The Contribution of Housing Markets to the Great Recession

Peter Pedroni[1]

and

Stephen Sheppard[2]

Please check with authors before quoting

November 21, 2010

[1]Department of Economics, Williams College, 24 Hopkins Hall Drive, Williamstown, MA 01267

[2]Department of Economics, Williams College, 24 Hopkins Hall Drive, Williamstown, MA 01267

Abstract

The recession of 2008-2009 was the most severe downturn experienced by the US since 1981-82 and, by some measures, the most severe since the 1930s. Most economists and policy makers attribute the recession to the collapse or near collapse of financial institutions that held large quantities of assets whose values derived from mortgages on residential property. The values of these assets dropped precipitously with the collapse of house prices in 2006 and 2007 leading to sharp reductions in investment and lending. Household wealth was also reduced by the collapse of house prices leading to reductions in household spending that exacerbated the contraction.

In this paper, we employ a novel panel-based structurally-identified VAR approach to examine the evidence for the impact of housing markets on both local and national unemployment and output during the crisis. We use quarterly data on house prices, unemployment and output in 384 metro areas. The approach allows us to identify and isolate shocks originating in the local housing markets, and to distinguish these from other economics shocks originating at either the local or national level. The approach also accounts for heterogeneity in the dynamic responses to these shocks at the local level, and investigates the relationship between the size of the local responses and the degree of regulation of housing supply present at the local level.

Our approach is able to better quantify the extent to which the housing market contributed to the recent national economic downturn and illustrates the contribution that is possible when the economies of the regions within the national economy are used to test and deepen our understanding of aggregate economic behavior.

1 Introduction

The recession that began in December 2007 and officially ended in June of 2009 has been dubbed the "Great Recession" in both the popular press and by at least some academics. It is certainly the most severe recession endured by the US economy since the recession of 1981-82, and by some measures it is the most severe economic contraction since the 1930s.

During the recession itself and during the somewhat tepid recovery that followed, an analytic narrative has captured the imagination of policy makers, the popular press and many economists to explain the recession. This narrative runs roughly as follows: sometime after 2000 or 2001 a housing bubble developed in the US. The bubble was fueled by excessive and perhaps fraudulent lending by financial institutions to buyers whose only hope of meeting their repayment obligations rested on hoped-for capital gains derived from continuation of the bubble. The financial institutions securitized the resulting mortgages into assets that were so complex that the risk associated with the assets was difficult or impossible to determine, and these derivative assets were widely sold to investors and to other financial organizations who in turn built highly leveraged operations that assumed continued and secure returns from these assets. The bubble burst in 2006, and the subsequent decline in house prices led to increased foreclosures and threatened the stability of major financial institutions. This instability led to severe restrictions in consumer spending that had been encouraged by expected capital gains from residential property and in credit that could no longer be extended to businesses by overleveraged financial institutions. These restrictions sent the economy into recession beginning in December 2007.

The role of housing in this story is central and the extent of the bubble that affected the US housing market is illustrated in Figure 1, adapted from data in Shiller (2005). These data or figures similar to this have been widely reproduced in the popular and financial press and are presented as part of the evidence supporting the narrative described above. The data support the idea of "the" US housing bubble of 2001-2006, and may reinforce an impression that a graph similar to Figure 1 describes most US urban areas. This is certainly not correct, as shown in Figure 2.

Figure 2 shows the FHFA house price index for 384 US metro areas, normalized to the same year (1995) as the overall US house price index shown in Figure 1. We see that the actual experience of US housing markets has been highly varied, with the most extreme areas like Naples, Florida peaking not at 180 percent of their 1995 levels but at 365 percent. At the same time many metro areas experienced little or no "bubble" in residential house prices. For example Lafayette, Indiana peaked at 134 percent of 1995

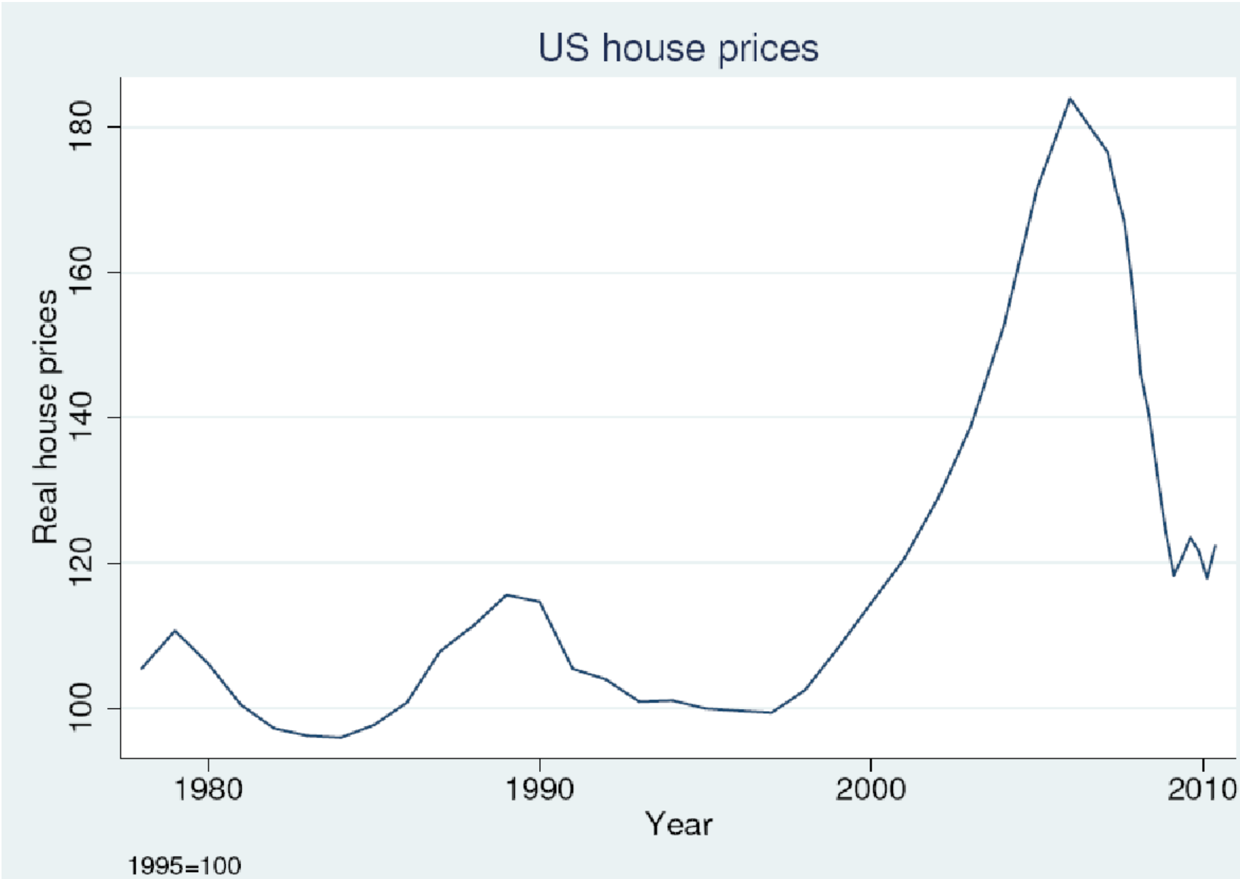

Figure 1: US house prices: an aggregate view of the bubble

prices at the beginning of 2007, and declined to 133 percent by first quarter of 2010.

It is clear from Figure 2 that the behavior of house prices prior to, during and after the recession of 2007-09 has exhibited considerable variation. To ignore this variation is to ignore an important opportunity to learn about the role of housing in causing the recession and about impacts of policies designed to prevent future dislocations attributable to the housing or financial markets.

Once we recognize the variation in actual US metro house price dynamics, it is natural to inquire whether there is a spatial structure to this variation. Are cities that experienced more severe house price fluctuations concentrated in particular regions or are they interspersed with those cities where the magnitude of house price fluctuations was less severe? Figure 3 provides an initial look at this question.

Figure 3 shows the distribution of house price variability in two ways. The shading of individual metro areas represents the difference between the average annual rate of growth in house prices from 1995 through the peak of the local market and the annual rate of decline from the peak of the local market to the post 2006 trough. The height of the superimposed bars is proportional to the standard deviation of the metro FHFA all sales house price index from 1995 through the first quarter of 2010.

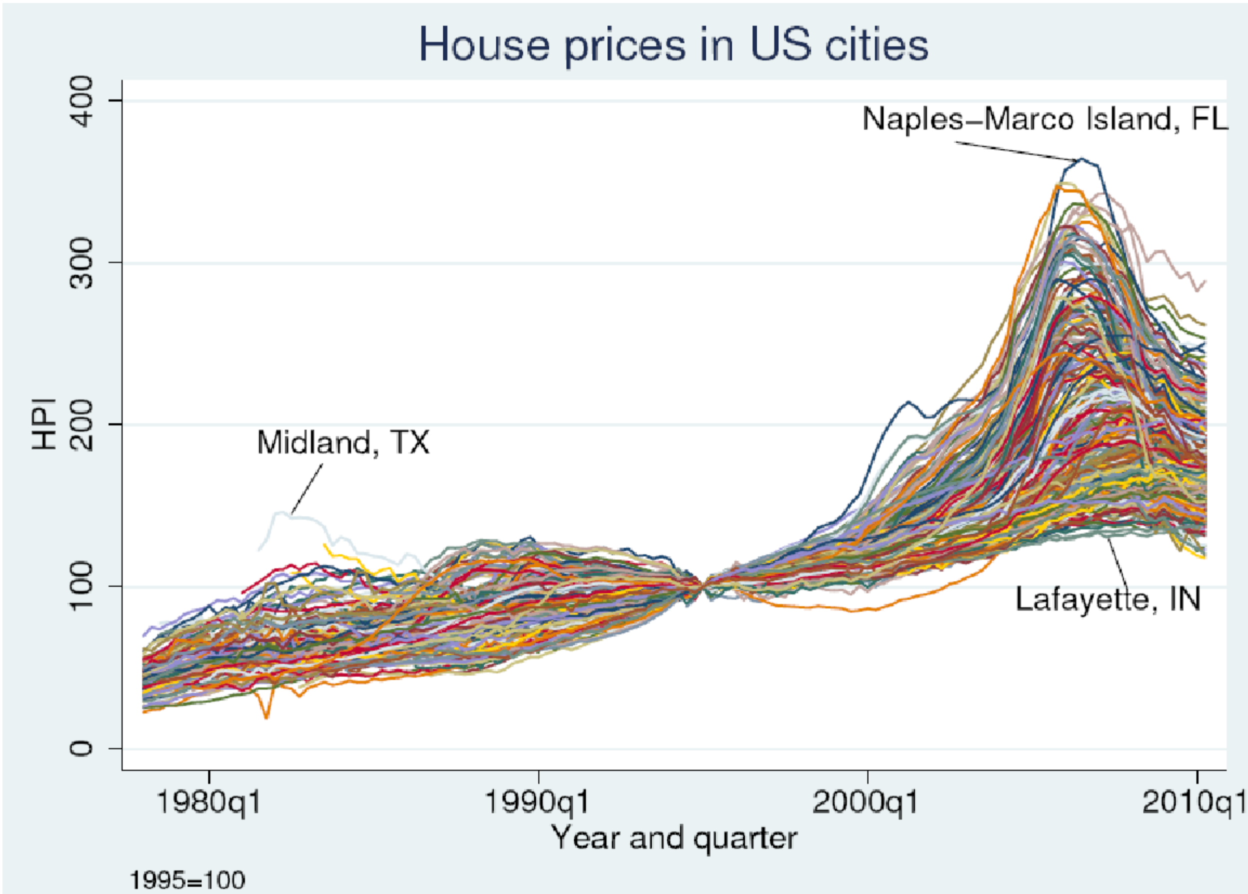

Figure 2: US house prices: 384 distinct markets

The data shown in Figure 3 suggest that the most extreme swings in house prices took place in cities along the Atlantic and Pacific coasts. While possible, it would be surprising if lending practices in these regions were uniquely irresponsible or fraudulent. The securitization of mortgages was a national phenomenon and while it is possible that mortgage lenders away from coastal states were less likely to be the source of mortgages for creation of derivative securities it would surely not follow that non-coastal lenders were insulated from risk or from broader dislocations in national credit markets.

Coleman, LaCour-Little & Vandell (2008) analyze a short panel of data from 20 metro areas and find no evidence that subprime lending practices were significant determinants of future house price returns, but do find some other shifts in credit markets that appear to be linked to the housing market bubble. This sort of co-determination should not be surprising and poses challenges for testing hypotheses about the role of housing markets in causing or contributing to the economic decline.

It is possible that there are significant differences between coastal and non-coastal regions in house price dynamics and/or linkages between the housing sector and the rest of the economy. There are several

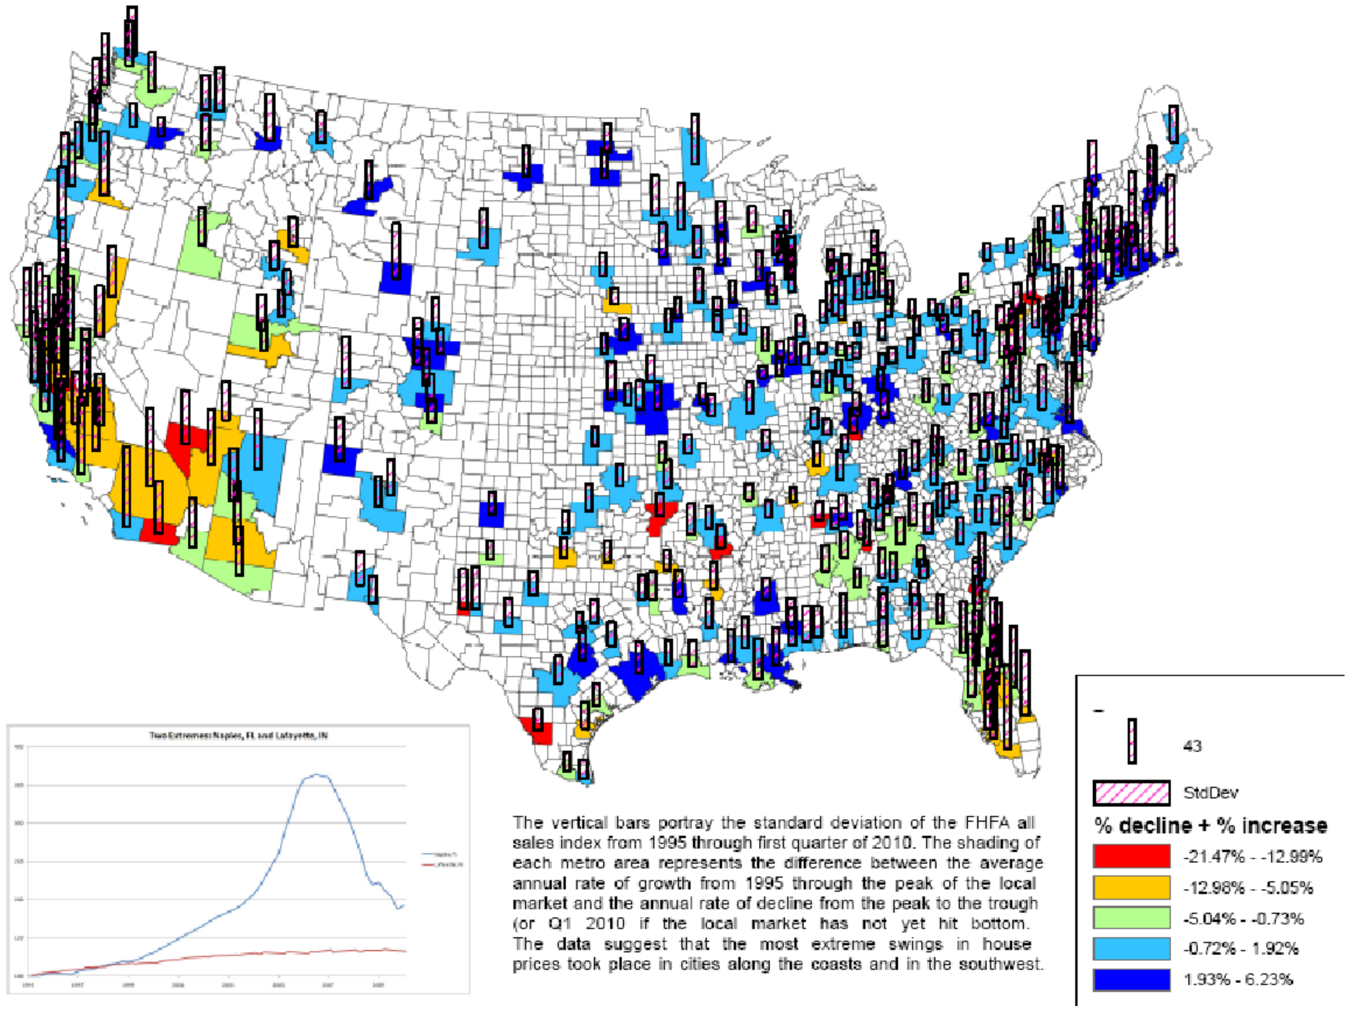

Figure 3: Spatial structure of housing bubble

possible reasons why such differences might exist. Differences in the industrial structure of the regions are clearly present, and these may give rise to different house price dynamics.

There are also general differences in land use and housing regulation between the regions. These regulatory regimes can have large impacts on the ability of housing suppliers to respond to changes in house prices, and such differences in supply elasticity might offer an explanation. Glaeser, Gyourko & Saiz (2008) present a model of housing supply and housing bubbles that incorporates differences in the elasticity of housing supply as an explanation for why some urban areas may be more prone to housing bubbles. Their examination of the data provides some support for their model.

Many writers attribute the house price bubble to irrational exuberance (Shiller (2005)) or to some non-linearity in discounting of rents and the resultant hyper-sensitivity to interest rates (Himmelberg, Mayer & Sinai (2005)). While cultural and other differences between coastal and non-coastal regions in the US clearly exist, it seems that if such arguments are to provide an explanation of what the US economy has actually experienced they will probably have to explain why market participants in coastal cities are more irrational or more non-linear than their counterparts in the central part of the country.

In this paper we provide an analysis of linkages between the local economy, house prices and unemployment based on a structural vector auto-regression (SVAR) analysis of the panel of 384 US metropolitan areas for which FHFA house price data are available. This approach allows us to examine some of the possible explanations of how house price swings might have contributed to the recession of 2007-09, and to explore in greater detail the structure of house price dynamics in US metro areas.

2 Empirical methodology

We use a structural panel VAR methodology as developed in Pedroni (2008). The approach enables us to address several challenging aspects of our empirical analysis. For example, the approach is able to accommodate the fact that the true economic shocks of interest are unobserved, and the observed variables are highly endogenous. It also accounts for the fact that the economic activity of different metropolitan areas are interdependent, and that some shocks of interest originate at the local level while others do not. Importantly, the approach accounts for the fact that metropolitan areas are very different from one another, so that the local responses to both local and national economic shocks may be very heterogeneous among metropolitan areas.

The panel structural VAR approach can be viewed as a generalization of time series based structural VAR approaches such as Blanchard & Quah (1989) to a panel data framework. Analogous to the time series based approach, identification of the unobserved structural economic shocks is accomplished by a set of identifying restrictions. These typically include the assumption that the true unobserved structural shocks of interest correspond to the forcing process for the economy being modeled, and therefore are naturally thought of as orthogonal to one another, and arbitrarily normalizable. Additionally, assumptions are typically made regarding the admissible dynamic responses of the observed variables to the unobserved shocks. In the case of Blanchard and Quah, these additional identifying restrictions come in the form of restrictions on the steady response of certain variables to certain shocks. Accordingly, the contemporaneous and short run dynamic responses of all variables are considered completely endogenous and unrestricted, with only the long run steady state response of some variables restricted.

The challenge for panels such as ours is to recognize that individual metropolitan areas are very different from one another, but also interdependent. Toward this end, the panel structural VAR approach allows for complete heterogeneity in the responses of the metropolitan level variables to the unobserved shocks. Confidence intervals for the median responses are then based on the spatial variation in the estimated responses among metropolitan areas. Furthermore, in light of the fact that metropolitan areas are interdependent and likely responding as a group to common shocks, we also identify the metropolitan responses to economic shocks originating nationally. This allows us to construct confidence intervals that are valid despite the fact that metropolitan area responses are cross-sectionally dependent. Finally, this approach allows us to relate the pattern of estimated dynamic responses among metropolitan areas and to relate these to observed characteristics of the metro areas.

Since we are primarily interested in knowing the consequences of structural shocks arising in the housing market, we take this structural economic shock to be any shock that leads to a long run effect on housing price, and refer to it as a "house" shock. We want to examine the dynamic responses of output and unemployment to such shocks. However, we recognize that many other shocks also impact output and employment, and so we must control for these in a way that allows for complete endogeneity. In particular, it will be important to control for shocks that may have a permanent effect on housing prices, but which are not fundamentally housing market shocks. We posit that one such class of shocks, which we refer to as economic *supply* shocks, can be distinguished from *housing market* shocks by virtue of the fact that they also permanently alter local economic activity as measured by metropolitan GDP in the long run. Specifically, we assume that while both the supply shocks and housing market shocks may have a permanent long run effect on housing prices, only the supply shocks have a permanent effect on output. Our identification treats these two classes of shocks as the only shocks that permanently alter housing prices, and allow both shocks to permanently alter unemployment. Of course, these are not the only classes of shocks that might permanently alter unemployment. Therefore, we also control for a third class of shocks which are permitted to have a permanent effect on unemployment, but no long run effect on output or housing prices. We refer to these as *other* shocks.

To implement this identification scheme while allowing for heterogeneity and interdependence among metropolitan areas, we take the following steps. First, we estimate a dynamic vector autoregression for three variables, namely metropolitan area output, housing prices and unemployment for each metropolitan area in our sample. The lag length for each metro area VAR is chosen by a standard Akaike Information Criterion, and is allowed to differ among metro areas. Next, means for each variable are computed across metropolitan areas for each time period, and a separate VAR is fitted for these averages. The reduced form impulse responses are computed for both the estimated individual and averaged VARs, and the triangular steady state identification scheme is applied to each of these in order to extract structural shocks. The structural shocks identified from the metropolitan area VAR are composed of both idiosyncratic local metro specific shocks and common national shocks. However, since the structural shocks are orthogonal, these can be further decomposed into region specific and national shocks by OLS regression of the composite shocks on the national shocks. Once these regional and national shocks have been identified, the structural impulse responses and variance decompositions are computed. For more details, we refer readers to Pedroni (2008).

3 Theoretical Approach

Our analysis is focused on an urban economy where total output (gross metropolitan product), real house prices, and unemployment are interdependent and mutually determined. As noted above, the analysis is based on a dynamic model in which white noise shocks affect the economy and generate a series of general equilibrium responses. The supply shocks and the housing shocks are the primary shocks of interest whose effects we seek to identify.

Our interest is in developing an improved understanding of the metropolitan economy, including the dynamic structure of its responses to shocks and the long-run implications of shocks. Our primary motivation is in better understanding the implications of housing shocks so much of the discussion below develops with that objective in mind.

4 The data

There are 384 metropolitan areas in the US for which FHFA produces a quarterly all sales house price index. These indices use data from the MSA defined in 2009 by the Office of Management and Budget. Some of the house price indices are available beginning in the first quarter of 1978, most are available throughout all of the 1990s and some do not become available until the last few years. Further details on the mechanism of constructing this index are provided in Calhoun (1996).

We combine the price data with data from the Local Area Unemployment series provided by the Bureau of Labor Statistics for the same geographic areas. These unemployment rates are provided monthly and we produce quarterly averages to match our other data.

Total metropolitan GDP is made available for all metropolitan areas in the US by the Bureau of

Economic Analysis, but only for years after 2000. Moodys Analytics has used the BEA methodology to produce a quarterly series of real GDP estimates for all 384 metro areas, and we utilize these for our measure of total output.

These are the only variables required for estimation of our SVAR model. In what follows we will examine the pattern of impacts that is estimated by the analysis, and relate these impacts to the industrial structure in each city. Because we will want to use location quotients as an indicator of industrial structure in each metropolitan area, we use data from the Quarterly Census of Employment and Wages. Unfortunately, QCEW data for metro areas are based on the 379 areas defined in 2004. We adapt their definitions to match our other data, but drop the 5 newest metro areas from the data used in our analysis when examining the impact of industrial structure.

The result is an unbalanced panel of 384 areas for which we have up to 130 complete quarterly observations of real output, house prices and unemployment for each city beginning with first quarter of 1978 and running through the first quarter of 2010. These are combined with data identifying each city as coming from a "coastal" state (located on the Atlantic Coast from Maine to Florida, on the Pacific Coast, and including Nevada and Arizona) or non-coastal state (all other states). Also included are location quotients for 2-digit NAICS industries for all cities in the sample based on industrial composition in 2009.

5 Results

The structural panel VAR approach described above focuses our attention on three different types of shock, each of which can be considered as applying only to an individual metro area or to the national urban system as a whole. To present the results we focus attention on the variance decompositions (given in the appendix) and the impulse-response functions for each of the six resulting types of shock.

The variance decompositions are given in the appendix, and show that the variance in MSA output is largely explained by shocks to regional supply, with some impact in first few years from regional house shocks, but this impact is largely dissipated after 5 years. The variance in local house prices is predominantly explained by regional supply shocks and regional house shocks. The variance in unemployment is explained by regional shocks of all three varieties. While national shocks do have some explanatory power the share of variance explained is much less than for regional shocks, as would be expected.

Tables 1, 2 and 3 present results in the form of plots of the impulse response functions for metropolitan output, house prices, and unemployment respectively. These impulse-response functions are derived from

Table 1: Impulse-response functions for output

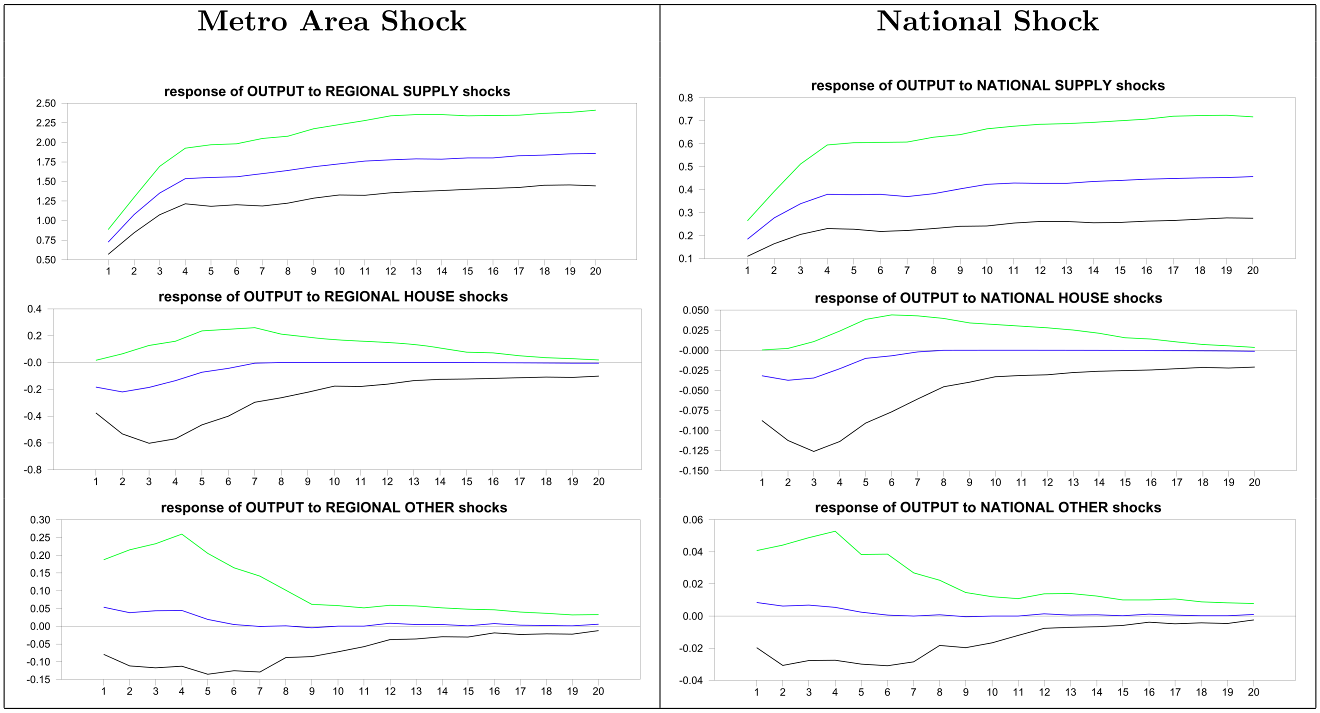

our model estimates, and summarize the dynamic response of the variables to the shocks that affect the urban system. In the left hand column of each table are the responses to regional limited to the metropolitan area supply, housing and other shocks. In the right hand column are the responses to national shocks of the same type.

For each type of shock and each variable that responds, we present the response over 20 quarters. The central (blue) line shows the median response of our panel of cities at each lag from the shock. The upper (green) and lower (black) lines show the 75th and 25th percentiles of our panel. This provides an indication of the range of responses to each shock. The vertical scales are adjusted to accommodate the range of responses to each shock, but show the same variable within each figure. It is tempting to think of the upper line as identifying the impulse-response of the "most responsive" MSA and the lower line as the dynamic behavior of the "least responsive" but the reader should keep in mind that these orderings will generally change from quarter to quarter so a particular MSA may be in the most responsive quartile at one particular lag but closer to the median one or two lags later.

There are two general observations that can be offered that apply to all three of the figures:

- The regional and the national shocks generate qualitatively similar dynamic patterns of response

- The national shocks generate responses that are much smaller than the regional shocks, with national shock responses ranging from one tenth to one third of the magnitude of the regional shocks

In the graphs shown in Table 1, the vertical axis presents a response of local output measured in units, so that the numerical quantities indicated can be interpreted as percentage changes in metropolitan GDP that occur at various points in time as a result of the shock. Per our identification restrictions, the impact on local output of a house or other shock converges to zero as the metropolitan economy returns to steady state. The response of output to a regional supply shock converges to about 1.75% of local GDP for the median metropolitan area, with most of the impact taking place in the first four quarters.

In the graphs shown in Table 2, the vertical axis presents a response of local house prices measured in units, so that the numerical quantities indicated can be interpreted as percentage changes in the house price index for the metropolitan area that occur over time as a result of the shock. Both house shocks and supply shocks are seen to generate noticeable changes in metropolitan region house prices, with the median metropolitan area converging to 1.5%-2% increase in house prices in response to a regional house shock. It is worth noting that there is considerable heterogeneity across the US urban system in the magnitude of response to these shocks, with the most responsive quartile of MSAs eventually realizing more than 4% increase in house prices as a result of a shock, and the least responsive quartile experiencing barely 1%. This heterogeneity in response will be discussed at greater length below.

In the graphs presented in Table 3, the vertical axis presents a response of the local unemployment rate and the numerical quantities indicated are changes in the this rate for the metropolitan area that occur at various points in time as a result of the shock. As might be expected, a supply shock has a negative impact on the MSA unemployment rate. A regional supply shock that for the median MSA increases output and MSA house prices by somewhat more than 1.75% will decrease the MSA unemployment rate by about 0.425.

The graphs shown in the second row of the table are of particular interest. They indicate that a house shock, that will increase house prices by somewhat more than 2%, will actually reduce the MSA unemployment rate by nearly 0.10 percentage points. Why might this effect arise? At first one might think that the house shock is stimulating output, and the increase in MSA output increases labor supply. Review of the data in Table 1 above, however, shows that in general this is not the case. For more than half of the metro areas, a regional house shock alone has a negative impact on output. As the economy converges towards the steady-state, the impact of a house shock on output goes to zero. What the analysis shown in

Table 2: Impulse-response functions for house prices

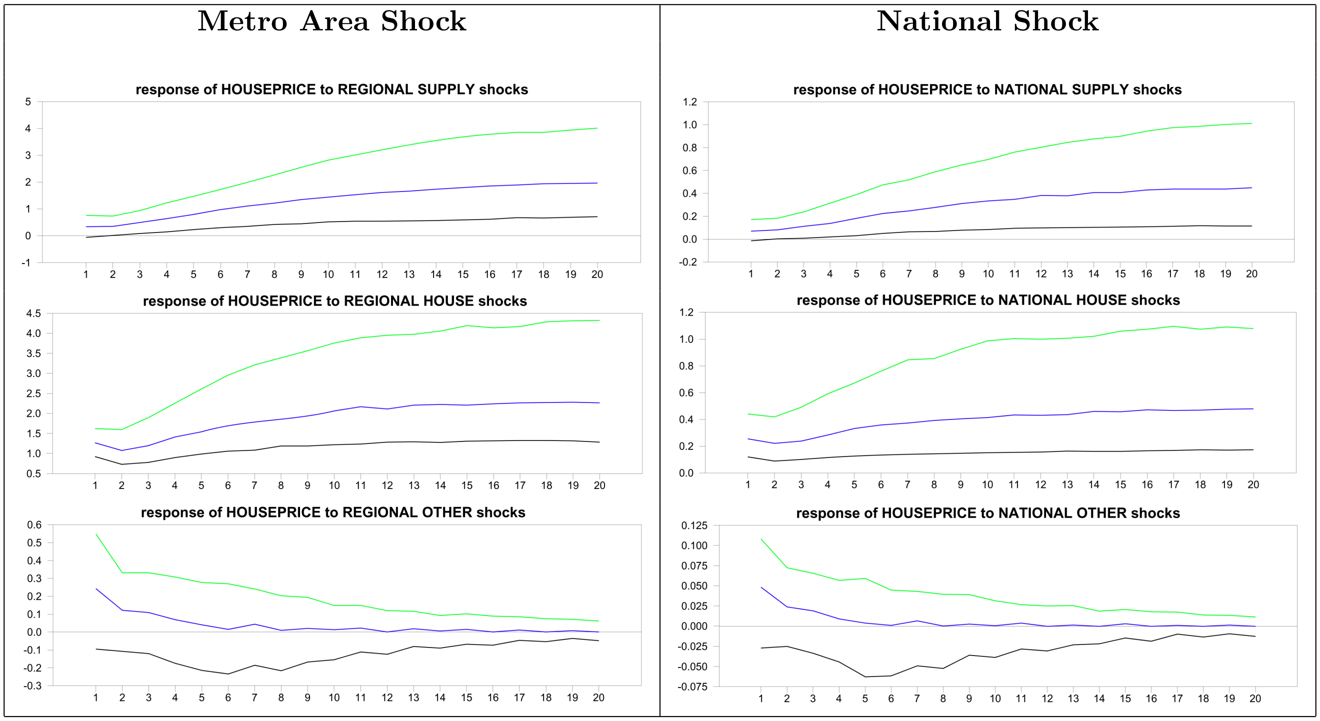

Table 3 reveals is that a house shock that increases house prices in the steady state lowers unemployment for most cities.

This finding is consistent with a perspective suggested by Oswald but as yet unpublished. Oswald's hypothesis was in fact inspired by the empirical analysis presented in Hughes & McCormick (1981) and Hughes & McCormick (1987) that demonstrated that tenure in council housing would discourage internal migration and potentially retard labor market adjustment, leading to higher rates of unemployment. The issue has subsequently been investigated by Munch, Rosholm & Svarer (2008), Coulson & Fisher (2009) and others linking local labor market congestion to local housing markets. A shock that increases house prices will plausibly be associated with easier real estate transactions and reduced time on the market for houses being sold. This in turn facilitates geographic labor mobility, reducing local labor market congestion and reducing unemployment. This finding is important because micro-data studies of US metropolitan markets have generally found mixed results and at best weak support for the "Oswald hypothesis". The studies however are made difficult because of endogeneity in labor search decisions and housing tenure decisions. Under the identification assumptions used in our analysis, the estimates are robust to general endogeneity and can be expected to more consistently measure the general equilibrium impacts.

As noted above, there appears to be considerable heterogeneity amongst US metropolitan areas in the

Table 3: Impulse-response functions for unemployment

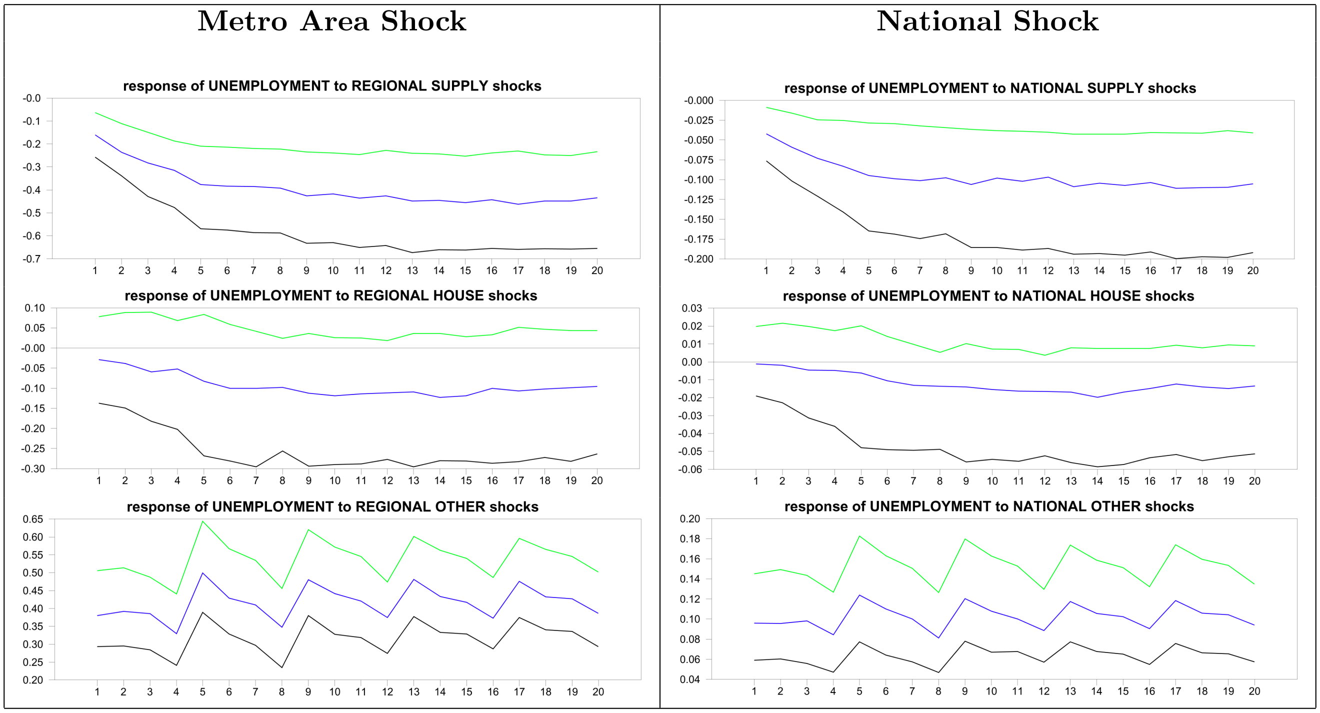

responses of local housing prices to house shocks. The introductory discussion and the presentation of data in Table 3 furthermore suggests that part of the heterogeneity in house price dynamics may exhibit spatial structure, with cities located in states along the Atlantic and Pacific coasts behaving differently from those in the central portion of the US.

In order to explore this possibility, we divide our panel roughly in half, with 194 MSAs located in Nevada, Arizona, and states along the Atlantic and Pacific coasts, and 190 located in all other states. We characterize these subsamples as coastal and non-coastal respectively, recognizing that the division is chosen because it conveniently divides the sample of US metropolitan areas almost exactly in half and is suggested by what appear to be important differences in the magnitude of recently observed fluctuations in house prices.

Tables 4, 5 and 6 present a graphical comparison of the impulse response functions for MSA output, house prices and unemployment rates in the two subsamples. For each figure, the left column presents the response of cities in coastal states and the right hand column presents the response of cities in non-coastal states. As in the figures above, the first row shows the impulse-response function for regional supply shocks, the second row for regional house shocks and the final row presents the response for other regional shocks.

Table 4 immediately presents us with an interesting finding. While the response of output to a supply shock or to a regional other shock is dynamically and quantitatively similar, there are clear and significant differences between the two subsamples in their dynamic response to a regional house shock. A house shock that increases house prices in both subsamples, and as assumed for identification has no impact on steady state output in the MSA.

Table 4: Comparison of output responses in coastal versus non-coastal regions

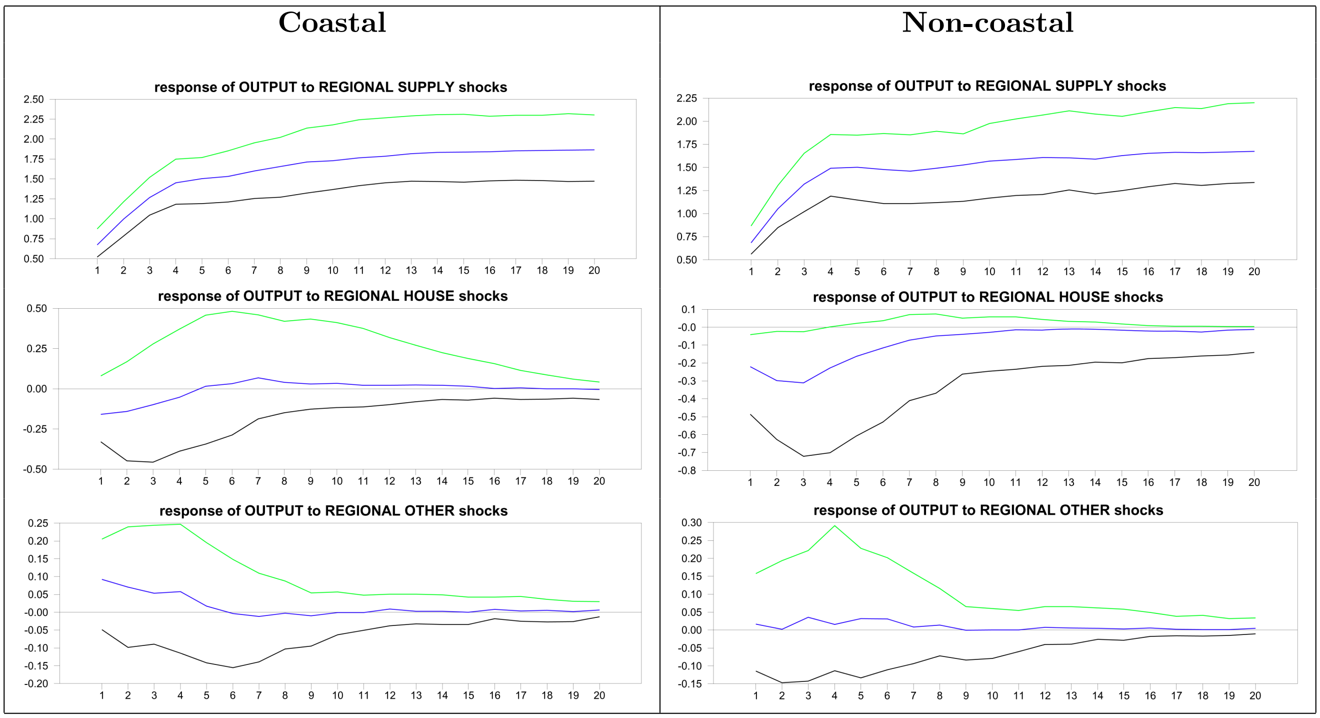

The dynamic response of output to the house shock, however, shows that in the non-coastal regions MSA output responds negatively or very close to negatively for 75% of the urban areas in the subsample. In the coastal sample, by contrast, more than 25% of the MSAs exhibit only positive output response to a regional house shock and half of the cities are exhibiting positive response to a regional house shock from the fifth quarter onwards.

Table 5 continues the comparison to the response of MSA house prices. Here we note that the coastal and non-coastal responses to other regional shocks are dynamically and quantitatively similar. Important differences emerge, however, when comparing the response of house prices to supply shocks and to house shocks. The coastal region consistently exhibits greater responsiveness of house prices to both supply shocks and to house shocks. For the median MSA, a supply shock that generates about 1.75% increase in metropolitan GDP in either coastal or non-coastal regions (first row of Table 4) will generate a 2.5% increase in house prices in the coastal region and about a 1.5% increase in house prices for the non-coastal region. Similarly, a house shock that generates about a 4% steady-state change in house prices in the coastal region will generate slightly more than 1.5% change in house prices in the non-coastal region. There is some factor of housing markets in coastal cities that makes house prices more responsive to shocks.

Table 5: Comparison of house price responses in coastal versus non-coastal regions

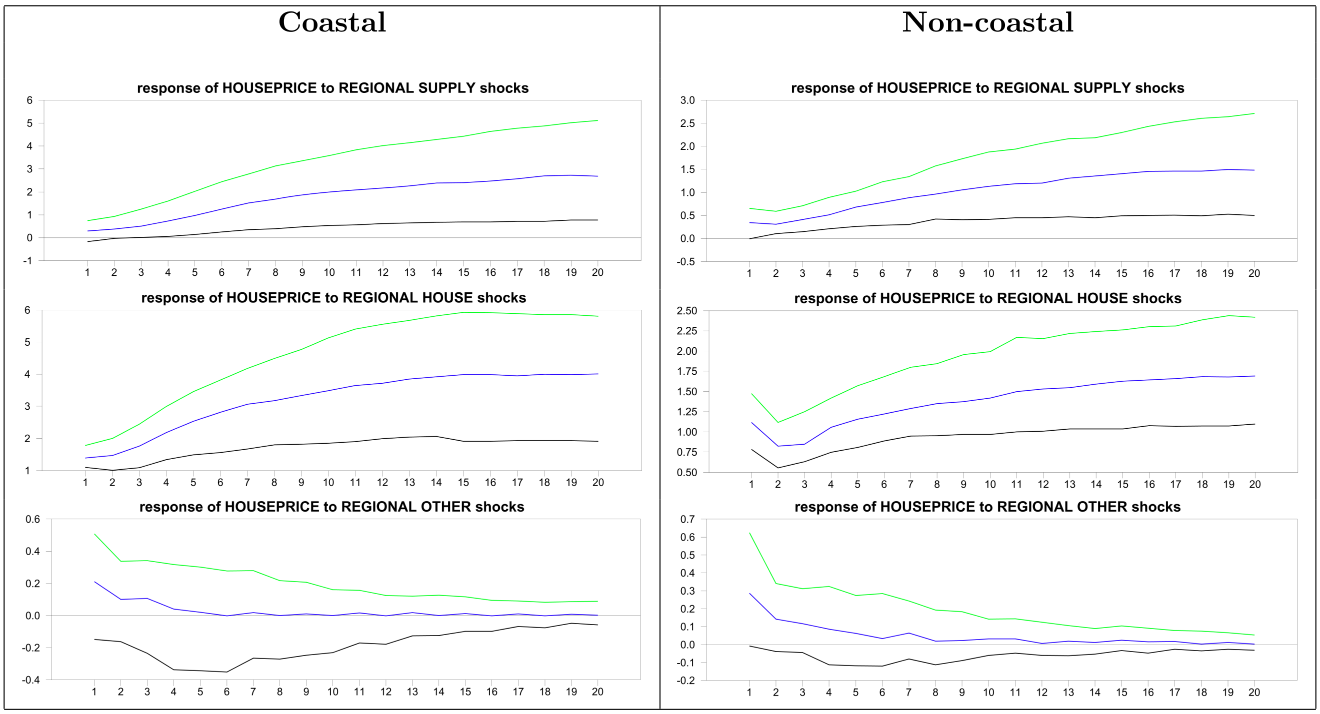

Table 6 presents a comparison of responsiveness of unemployment. The dynamic response and steady state levels of unemployment in response to supply shocks are very similar in the two regions. Similarly, the response of unemployment to other regional shocks exhibits very similar patterns in the coastal and the non-coastal regions.

There are very modest differences between the two regions in how unemployment responds to house prices shocks. The magnitudes of the responses and the values to which the MSAs appear to converge in the steady state are similar, as are the dynamic structure of the responses. Both regions continue to exhibit behavior consistent with the Oswald hypothesis noted for the combined data above. In the coastal region it appears that about 75% of metropolitan areas experience long-run declines in unemployment in response to a house shock that increases house prices. In the non-coastal region this share is on the order of 65%. This still seems to provide general support for the hypothesis.

Next, we turn attention to another possible explanation for the differences in observed dynamics of US metropolitan economies: differences in industrial structure. In order to investigate this we focus attention

Table 6: Comparison of unemployment responses in coastal versus non-coastal regions

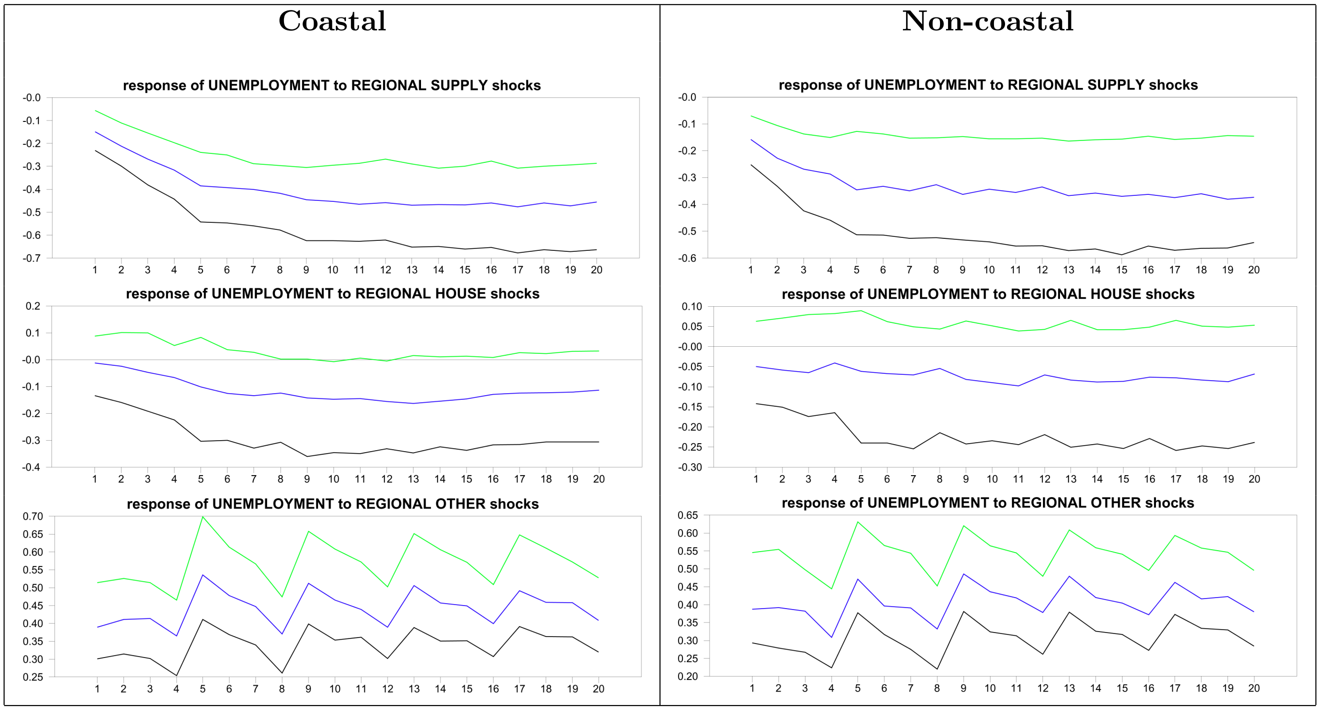

on the variation across panel members in their fourth quarter responsiveness to various types of shocks. This must be considered a preliminary check because responses might vary with respect to a great many factors.

We matched our data and estimated responses to shocks with the QCEW industrial data on establishments, employment and wages paid. We calculated location quotients for two-digit NAICS industry groups and examined the relationship between magnitudes of location quotients and response to shock. In order to encompass the largest possible group of cities we use location quotients calculated for establishments because these numbers are not suppressed for confidentiality reasons in any MSA.

The three scatter plots with linear prediction overlay illustrate the point. Here we see the impact on regional house prices or regional unemployment of house shocks or supply shocks. In Figure 4 we see that as the location quotient for finance and insurance establishments increases, the response of house prices to a house shock tends to decrease. Figure 5 indicates by contrast that as the MSA exhibits greater concentration of real estate and property management establishments, the response of house prices to a house shock tends to increase. Finally, Figure 6 reveals that as the location quotient for construction industry establishments increases, the fourth quarter response of unemployment to supply shocks increases.

We can of course move beyond this bivariate analysis and examine the relationship of location quotients

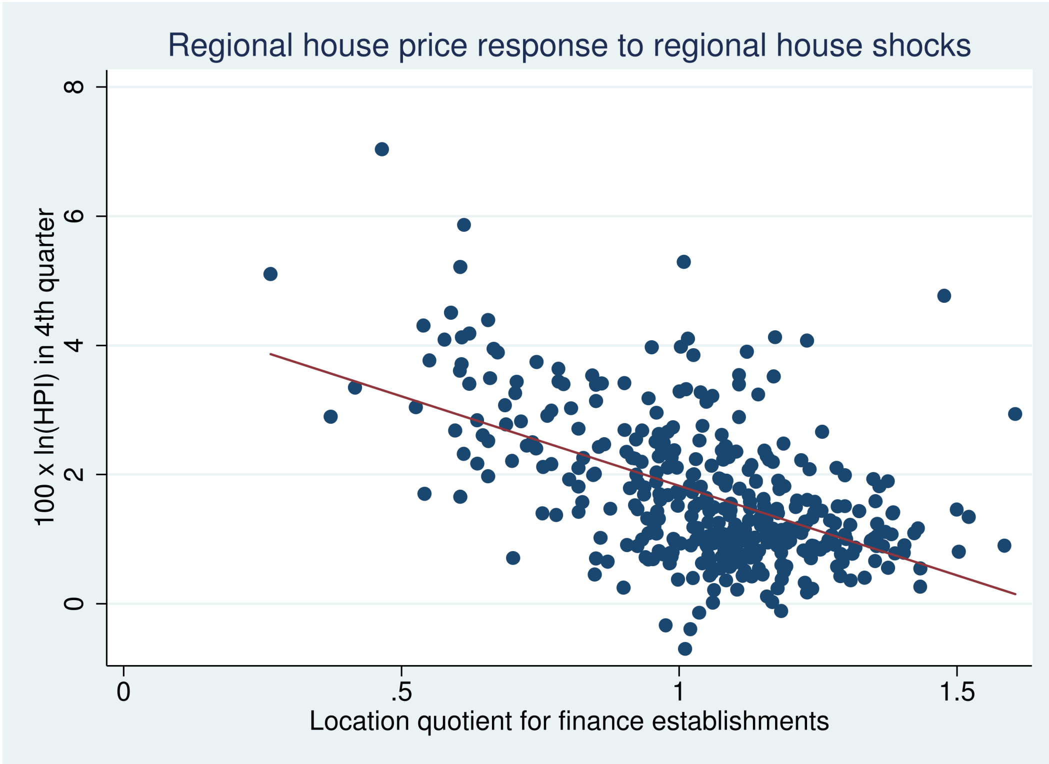

Figure 4: Impact of finance and insurance on Q4 response to house shocks

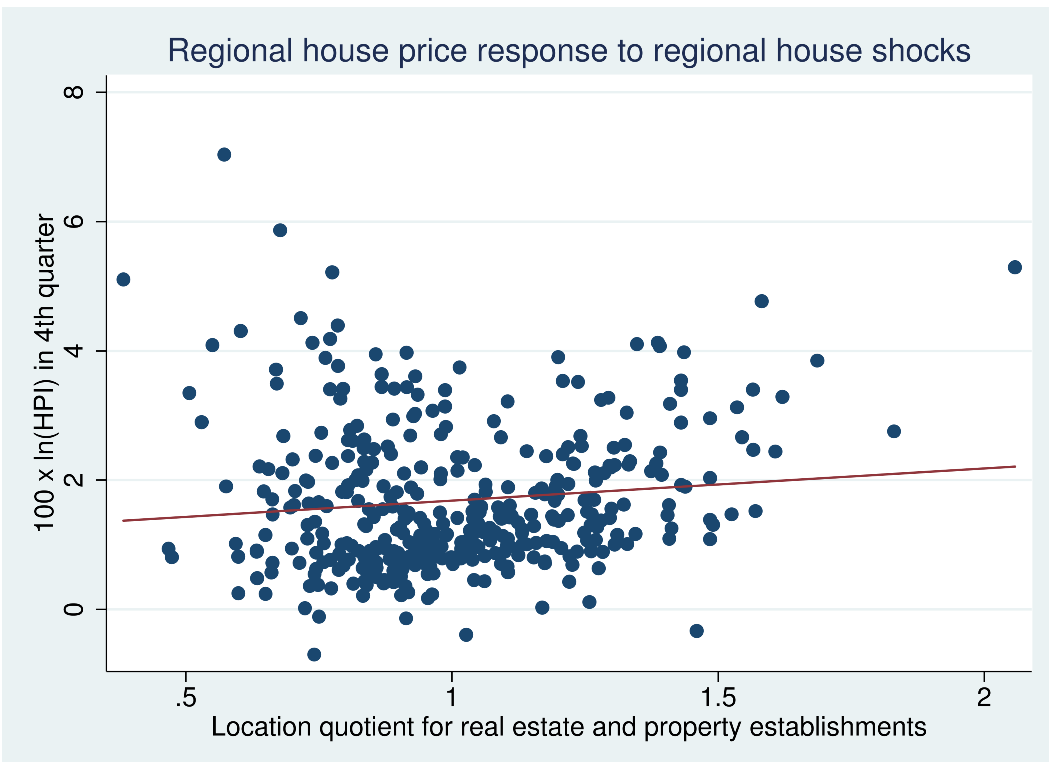

Figure 5: Impact of real estate and property management on Q4 response to house shocks

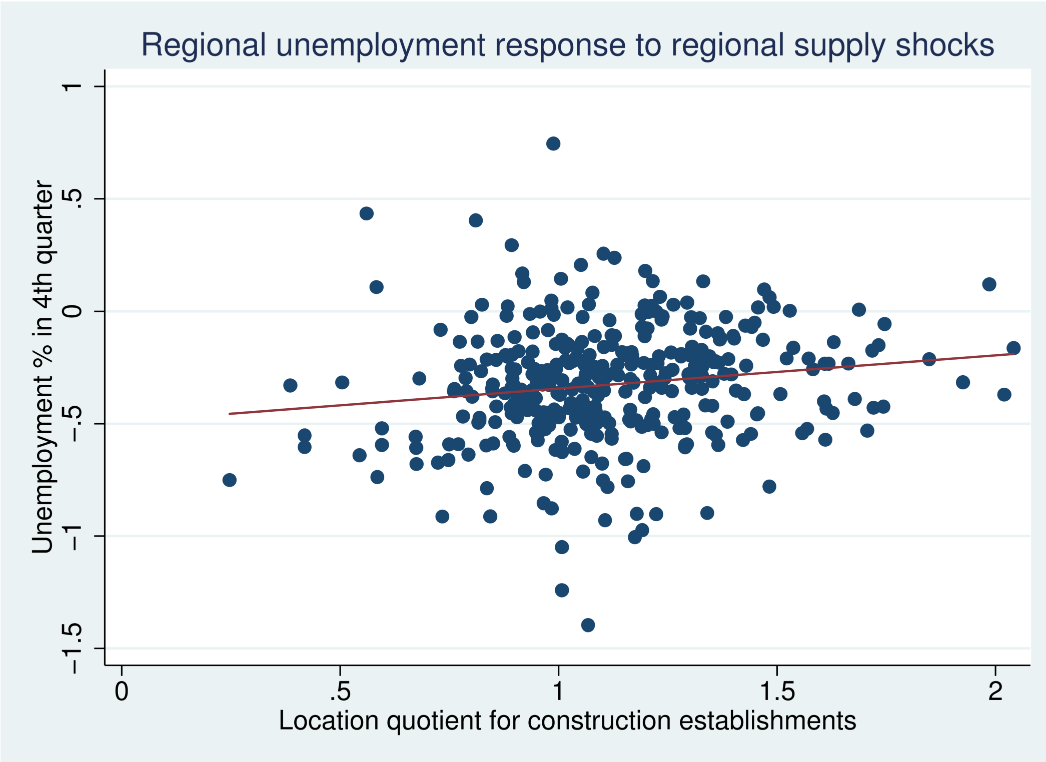

Figure 6: Impact of construction on Q4 response to supply shocks

	Industry	**Impacts on Y** β	**Impacts on HPI** β	**Impacts on UE** β
Supply	Agriculture	0.029**	0.076***	0.008
Shock	Mining	0.044***	-0.005	-0.001
	Utilities	-0.005	-0.381***	0.115***
	Construction	0.067	-0.03	0.148***
	Manufact	0.079	-0.234**	-0.137***
	Wholesale	0.192*	0.144	-0.243***
	Retail	-0.257**	-1.434***	0.15***
	Transport	-0.118*	-0.473***	0.092***
	Information	-0.181	-0.12	0.106**
	Fin/Ins	-0.204	-1.078***	0.01
	RealEstate	0.156	0.302*	0.111**
	ProfServ	0.064	0.508***	-0.028
	Management	-0.108**	0.11	-0.028
	AdminServ	-0.056	0.663***	-0.009
	Education	-0.151	0.862***	0.048
	Health	-0.341***	-0.685***	0.089**
	Arts/Cult	-0.112	0.018	0.049
	Hospitality	-0.307**	-1.012***	0.149***
	OtherServ	0.045	0.435***	-0.061***

*** - significant at 1%, ** - significant at 5%, * - significant at 10%

Table 7: Impact of industry structure on response to supply shock

	Industry	Impacts on Y β	Impacts on HPI β	Impacts on UE β
House	Agriculture	-0.012	0.162***	0.001
Shock	Mining	-0.005	-0.015	0.001
	Utilities	0.029	-0.445***	-0.026
	Construction	0.154	-0.155	-0.069**
	Manufact	-0.313***	-0.952***	0.041**
	Wholesale	-0.276**	-1.459***	0.049
	Retail	-0.072	-2.34***	-0.12***
	Transport	-0.052	-0.732***	-0.038*
	Information	-0.075	-0.594***	0.064*
	Fin/Ins	-0.37***	-2.772***	0.015
	RealEstate	0.362***	0.498**	-0.048
	ProfServ	0.069	0.191	0.027
	Management	-0.03	-0.199**	0.037**
	AdminServ	0.044	-0.37*	0.047
	Education	-0.231	0.208	0.097*
	Health	0.039	-1.025***	0.01
	Arts/Cult	-0.014	-0.109	0.029
	Hospitality	0.071	-1.321***	-0.032
	OtherServ	0.063	1.098***	0.017

*** - significant at 1%, ** - significant at 5%, * - significant at 10%

Table 8: Impact of industry structure on response to housing market shock

	Industry	Impacts on Y β	Impacts on HPI β	Impacts on UE β
Other	Agriculture	-0.008	-0.011	0.017***
Shock	Mining	0.004	0.019***	0.01***
	Utilities	-0.011	0.087**	0.022
	Construction	0.027	0.333***	-0.111***
	Manufact	0.012	0.104**	-0.018
	Wholesale	0.131**	0.165**	-0.088***
	Retail	-0.157**	0.315***	-0.009
	Transport	-0.065**	0.162***	-0.008
	Information	0.212***	0.014	-0.238***
	Fin/Ins	0.126**	0.34***	-0.17***
	RealEstate	0.013	0.2**	-0.072**
	ProfServ	0.075	-0.041	-0.124***
	Management	0.06**	-0.014	-0.05***
	AdminServ	0.09	0.118	-0.183***
	Education	0.157**	-0.243**	-0.234***
	Health	-0.052	0.075	-0.042
	Arts/Cult	0.032	0.001	-0.08***
	Hospitality	-0.02	0.063	-0.103***
	OtherServ	0.001	-0.192***	0.074***

*** - significant at 1%, ** - significant at 5%, * - significant at 10%

Table 9: Impact of industry structure on response to other shock

in all of the industries to output, house price, and unemployment responses to supply, house and other shocks. The results are summarized in Tables 7, 8 and 9 .

6 Interpretation and comments

Our analysis has presented the implications of a structured panel VAR approach to understanding house price and output dynamics in US metropolitan areas. We use quarterly data from all 384 MSAs covering a period that for some cities extends for 32 years back to the beginning of 1978. We began by observing that contrary to the impression that might be formed from literature that deals with "the US housing bubble" the actual experience of us cities is highly varied. Some have experienced tremendous increases in house prices followed by spectacular declines during 2007-2009. Others experienced very modest increases and have seen only modest declines or even growth in prices during the "great recession." This pattern of difference in house price dynamics appears to have a clear regional structure, and its underlying causes are interesting.

Our analysis has established that there are clear differences between cities in coastal states and cities in non-coastal states. These differences can be characterized as:

1. While steady-state impact on output from house shocks is zero in both regions, the dynamic path of adjustment is different, with more cities exhibiting short run positive output response to a house shock in the coastal region than in the non-coastal region.

2. House prices in the coastal region are generally more responsive to both supply shocks and house shocks. While the dynamic patterns of house price response are similar, the magnitude of response is 1.5 to 2 times larger in the coastal region than in the non-coastal region.

In addition, our analysis has revealed general support in both regions for labor market linkages with housing markets. A house shock that increases house prices reduces unemployment in the majority of cities. This does not appear to be a transitory phenomenon but remains evident five years after the shock has occurred. This provides general support for the so-called Oswald hypothesis in US cities.

Finally, we have shown that the short run responsiveness to shocks is dependent on the industrial structure of the urban area. House price responses and unemployment responses, in particular, seem to depend significantly on industrial concentration.

The analysis presented here leaves many questions unanswered. What are the specific differences between coastal and non-coastal regions that might explain the clear differences in dynamics that exist? Are there demographic, regulatory, or other differences that explain why a house shock in the coastal region generates a much larger house price response? Finally, what structural or other MSA characteristics influence the strength of linkages between housing and labor markets? Can we identify the features that determine whether a house shock in a city is likely to induce a steady-state drop in unemployment?

These and others are questions with important policy implications. It is clear that the structured panel VAR technique applied in this analysis can be an important tool for finding the answers.

References

Blanchard, O. J. & Quah, D. (1989), 'The dynamic effects of aggregate demand and supply disturbances.', *American Economic Review* **79**(4), 655–673.

Calhoun, C. A. (1996), Ofheo house price indexes: Hpi technical description, Technical report, Office of Federal Housing Enterprise Oversight, 1700 G Street, NW, Washington, DC 20552. Downloaded on November 5, 2010 from http://www.fhfa.gov/webfiles/896/hpi_tech.pdf.

Coleman, M., LaCour-Little, M. & Vandell, K. D. (2008), 'Subprime lending and the housing bubble: Tail wags dog?.', *Journal of Housing Economics* **17**(4), 272–290.

Coulson, N. E. & Fisher, L. M. (2009), 'Housing tenure and labor market impacts: The search goes on.', *Journal of Urban Economics* **65**(3), 252–264.

Glaeser, E. L., Gyourko, J. & Saiz, A. (2008), 'Symposium: Mortgages and the housing crash: Housing supply and housing bubbles.', *Journal of Urban Economics* **64**(2), 198–217.

Himmelberg, C., Mayer, C. & Sinai, T. (2005), 'Assessing high house prices: Bubbles, fundamentals and misperceptions.', *Journal of Economic Perspectives* **19**(4), 67–92.

Hughes, G. A. & McCormick, B. (1987), 'Housing markets, unemployment and labour market flexibility in the uk.', *European Economic Review* **31**(3), 615–641.

Hughes, G. & McCormick, B. (1981), 'Do council housing policies reduce migration between regions?.', *Economic Journal* **91**(364), 919–937.

Munch, J. R., Rosholm, M. & Svarer, M. (2008), 'Home ownership, job duration, and wages.', *Journal of Urban Economics* **63**(1), 130–145.

Pedroni, P. (2008), The econometrics of a panel svar. Williams College Working Paper.

Shiller, R. J. (2005), *Irrational Exuberance*, second edn, Princeton University Press, Princeton, NJ.

7 Appendix

U.S. Complete Metro Area Variance Decompositions

*LR Triangular SVAR: [lnY*100, lnP*100, Unemp %]` = A(1)*[SUPPLY, HOUSE, OTHER]`*

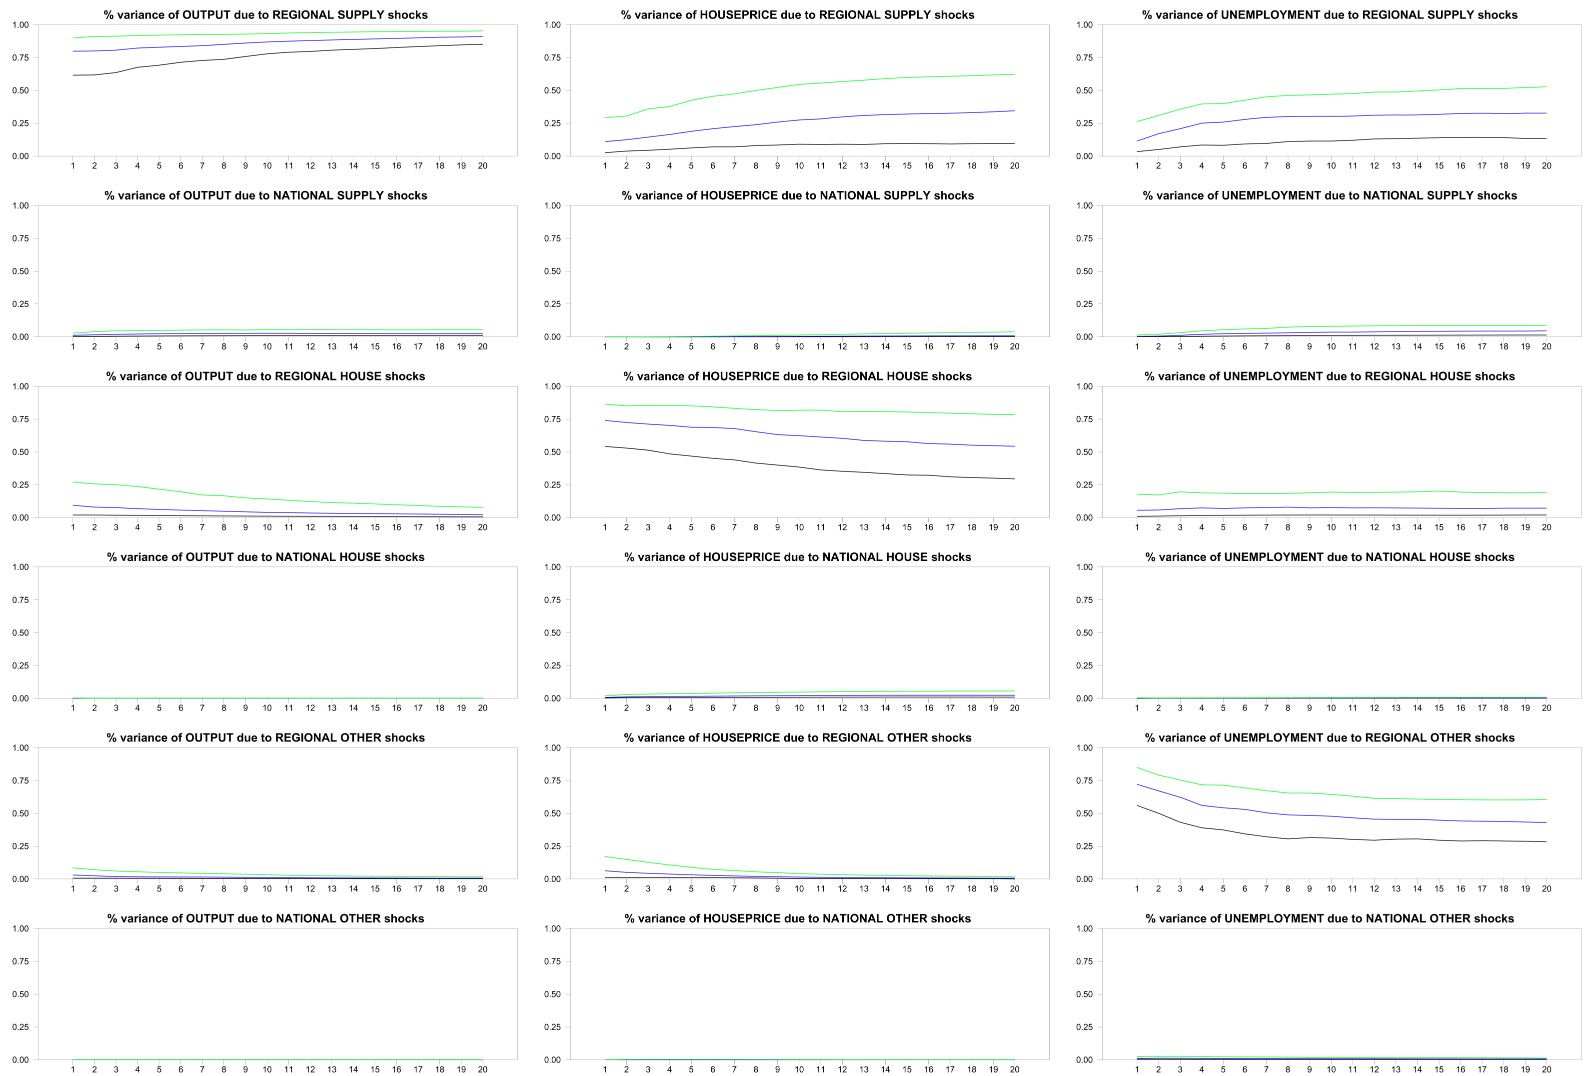

U.S. Coastal Metro Area Variance Decompositions

*LR Triangular SVAR: [lnY*100, lnP*100, Unemp %]` = A(1)*[SUPPLY, HOUSE, OTHER]`*

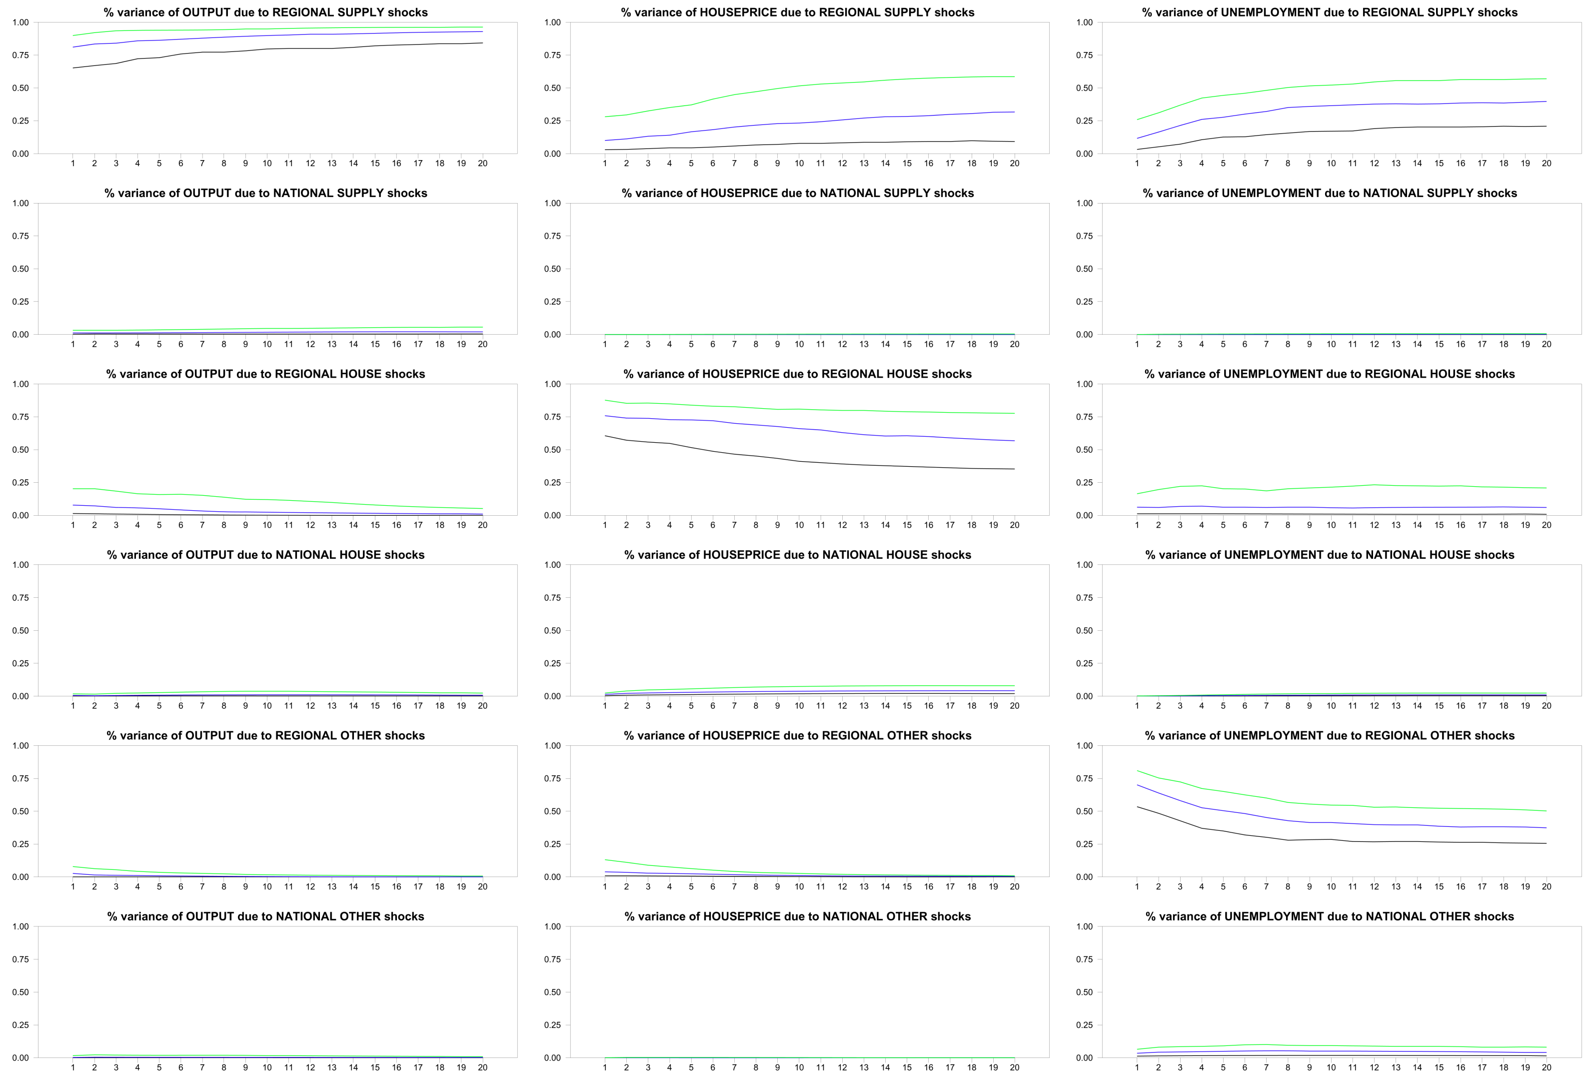

U.S. NonCoastal Metro Area Variance Decompositions

*LR Triangular SVAR: [lnY*100, lnP*100, Unemp %]` = A(1)*[SUPPLY, HOUSE, OTHER]`*