supply elasticity, we should see a significantly negative correlation between supply growth and cycle severity. In fact, the cross sectional correlations between supply growth and the three cycle measures b_1, b_2, and b_3 are all positive, respectively: .18, .18, and .26.[6] For all three measures, assuming independence of residuals across metropolitan areas, we can reject a zero or negative relationship between supply elasticity and cycle magnitude. The data thus provide no support for the joint hypotheses of a shared national demand shock combined and a causal role for supply inelasticity in generating cross-metropolitan variation in cycle amplitude.

2.4 Proxies for Supply Elasticity and Demand Variability Derived From 1980s Price and Quantity Growth

That changes in log owner housing demand were identical across metropolitan areas in the 2000s, as assumed in the first regression specification, is not literally true. Presumably the positive correlation between 2000s supply growth and cycle amplitude is biased upward by the omission of a suitable control for demand growth during the boom. Maintaining the assumption that supply curves have not changed over time, it is possible to use historical price and quantity data to infer historical supply elasticities and demand pressure.

In an early assessment that housing prices were too high, Shiller (2003) argues that the markets experiencing high price growth in the 2000s were "glamor" markets that had experienced high historical demand growth. The lending and speculative frenzy in the Sand States and elsewhere may have been results of historical demand growth (see, e.g. Olesiuk and Kalser (2009)). 1980s demand growth might therefore be a reasonable proxy for fluctuations in demand in the 2000s.[7] Figure 3, described above and below, supports this approach.

For demand growth in the 1980s to be a satisfactory proxy for demand volatility in the 2000s, it must be the case that metropolitan area supply curves do not change over time, and that quantity and price growth in the 1980s did not directly affect supply changes in the 2000s. This lack of lingering direct effect assumes long run supply had fully adjusted to 1980s price growth by the start of the 2000s. The steep decline in US building permits between 1989 and the early 1990s suggests that this is not too strong of an assumption.

Suppose that equation (2) is specialized to:

$$\ln q^s_{m1990} - \ln q^s_{m1980} = \eta_m \left[\ln p_{m\overline{1980s}} - \ln p_{m1980}\right] \tag{11}$$

$$\eta_m^{-1} = \frac{\ln p_{m\overline{1980s}} - \ln p_{m1980}}{\ln q^s_{m1990} - \ln q^s_{m1980}}. \tag{12}$$

$p_{m\overline{1980s}}$ is the maximal price attained in metropolitan area m in the 1980s. Equation (11) makes the moderately strong assumption that quantity growth between 1980 and 1990 was a long run supply response to maximal price in the 1980s.

Unlike the housing boom of the 2000s, in which the date at which prices hit their metropolitan area maximum was tightly distributed around 2007, the maximal value in the 1980s was widely dispersed, with a large number of metropolitan areas seeing a real maximum at the start of the 1980s, due to "stagflation" and the collapse of an oil boom. The standard deviation across metropolitan areas of breakpoint dates that minimize the within-metropolitan squared residuals in (7) is 1.6 for the 2000s and 3.2 for the 1980s. That metropolitan areas witnessing large quantity growth and negative price growth over the 1980s did so due to highly negative supply elasticities is implausible. It thus makes sense to allow variation across metropolitan areas in the date of the price maximum to which 1980s quantity growth was a response. The use of 1980 as a start date has the cost of sharply reducing the number of metropolitan areas for which FHFA price data is available, but the main results are not sensitive to the start date.

Allowing for heterogeneous demand growth α_m in the 1980s, but retaining a common demand elasticity γ, we obtain by transforming equation (1):

$$\ln \alpha_m = \ln q^d_{m1990} - \ln q^d_{m1980} + \gamma \left[p_{m\overline{1980s}} - p_{m1980}\right]. \tag{13}$$

Unfortunately, γ is unknown, even if a common national value exists. Based on estimates in the literature of rental housing demand (e.g. Hanushek and Quigley (1980), Davis and Ortalo-Magne (2011)), $\gamma \in (0, 1)$ is a sensible range.

Figure 3 provides visual evidence that strong historical demand, rather than inelastic supply, drove variation in housing cycle magnitude across metropolitan areas. Using Census housing unit counts and deflated prices as above, Figure 3 plots price growth $\frac{p_{m\overline{1980s}}}{p_{m1980}}$ against quantity growth $\frac{q_{m1990}}{q_{m1980}}$. "Bubbles" are proportional to 2000s cycle magnitude measure b_1. Two facts stand out: first, "low supply elasticity" markets that saw high price growth and low quantity growth in the 1980s almost uniformly saw larger than average housing cycles in the 2000s. Second, "high elasticity" markets that saw high quantity growth and low price growth also saw large price cycles, although with greater variation in outcomes than for the low elasticity markets. Weak cycles in the 2000s were concentrated in markets that saw low rates of quantity and price growth in the 1980s. There is no visual evidence that supply elasticity in the 1980s, which should have pushed demand pressure away from price growth to quantity growth, was associated with less severe cycles in the 2000s.

Quantitatively, metropolitan areas in the lower right quadrant of Figure 3, with greater than median quantity growth and less than median price growth in the 1980s saw median and mean values of 2000s cycle magnitude b_1 of .11 and .12. Metropolitan areas in the upper left quadrant, with greater than median price growth and less than median quantity growth in the 1980s saw median and mean values of .11 and .11. For measures b_2 and b_3, the evidently elastically supplied markets experiencing high demand in the 1980s in the lower right of Figure 3 saw both higher means and medians than the "inelastic" markets in the upper left quadrant.

Markets in the lower left of Figure 3, which saw less than median price growth and less than median quantity growth in the 1980s had much lower values for all three measures; for b_1, the median and mean values in this low growth quadrant were .04 and .06. Statistically, we cannot reject that markets in the "inelastic supply" upper left quadrant and markets in the "elastic supply" lower right quadrants had values for b_1, b_2, and b_3 drawn from distributions with identical means. We can, however, reject that either the high price growth, low quantity growth or the high quantity and low price growth group had the same mean for the measures of cycle amplitude b_1 and b_2 as the low price growth, low quantity growth lower left quadrant. The hypothesis of equal means for the crash measure b_3 is rejected for the lower right versus lower left comparison, but not for the upper left versus lower left comparison.

That markets with historically high price growth and low quantity growth saw no more severe cycles in the 2000s than markets with historically low price and high quantity growth suggests that conditional on a moderately high level of demand growth, supply elasticity does not soften cycle severity. However, even taking equations (11) and (13) literally, a problem of interpretation arises: the relative importance of price and quantity growth in "demand growth" α is unknown as long as demand elasticity γ is unknown. Depending on γ, average demand growth in the northwest quadrant of Figure 3 may have been less than, equal to, or greater than demand growth in the southeast corner. While it is hard to believe that upper left markets such as San Francisco and Boston did not enjoy substantial demand growth in the 1980s or 2000s, this could have been the case if demand is sufficiently price inelastic that high price growth need not reflect strong demand pressure. Theoretically, it could then be the case that conditional on a proper measure of demand, elastic supply *was* associated with less severe cycles in the 2000s.

To address this possibility, Table 3 quantifies Figure 3 by estimating equation (5) with the inverse elasticity measure described in equation (12), and using values of γ of 0, .25, and .75 to calculate growth α_m from (13). At $\gamma = 0$, the demand control is a control for log quantity growth in the 1980s, so that conditional on demand, the inverse supply elasticity is essentially a measure of price growth. As γ grows in the conditioning demand variable, more weight is put on the ratio of price to quantity growth and less on price growth in the estimated inverse supply elasticity η_m^{-1}. A significant negative relationship between estimated supply elasticity and 2000s cycle severity would be indicated by a significantly *positive* coefficient on inverse elasticity.

We find in column (1) that for all three cycle measures, there is no significant unconditional relationship with the "inverse elasticity" ratio of price to quantity growth in the 1980s.

In column (2), we find that conditional on quantity growth only ($\gamma = 0$), the ratio of price to quantity growth has a small and insignificantly positive association with b_1, b_2 and b_3. With positive γ, such that both price and quantity growth indicate demand growth, there is never a significantly positive relationship between inverse supply elasticity $\frac{1}{\eta}$ and 2000s cycle intensity. Not surprisingly as the demand proxy puts more weight on price growth relative to quantity growth in the 1980s, inverse supply elasticity has a more negative coefficient. In all specifications, estimated demand growth in the 1980s is significantly positively associated with more intense cycles in the 2000s. Both price and quantity growth in the 1980s are associated with severe cycles in the 2000s. There is thus evidence that markets subject to historically high demand pressure saw relatively severe price cycles in the 2000s, but no compelling evidence to support the notion that supply inelasticity caused severe price cycles.

2.5 Land Availability and State Fixed Effects as Supply and Demand Measures

The assumptions required for the second identification approach are jointly strong: a particular value for the demand parameter γ, demand elasticity that is both common and meaningful in a speculative growth environment, and historically constant demand and supply elasticities. We now turn to estimation that does not involve converting historical price and quantity statistics into supply and demand parameters.

Ideally a measure of supply elasticity on the right hand side of regression (5) would have no mechanical correlation with price movements. Measures of regulation or lack of buildable land are attractive in this way. As emphasized by Kolko (2008) and Saiz (2008), however, these measures are correlated with demand, and hence likely with demand-side sources of variation in cycle amplitude. Identifying a role for elasticity with a measure of supply constraints as a proxy, given these measures would fail the exogeneity requirement of an instrumental variable, thus requires a satisfactory control variable approach. As discussed above, state fixed effects appear to capture a lot of variation in a critical source of demand heterogeneity: fluctuations in mortgage credit conditions over the 2000s. Importantly, the fact that state averages are correlated with supply constraints does not compromise identification as long as the supply measure remains negatively correlated with elasticity conditional on the demeaning. Given the direction of results, a challenge to identification would have to come from a reason for supply constraints to be negatively correlated with demand variability within states; the source of such a correlation is not easy to imagine.

Three common measures of supply constraints are: the Rappaport and Sachs (2003) status as a Coastal market based on a threshold distance to an ocean, the Gulf Coast, or a Great Lake; the Saiz (2008) calculation of the fraction of land lost to steep slopes and water; and the Gyourko et al. (2006) measure of local regulatory intensity. I aggregate each of these measures to the census CBSA (or Metropolitan Division if applicable) level. In the case of the Coastal indicator, the market value is the maximum among component counties.

Two desirable features of a proxy for supply inelasticity are: a negative unconditional correlation with supply, and a correlation that is not significantly diminished when the measure and supply growth are both demeaned at the state level to purge important sources of variation in demand. The top panel of Table 4 presents correlations of the three measures "Coastal", "unavailable", and "regulation" with each other and with quantity growth in the 2000s across markets. The bottom panel presents these correlations when all variables have been demeaned at the state level. We find that the Saiz (2008) unavailability measure is most attractive because of a significant negative unconditional correlation with supply growth in the 2000s that becomes more negative conditional on state fixed effects.

Table 5 presents results of estimating regression (5) with each of the three proxies for inelasticity in a separate panel and state fixed effects present in some specifications. The results are remarkably consistent across specifications. Unconditionally, measures of supply constraints are all significantly positively correlated with all three cycle severity measures. In all cases, state fixed effects explain a very large fraction of variation in all three of the cycle severity measures b_1, b_2, and b_3 conditional on the supply proxy. The adjusted R-square always increases by at least .5 when the state fixed effects are included. Moreover, conditional on state fixed effects, there is always a significant decrease in the estimated coefficient on the supply proxy, and there is never a significant positive conditional relationship between supply constraints and cycle amplitude. In the case of the Saiz (2008) unavailable land measure, conditional on state averages there is a significantly negative association with cycle amplitude measure b_1 and crash measure b_3.[8]

3 Conclusions

More than two-thirds of the variation in housing cycle severity in the 2000s across US markets can be explained by state averages, with most of the variation across states explained by the four "Sand States" of Arizona, California, Florida, and Nevada. While there is significant variation in supply conditions within states, particularly California, this residual variation in supply conditions is not associated with increased residual cycle severity. This fact is consistent with two other findings: first, that supply growth in the 2000s is positively, not negatively, correlated with cycle severity; and second that the ratio of price to quantity growth in the 1980s is not significantly positively correlated with 2000s cycle severity either unconditionally, or conditional on demand growth measures that put positive weight on either 1980s quantity growth or 1980s price growth. While there is theoretical reason to suspect that constraints on supply growth may trigger price volatility, the housing cycle of the 2000s does not provide empirical support.

One natural direction for future research would be to build on the work of Pavlov and Wachter (forthcoming) and establish a tight causal link between state-level credit conditions and average price movements. Any such project faces the difficulty that credit conditions are presumably determined in part by expectations of future price movements.

A second direction would be to think seriously about standard errors in analyses like those above. California is important to the results in every sense. The correlation between measures of supply constraints and cycle amplitude is driven by the housing markets around San Francisco and Los Angeles. The significantly negative within-states relationship between cycle amplitude and the Saiz (2008) land unavailability measure is driven by the fact that more elastically supplied markets within California had even more severe cycles than the Coastal markets. Anecdotally, one heard arguments during the boom that markets such as Bakersfield and Fresno, while far removed, were on the way to becoming suburbs of Los Angeles or San Francisco. The correct theoretical error correlation structure among markets depends would rely on an economic model of market integration that has yet to be developed.

Future analysis of the role of supply elasticity in price movements could usefully distinguish differences between short-run and long-run elasticity. As emphasized above, the proxies we have for supply elasticity are imperfect, and they are certainly not sufficient to separately identify short- versus long-run elasticities.

Finally, it would be useful to explore the role of supply elasticity in other countries and other US housing cycles. A glance at the distribution of cycle severity across markets in Spain suggests that supply elasticity was no protection against large price swings in the 2000s. By contrast, Glaeser et al. (2008) emphasize that there were few large cycles in the 1980s in elastically supplied US markets.

Notes

[1]Any comparative statics would be even more difficult to obtain in a model that acknowledges the irreversibility and real optionality of supply, as in Capozza and Helsley (1990) and Novy-Marx (2007).

[2]While Gyourko et al. (2004) show that quantity growth slowed in the Sand States between the 1980s and 2000s, it is difficult to argue based on the box plots that limited short run or anticipated long run quantity growth caused their relatively large price increases and decreases in the mid- to late-2000s.

[3]Whether this fact reflects a causal role for Washington Mutual and its abnormally high loan to value mortgages, is the subject of research underway. Hudson (2010) and Katz (2009) document an important role in the evaporation of lending standards for Long Beach Mortgage, acquired by Washington Mutual in 1999.

[4]Dynamics in this setting are described by Wheaton (1999).

[5]The correlation between supply growth between (a) 2000 and 2007 and (b) 2000 and 2009 across US counties is over .99, so this choice has little consequence.

[6]When individual metropolitan area cycle peaks are allowed, the correlations with b_1 and b_3 are .25 and .13 respectively.

[7]In principle, the approach below could be used with data from the 2000s, but the price cycle measures are mechanically correlated with price growth between 2000 and 2007, hence the use of historical data.

[8]Minorities and foreign-born residents are over-represented in Coastal and Sand states. Not surprisingly, controls for foreign-born share, suggested by Pavlov and Wachter (forthcoming) and Saiz (2008), not reported below, do not alter the results of Table 5.

References

Jushan Bai and Pierre Perron. Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1):47–78, 1998.

Dennis Capozza and Robert Helsley. The stochastic city. *Journal of Urban Economics*, 28: 187–203, 1990.

Morris A. Davis and Francois Ortalo-Magne. Househould expenditures, wages, rents. *Review of Economic Dynamics*, 14(2):248–261, 2011.

Andra Ghent and Marianna Kudlyak. Recourse and residential mortgage default: Theory and evidence from us states. working paper, Federal Reserve Bank of Richmond, 2010.

Edward L. Glaeser, Joseph Gyourko, and Albert Saiz. Housing supply and housing bubbles. *Journal of Urban Economics*, 64(2):198 – 217, 2008.

Joseph Gyourko, Christopher J. Mayer, and Todd Sinai. Superstar cities. Working Paper, Wharton, 2004.

Joseph E. Gyourko, Albert Saiz, and Anita A. Summers. A New Measure of the Local Regulatory Environment for Housing Markets: Wharton Residential Land Use Regulatory Index. *SSRN eLibrary*, 2006.

Eric A. Hanushek and John M. Quigley. What is the price elasticity of housing demand? *The Review of Economics and Statistics*, 62:449–454, August 1980.

Charles Himmelberg, Christopher Mayer, and Todd Sinai. Assessing high house prices: Bubbles, fundamentals and misperceptions. *Journal of Economic Perspectives*, 19(4):67–92, Fall 2005.

Michael W. Hudson. *The Monster*. Times Books, New York, 2010.

Alyssa Katz. *Our Lot*. Bloomsbury USA, New York, 2009.

Jed Kolko. Did God make San Francisco expensive? Topography, regulation, and housing prices. Working paper, Public Policy Institute of California, 2008.

Paul Krugman. That hissing sound. *New York Times*, 2005. Op-Ed, August 8.

Adam J. Levitin and Susan M. Wachter. Explaining the Housing Bubble. *SSRN eLibrary*, 2010.

Atif Mian and Amir Sufi. The consequences of mortgage credit expansion: Evidence from the u.s. mortgage default crisis. *Quarterly Journal of Economics*, 124(4), 2009.

Robert Novy-Marx. An equilibrium model of investment under uncertainty. *Review of Financial Studies*, 20(5):1461–1502, 2007.

Shayna M. Olesiuk and Kathy R. Kalser. The sand states: Anatomy of a perfect housing-market storm. *FDIC Quarterly*, 3(1):30–32, 2009.

Andrey D. Pavlov and Susan M. Wachter. Subprime lending and real estate prices. *Real Estate Economics*, forthcoming.

Karen M. Pence. Foreclosing on opportunity: State laws and mortgage credit. *Review of Economics and Statistics*, 88(1):177–182, 2006.

Jordan Rappaport and Jeffrey D. Sachs. The united states as a coastal nation. *Journal of Economic Growth*, 8:5–46, 2003.

Albert Saiz. On local housing supply elasticity. Working paper, Wharton, 2008.

Robert J. Shiller. Is there a bubble in the housing market? *Brookings Papers on Economic Activity*, 2:299–362, 2003.

Stijn Van Nieuwerburgh and Pierre-Olivier Weill. Why has house price dispersion gone up? working paper, NYU and UCLA, 2006.

William Wheaton. Real estate cycles, some fundamentals. *Real Estate Economics*, 27(2), 1999.

Figure 1: Distribution of Housing Cycle Amplitude measure b_1 among four types of markets. Vertical axis: annualized metropolitan real FHFA repeated sale home price index growth 1999 to 2007 minus annualized growth 2007 to 2010. "Sand States" are Arizona, California, Florida, and Nevada. "Coastal" markets abut a Great Lake or ocean or gulf, as defined by Rappaport and Sachs (2003).

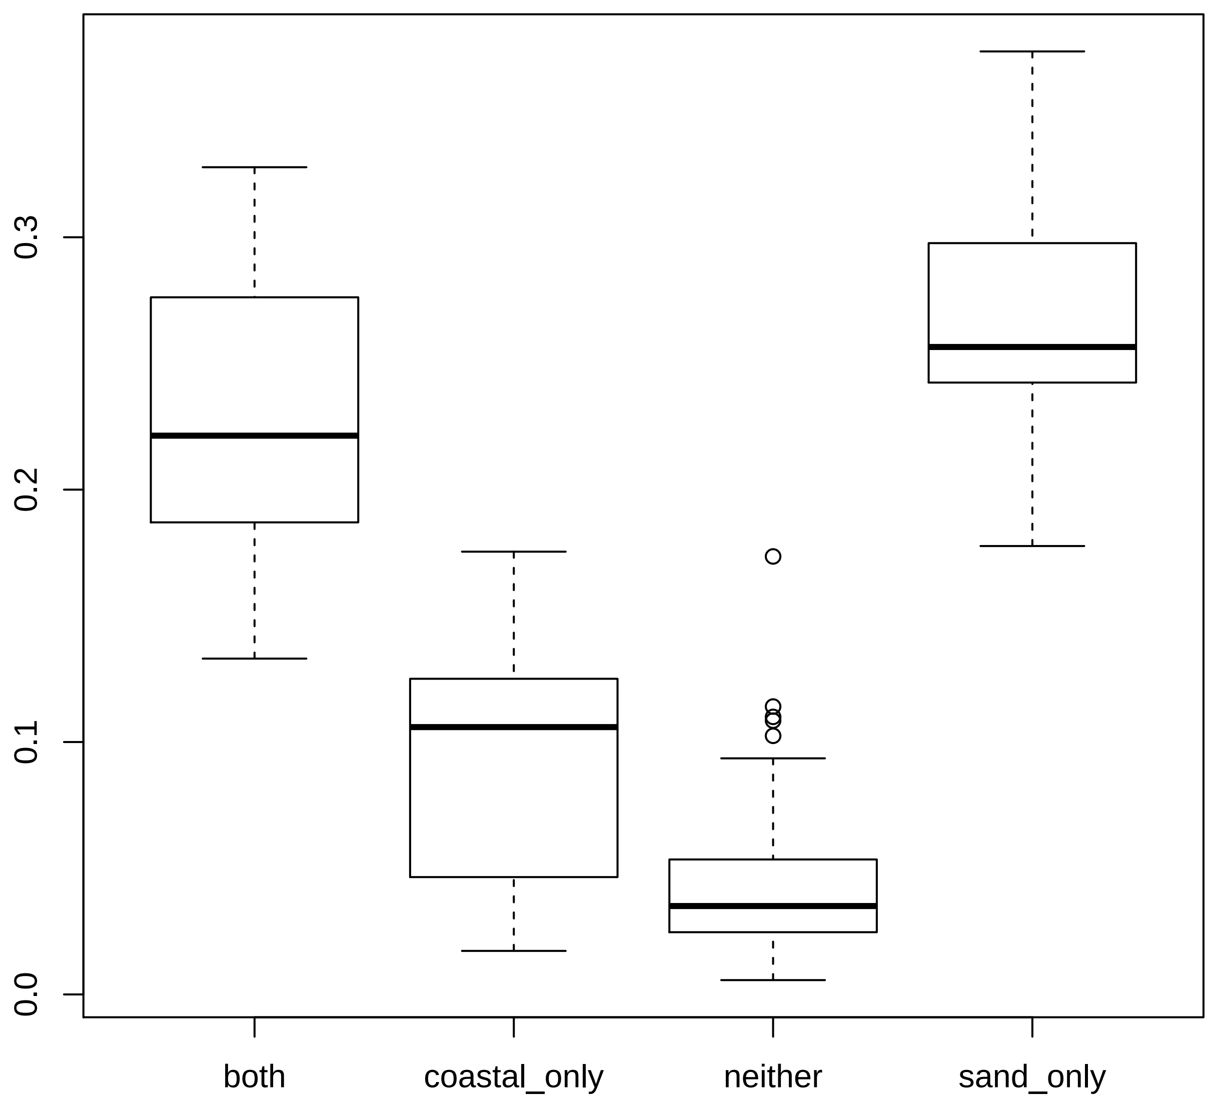

Figure 2: Distribution of log $\frac{\text{housing units in 2009}}{\text{log housing units in 2000}}$ by market type. Units as estimated by US Census.

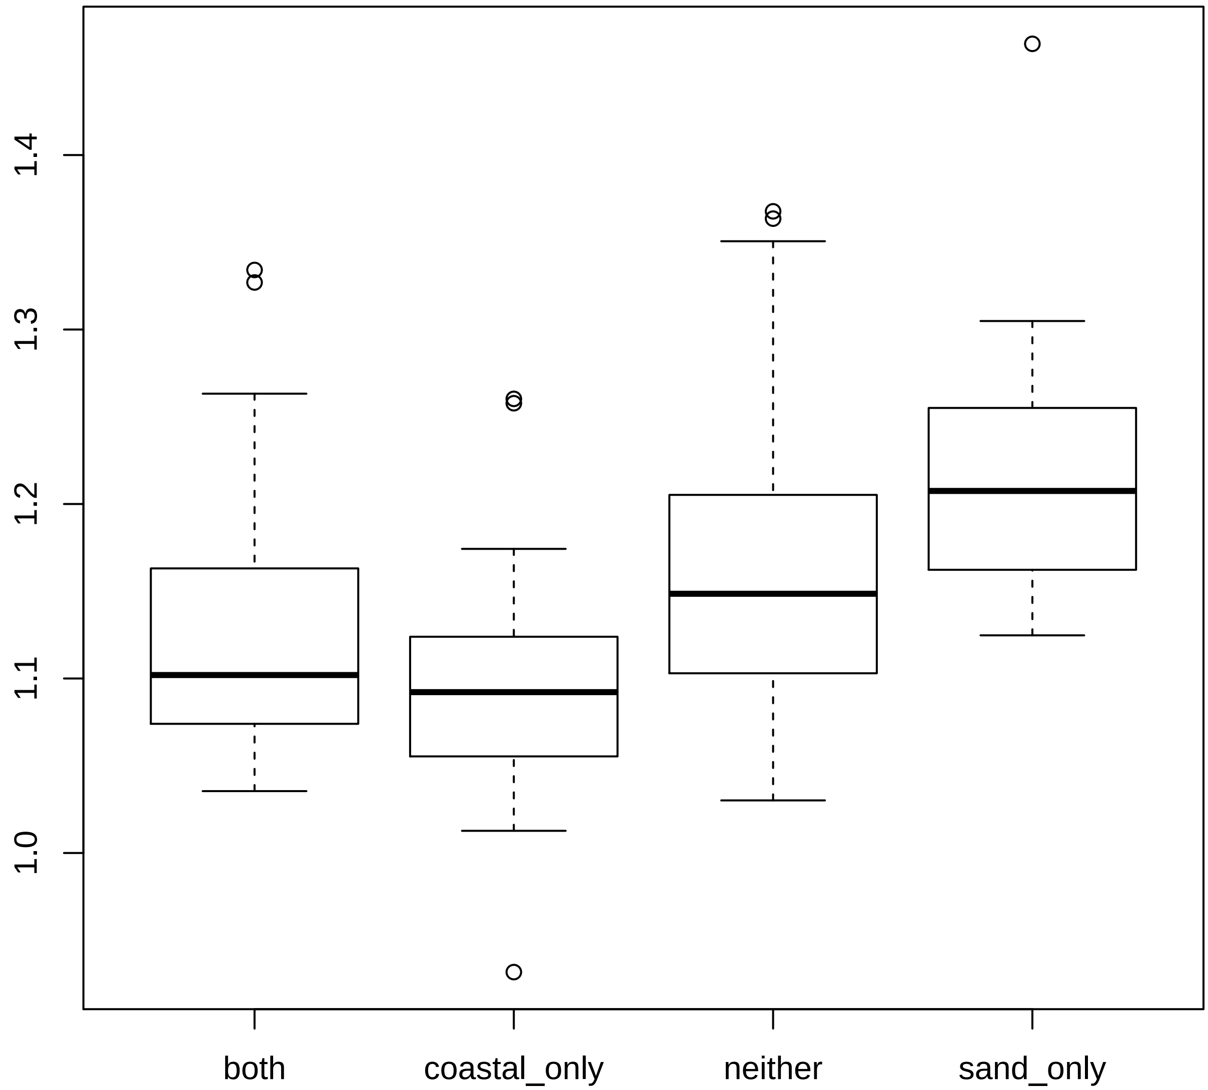

Notes: In ambiguous cases, I define membership of a housing market in a state based on the first city in the metropolitan area's name. I define markets (CBSA or Metropolitan Division, if applicable) as "Coastal" if they have at least one county that Rappaport and Sachs (2003) deem adjacent to either the Atlantic or Pacific ocean, the Gulf of Mexico, or a Great Lake. Similar results were obtained in an earlier version of this paper that considered only the California and Northeast Coastal markets.

Figure 3: 1980s quantity growth and maximal price growth. "Bubbles" are 2000s housing price cycle magnitude measure b_1.

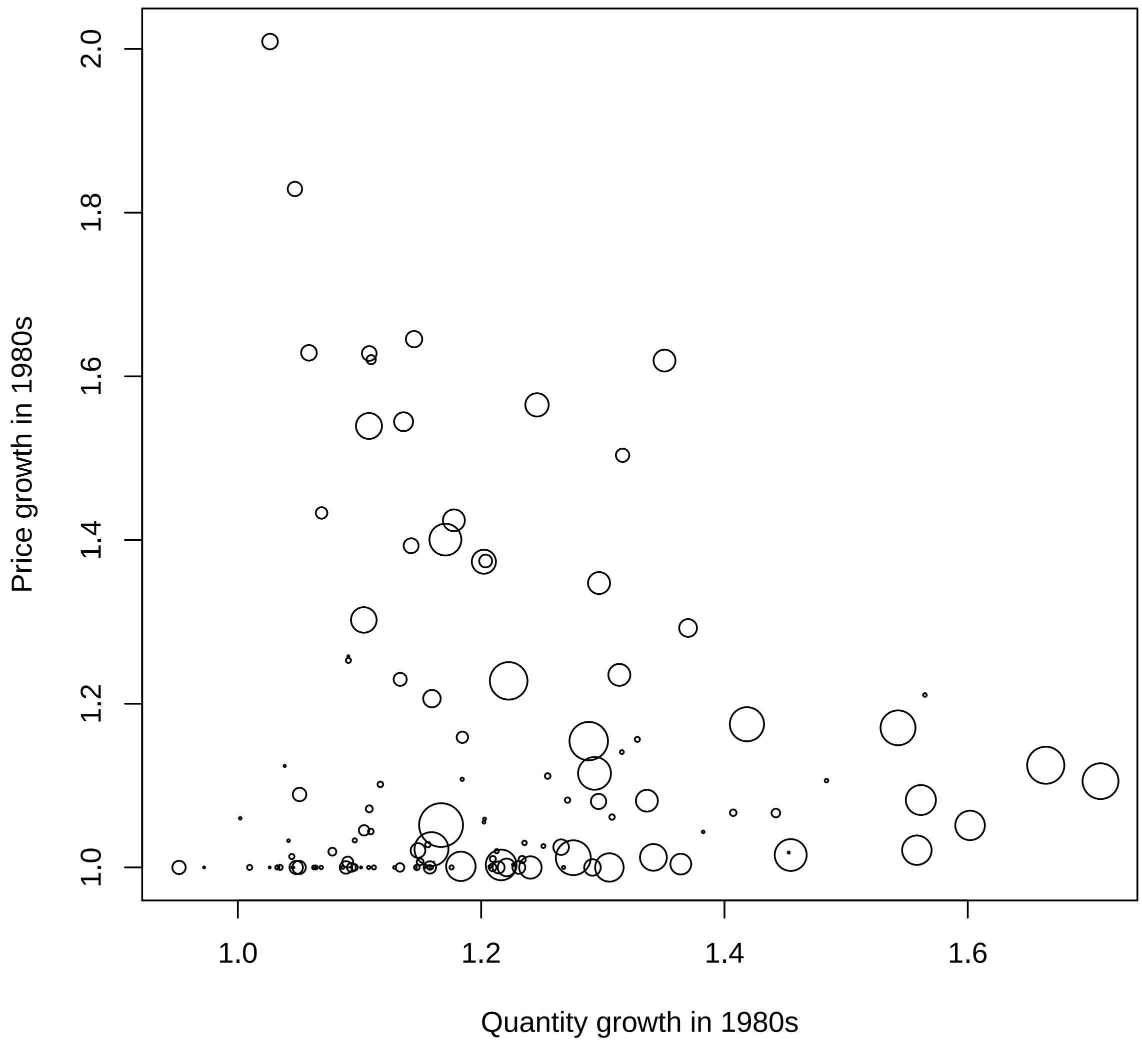

Table 1: Summary Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
$\frac{\text{Units 1990}}{\text{Units 1980}}$	121	1.203	0.149	0.952	1.709
$\frac{\text{Units 2009}}{\text{Units 2000}}$	121	1.139	0.086	0.932	1.464
$\frac{\text{Units 2009}}{\text{Units 2000}} / \frac{p_{2007q2}}{p_{2000q1}}$	121	0.648	0.712	-0.233[a]	3.155
1980s price growth (see text)	121	1.128	0.201	1	2.009
b_1: growth rate difference	121	0.11	0.093	0.006	0.374
b_2: standard deviation	121	0.067	0.05	0.016	0.223
b_3: $\frac{p_{2007}}{p_{2010}}$	121	1.278	0.299	0.982	2.454
Unavailable Land	121	0.279	0.233	0.01	0.86
Coastal	121	0.504	0.502	0	1
Regulations	121	0.157	0.701	-1.345	2.753
Sand State	121	0.273	0.447	0	1
1980s inverse elasticity (see text)	121	0.013	0.039	0	0.299

Note a: New Orleans has a negative elasticity estimate. All other metropolitan areas have positive values, with Nassau-Suffolk the lowest at .04.

Table 2: Cycle amplitude measures: pairwise correlations

	b_1	b_2	b_3
b_1	1	.97	.91
b_2	0.97	1	.96
b_3	0.91	0.96	1

Table 3: Regressions of cycle amplitude measures $b1$, $b2$, and $b3$ on estimated inverse supply elasticity: $\frac{\text{Log real price growth 1980 to maximum 1980s}}{\text{Log quantity growth 1980 to 1990}}$ and demand growth: log quantity growth 1980s + $\gamma\times$ log real price growth 1980 to maximum 1980s. γ varies from 0 to .75 across columns (2) through (4). Inverse elasticity divided by 100 for readability.

Dependent variable: b_1: difference in growth rates boom - bust

	(1)	(2)	(3)	(4)
		$\gamma = 0$	$\gamma = .25$	$\gamma = .75$
constant	0.1104**	-0.2551**	-0.3509**	-0.3782**
	(0.0089)	(0.0639)	(0.0705)	(0.0739)
inverse elasticity	-0.001	0.2752	0.0627	-0.2941
	(0.217)	(0.1983)	(0.1866)	(0.191)
demand		0.3008**	0.3101**	0.2403**
		(0.0522)	(0.0471)	(0.0362)
Adj. R-sq.	-0.01	0.21	0.26	0.26
degrees.freedom	119	118	118	118

Dependent variable: b_2 standard deviation of real price growth

	(1)	(2)	(3)	(4)
		$\gamma = 0$	$\gamma = .25$	$\gamma = .75$
constant	0.0668**	-0.1203**	-0.1747**	-0.1965**
	(0.0048)	(0.0347)	(0.0382)	(0.0396)
inverse elasticity	-0.0062	0.1351	0.0272	-0.1641
	(0.1165)	(0.1078)	(0.1011)	(0.1024)
demand		0.1539**	0.1623**	0.1295**
		(0.0284)	(0.0255)	(0.0194)
Adj. R-sq.	-0.01	0.19	0.24	0.26
degrees.freedom	119	118	118	118

Dependent variable: b_3: log price 2010 minus log price 2007

	(1)	(2)	(3)	(4)
		$\gamma = 0$	$\gamma = .25$	$\gamma = .75$
constant	1.2844**	0.1897	-0.0136	0.0241
	(0.0287)	(0.209)	(0.2362)	(0.2528)
inverse elasticity	-0.5115	0.3158	-0.3322	-1.2674*
	(0.6971)	(0.6489)	(0.6248)	(0.6533)
demand		0.9011**	0.8725**	0.6198**
		(0.1707)	(0.1578)	(0.1237)
Adj. R-sq.	0	0.18	0.2	0.17
degrees.freedom	119	118	118	118

Table 4: Correlation of unit growth 2000 to 2009 different measures of supply constraints. Top panel: unconditional; bottom panel all variables demeaned at state level.

	$\frac{\text{units 2009}}{\text{units2000}}$	Coastal	unavailable	regulations
$\frac{\text{units 2009}}{\text{units2000}}$	1.00	-0.36	-0.13	-0.02
Coastal	-0.36	1.00	0.50	0.21
unavailable	-0.13	0.50	1.00	0.34
regulations	-0.02	0.21	0.34	1.00
$\frac{\text{units 2009}}{\text{units2000}}$	1.00	-0.26	-0.50	-0.03
Coastal	-0.26	1.00	0.51	0.07
unavailable	-0.50	0.51	1.00	0.18
regulations	-0.03	0.07	0.18	1.00

Table 5: Regressions of cycle amplitude measures on Saiz (2008) measure of land unavailability, Rappaport and Sachs (2003) measure of Coastal status, and Gyourko et al. (2006) regulatory index. State dummy variables are present in even columns only. OLS standard errors in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var.	b_1	b_1	b_2	b_2	b_3	b_3
constant	0.0535**	0.0574	0.035**	0.0316	1.1383**	1.1487**
	(0.0108)	(0.0408)	(0.0058)	(0.0206)	(0.0375)	(0.1624)
Unavailable	0.2068**	-0.0458*	0.1144**	-0.0201	0.4978**	-0.3547**
	(0.0305)	(0.0258)	(0.0162)	(0.013)	(0.1054)	(0.1027)
State dummies?	No	Yes	No	Yes	No	Yes
Adj. R-sq.	0.26	0.8	0.28	0.82	0.14	0.7
degrees.freedom	126	90	126	90	126	90

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var.	b_1	b_1	b_2	b_2	b_3	b_3
constant	0.0795**	0.0508	0.0505**	0.0288	1.2034**	1.0978**
	(0.0107)	(0.0413)	(0.0058)	(0.0208)	(0.0353)	(0.1719)
Coastal	0.0623**	0.0025	0.0321**	0.0012	0.1446**	-0.0215
	(0.0152)	(0.0106)	(0.0082)	(0.0054)	(0.0504)	(0.0442)
State dummies?	No	Yes	No	Yes	No	Yes
Adj. R-sq.	0.11	0.79	0.1	0.82	0.05	0.66
degrees.freedom	126	90	126	90	126	90

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Var.	b_1	b_1	b_2	b_2	b_3	b_3
constant	0.1018**	0.0508	0.0616**	0.0288	1.2542**	1.0961**
	(0.0075)	(0.0414)	(0.004)	(0.0208)	(0.0253)	(0.1717)
Regulation	0.055**	-6e-04	0.0312**	8e-04	0.1336**	-0.0276
	(0.0106)	(0.01)	(0.0056)	(0.005)	(0.0357)	(0.0415)
State dummies?	No	Yes	No	Yes	No	Yes
Adj. R-sq.	0.17	0.79	0.19	0.82	0.09	0.66
degrees.freedom	126	90	126	90	126	90