

ASSET PRICES, LOCAL PROSPECTS AND THE GEOGRAPHY OF HOUSING
DYNAMICS

Preetesh Anand Katak

A dissertation submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Finance at the Kenan-Flagler Business School.

Chapel Hill
2017

Approved by:

Gregory Brown

Riccardo Colacito

Wayne Landsman

Christian Lundblad

Jacob Sagi

©2017
Preetesh Anand Katak
ALL RIGHTS RESERVED

ABSTRACT

Preetesh Anand Kantak: Asset Prices, Local Prospects and the Geography of Housing Dynamics
(Under the direction of Riccardo Colacito)

I uncover a novel link between economic fundamentals and real estate returns. First, using a measure of “local” (i.e. MSA) growth prospects from both industry level employment and financial data, I am able to explain a substantial portion of variability in excess real estate returns and price-to-rent ratios. A one standard deviation decline in local growth prospects is associated with a 120bp drop in the housing risk premia and a 500bp increase in the price to rent ratio. Second, I establish the underlying dynamics of relative, housing versus non-housing, consumption in the MSA-level cross section of economic agents. I document that a one standard deviation decline in local growth prospects is associated with a 1000bp rise in level and 700bp decrease in variance of relative consumption. In contrast, a one standard deviation increase in local growth prospects is also associated with a rise in relative consumption. This suggests that housing acts as a hedge against long-run economic growth. I then investigate these empirical facts through the lens of a consumption based, asset pricing model. I combine a persistent component in consumption, Epstein and Zin (1989) preferences, and nonseparable consumption of non-housing goods and housing services. Consumption of housing services provides a hedge to long-run growth prospects. The model is able to replicate the excess returns and price to rent ratio regression results from the data.

To my wife and best friend. You may not believe it, but I couldn't have done any of this without you.
Thank you for your support and understanding along the way.

ACKNOWLEDGEMENTS

Invaluable suggestions and advice were provided by all members of my committee. Without them, I would not have completed this work. Additional important insight was provided by Robert Connolly, Jennifer Conrad, Mariano Croce, Eric Ghysels and Adam Reed. I'd also like to thank seminar participants at Kenan-Flagler Business School and the 2016 IPC Real Estate Conference at the University of North Carolina. All errors are my own.

TABLE OF CONTENTS

LIST OF TABLES	viii
LIST OF FIGURES	ix
1 INTRODUCTION	1
2 MOTIVATIONAL EVIDENCE	8
2.1 Low Frequency Cross-sectional Evidence	8
2.2 Time-series Evidence	9
3 LITERATURE	11
4 DATA AND VARIABLE CONSTRUCTION	13
4.1 National Data Computation	13
4.2 Price-to-dividend Computation	14
4.3 Validity of Price-to-dividend	16
4.4 Other Independent Variables	19
4.5 Construction of Local Variables	20
5 HOUSING DYNAMICS ANALYSIS	23
5.1 Return Regression	23
5.2 Robustness	25
5.3 Capitalization Rate Regressions	27
6 ECONOMIC FOUNDATIONS	31
6.1 Setup of the Economy	31
6.2 Financial Markets and Equilibrium of the Economy	34

6.3	Calibration of the Model	36
6.4	Numerical Results	39
6.5	Further Empirical Implications	42
7	IDENTIFICATION OF DURABILITY	44
7.1	Level Regression	45
7.2	Level Regression Robustness	47
7.3	Variance Regression	49
8	CONCLUSION	53
	APPENDIX A Gauss-Newton Conversion Algorithm	55
	APPENDIX B Dividend and Repurchase Yield	58
	APPENDIX C Annual GMM	59
	APPENDIX D Extending Annual GMM to Quarterly	60
	APPENDIX E Campbell-Shiller Decomposition	63
	APPENDIX F Cash Flow Component Extraction	64
	APPENDIX G Solving Model's Equilibrium	66
	APPENDIX H Pricing of Aggregate Markets	68
	APPENDIX I Simulation Equations	70
	APPENDIX J Estimation of Composition Variance	71
	REFERENCES	72

LIST OF TABLES

Table 1	Predictive Regressions	10
Table 2	Pre-regression, Conversion Statistics	15
Table 3	GMM and Business-cycle Parameter Summary	18
Table 4	Panel Summary Statistics	22
Table 5	Return Regression Results	24
Table 6	Return Regression Robustness Results	26
Table 7	Capitalization Rate Regressions	29
Table 8	Model Calibration	36
Table 9	Aggregate Calibration Regressions	38
Table 10	Identification Regressions	46
Table 11	Identification Robustness Regressions	48
Table 12	Identification Auxiliary Regressions	51

LIST OF FIGURES

Figure 1	Industry Concentration and Cross-sectional Evidence	7
Figure 2	Equilibrium Prices and Empirical Replication	41
Figure 3	Estimated $Var_t(sc_{t+1})$ versus x_t	52

1 INTRODUCTION

This paper shows that the cross-section and time-series of long-run industry growth prospects is a key factor in explaining the geographic differences in home prices. I use a market based measure of growth prospect to quantify the *unique* relationship between prospects and housing returns. I document that a decline in local growth prospects is associated with a decrease in the expected housing risk premia and an increase in the price-to-rent ratio. I attribute these finding to the hedging characteristics of housing assets and replicate them in a consumption based asset pricing model in which (1) investors care about the temporal distribution of uncertainty, (2) consumption of non-housing goods and housing services are exposed to time-varying, long-run growth prospects, and (3) housing services provide a hedge against negative news about expected long-run growth. In addition, I show strong empirical evidence of these hedging characteristics in the geographic cross-section of consumption data.

Real estate is a significant portion of aggregate wealth in the United States. The total value of household real estate is approximately \$25 trillion, which is greater than the almost \$23 trillion households have in corporate equity. Even after netting this gross figure by the \$9.3 trillion in household mortgage liabilities, the portion of wealth concentrated in real estate is staggering. U.S. GDP figures show a similar picture; expenditures on shelter account for almost 20% of the nation's annual consumption (Federal Reserve Q4 2015 statistical release). These aggregate figures, however, belie the inherent geographic differences in real estate wealth across the country (see Gyourko, Mayer and Sinai (2013)).

Urban economists have shown that industries cluster for a wide array of reasons - e.g., close proximity to suppliers or purchasers (see Rosenthal and Strange (2004) for overview of agglomeration economies). Likewise, economists have studied how time variation in industry-specific productivity drive variation in long-term economic growth (see Hsieh and Klenow (2009) and Bartels-

man, Haltiwanger and Scarpetta (2013) for recent work). It follows from this evidence that growth prospects of the local industry mix should be reflected in geographic differences in economic growth and price dynamics of local (segmented) asset classes. Housing is the segmented asset I use for my empirical analysis; its value is intimately tied to its location. For example, the structural transformation of the United States from a largely manufacturing to service based economy has played out differently over long-horizons between the industrial Midwest and the globally integrated, finance and technology hubs on the coasts. My analysis shows that these differing economic fundamentals have had important implications for local housing returns. This paper seeks to understand how house prices move in the context of its role as both a *local* consumption asset (i.e. provides access to *local* growth opportunities) and the primary store of wealth for the median household. As such, my findings have important implications for, inter alia, the transmission of monetary and fiscal policy to inclusive and broad sectoral economic growth, and long-run income inequality.

I first generate a formal measure of heterogeneity in long-run expectations between different regions; its construction depends on two assumptions. First, industry level price-to-dividend (PD) ratios from a cross-section of 61 industries contain information about common and industry specific growth opportunities. As shown in Campbell and Shiller (1987), with all else equal, dividends of high PD-ratio equities are either expected to grow faster or be less uncertain than those of low PD-ratio equities. My results show there is valuable economic information in how these industry PD-ratios vary through time and in the cross-section. Second, local growth prospects should reflect the growth opportunities of the industries employing locally. Utilizing a procedure first developed by Bartik (1991), I weight industry prospects (PD-ratios) taken from the aggregate market by the industry's MSA (metropolitan statistical area) employment share to obtain a measure of local prospects. Many important industries are highly concentrated - e.g., computer design and finance (see figure (1A) for more detail). To the degree that an individual's wealth (human capital and real estate) is tied to their location, it is natural to think that the consumption choices and asset holdings of individuals living in areas where local industries have strong growth prospects will

be different from those that do not. These two assumptions are similar to those made by Bekaert, Harvey, Lundblad and Siegel (2007). In their study they compute measures of country-level growth opportunities by weighting industry price-to-earnings (PE) ratios to understand the cross-section of emerging market returns. In the same spirit, I weight industry PDs by local employment share to understand the cross-section of MSA-level asset returns.

I run predictive regressions to show that my measure of local prospects explain a substantial portion of variability in the cross-section and time-series of excess housing returns. I find that a one standard deviation increase in my proxy of local prospects leads to a statistically significant 120bp increase in annual excess housing returns. In addition, I decompose my measure of local prospects into a common (all-MSA or global) and orthogonal (local) component. A one standard deviation increase in the global, and local component leads to an statistically and economically significant 150bp and 35bp increase in annual excess returns, respectively. Surprisingly, although my industry-level measures are taken from the aggregate markets, their impact is geographically heterogeneous due to differences in MSA-level employment composition. These return predictability findings are robust to various business cycle controls from the literature (see, e.g., Abraham and Hendershott (1994); Favilukis, Kohn, Ludvigson and Van Nieuwerburgh (2013); and Tuzel and Zhang (2016)).

A present value decomposition of the price-to-rent (PS) ratio shows that PS-ratios are directly proportional to expected growth in rent and inversely proportional to expected returns (similar to Campbell and Shiller (1987) PD decomposition). The urban economics literature finds a positive link between expected economic growth and rent (see Glaeser and Gyourko (2005)). My predictive regression results also indicate a positive link between expected economic growth and returns. Given the opposite sensitivities of PS-ratios to each component, I test which component dominates by contemporaneously regressing MSA-level PS ratios onto my measure of local prospects. I estimate the PS ratio as the inverse of surveyed cap-rates (expected operating income dividend by net asset value) from local (62 MSAs) developers of multi-family residential properties from Integra Realty Resources (IRR). I find that a one standard deviation increase in

local prospects lowers the log price-to-rent by more than 500bps. The negative coefficient implies that the expected return dominates the expected rent channel.

These risk based results stand in contrast to the findings of traditional predictive regressions. In the equity market better prospects, which are typically associated with high PD-ratios, usually imply less risk and therefore lower expected excess returns going forward (see Campbell and Shiller (1988); Cochrane (2008); Lewellen (2004)). The key difference between housing and other consumption assets, however, is the relative durability of the services that it provides. Housing is an *extremely* durable asset: when local prospects are good they strongly influence plans to build infrastructure and locate close to suppliers of capital and labor (i.e. agglomeration economies). Once in place these pieces are both permanent in nature and geographically specific. For example, New York City and San Francisco will likely continue to be centers of finance and technology regardless of higher frequency economic fluctuations. Thus, in the context of risk to housing, this permanence acts as a hedge against negative shocks to long-run local growth prospects. Urban economists have argued that the durability of *physical structures* is key to understanding many multi-decade phenomena in urban and labor economics (see Glaeser and Gyourko (2005) and Notowidigdo (2013)). My argument is similar regarding the durability of the location specific *services* housing provides.

To formalize how this characteristic impacts the profile of housing returns, I propose a consumption based model that is centered around the problem of a representative MSA. The MSA is inhabited by an investor with Epstein and Zin (1989) recursive preferences defined over housing and non-housing consumption goods. I specify endowment processes for the growth rate of consumption of non-housing and housing services of the representative MSA. Both endowments are exposed to a highly persistent state variable. However, while the growth of the non-housing endowment is symmetrically exposed to growth prospects - i.e. the endowment both conditionally expands and contracts with the state variable - the growth rate of the endowment of housing services is durably exposed. That is, housing services expand when prospects are good, but do not contract when prospects are bad.

First, recursive preferences are required to price persistent dynamics, such as those of my measure of local prospects (see inter alia Bansal and Yaron (2004), Colacito and Croce (2011, 2013)). Second, the persistence of my state variable contains important information in that its fluctuations embed long-horizon implications. This feature has immediate intuitive appeal when it comes to the pricing of long-lived assets such as housing. The asymmetric exposure to prospects reflects the hedging quality of housing services; once they are “built” they are difficult to take away. In addition, both non-housing and housing endowments have short-run i.i.d. shocks. These shocks reflect higher frequency business-cycle or MSA-level idiosyncratic shocks.

This simple setup means that the endowment of housing is exposed to two sources of risk when growth prospects are high (short- and long-run), but only one when growth prospects are low (short-run). The *expected* risk to housing thus becomes a function of the conditional probability of being in a high expected growth state tomorrow. In addition, as growth prospects are highly persistent, this probability is increasing in the expected growth rate of consumption, giving the return profile characteristics of a regime switching model. I then replicate my empirical results by simulating 375 MSAs over 40yrs and conduct both the predictive excess return and contemporaneous PS-ratio regressions on generated data. My average coefficients are close to their empirical counterparts. A one standard deviation higher growth prospect leads to a 110bp increase in expected return. The model illustrates how durability of housing services with respect to long-run growth prospects generates time varying volatility in the ratio of housing services versus non-housing consumption (henceforth relative consumption), which has immediate implications for return and pricing dynamics.

Finally, I provide empirical support for the key assumption of my model: the asymmetric exposure of housing services to long-run growth prospects. I use the BLS consumer expenditure survey (CEX) data to test for these dynamics in the cross-section of MSAs. The evidence of asymmetry is statistically significant and economically large. During bad economic times (i.e. growth prospects are below their long-run mean), a one standard deviation fall in prospects leads to a 1000bp rise in relative consumption. During good economic times, a one standard deviation rise

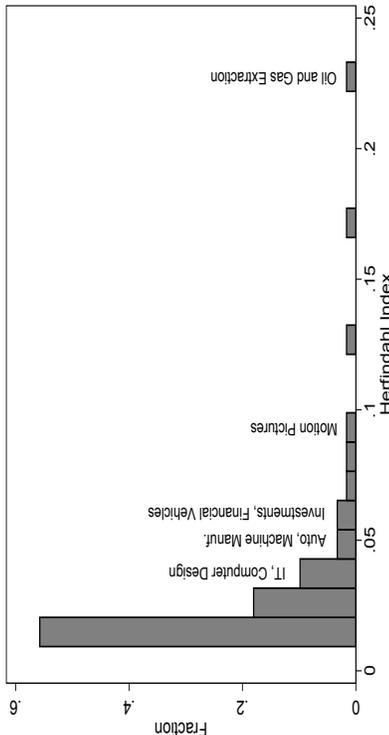
in prospects leads to a much smaller rise in relative consumption. This suggests that the local households raise consumption of housing services in lock step with non-housing consumption during times of good long-run economic prospects, but do not drop them as much as non-housing consumption during bad times. This evidence is not present in renter consumption data; I find that a one standard deviation fall and rise in long-run growth prospects lead to an equally large rise and *fall* of relative consumption. This suggests that “sticky” contracting issues are not the primary driver of my results. I also confirm the profile of time varying volatility in the data. The variance of relative consumption increases by more than 700bps from the low to high variance regime. Given the high persistence of the underlying economic state variable, this increase in variance has a long-horizon impact on housing consumption decisions. This verifies the model’s primary risk channel.

The paper is organized as follows. Section (2) develops my research question using conditional correlations between local economic growth prospects, output and housing returns. Section (3) describes the contribution of my work to existing literature. Section (4) describes the data used and variables constructed for my analysis. Section (5) has my first set of empirical results, which motivates the connection between geographic heterogeneity in long-run prospects and excess housing returns. Section (6) presents my model and quantitative results. Section (7) identifies durability and the time-varying volatility channel in the cross-section of consumption data. Section (8) concludes.

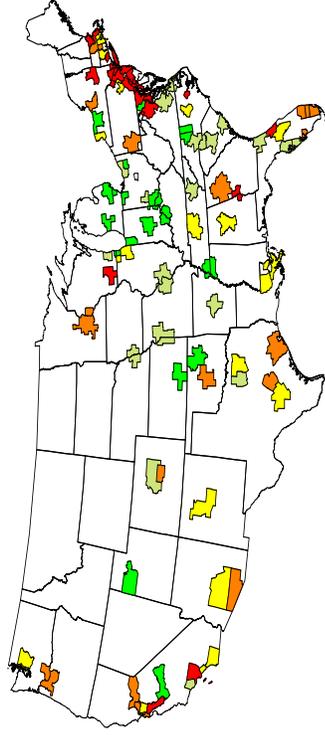
Figure 1: Industry Concentration and Cross-sectional Evidence

Panel 1A captures the distribution of Herfindahl-Hirschman Indices (HHI) across 61 industries and 375 MSAs on December 31st, 2011. Employment share for each industry i and MSA a is defined as $s_{i,a}$. Panel 1B projects the average MSA-level employment share weighted PD-ratios from 1991-2000 for the 100 most populated MSAs. Panel 1C projects MSA-level real per-capita GDP from 2001-2010. Data is from The Bureau of Economic Analysis (BEA). Panel 1D projects MSA-level average housing excess price appreciation from 2001-2010. Data is from the FHFA. The colors (green-yellow-red scale) correspond to five equally sized buckets ranked from the lowest to highest 20% of each variable.

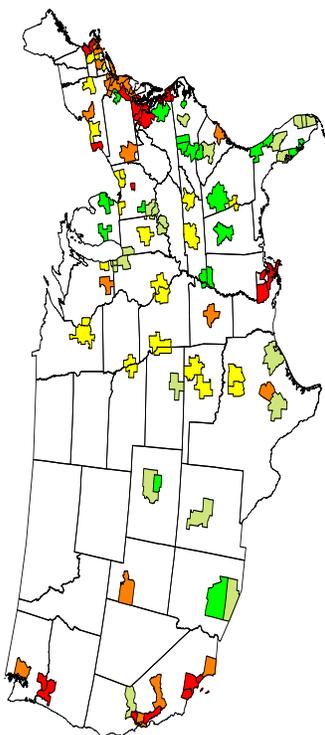
A: 2011 Herfindahl-Hirschman Indices



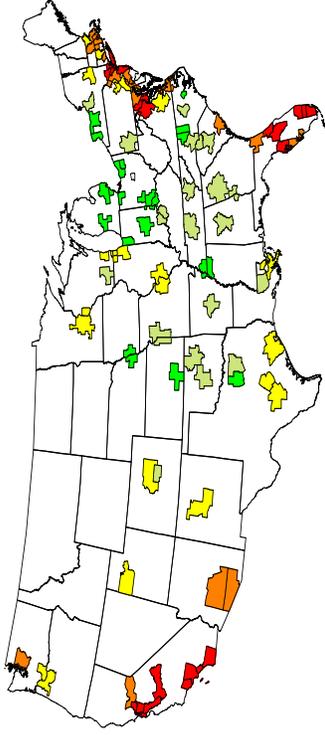
B: 1991-2000 Employment Weighted PD-ratio



C: 2001-2010 Real Per-capita Output



D: 2001-2010 Excess Price Appreciation



2 MOTIVATIONAL EVIDENCE

I use the aggregate industry PD ratio to obtain measures of industry-specific growth opportunities for two reasons. First, I would like my measure to reflect heterogeneity in prospects only because of geographic differences in industry clustering. Second, given its durable nature and role as both a consumption good and component of wealth, housing is exposed in a fundamentally different way than stocks, both in magnitude and direction, to long-run sources of economic risk. It's important that my measure captures this difference. In this section, I'll show evidence that my measure of long-run growth prospects satisfies both requirements.

2.1 Low Frequency Cross-sectional Evidence

Location specific measures of prospects (e.g., price-to-rent) are 1) limited in availability and quality, and 2) confounded by geographic-specific information orthogonal to long-run growth prospects (e.g., weather or regulatory factors). To better motivate the specific type of geographic risk in which I am interested, I conduct a few conditional correlation tests using my main measure of local prospects. In figure (1B) I project the average (1991-2000) local, employment share weighted PD-ratios onto the map of the continental United States. I split their values into five equally sized buckets (green-yellow-red scale). The highest PD-ratios occur on the technology and banking heavy coastal areas, whereas the lowest PD-ratios occur in the industrial Midwest. In figure (1C), I show the subsequent real per-capita output in the same MSAs over the subsequent decade (2001-2010). The intra-decade spatial correlation of the two cross-sectional samples is almost 50% and significant at a 5% level. This suggests that my localized measure of prospects does a good job of predicting long-run future output growth. In figure (1D), I do the same for average MSA-level housing excess price appreciation for 2001-2010. The coasts and Midwest also have the highest and

lowest excess price appreciation, respectively. The intra-decade spatial correlation between these returns and the local average PD from 1991-2000 is more than 60% and significant at a 5% level. There seems to be strong association between geographic heterogeneity in output and housing returns and differences in industry prospects and employment composition.

2.2 Time-series Evidence

Table (1) highlights the difference between housing and stocks in their exposure to long-run growth prospects using times-series regressions on aggregate indices on a quarterly basis. Column (1) is the classic Campbell and Shiller (1987) result using data from 1953-2014; the excess return of the stock market is regressed onto my economic state variable, the lagged market PD ratio (x_{t-1}^{global}). The state variable predicts excess returns on housing with a statistically significant *negative* coefficient. This results fits the intuition that high PD-ratios reflect a conditionally better economic environment, signifying lower economic risk. The marginal agent thus demands a lower risk premia in such circumstances. Cochrane (1994), Bansal and Yaron (2004), Lettau and Ludvigson (2005) and others have explored the underlying economic drivers of this result. In column (2) I regress housing excess returns onto my state variable. Aggregate economic prospects predict aggregate housing returns with a *positive* and statistically significant coefficient.

The balance of columns in table (1) are for robustness. As Ghysels, Plazzi, Torous, Valkanov (2012) point out, returns estimated from repeat sales indices are by construction highly persistent. In column (3), I add the lagged excess return to the regression in order to pre-whiten the data. The excess return measure is still positive and statistically significant. Finally, in the last two columns, I perform the same predictive regression on the REIT total return indices from the CRSP Ziman real estate data series. Columns (4) and (5) are the result of regressing excess returns from the general REIT universe and multi-family home REITS, respectively, onto my state variable. The sign remains the same although the statistical significance falls as REIT data is only available starting in the early 1980s. The key difference in risk dynamics between stocks and housing seems to be that the risk premia to housing falls when *long-run* economic prospects are bad. That

Table 1: Predictive Regressions

This table reports results from predictive regressions of excess returns for the stock and housing markets onto the lagged market price-to-dividend, my state variable of interest (x_{t-1}^{global}). The regressions are on quarterly data. Following the methodology of Boudoukh, Michaely, Richardson and Roberts (2007), my measure of dividend includes share repurchases. Model (1) is the classic Campbell and Shiller (1987) result using data from 1953-today. Model (2) regresses housing excess returns onto my state variable. Housing returns are computed from the Shiller repeat sales index available online at <http://www.econ.yale.edu/shiller/data.htm>. A national estimate of rent is then added to obtain a proper return measure (see <http://www.lincolninst.edu/>). The 3m T-bill rate is then subtracted from this measure. Repeat sales indices are highly persistent and demonstrate seasonality by construction; in model (3) I include the excess return quarter and annual lag as regressors. Columns (4) and (5) are the result from regressing excess returns on the general REIT universe and multi-family home REITS, respectively, onto the lagged market price-to-dividend ratio (x_{t-1}^{global}). The REIT indices are from the CRSP Ziman real estate total return data series. The sample size falls from the Campbell-Shiller to housing return regressions because rent and REIT data is only available after the 1980s.

	Stocks		Housing		
	(1)	(2)	(3)	(4)	(5)
x_{t-1}^{global}	-0.0347* [-1.92]	0.0241*** [3.77]	0.0081** [2.26]	0.0435 [1.39]	0.0567* [1.81]
F-statistic	3.673	14.189	67.561	1.927	3.293
p-value	0.057	0.000	0.000	0.168	0.072
Observations	236	132	128	128	128

is housing acts as a hedge against downside risks to long-run economic prospects. This result corroborates the findings from the low frequency cross-sectional correlations in the previous section, but at a much higher frequency.

3 LITERATURE

This paper extends work done in the long-run risk asset pricing literature. Bansal and Yaron (2004) show that highly persistent, low-frequency dynamics in consumption and recursive preferences (collectively known as long-run risk) are important in explaining the equity risk premia puzzle. Colacito and Croce (2011, 2013) utilize these same dynamics and preferences to better understand key international asset pricing dynamics such as the currency risk premia. My empirical specification is similar to that of Gomes, Kogan, and Yogo (2009) and Eraker, Shaliastovich and Wang (2016) in that I attribute heterogeneity in asset price dynamics to heterogeneity in long-run dynamics of different firms or industries.

This paper also relates to a strand of the macroeconomic literature that links housing to asset prices and consumption behavior. Piazzesi, Schneider and Tuzel (2007) show that non-separable composition risk between housing and non-housing consumption can explain a significant portion of the equity premium puzzle. Lustig and Van Nieuwerburgh (2005, 2009) find connections between housing collateral, consumption and asset prices. My objective is somewhat more fundamental in that I seek to understand how persistent local growth prospects impact the risk premia of home prices because of their inherent geographic heterogeneity rather than vice versa.

My geography-centric approach is largely related to work done in urban economics. Beginning with Rosen (1979) and Roback (1982), urban economists have used spatial general equilibrium models to understand geographic heterogeneity in home prices and productivity.¹ In terms of higher frequency home fluctuations, Glaeser, Gyourko, Morales and Nathanson (2014), show that a spatial equilibrium model with durable housing and lagged investment, explain certain housing return “puzzles.” Their model and calibration, however, largely ignores the importance of time-

¹See Glaeser and Gyourko (2005); Glaeser, Gyourko, and Sacks (2005); Van Nieuwerburgh and Weill (2010); Kline and Moretti (2016); and others.

varying expectations and risk aversion in prices. Hizmo (2015) blends a spatial equilibrium with a real business cycle model to estimate the housing risk-premia in different cities. While risk aversion plays an important role, agents ignore conditional dynamics of expectations when making decisions on how much housing services to consume.

Finally, my empirical analysis is related to the growing literature connecting location to economic activity and fundamentals. Dougal, Parsons and Titman (2015) illustrate the strong relationship in investments between firms in close geographic proximity. Engleberg, Ozoguz and Wang (2010) show strong comovement in fundamentals of firms in the same industry cluster. More recently Tuzel and Zhang (2016) explore how the cyclical nature of the local economy influences real estate and asset prices. My work differs from theirs in the “frequency” of fluctuations in which I am interested. While they primarily explore business-cycle dynamics, I attempt to capture the effects of much longer economic cycles on house prices.

4 DATA AND VARIABLE CONSTRUCTION

In this section I present the sources of my data and the methods used to validate my estimate of long-run prospects. I first introduce my aggregate measure of industry output. I then extract my industry-level measure of prospects and show that it strongly predicts the industrial cross-section of real per-capita output. I weight this measure of prospects using MSA-level employment composition data to obtain a local measure of long-term growth prospects. My methodology follows that of Bekaert, Harvey, Lundblad and Siegel (2007). In keeping with an established literature, I occasionally refer to these prospects as long-run risk (see Bansal, Kiku and Yaron (2012) and Colacito and Croce (2011) for references).

4.1 National Data Computation

My measure of aggregate industry output is from The Bureau of Economic Analysis (BEA) - US-wide (aggregate) value-added and employment data for approximately 60 broad industries running from 1947 to today. The data is primarily at the 3 digit NAICS (North American Industry Classification System) or 2 digit SIC (Standard Industry Classification) level. The NAICS data is available from 1977 to today and SIC data from 1947 to 1997. I convert this data into a single classification system to properly identify the strengths of my measure of long-run industry prospects (see section (4.3)). In order to best capture the more contemporary trends in GDP and employment away from manufacturing towards information technology and services, I convert the older (SIC) to the newer (NAICS) system.

Unfortunately, there is not a 1-to-1 conversion between BEA-available SIC and NAICS accounts. I utilize the 1-to-N crosswalk between 3-digit SIC and 6-digit NAICS codes available from the Bureau of Labor Statistics (BLS) to conduct my conversion for non-overlapping years for both

output and employment (see APPENDIX A for details on my conversion algorithm). Table (2A) provides the summary statistics of the NAICS to SIC conversion weights. I now have an estimate for value-added GDP and employment from 1947-2012 for a single industry classification. The value-added GDP data is deflated by industry specific deflators. A *per-capita* measure is computed by dividing each industry's value-added GDP by its employment figure.

4.2 Price-to-dividend Computation

The regressions presented in section (5) require an estimate of the price-dividend (PD) ratios for the same 61 NAICS industries for which I now have value-added GDP. I compute PD ratios using data from the Center of Research in Security Prices (CRSP).² Following Boudoukh, Michaely, Richardson and Roberts (2007), I also include share buybacks in my measure of dividend (see APPENDIX B for details). Including this payout is critical in comparing industry prospects due to heterogeneity in repurchases yields (RP_t) across industries. This is similar to what is done by Gomes, Kogan and Yogo (2009), although on a more disaggregate industry basis.

For my analysis, I require both annual and quarterly PD ratios. Given its monthly frequency, the dividends from CRSP are simply summed on both basis. For example, if there are N companies within industry i , dividends for annual or quarterly t are summed if they are in $(t - m, t)$, where m is 11 and 2 for each frequency, respectively. The dividend yield is thus

$$DP_t^i = \frac{\sum_{k=t-m}^t \sum_{j=1}^N (ShrOut_{j,k}) (Div_{j,k})}{\sum_{j=1}^N (ShrOut_{j,t}) (Px_{j,t})}.$$

For CRSP data, quarterly seasonality is minimal so this simple change in index is sufficient to obtain accurate yields at both frequencies. This is not the case with repurchases, where taxes and regulatory considerations affect timing of company actions. For my quarterly projections, I thus evenly divide a given year-industry's annual repurchases over the four quarters. My PD ratio measure is the inverse of dividend and repurchase yield, $1 / (DP_t + RP_t)$.

²See van Binsbergen and Kojen (2010) for references. CRSP distribution codes are restricted to all cash ordinary (code 1xxx) and liquidation dividends (codes 2xx2/2xx8).

Table 2: Pre-regression, Conversion Statistics

I use the crosswalk between 3-digit SIC and 6-digit NAICS codes from the Bureau of Labor Statistics (BLS) to convert SIC to NAICS GDP. Indicators are given to each correspondence and then rolled-up to the BEA-available primarily 3-digit SIC level. The first line of Panel A provides summary statistics of how many NAICS industries each SIC account converts. The BEA SIC and NAICS GDP and employment data also have common years from 1977 to 1997. As detailed in APPENDIX A, using this 1-to-N correspondence and the common years, I use an iterative algorithm to estimate SIC-to-NAICS conversion weights for the years NAICS data is not available. The balance of Panel A provides the summary statistics of these estimated weights over various subsamples. Panel B are the summary statistics of my price-to-dividend ratio from 1947-2011. Cash dividend (distribution codes 1xxx and 2xx2/2) yields are estimated from the full CRSP database - I sum the dividends paid by a given industry (for PD_t^{ind}) or the "market" (for PD_t^{mkt}) over either the quarter or year and then divide by the total market capitalization on the last month of said period. Intra-period dividends are reinvested at the risk-free rate. Repurchase yields are computed from COMPUSTAT data. I compute repurchases following Boudoukh, Michaely, Richardson and Roberts (2007). In order to remain internally consistent, I remove each industry's price and dividend information in the measure of its market PD-ratio. PD-ratios are the inverse of the cash dividend plus repurchase yields.

A: SIC-to-NAICS Statistics

	N	Mean	p25	p50	p75	max/min
SIC-to-NAICS Count	60	4.13	2.00	3.50	5.00	19.00
<i>1977-1997 weights:</i>						
Employment	247	0.24	0.02	0.12	0.35	0.00
Value-added GDP	247	0.24	0.02	0.12	0.33	0.00
<i>1977-only weights:</i>						
Employment	247	0.24	0.02	0.12	0.33	0.00
Value-added GDP	247	0.24	0.02	0.11	0.32	0.00
<i>1997-only weights:</i>						
Employment	247	0.24	0.02	0.11	0.35	0.00
Value-added GDP	247	0.24	0.01	0.11	0.32	0.00

B: Price-to-Dividend Statistics

	Composition		PD statistics			
	n(max)	n(min)	mean	sd	skewness	kurtosis
$\log(PD_t^{mkt})$:	8806	721	4.62	0.30	-0.11	2.49
$\log(PD_t^{ind})$:						
p25	49	2	4.62	0.41	-0.20	2.75
p50	106	4	4.75	0.48	0.13	3.38
p75	219	15	5.02	0.76	0.71	5.14

Summary statistics of my PD measure are presented in table (2B). As can be seen, over my sample there is sharp increase in the number of companies at both the market and industry-specific level over time. For example, during the early years almost 25% of industries have 3 or less characteristic firms in my PD-ratio estimate. The sharpest increase in firm composition comes between the years 1960 to 1975, during which the AMEX and NASDAQ were introduced to CRSP.

4.3 Validity of Price-to-dividend

The premise of my analysis is that industry PD-ratios embed realistic expectations of an industry's future cash flow and riskiness. In this section, I use a predictive regression approach, initially used in work by Harvey (1988), to show that this is in fact the case. More recently, Bansal, Kiku and Yaron (2012) and Colacito and Croce (2011, 2013) use a similar approach in the long-run risk context. Specifically, I show that the past industry and market PD-ratios computed in section (4.2) project next period industry value-added GDP growth computed in section (4.1). Assuming that expectations of GDP growth are, on average, correct, this intuition is captured in a simple regression for industry i ,

$$\Delta y_{i,t+1} = \beta pd_{i,t} + \varepsilon_{i,t+1},$$

where lower case letters signify the log of each variable.

I can take this intuition a step further by decomposing the information embedded in a given industry's PD-ratio into two orthogonal components. Either industry i 's cash flow is directly tied to broader underlying economic activity (e.g., shift from manufacturing to information technology) or its distributional characteristics provide diversification benefits to the agent. In both instances these dynamics are appropriately captured by the industry PD-ratio's co-movement with the market-wide PD-ratio. Industry i 's PD-ratio can also reflect *industry-specific* considerations. This dynamic is captured by the residual from the decomposition regression. In addition, the risk-free rate embeds information about market-wide prospects or long-run risk (see Bansal and Yaron

(2004)). I thus include the annualized 3m treasury bill rate in the projection. Finally, these long-run expectations of output and riskiness tend to be extremely persistence. I capture this persistence as an AR(1) process. The motivating idea is that investment decisions, especially those of large scale, entail projecting economic growth substantially in the future and are usually made ignoring higher frequency, “temporary” economic (e.g., business) cycles. For example, Croce (2014) and Colacito, Croce, Ho and Howard (2013) use persistent AR(1) process in TFP growth to capture previously unexplained production and asset price puzzles.

My primary system of equations for any industry i is thus,

$$\begin{aligned}
 pd_{i,t} &= \gamma_i pd_{-i,t} + \eta_{i,t}, & (1) \\
 \Delta y_{i,t+1} &= \beta_i pd_{-i,t} + \alpha_i \eta_{i,t} + \delta_i r_t^f + e_{i,t+1}, \\
 x_{i,t} &= \beta_i pd_{-i,t} + \alpha_i \eta_{i,t} + \delta_i r_t^f = \rho_i x_{i,t-1} + \varepsilon_{i,t}.
 \end{aligned}$$

The first equation of system (1) captures the decomposition of the industry PD-ratio onto its market (captured by γ) and industry-specific (η_i) components. In order to remain internally consistent, I remove each industry’s price and dividend information in the measure of its market PD-ratio (i.e. $pd_{-i,t}$). The second equation of system (1) is the projection or predictive regression. If the PD-ratio embed information about long-run expectations of industry cash flows, a negative β and α implies that agents ascribe “safety” to the industry’s cash flows. Intuitively, one would assume non-durable industries, such as food-manufacturing, would exhibit negative β s, whereas technology related industries, such as computer-manufacturing, would exhibit positive α s. The third equation of system (1) parameterizes the persistence of these expectations.

For statistical inference, I jointly estimate each industry’s α , β , γ , and δ coefficients using GMM (see APPENDIX C for orthogonality conditions). Separate GMM regressions are run for each of my 61 NAICS industries from 1948 to 2012 with annual data. Table (3A) gives the results of this estimation for the main parameters at the 25th, 50th, and 75th percentiles. The β coefficients and volatility of the estimated conditional component, x_t , match closely to those in Colacito and

Table 3: GMM and Business-cycle Parameter Summary

This table reports the GMM parameter statistics from section (4.3). Separate GMM regressions are run for each of my 61 NAICS industries from 1948 to 2012 with either annual (Panel A) or quarterly (Panel B) data. This allows me to jointly estimate my projection equations. The moment conditions for the quarterly GMM require minor modifications due to the mixed-frequency data. Those details are in APPENDIX D. In addition, I provide results of a Wald statistic that test whether the main projection coefficients are jointly different than zero, and estimates of the business-cycle beta from equation (2). Results are given at the 25th, 50th, and 75th percentiles.

A: Annual Parameters and Statistics

	count	mean	sd	p25	p50	p75
Idiosyncratic P/D, α	61	-0.001	0.037	-0.016	0.000	0.007
Market P/D, β	61	0.019	0.032	0.001	0.009	0.037
Decomp. Coef., γ	61	1.014	0.467	0.804	0.937	1.116
Risk-free Rate, δ	61	0.021	0.274	-0.089	0.010	0.097
LRR Persistence, ρ	61	0.630	0.206	0.519	0.681	0.774
Wald p-values, $\alpha = \beta = \delta = 0$	61	0.319	0.356	0.007	0.111	0.625
Variance ratio, $\sigma_{Irr}^2/\sigma_{\Delta y}^2$	61	0.032	0.037	0.005	0.019	0.044
Business-cycle, β_{bc}	61	0.550	0.665	0.200	0.447	0.697

B: Quarterly Parameters and Statistics

	count	mean	sd	p25	p50	p75
Idiosyncratic P/D, α	61	-0.001	0.005	-0.002	-0.000	0.001
Market P/D, β	61	0.002	0.007	-0.002	0.001	0.004
Decomp. Coef., γ	61	0.982	0.524	0.652	0.878	1.145
Risk-free Rate, δ	61	0.001	0.084	-0.035	-0.000	0.037
LRR Persistence, ρ	61	0.971	0.024	0.961	0.979	0.989
Wald p-values, $\alpha = \beta = \delta = 0$	61	0.170	0.259	0.000	0.026	0.226
Variance ratio, $\sigma_{Irr}^2/\sigma_{\Delta y}^2$	61	0.003	0.004	0.000	0.001	0.003
Business-cycle, β_{bc}	61	0.550	0.665	0.200	0.447	0.697

Croce (2011) and Bansal, Kiku, Yaron (2012). In addition, high and low β industries are largely composed of durables (e.g., autos) and non-durables (e.g., food), respectively. In addition, the null for the Wald test $\alpha = \beta = \delta = 0$ is rejected for nearly 50% of the industries at a 5% level, providing evidence that past PD-ratio provide insight into future industry output. Much of the data I use can also be extracted at a quarterly frequency. Table (3B) shows that higher frequency data yields even stronger evidence of this predictive relationship (see APPENDIX D for methodology). My result validate the usefulness of industry price-to-dividend as a measure of prospects. For the specifications where an industry's PD-ratio is split into its market and industry components, I simply use $\gamma_i pd_{-i,t}$ and $\eta_{i,t}$ from system (1), respectively. Given the nomenclature used in system (1), henceforth $pd_{i,t}$, $\gamma_i pd_{-i,t}$, $\eta_{i,t}$ will be referred to as $x_{i,t}$, $x_{i,t}^{global}$ and $x_{i,t}^{local}$, respectively.

4.4 Other Independent Variables

From the Campbell-Shiller decomposition (see APPENDIX E for derivation), one can see that the PD-ratio is a convolution of expected dividend growth and expected returns. Given the economic story I developed in section (2), it would seem that the cash flow rather than the return component of the PD-ratio would be the primary driver of real estate dynamic. I exploit the standard ICAPM risk and return relationship to derive a measure of long-run industry prospect that seeks to extract only the cash-flow portion from the industry price-to-dividend ratios (see APPENDIX F for details). I refer to the alternative measure of industry prospect as x_t^{cf} .

In addition, it is important that I isolate the impact of long-run industry prospects on housing returns from those of higher frequency (henceforth business cycle) dynamics. Previous literature has shown that housing returns are strongly effected by cycles of higher frequency (see Davis and Heathcote (2005) for literature overview). For my analysis in section (5), I use three measures of business cycle risks that are motivated in Tuzel and Zhang (2016). Two utilize the beta estimate, β_{bc} , from regressing my annual industry value-added GDP onto annual national GDP growth, i.e.

$$\Delta y_t^i = \beta_{bc} \sum_{q=1}^4 \Delta y_{q,t}. \quad (2)$$

Details of the construction of these three measure are provided in the next section. The summary statistics from the cross-section of β s from this regression are also provided in tables (3A) and (3B).

In summary, given consistent parameter estimates at two frequencies, I am comfortable constructing point estimates of the long-run industry prospects, x_t , for all 61 aggregate NAICS industries. In addition, following Tuzel and Zhang, I estimate business cycle variables for these same industries. The next step is to “localize” these aggregate measures.

4.5 Construction of Local Variables

I weight each industry variable estimated at the aggregate level with the industry’s employment share at the local MSA level and then sum to obtain my local variable. In the urban and labor economics literature this methodology is called the Bartik (1991) procedure and is used to generate measures of local “demand” from aggregate market measures (see Bound and Holzer (2000), Autor and Duggan (2002) and Luttmer (2005) for applications). In the financial context, Bekaert, Harvey, Lundblad and Siegel (2007), Davidoff (2015), and Tuzel and Zhang (2016) have motivated similar approaches in their analysis. I use the procedure to connect aggregate industry prospects to growth prospects of local economies.

For example, my MSA-level long-run prospect measure for any period, t , and MSA, msa , is

$$x_{msa,t} = \sum_{i=1}^I s_{i,msa,t} x_{i,t}, \quad (3)$$

where $s_{i,msa,t}$ is the fraction or share of total employment of industry i in msa at time t . My source of MSA-level employment data is the BLS quarterly census of employment and wages (QCEW) data. Applying the Bartik procedure, even dynamics that vary across only industry (e.g., β_{bc}) will now vary in both space and time due to changes in employment composition - e.g., β_{bc} becomes $\beta_{msa,t}^{bc}$.

Given that geographic definition of MSAs are periodically updated with the addition and dele-

tion of counties, it is important that my employment shares are consistent with my main dependent variable, the MSA-level FHFA excess housing returns. I thus download county level employment data and aggregate up to the MSA using the most recent BLS county-to-MSA crosswalk. My regressions are also supplemented with four measure of business-cycle variation for each MSA. The first is simply $\beta_{msa,t}^{bc}$. The second is the aggregate GDP growth, Δy_t . The third is the expectation of industry GDP growth given its beta exposure to the GDP shock, i.e. $\beta_{msa,t-1}^{bc} \times \Delta y_t$. Finally, employment growth itself may be an important predictor of real estate returns. The fourth addition is thus the MSA Bartik-localized employment growth, $\Delta emp_{msa,t}$.

Tables (4) provide summary statistics for my primary MSA-level dependent and independent variables of interest. Most of my independent variables show significant non-normal behavior. This is suggestive of persistent dynamics and correlated errors across MSAs. In my regressions, I thus control for correlations in errors by clustering standard errors along either the MSA or time dimension and add MSA-level and date fixed effects. Additionally, my empirical analysis explains fluctuations in the housing risk premium not the gross housing returns; this is an important distinction versus some of the work listed in section (3). Given that the risk-free rate is expected to be higher during good economic times, it would be unsurprising to find that x_t , my local PD ratio, predicts the gross real return to housing as well. In other words, housing returns, asset returns and the risk-free may have common cycle components; by using the housing risk premia rather than real returns as my independent variable, I should be able to better identify the main drivers of *long-run* heterogeneity in house price dynamics. The lagged 3m risk-free security yield is subtracted to obtain my MSA-level housing excess price appreciation, r_t^p .

Table 4: Panel Summary Statistics

This table reports summary statistics of my estimated MSA-level dependent variable and regressors using the annual and quarterly frequency data. I localize these regressors by following the Bartik (1991) procedure - each industry variable I estimate at the aggregate level is interacted with the MSA-level employment share and then summed at a MSA-level. Summary statistics are also given on samples collapsed either cross-sectionally (375 MSAs) or by time-series (35yrs). r_t^p is the excess price appreciation estimated by the MSA-level FHFA repeat sales indices. \hat{x}_t is my industry specific measure of long-run risk - the demeaned industry pd-ratios, $\hat{\varepsilon}_t$ is the residual from system (1), $\hat{\beta}^{bc}$ is the business cycle beta computed from equation (2) and Δemp_t is the employment share weighted average aggregate industry employment growth.

	Annual Panel				Quarterly Panel			
	mean	sd	skewness	kurtosis	mean	sd	skewness	kurtosis
<i>Full panel:</i>								
r_t^p	-0.011	0.065	-0.636	13.005	-0.003	0.026	-0.477	39.929
x_t	0.074	0.396	-0.309	1.884	0.051	0.342	-0.158	1.893
$\hat{\varepsilon}_t$	0.020	0.269	0.441	3.043	0.009	0.146	0.082	4.481
$\hat{\beta}_t^{bc}$	0.462	0.058	1.232	11.129	0.462	0.059	1.150	12.074
Δemp_t	0.012	0.021	-1.100	5.420	0.003	0.004	-0.509	2.925
<i>Time-series:</i>								
r_t^p	-0.010	0.057	-0.214	3.998	-0.002	0.024	-0.474	7.116
x_t	0.086	0.395	-0.390	2.086	0.061	0.338	-0.215	1.998
$\hat{\varepsilon}_t$	0.022	0.273	0.431	2.932	0.010	0.146	0.105	3.829
$\hat{\beta}_t^{bc}$	0.462	0.032	0.029	2.576	0.462	0.032	0.046	3.005
Δemp_t	0.012	0.021	-1.209	5.164	0.003	0.004	-0.578	2.677
<i>Cross-section:</i>								
r_t^p	-0.011	0.061	-0.295	4.213	-0.003	0.025	-0.517	7.812
x_t	0.061	0.396	-0.266	1.941	0.040	0.337	-0.119	1.925
$\hat{\varepsilon}_t$	0.016	0.270	0.446	3.003	0.007	0.145	0.112	3.820
$\hat{\beta}_t^{bc}$	0.463	0.032	0.036	2.569	0.463	0.032	0.056	3.112
Δemp_t	0.013	0.021	-1.084	5.002	0.004	0.004	-0.530	2.712

5 HOUSING DYNAMICS ANALYSIS

5.1 Return Regression

My full regression specification is,

$$r_{msa,t}^p = b_x x_{msa,t-1} + \text{Time FE} + \text{MSA FE} + \text{Controls} + \varepsilon_{msa,t}. \quad (4)$$

My business cycle control variables from section (4) are $\beta_{msa,t-1}^{bc}$, which captures the exposure or β of the local industry mix to changes in aggregate gdp; the interaction of this β and aggregate GDP growth, $\beta_{msa,t-1}^{bc} \times \Delta y_t$; and the five year average local industry employment growth, $\Delta emp_{msa,t-1}$. Given that I have time fixed-effects, variation in credit supply and contemporaneous shocks to GDP, will not drive my results (see Favilukis, Kohn, Ludvigson and Van Nieuwerburgh (2013) for discussion). This is especially important due to the credit boom-bust cycle from 2006-2009. In addition, I include the lagged MSA return premia, $r_{msa,t-1}^p$, as a control due to the highly degree of autocorrelation in return measures computed from repeat-sale indices (see Ghysels, Plazzi, Torous, Valkanov (2012)).

My primary coefficient of interest is b_x . In section (2) I showed that a geographic concentration of high PD-ratio industries today seems to imply greater excess price appreciation going forward. $x_{msa,t}$ is a employment weighted average of the PD-ratios of industries within a particular MSA. I would thus expect that b_x is positive and significant. I conduct this analysis for annual $x_{msa,t}$; the results are presented in table (5). All variables are standardized, representing the marginal effect from a one standard deviation move in the regressor.

b_x is significant in the expected direction. A one standard deviation higher measure of prospects implies a greater than 120bp higher price appreciation in the next period. In column (2) I split my

Table 5: Return Regression Results

This table reports the regression results of equation (4) given annual data. I cluster standard errors along MSA. The repeat-sales method used to compute the FHFA price indices has a high degree of autocorrelation in returns by construction. I thus include the lagged MSA return premia, r_{t-1}^p , as a control. Model (1) is the regression on the full sample. Model (2) splits my measure of long-run local growth prospects into its market-wide (global) and MSA-specific (local) components. Model (3) adds the Tuzel and Zhang (2016) motivated business cycle variables. Model (4) adds the estimated local expected employment growth. Model (5) uses the MSA-level rental data from Campbell, Davis, Gallin and Martin (2009) for a cross section of 23 MSAs from 1975-2007. As this regression limits the time series and cross-section, I do not include MSA and time fixed effects, but rather cluster along both dimensions. I also add the credit supply variable from Favilukis, Kohn, Ludvigson and Van Nieuwerburgh (2013). *, **, and *** denote significance at the 10%, 5%, and 1% levels.

	(1)	(2)	(3)	(4)	(5)
	β	β	β	β	β
	[t-stat]	[t-stat]	[t-stat]	[t-stat]	[t-stat]
$x_{msa,t-1}$	0.0126*** [6.29]				0.0120** [2.32]
$\hat{x}_{msa,t-1}^{global}$		0.0155*** [10.42]	0.0141*** [8.43]	0.0175*** [8.73]	
$\hat{x}_{msa,t-1}^{local}$		0.0036** [2.21]	0.0038** [2.34]	0.0036** [2.24]	
Inc. r_{t-1}^p	+	+	+	+	+
Inc. $\hat{\epsilon}_t$	+	+	+	+	+
BC Controls			+	+	+
Inc. Δemp_{t-1}				+	+
FKLV _{cs}					+
Observations	10,003	10,003	10,003	10,003	586
R^2	0.5078	0.5084	0.5089	0.5100	0.5349
Wald Test		0.0000	0.0733	0.0001	

measure of industry long-run prospects into its market-wide ($\hat{x}_{i,t}^{global} = \gamma_i pd_{-i,t}$) and orthogonal industry-specific ($\hat{x}_{i,t}^{local} = \eta_{i,t}$) components in the regression (see equation (1)). These are employment share weighted according to equation (3) to obtain a common, $x_{msa,t}^{global}$ and local, $x_{msa,t}^{local}$, portion of long-run prospects across MSAs. Both parts of long-run prospects load significantly and positively on expected future price appreciation. In addition, I test if disentangling these orthogonal components give us any additional predictive power. The null that the coefficients are equal is

rejected as seen by the Wald Test p-value. Column (3) and (4) add the business cycle variables. While the business cycle regressors contribute to explaining the variation in excess returns (see Wald tests p-values), they explain substantially less than the long-run prospect derived variable ($x_{msa,t}$). Finally, our measure of price appreciation is not a proper total return measure as it does not include rent. Campbell, Davis, Gallin and Martin (2009) derive a measure of rent for a cross section of 23 MSAs from 1975-2007 using Census data. I add rent to the price appreciation and limit the regressions to these years and MSAs. The results are in column (5). As the cross-section and time series are now limited, I no longer include MSA and time fixed effects, but cluster along both dimensions following Petersen (2009). All results point to the same conclusion - the better the conditional prospects the higher the expected returns next period.

5.2 Robustness

My results are robust to a variety of tests which I present in table (6). The first two columns remove observations from the full-sample due to limitations of my data. First, as described in section (4.5), employment data between 1975-1989 was converted using the national employment conversion matrix estimated by my Gauss-Newton algorithm in section (4.1). I thus run equation (4) on only the > 1989 sample, which reduces the number of observations by nearly 25%. The significance of the regressors remains. Second, the QCEW database suppresses a substantial number of county-industry employment figures due to disclosure restrictions. This is more prevalent in less populated counties; during the computation of employment weights, I thus define a variable that captures the percentage of industry-county employment information that is suppressed within an MSA. Column (2) restricts the sample over which I run regression (4) on only those MSAs where $< 10\%$ of county-industries are suppressed. Although the sample is cut by more than half, the statistical and economic significance of the coefficient of interest remains. Column (3) limits the sample to the largest MSAs. Major cities tend to be the main source of a region's economic dynamism. From an economic perspective, my findings would carry much less importance if they were limited to smaller MSAs. Restricting the sample to MSA-level population

Table 6: Return Regression Robustness Results

This table reports robustness results for the regressions presented in table (5). Details on standard errors and controls are identical to those applied in that table. Model (1) uses data from 1990-2011, which removes employment share estimates using converted employment data. Suppression of certain count-year employment data from the QCEW may be a problem in estimating my coefficients. I specify an indicator where each county-*msa* employment data is not disclosed. I then compute the percentage of county-industries within an MSA with suppressed data. Model (2) restricts the sample over which I run regression (4) to only those MSA-years where $< 10\%$ of its county-industries are suppressed. Model (3) restricts the sample to MSA-level populations of $> 680,000$, which represents the top 25% of MSA size. Model (4)- (5) tweaks my employment share measure to include lagged employment growth to varying degrees; the tilt parameter, λ , is 1 and 5 for (5) and (6), respectively (see section (5.2) for details). *, **, and *** denote significance at the 10%, 5%, and 1% levels.

	(1)	(2)	(3)	(4)	(5)
	β	β	β	β	β
	[t-stat]	[t-stat]	[t-stat]	[t-stat]	[t-stat]
$x_{msa,t-1}$	0.0081*** [4.20]	0.0161*** [6.19]	0.0132* [1.89]	0.0079*** [2.83]	0.0080*** [2.93]
Inc. r_{t-1}^p	+	+	+	+	+
Inc. $\hat{\epsilon}_t$	+	+	+	+	+
BC Controls	+	+	+	+	+
Inc. Δemp_{t-1}	+	+	+	+	+
Observations	8,087	4,589	2,565	10,003	10,003
R^2	0.6213	0.4898	0.5855	0.5076	0.5073

of $> 680,000$, which is the 25th percentile of MSA size across my panel, the coefficient of interest shows similar economic and statistical significance.

The area of greatest overlap between my paper and urban economics is the literature linking local labor shocks, city population size and house prices. Glaeser and Gyourko (2005), in particular, explore how negative shocks to growth are linked to multi-decade employment growth and house price movements. The proxy traditionally used for labor shocks is variations in employment growth. Thus, another possible issue with my methodology is that while I weight my variables by employment share, I ignore expectations in employment growth within my measure of industry prospects. As an extreme example, assume an industry with high long-run prospects has a current share of employment of 0% in a particular region. Expectations, however, are that em-

ployment in that industry will grow to 10% share in the next year. One would assume that given this scenario the industry's prospects would contribute to the local long-run growth prospects far beyond its current employment share. Columns (4) and (5) attempt to answer this question by shifting MSA industry employment weights by relative average past 5-year industry employment growth such that for each industry, i , in msa ,

$$w_{msa,t}^i = \frac{emp_{msa,t}^i \times \exp\left(\lambda \times \Delta emp_{msa,t-5,t-1}^i\right)}{\sum_{k=1}^N emp_{msa,t}^k \times \exp\left(\lambda \times \Delta emp_{msa,t-5,t-1}^k\right)}.$$

This setup effectively shifts the weights on all variables towards industries with higher employment growth. Given that an industry's employment growth is also an extremely persistent variable, this analysis assumes that past growth is a good indicator of expectations of employment growth. The quarterly regression specification was re-run with arbitrarily chosen shift parameters, λ , values of 1 (column (4)) and 5 (column (5)). These new weights were then used to localize the aggregate industry dynamics to the MSA level. As seen in table (6) the results are robust to these shifts as well.

5.3 Capitalization Rate Regressions

The motivating assumption of my empirical approach is that industry clustering generates geographic heterogeneity in local prospects. I proxy for these "local" prospects by weighting PD-ratios of different industries by their employment share within a cross section of MSAs. A lower local long-run prospect measure, $x_{msa,t}$ signifies either lower expected cash flows, $\Delta y_{msa,t+1+j}$ and/or higher expected returns, $r_{msa,t+1+j}$, for the local industry mix.

Housing assets, however, seem to include a hedge on $x_{msa,t}$ in that expected excess price appreciation is lower in bad than good times. This is implied by the robust positive coefficient when running excess price appreciation onto the local economic state variable. This has immediate implications on the price-to-rent ratios of real estate. If real estate income or rent (henceforth s_t) is a

valid proxy for the utility provided by housing, then from Campbell and Shiller (1987),

$$p_t - s_t = \frac{k}{1 - \rho} + \sum_{j=1}^{\infty} \rho^j E_t [\Delta s_{t+1+j}] - E_t [r_{t+1+j}], \quad (5)$$

where lower case letters signify the logarithm of variables. In the spatial equilibrium models from urban economics, housing acts as a wedge (Rosen (1979) and Roback (1982)). In order to maintain a no-arbitrage condition between cities, dynamics in income growth filter to both housing services and non-housing consumption. Given that high local prospects should lead to future expected income gains, this implies that $E_t [\Delta s_{msa,t+1+j}]$ increases in $x_{msa,t}$. Additionally, from the regressions above, expectations of future returns also increase in x_t . As both cash flows and expected returns to housing assets are increasing in $x_{msa,t}$, a simple regression of the price-to-rent ratio onto my measure of local growth prospects will clarify which channel dominates house price dynamics.

To empirically tackle this question, I use the inverse of operating income-to-price or capitalization rate (henceforth cap-rates) data from Integra Realty Resources (IRR). IRR is a large real estate valuation and advisory service provider with offices in 62 different MSAs (www.irr.com). Annually from 1995-today IRR has been collecting cap-rate projections from major developers of commercial property across their regional offices. I then run a simple contemporaneous regression, projecting the MSA-level cap-rates onto the local x_t ,

$$\ln \frac{1}{CapRate_{msa,t}} = b_x x_t + \text{MSA FE} + \varepsilon_{msa,t}. \quad (6)$$

It is important to note that besides zeroing in on the the specific channel that drives the pricing results of section (5), this regression also confers a robustness benefit. First, expected returns computed via repeat sales indices (as in section (5)), tend to conflate past with future returns if transaction volumes are extremely low as is the case in small MSAs and during economic downturns (see Ghysels, Plazzi, Torous, Valkanov (2012)). In regression (4) I attempted to control for this with one-period lagged excess appreciation. The IRR data is survey rather than transactional

Table 7: Capitalization Rate Regressions

The dynamics of long-run prospects seem to be a key driver of real estate dynamics. As seen in sections (5) and (7), the link has both a expected return and cash flow component. This table reports the results of regression (6), which empirically answers which process dominates. The capitalization rate data is from Integra Realty Resources. I cluster along both the time and MSA dimension. Model (1) runs the regression on the full sample of capitalization rates for urban multi-family home property types. Model (2) adds the contemporaneous business cycle variables from section (4). Model (3) splits my measure of long-run prospects into its global and local components. Model (4) and (5) seek to mitigate spurious correlation issues as both the dependent and independent variables are highly persistent. I first difference the PS-ratios and measure of long-run growth prospects. That is, I regress the annual change in capitalization rate onto the change in local prospects (see equation (7)). Models (4) does this regression on urban multi-family home capitalization rates; model (5) on suburban multi-family home capitalization rates. Both regressions are still clustered along the time dimension.

	(1)	(2)	(3)	(4)	(5)
	β	β	β	β	β
	[t-stat]	[t-stat]	[t-stat]	[t-stat]	[t-stat]
x_t	-0.1274*** [-6.71]	-0.1214*** [-6.92]			
\hat{x}_t^{global}			-0.0916*** [-5.03]	-0.0462*** [-4.44]	-0.0453*** [-4.85]
\hat{x}_t^{local}			-0.0745*** [-2.97]	-0.0356*** [-3.11]	-0.0375** [-2.90]
T-Z Beta		+	+		
Δemp_t		+	+		
FKLV _{cs}		+	+		
Observations	753	753	753	658	704
R^2	0.489	0.536	0.540	0.213	0.211

data; to the degree that expectations will be immediately embedded in the surveyed price-to-rent ratios, the contemporaneous regression presented in equation (6) will even better mitigated this possible issue. Second, my return measure from earlier does not include rental income and is thus not a proper martingale. The capitalization rate data implicitly includes a measure of rent.

Column (1) of table (7) gives the results of this regression over the sample of urban multi-family home property types. I cluster standard errors along time and MSA and include MSA-level fixed effects. When times are good (i.e. high x_t), the negative and statistically significant coefficient implies that the expected return impact dominates. That is in good times, the PS-ratio falls. Column (2) adds the contemporaneous business cycle variables from section (4) - i.e. employ-

ment growth and business cycle beta, and credit supply. As with the return regressions, both load on these regressions, but their economic impact is smaller than that of long-run prospects. Column (3) splits my measure of long-run prospects into its global and local components. Similar to my return regressions, both components load significantly in the expected direction. As Plazzi, Torous and Valkanov (2010), highlight from their capitalization rate data, the cross-sectional variation in cap rates is much greater than the time series variation. Given the high degree of persistence, especially in panel, of both the right and left sides of equation (6) there is concern that these results could be spurious. As robustness, I first difference the MSA-level PS-ratios and measure of long-run local growth prospects. That is, I'm seeing how the price-to-rent (ps_t) ratios change with local prospects,

$$\Delta ps_{msa,t} = b_{global} \Delta x_{msa,t}^{global} + b_{local} \Delta x_{msa,t}^{local} + \varepsilon_{msa,t}. \quad (7)$$

Column (2) looks at this regression for urban multi-family homes whereas column (3) focuses on *suburban* multi-family homes. Standard errors for both regressions are clustered in the time dimension. All results seem to point to the same conclusion from my initial regression - the risk premia of housing is higher in good than bad times.

6 ECONOMIC FOUNDATIONS

In this section, I develop a representative agent model that helps account for my empirical findings. First, the decision to build and the process of consuming local specific housing services involve horizons that are longer than the duration of a typical business cycle fluctuation. My agent's post-trade endowment is thus exposed to a low volatility, highly persistent state variable that seeks to reflect the dynamics of my local measure of long-run growth prospects. To properly internalize the risk to their wealth from the dynamics in long-run prospects, my agent has recursive preferences (see Bansal and Yaron (2004)).

Second, my empirical findings in section (5) suggest that when the economic environment is bad - i.e. expectations of future cash flow is low - for firms located within an MSA, the expected risk premia of housing decreases. This implies that housing acts as a hedge to poor long-run growth prospects. In my model, the location specific consumption provided by housing is thus highly durable to the downside. This downside protection generates a regime switching dynamic in the expected excess returns to housing. The transition between these states is gradually increasing in my state variable as it is a function of the probability of being in a high growth state tomorrow given the growth state today. This setup allows me to replicate my empirical findings and then motivates a more in-depth empirical analysis of underlying consumption behavior in section (7).

6.1 Setup of the Economy

The MSA is populated by a representative agent with Epstein and Zin (1989) recursive preferences with unit intertemporal elasticity of substitution (IES), defined over a CES aggregate of

housing (S_t) and non-housing (C_t) goods

$$U_t = (1 - \beta) \ln(u_t) + \beta \ln E_t [\exp(\theta U_{t+1})]^\theta, \quad \text{where} \quad (8)$$

$$u_t = \left[(1 - \alpha) C_t^{\frac{\epsilon-1}{\epsilon}} + \alpha S_t^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}},$$

where $\theta = \frac{1}{1-\gamma}$ and is a function of risk aversion, γ , ϵ is the intra-temporal elasticity of substitution, and α captures the agents relative preference for housing versus non-housing consumption. I choose to limit my analysis to unitary IES as this corresponds to the average number estimated in the empirical literature and several theoretical studies have already employed this specification (e.g., Tallarini (2000); Colacito, Ghysels, Meng and Siwasarit (2016)). Epstein and Zin preferences have been used extensively in RBC models to explain asset pricing puzzles although the analysis as been largely limited to equities, bonds, and their derivatives (e.g., Bansal and Yaron (2004); Colacito and Croce (2011); Gomes, Kogan, Yogo (2009); Eraker, Shaliastovich, and Wang (2016)). Fillat (2008) uses a similar setup as mine with housing assets being priced, but largely focuses on its impact on equity and bond prices.

By separating the risk (captured by θ) from inter-temporal consumption decisions (captured by IES), EZ-preferences provide a channel through which a MSA's long-run industry prospects play an important role in an agent's consumption behavior. This is illustrated by the additional variance term if I assume the agents wealth follows a log-normal distribution,

$$U_t = (1 - \beta) \ln(u_t) + \beta E_t [U_{t+1}] + \beta \frac{1}{2\theta} V_t [U_{t+1}]. \quad (9)$$

In the case of housing returns, this framework has immediate intuitive appeal. To the degree housing decisions factor into an agent's wealth, one would think that even business frequency cycles should not rest as heavily on their mind as say 15+ year dynamics. As discussed in section (4.3), even small shocks to my prospects carry over long-horizons due to its persistent. Its impact is felt through the riskiness of future utility (wealth) term in equation (9).

I complete my local economy by specifying an endowment process for the MSA's non-housing good and housing service consumption growth. Both contain exposure to a AR(1) predictive component, x_t . In my monthly calibration the persistence coefficient, ρ , is close to but still less than one, and matches the average persistence of my empirical measure of local long-run growth prospects. I assume that the endowment of both goods is a function of the same underlying state variable, x_t . This is similar to Gomes, Kogan and Yogo (2009), who assume a single aggregate productivity measure drives production for both non-durable and durable goods manufacturing firms.

$$\begin{aligned}\Delta c_{t+1} &= \mu_c + \phi_c x_t + \frac{\lambda}{2}(s_t - c_t) + \sigma_c e_{c,t+1}, \\ \Delta s_{t+1} &= \mu_s + \mathbb{I}_{x>0} \phi_s x_t - \frac{\lambda}{2}(s_t - c_t) + \phi_s \sigma_c e_{s,t+1}, \quad \text{where} \\ x_t &= \rho x_{t-1} + \phi_x \sigma_c \varepsilon_{x,t}.\end{aligned}\tag{10}$$

μ_c and μ_s are the unconditional growth non-housing consumption and housing service endowments, ϕ_c and ϕ_s are leverage exposures to underlying local prospects. The first and second equations in system (10) also include a co-integration term, $\frac{\lambda}{2}(s_t - c_t)$. As highlighted by Piazzesi, Schneider and Tuzel (2007), housing as a share of total expenditures has stationarity properties over my sample horizon; this is not the case, for example, with durable versus non-durable goods consumption. My no-bubble condition links non-housing and housing services consumption across time to maintain this property (see Giglio, Maggiori and Stroebel (2016) for further discussion on the existence of bubbles in housing markets). The parameter λ determines the rate of convergence of the relative quantities. In addition, I assume that the shocks in system (10) are standard normal and orthogonal.

Housing services are broadly defined - they include everything from housing capital to agglomeration economies. The key difference between it and non-housing goods, is that the production of the former is highly durable and location dependent. Both tangible and intangible infrastructure are built so that people and companies within a particular industry can be in close proximity. However, once built, these accumulated "benefits" cannot be easily shifted given the externalities of agglom-

eration. Thus, in my model, while the endowment of non-housing goods is symmetrically exposed to x_t (following Bansal and Yaron (2004)), housing services can expand in “good” times, which is activated by the indicator variable $\mathbb{I}_{x>0}$, but not contract in “bad.” Given that x_t is a mean zero process, the kink resides at $x_t = 0$. The assumption of a hard kink is admittedly extreme; however, the model captures the more general idea that in bad times housing services cannot contract at the same rate as non-housing consumption. That is the relative, not absolute, durability of housing services is what is important from a pricing perspective.

6.2 Financial Markets and Equilibrium of the Economy

My empirical setup effectively assumes that each MSA acts in autarky for both consumption goods - i.e. the representative agent in each MSA consumes only the good with which he is endowed. If the implied prospects of the local industry mix is high, this is immediately reflected in the consumption endowments locally for both non-housing goods and housing services. This setup can be motivated as a segmented or post-trade economy.

As I allow long-run prospects and thus the composition of consumption ($\ln \frac{S_t}{C_t} = sc_t$) to be conditioned on past values, both x_t and sc_t are state variables. I solve for equilibrium of my utility function and price-to-consumption ratios using standard iteration methods along a 3σ grid of x_t and sc_t . Expectations are computed across a 99-point gaussian quadrature, which randomly assigns shocks for e_c , e_s , and ε_x (see details of algorithm and derivations in APPENDIX G).

Using the non-housing good as the numéraire, the stochastic discount factor is characterized by

$$M_{t+1} = \left[\frac{\beta \exp(U_{t+1}/\theta)}{E_t \exp(U_{t+1}/\theta)} \right] \left[\frac{u_{t+1}}{u_t} \right]^{\frac{1}{\varepsilon}-1} \left[\frac{C_{t+1}}{C_t} \right]^{-\frac{1}{\varepsilon}}. \quad (11)$$

Equation (11) illustrates the components that will potentially impact an asset’s risk premia. The last bracketed term is related to the consumption growth risk seen in traditional asset pricing models (e.g., Mehra and Prescott (1985)). The middle bracketed term is a function of the growth of

aggregated bundle of non-housing goods and housing services and is thus a composition growth risk term. Given that both non-housing and housing service endowment growth is a function of highly persistent terms, the continuation utility (first) term is impacted by both risks in the SDF.

I then use the Euler equation, $E_t [M_{t+1} R_{s,t+1}] = 1$, to pin-down expected excess returns for the local housing. For the housing asset, the price-to-rent, $\ln \frac{P_{s,t}}{s_t}$, is

$$\ln \frac{P_{s,t}}{s_t} = ps_t = \ln E_t [\exp(m_{t+1} + \log(\exp(ps_{t+1}) + 1) + \Delta s_{t+1})]$$

where lower cases signify logs (see APPENDIX G for explicit recursive formulations). Given the value functions of the price-to-rent, both the excess return, $r_{s,t+1} = \ln(\exp(ps_{t+1}) + 1) - ps_t + \Delta s_{t+1}$, and excess price appreciation, $r_{s,t+1} = ps_{t+1} - ps_t + \Delta s_{t+1}$, can be computed for local housing.

I assume that the aggregate, or non-segmented markets, are priced by a subset of people that are not tied down to the location of housing services. That is they can allocate or trade housing between multiple MSAs. One could think of these individuals as represented by representative institutional money manager or a wealthy individual. I show in APPENDIX H that this representative agent then prices aggregate markets as in Bansal and Yaron (2004) with long-run risk and stochastic volatility. Given this result, my empirical strategy immediately follows. I can identify x_t through information in the aggregate market price-to-dividend ratios to elucidate their effects on my local asset, housing. There is empirical evidence of the existence of different marginal investors in aggregate versus local markets. As documented by Chen and Stafford (2016) historically only around 25-35% of households participate directly in the stock market. This is inverse the home ownership rate of approximately 65%. In addition, as Babkin, Glover and Levine (2016) show companies and wealthy executives readily move when local conditions do not suite (e.g., corporate inversion - Pfizer, Medtronic, Tyco; or Facebook's Eduardo Saverin). The median homeowner on the other hand does not have this option.

Table 8: Model Calibration

This table reports the calibration used in my model. β , ρ , φ_x , and γ are taken from Bansal and Yaron (2004). The utility weight parameter, α , on housing services is estimated as in Gomes, Kogan and Yogo (2009). The elasticity of substitution between non-housing goods and housing service consumption, ε , is estimated from the β when regressing expenditures of housing onto expenditures of non-housing goods. μ_c , μ_s , φ_s , and λ are estimated from aggregate housing to non-housing consumption data. ϕ_s is a free variable that is calibrated to match my empirical results from section (5).

Parameter	Description	Value	
ϕ_s	Housing service leverage to long-run consumption growth	5.0000	Match Return Regression β
α	Housing service versus non-housing consumption preference parameter	0.1599	Simulated \bar{sc}
ε	Intra-period elasticity of substitution between housing services and non-housing consumption	2.2070	Expenditure Correlations
μ_c	Average consumption growth	0.0023	Using National Data
σ_c	Short-run consumption volatility	0.0057	...
φ_s	Ratio short-run housing service volatility to short-run consumption volatility	1.4200	Implied from ε and National Data
λ	Mean reversion coefficient on housing service to consumption ratio	0.0150	Predictive Regression
μ_s	Average housing service growth	-0.0021	Harding, et al (2007)
β	Subjective discount factor	0.9980	Bansal & Yaron (2004)
φ_x	Ratio long-run growth volatility to short-run consumption volatility	0.0602	...
ρ	Persistence coefficient on long-run consumption growth	0.9790	...
γ	Risk aversion	15.000	...

6.3 Calibration of the Model

Table (8) provides the monthly parameter values for my model. The parameters are chosen at a monthly frequency in the spirit of work done by Bansal, Kiku, and Yaron (2012), who estimate the decision frequency of agents in their economy to be roughly one month. The subjective discount factor (β), persistence on long-run prospects (ρ), relative variance of long-run to non-housing consumption shocks (φ_x) and risk aversion coefficient (γ) are taken directly from Bansal and Yaron (2004). I equate the average housing service growth (μ_s) to depreciation; Harding, Rosenthal and Sirmans (2007) estimates that housing depreciates at roughly -2.5% per year.

The relative dynamics of housing services to non-housing goods consumption is the compo-

ment of the local stochastic discount factor (SDF) that determines excess housing returns. Coming directly from the endowment equations (10), I can compute the level of this quantity,

$$\log \frac{S_{t+1}}{C_{t+1}} = sc_{t+1} = (1 - \lambda)sc_t + \mu + (\phi_s - \phi_c) \mathbb{I}_{x>0}x_t - \phi_c \mathbb{I}_{x<0}x_t + e_{sc,t+1}. \quad (12)$$

NIPA provides quantity estimates which I can use to estimate this regression; however, the Bureau of Labor Statistics' (BLS) process of splitting housing expenditures into quantities and prices is fraught with error. For example, Boskin, Dulberger, Gordon, Griliches, and Jorgenson (1998) discuss the historical inconsistencies in how the BLS estimates housing services (quantities) and rent (prices) due to changing surroundings and technology - e.g., agglomeration economies, pollution and widespread use of electricity. To avoid these issues tainting my calibrations, I follow Piazzesi, Schneider and Tuzel (2007) in using expenditure rather than price or quantity data to estimate parameters.

This is straightforward given that prices and quantities of my assets are linked via the intra-period equilibrium. Assuming that non-housing is the numéraire,

$$\frac{u_{s,t}}{u_{c,t}} = \frac{p_t^s}{p_t^c} = \frac{\alpha}{1 - \alpha} \left(\frac{S_t}{C_t} \right)^{-\frac{1}{\varepsilon}}.$$

I define relative expenditures of housing services and non-housing goods as

$$\omega_t = \ln \frac{p_t^s S_t}{p_t^c C_t} = \ln \frac{\text{Shelter}}{\text{Total Expenditure} - \text{Shelter}}, \quad (13)$$

where p_t^s and p_t^c are the equilibrium price of housing services and non-housing goods. The numerator is designated expenditure on "shelter." The denominator is then the total expenditures minus expenditures on shelter. Using this relationship and equation (12), I obtain a simple expres-

Table 9: Aggregate Calibration Regressions

This table reports the regression results of equation (14) using aggregate expenditure NIPA data on an annual basis. The dependent variable is log housing to non-housing expenditures. All independent variables are standardized so that expressions represent a one standard deviation change. Model (1) uses all available data. Model (2) removes the boom-bust of the financial crisis, i.e. data post 2005. Model (3) uses data from 1960 to 2005. Standard errors are Newey-West corrected with a lag of 4. Given the annual frequency of the data, coefficients must be scaled when used as parameter estimates for my monthly calibration. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

	(1)	(2)	(3)
	β	β	β
	(t-stat)	(t-stat)	(t-stat)
ω_{t-1}	0.9554*** [65.50]	0.9550*** [62.91]	0.8555*** [21.46]
$I_{x>0}x_{t-1}$	0.0001 [0.07]	-0.0000 [-0.01]	-0.0012 [-0.65]
$I_{x<0}x_{t-1}$	-0.0041*** [-2.35]	-0.0033** [-1.78]	-0.0027* [-1.63]
Observations	211	190	118

sion for an OLS regression I run on national NIPA data at an annual frequency,

$$\omega_t = \mu + (1 - \lambda) \omega_{t-1} + \left(1 - \frac{1}{\varepsilon}\right) (\phi_s - \phi_c) \mathbb{I}_{x>0}x_{t-1} - \left(1 - \frac{1}{\varepsilon}\right) \phi_c \mathbb{I}_{x<0}x_{t-1} + \left(1 - \frac{1}{\varepsilon}\right) e_{sc,t}. \quad (14)$$

Table (9) reports the regression results. All right-hand-side variables are standardized; coefficients thus represent the effects of a one standard deviation change. Standard errors are Newey-West corrected with a lag of 4. Column (1) is on the full sample, 1947-2011. Column 2 removes the boom-bust of the recent financial crisis, i.e. post-2005 data. Column 3 runs the regression on data from 1960-2005 as the data is a bit more volatile in the early sample. $(1 - \lambda)$ in particular is extremely well identified in all cases. Given the annual frequency of the data, my monthly λ is estimated as 0.0150. The kink is statistically identified, but economically weak in the national data. As I will show in section (7), it is much stronger using MSA-level expenditure data.

Given the extremely low λ , housing to non-housing consumption is also extremely persistent. Correlation in expenditures of housing services and non-housing goods consumption should be dominated by the higher volatility i.i.d. shocks in system (10). I can thus estimate the elasticity of substitution (ε) via the covariance of growth in housing to non-housing expenditures in the data. From the intra-period CES aggregator,

$$\ln \left(\frac{p_{t+1}^s S_{t+1}}{p_t^s S_t} \right) = \left(1 - \frac{1}{\varepsilon} \right) \Delta s_{t+1} + \frac{1}{\varepsilon} \Delta c_{t+1}$$

The inverse of the β coefficient of expenditures will thus be my estimated value of ε . My estimate of 2.2070 is close to that estimated in the literature (see Piazzesi, Schneider and Tuzel (2007)).

Given that non-housing consumption is my numéraire and the orthogonality of short-run shocks, I extract estimates of non-housing and housing consumption dynamic parameters (μ_c , σ_c and ϕ_s) directly from the data. I compute an estimate of the ratio of short-run housing service to short-run non-housing consumption volatility from the volatility of the expenditures of sc_t . My model assumes that $e_{s,t+1}$ is orthogonal to $e_{c,t+1}$; ϕ_s is therefore $\sqrt{\sigma_{sc}^2 - \sigma_c^2} / \sigma_c$. Finally, using the approach of Gomes, Kogan and Yogo (2009), I estimate the housing to non-housing consumption preference parameter (α) from the mean aggregate expenditure ratio in the data and the mean housing to non-housing quantities from my model. The remaining free-parameter, ϕ_s , is calibrated by matching the slope coefficients from my price appreciation empirical analysis in section (5).

6.4 Numerical Results

Given that $\varepsilon > 1$, housing services and non-housing goods are substitutes. The embedded put in housing services to long-run local growth prospects will thus act as a hedge in bad times. Figure (2A) illustrates the resulting equilibrium expected excess return from my model. As the empirical regressions in section (5) are on excess price appreciation not returns, it's important that the equilibrium relationship of expected price appreciation with x_t has a similar shape. Figure (2B)

shows that this is in fact the case.

Both equilibrium figures also show the contours of the two risk regimes. When growth prospects are extremely high housing is exposed to two sources of risk (short- and long-run), but only one when prospects are extremely low. At the two extremes, expected excess returns and price appreciation are effectively constant. Without recursive preferences the transition between regimes would be immediate. The elongated transition zone from the low return to high return regime thus allows for empirical identification of the positive relationship between long-run growth prospects and expected excess price appreciation. During the time frame of our regressions, roughly 1975-2011, local long-run prospect, across all MSAs, were roughly 0.45σ above the mean over the longer CRSP industry PD-ratio sample. In addition, the shocks to local prospects, assuming the AR(1) construct from my model, are approximately 80% correlated across MSAs. This implies that the relationship between r_{t+1}^p and x_t is squarely in this elongated zone during the sample from my empirical exercise. In addition the relationship is well identified because of the large cross section of MSAs.

To replicate the regressions in section (5), I simulate 375 MSAs over 40years on a monthly basis at the 1975-2011 mean of long-run growth prospects and correlation of shocks. I then aggregate the simulated $x_{msa,t}$, $ps_{msa,t}$ and $r_{msa,t+1}^p$ on an annual basis and run the same pooled coefficient regressions from section (5.1) and (5.3) on the simulated data, respectively

$$r_{msa,t}^p = b_x x_{msa,t-1} + \varepsilon_{msa,t}, \quad \text{and}$$

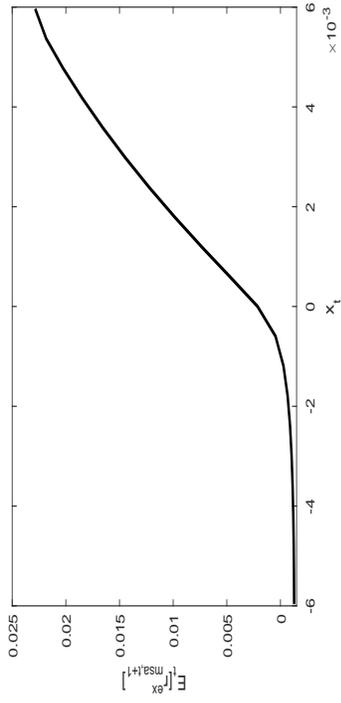
$$p_t - s_t = \beta_x x_t + \varepsilon_{msa,t}.$$

Figure (2C) illustrates the distribution of individual MSA estimated b_x . The pooled b_x is represented by the black line. My simulated b_x is on average 0.011, which is very close to my empirical b_x of 0.012. For my price-to-rent ratios regression I further restrict the regression to a random sample similar in cross-section and time-series to my empirical work (62 MSAs over 16 years). The simulated β_x is -0.037, close to my empirical β_x of -0.045 from the more robust first difference

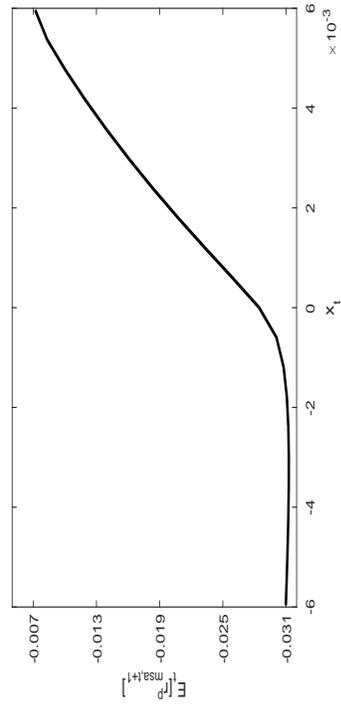
Figure 2: Equilibrium Prices and Empirical Replication

Figure (2A) is the equilibrium expected return of housing given the state variable, x_t , from my model. My regressions from section (5) were on a price appreciation measure; figure (2B) presents the equilibrium price appreciation. During the period of my sample of housing returns, x_t was on average 0.45σ higher than average over the full, 1947-2011, CRSP sample. In addition, assuming the AR(1) process from my model, shocks to long-run prospects are on average 80% correlated in my sample. I simulate 375 MSAs over 40 years on a monthly basis around this historically high mean and correlation. I then aggregated the simulated $x_{msa,t}$, $p_{msa,t}$ and $r_{msa,t}^p$ on an annual basis and run predictive regressions similar to those in my empirical exercise. Figure (2C) is the distribution of b_x for each MSA. The pooled coefficient is represented by the black line. Table (2D) are the results from the tercile (model (1)) and supply elasticity (model (2)) regressions from section (6.5).

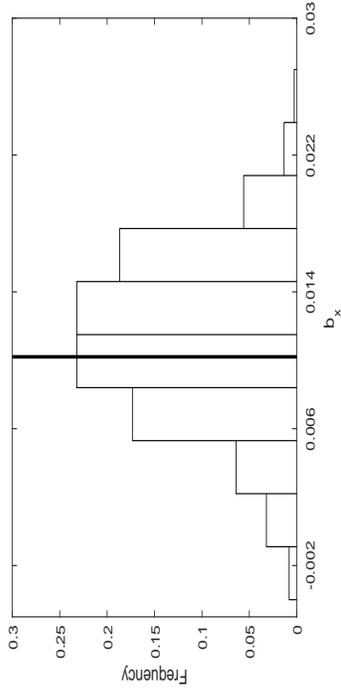
A: Expected Excess Return



B: Expected Excess Appreciation



C: Empirical Replication



D: Further Implications

$\mathbb{I}_{tr1} \cdot x_{t-1}$	0.0127***	0.0122***
	[2.83]	[4.49]
$\mathbb{I}_{tr2} \cdot x_{t-1}$	0.0183***	-0.0056***
	[5.23]	[-10.43]
$\mathbb{I}_{tr3} \cdot x_{t-1}$	0.0143***	
	[5.61]	
x_{t-1}		
$\log(\eta_{msa}) \cdot x_{t-1}$		

regressions in section (5.3). With this simple model I am able to closely replicate my empirical findings.

6.5 Further Empirical Implications

There are two additional empirical implications of my model. First, the relationship between long-run growth prospects and returns should not be linear. The sensitivity should be greatest when prospects are good, but low when prospects are either extremely good or extremely bad. To test this hypothesis, I run regression (4) with the full-set of controls described in section (5), but split the sensitivity of returns, b_x , into terciles of x_t ,

$$r_{msa,t}^p = b_1 \mathbb{I}_{tr1} x_{msa,t-1} + b_2 \mathbb{I}_{tr2} x_{msa,t-1} + b_3 \mathbb{I}_{tr3} x_{msa,t-1} + \text{Time FE} + \text{MSA FE} + \text{Controls} + \varepsilon_{msa,t}.$$

The results of the regression are in the first column of table (2D). Visually, although their economic difference is small, the difference in coefficient values fits the intuition developed from my model. Statistically, the differences in sensitivity is significant. A Wald test of equality between all three is rejected at a 10% level. The test statistic is largely driven by the rejection of equality of the coefficients of the second and third terciles, both of which are extremely well identified.

Second, as discussed in section (6.1), my model results are a function of the relative, not absolute, durability of housing services versus non-housing consumption. The degree of durability could vary across MSAs. Elasticity of housing is generally a function of two variables: physical limitations to the expansion of housing, and regulator restrictions on land-use (see Gyourko and Molloy (2014) for references). For example, in weak zoning MSAs, such as Houston, poor local growth prospects should lead to quicker shifts in industry composition. This implies that housing services are not *explicitly* durable when prospects are bad and expected excess returns and price appreciation should therefore display less sensitivity to the level of long-run growth prospects in these cities. That is, cities that are less constrained would provide agents less of a hedge against long-run growth prospects. Han (2014) makes a similar argument in the context of her empirical

work.

In urban economics there is a rich literature relating house price *levels* to elasticity. For example, Saiz (2010) generates empirical measures of MSA-level elasticity that embed both the physical and regulatory components of elasticity. As my analysis is looking at the risk-return tradeoff, I'm interested in changes to the degree of predictability of long-run prospects on expected housing returns given different levels of housing elasticity. I thus add the interaction of the Saiz measure, η_{msa} , and long-run local prospects to regression (4) to test my hypothesis.

$$r_{msa,t}^p = b_x x_{msa,t-1} + b_{saiz} \log(\eta_{msa}) x_{msa,t-1} + \text{Time FE} + \text{MSA FE} + \text{Controls} + \varepsilon_{msa,t}$$

The results of this regression are in the second column of table (2D). As hypothesized, the sensitivity of expected excess returns to long-run growth prospects is lower in more elastic MSAs. A one standard deviation higher elasticity reduces the sensitivity by nearly 60bps per year.

7 IDENTIFICATION OF DURABILITY

Any result from the model presented in section (6) hinges on the durability of housing services. Identifying the kink is thus key to understanding the driver of the pricing results seen in section (5). I do this using the cross-section of consumer units (households) from the Consumer Expenditure Survey (CEX) public-use microdata from 1990- 2011. The CEX data, which has been used in previous research (e.g., Piazzesi, Schneider, and Tuzel (2007), Meyer and Sullivan (2013)), provides households' expenditures for various goods on a quarterly basis. I utilize entries for housing services (shelter) and total expenditures over all households within a given principal-sample-unit (henceforth PSU) to extract a cross-sectional estimate of the housing to non-housing expenditure ratio (equation (13)) through time.

The survey also includes various characteristics of the individual households. The characteristics I use to isolate "local" consumption behavior are the population bucket of the PSU in which households reside (popsize), the state in which they reside (state), an indicator of whether they reside in a MSA or not (smsastat), and whether they are owner-occupiers or renters (cutenure). PSUs are composed of 1 or more (2.9 on average) MSAs and defined in a crosswalk available on the Consumer Expenditure BLS website. Specific PSU identifiers are also available 2007 onward for a sample of 22 PSUs with population greater than 1.25mm ($\text{popsize} \leq 2$). The larger PSUs compose nearly 50% of the households surveyed over these years. From 1990-2006, I identify the PSU of residence by linking households within MSAs of the same popsize and state. Given the time-series of PSU populations, which I can estimate from county-level population data, I can produce accurate expenditure ratios for the 22 largest PSUs (see <http://www.census.gov/popest/>). By limiting my analysis to just the largest popsizes (≤ 2), I isolate, with a high degree of certainty, specific PSUs. Correlation between my generated measure and the CEX identified PSUs is a sta-

tistically significant 65% for overlapping years (2007-2011). Although the CEX data runs from 1984, I only use data 1990 and after as there are no state identifiers before that year. Finally, renters tend to lock in expenditures on an annual or greater basis through a leasing contracts; given my quarterly data, I therefore focus on the consumption behavior of only owner-occupiers ($\text{cutenure} \leq 3$).

7.1 Level Regression

In order to use the CEX data there are two modifications that need to be made to equation (12). First, the data is in expenditure rather than quantity terms. As discussed in section (6.3), using the link between prices and quantities in equilibrium, I obtain equation (14). Second, my model ignores seasonality in $e_{sc,t+1}$ by assuming i.i.d. shocks. For my empirical specification, I need to add quarterly dummies. The modified regression formulation is

$$\omega_{t+1} = \mu' + (1 - \lambda) \omega_t + \left(1 - \frac{1}{\epsilon}\right) (\phi_s - \phi_c) \mathbb{I}_{x>0} x_t - \left(1 - \frac{1}{\epsilon}\right) \phi_c \mathbb{I}_{x<0} x_t + \text{Quarter Dummies} + \left(1 - \frac{1}{\epsilon}\right) e_{sc,t+1}.$$

This reduces to a simple regression that I can run on the panel of CEX data,

$$\omega_{psu,t+1} = b_\lambda \omega_{psu,t} + b_{x>0} \mathbb{I}_{x>0} x_{psu,t} + b_{x<0} \mathbb{I}_{x<0} x_{psu,t} + \text{Quarter Dummies} + \text{PSU FE} + e_{psu,t+1}. \quad (15)$$

For my main results, as the number of PSU's is small, I only cluster along the time dimension, but include PSU-level fixed effects. First, given durability of housing services, I would expect the coefficient in front of $\mathbb{I}_{x<0} x_{psu,t}$ to be less than zero. That is, expenditures on housing will fall less than non-housing when times are bad. Second, given the model calibration, I would also expect the coefficient in front of $\mathbb{I}_{x>0} x_{psu,t}$, $b_{x>0} = \phi_s - 2\phi_c$, to be positive. I present the results of this regression on the full-sample in the first column of table (10). All variables are standardized, rep-

Table 10: Identification Regressions

This table reports the regression results of equation (15) and (16). As the number of principal-sampling-units (PSUs), which are an amalgamation of an average 2.9 close-in-proximity MSAs, is relatively small (22), I only cluster along the time dimension. I also include PSU-level fixed effects. All regressors are standardized. Model (1) is the regression (15) on the full sample. As shown in Piazzesi, Schneider and Tuzel (2007) using annual-national NIPA data, the expenditure share of housing exhibits almost unit root behavior. As a check, in model (2), I deviate from my model by removing the no-bubble condition (i.e. assume $\lambda = 0$ and $\mu_s = \mu_c$) in equation (15). Model (3) restricts the sample to < 2006 . Models (4)- (5) are results from the regression equation (16). Model (4) regresses (t+1) expected variance onto my time t regressors. Model (5) regresses time (t+2) my time t regressors.

	(1)	(2)	(3)	(4)	(5)
	β	β	β	β	β
	(t-stat)	(t-stat)	(t-stat)	(t-stat)	(t-stat)
$\mathbb{I}_{x_t > 0} x_t$	0.0053 (0.73)	0.0155** (2.27)	0.0137** (1.97)		
$\mathbb{I}_{x_t < 0} x_t$	-0.0258*** (-3.14)	-0.0166** (-1.78)	-0.0418*** (-5.22)		
$\hat{P}_t (\mathbb{I}_{x_t > 0} = 1)$				0.0010*** (6.12)	0.0009*** (5.17)
Inc. ω_{t-1}	+		+		
Quarter Indicators	+	+	+		
Inc. $x_t^2 (\hat{P}_t - \hat{P}_t^2)$				+	+
Observations	1,776	1,776	1,412	1,333	1,355
R^2	0.1217	0.0078	0.1173	0.0206	0.0169
Wald Test	0.0253	0.0095	0.0000		

representing the marginal effect from a one standard deviation move in the regressor. Our results unequivocally show that $b_{x < 0} < 0$. The evidence that $b_{x > 0} > 0$ is mixed, although it is harder to pinpoint empirically due to its dependence on two parameters, ϕ_s and ϕ_c . Ultimately, the strongest test of the kink is a rejection of the null hypothesis that the two coefficients are equal, $b_{x < 0} = b_{x > 0}$, which is rejected at 5% significance level. The p-values for the Wald statistic are provided at the bottom of table (10).

As shown in my calibration, the expenditure share of housing exhibits almost unit root behavior - i.e. b_λ close to 1. In the regressions presented in columns 1 of table (10), $b_\lambda \in (0.23, 0.27)$. The value of λ is significantly higher than those used in my model calibration and shown in my national

data regressions. This may be a function of us pooling coefficients across different PSUs on a quarterly frequency basis. There may for example be MSA-level variability in the coefficients on quarterly dummies or in λ itself. A mis-specification of this sort could lead to non-i.i.d. errors and perhaps some higher frequency correlations between ω_{t+1} and ω_t . As a check, I deviate from my model by removing the no-bubble condition (i.e. $\lambda = 0$ and $\mu_s = \mu_c$). My regression specification thus simplifies to

$$\omega_{psu,t+1} - \omega_{psu,t} = b_{x>0}\mathbb{I}_{x>0}x_{psu,t} + b_{x<0}\mathbb{I}_{x<0}x_{psu,t} + \text{Quarter Dummies} + \text{PSU FE} + e_{psu,t+1}.$$

In Column (2), I present the results of this specification. $b_{x<0}$ is still significantly negative, while $b_{x>0}$ is now positive and significant. Column (3) is the same regression as column (1), but restricts the sample to the pre-housing boom and bust cycle (< 2006). There is stronger evidence in this regression that $b_{x>0}$ is positive.

7.2 Level Regression Robustness

In table (11), I present various robustness specifications for the level regression. One concern is that the results are due to differing “stickiness” between prices of non-housing goods and housing services. If housing service prices were relatively sticky, a fall in x_t would mechanically lead to a greater fall in $p_t^c C_t$ than $p_t^s S_t$, regardless of whether the kink existed. Of course the opposite is also true; one would also expect the coefficient on $\mathbb{I}_{x>0}x_t$ to be negative. In addition, mortgage payments tend to be fixed for long durations and are a fairly large portion of expenditures on shelter. Although the rotating sampling done by the BLS when executing the survey will help mitigate this issue, the stickiness may be institutional in nature. Columns (1)-(3) of table (11) uses the renters data from the CEX panel dataset to estimate the expenditure shares, ω_t . The hypothesis would be that one should not see the kink in the rental sample; as rents tend to be fixed for a year or more, expenditures would be extremely sticky to economic fluctuations. The regression

Table 11: Identification Robustness Regressions

This table reports robustness results for regression (15). Details on standard errors, fixed effects and controls are identical to those applied in table (10). One concern of ours is that results are due to differing “stickiness” between prices of non-housing goods and housing services, primarily due to fixed long-duration contracts. Models (1)-(4) uses renter- only data to estimate the dependent variable, expenditure share (ω_t). Another concern is the possible dominance of data from the recent boom-crisis period. The bust in particular has been argued to be the result of an abrupt contraction of credit and financial liberalization. Models (2) and (4) restrict the sample to only the pre-boom-crisis subperiod (< 2006). Models (3)-(5) uses the cash flow only residual (estimation described in section (5.1)) as my measure of x_t .

	(1)	(2)	(3)	(4)	(5)
	β	β	β	β	β
	(t-stat)	(t-stat)	(t-stat)	(t-stat)	(t-stat)
$\mathbb{I}_{x_t > 0} x_t$	-0.0218* (-1.54)	-0.0115 (-0.76)	-0.0146 (-1.20)	-0.0081 (-0.62)	0.0018 (0.26)
$\mathbb{I}_{x_t < 0} x_t$	-0.0251** (-1.92)	-0.0374** (-2.29)	-0.0298*** (-2.87)	-0.0360*** (-2.75)	-0.0236*** (-3.69)
Owner sample					+
Rent-only sample	+	+	+	+	
Quarter Indicators	+	+	+	+	+
Observations	1,759	1,395	1,759	1,395	1,776
R^2	0.1379	0.1179	0.1362	0.1165	0.1227
Wald Test	0.8951	0.3718	0.4572	0.2320	0.0274

results reflect this stickiness in the rental consumption data - the Wald statistic for $b_{x < 0} = b_{x > 0}$ is not statistically rejected. The main issue of using this subsample, however, is that the sample size falls dramatically. On average there are approximately 100 homeowner households that make up each of my PSU-level ω_t estimates in table (10). My primary assumption is that over time these 100 households are representative of the general PSU population. As renters compose only about 35% of the broader sample, a firm statement on the owner vs. rental markets is difficult as the smaller sample size significantly increases the variance in the dependent variable.

Another obvious concern is that the housing boom-crisis period (post 2005) dominates the rest of the sample. The bust in particular has been argued to be the result of an abrupt contraction of credit and financial liberalization (see Favilukis, Kohn, Ludvigson and Van Nieuwerburgh (2013) for review of literature). As some households lost their homes during this period, in the context of my

model, the discontinuity in my sc_t process might be semi- rather than full-durability - i.e. the exposure of Δs_t to long-run prospects is positive when $x_t < 0$, but smaller than when $x_t > 0$. Columns (2), (4), and (6) restrict the sample to only the pre-boom-crisis subperiod (< 2006). If anything, the relationship becomes stronger, suggesting that additional dynamics (i.e. those outside my model) were in play during the crisis.

Finally, as discussed in section (1) and (2), my motivation and model are largely based on a cash-flow story. In section (4) I generated an alternative measure of x_t that attempts to isolate the cash-flow growth component of local prospects. Columns (3)-(6) use this measure of x_t in regression (15). Economically, this alternative measure seems to have little impact on my results.

7.3 Variance Regression

The second dynamic that I seek to identify is the relationship between the variance of relative consumption and local prospects due to the durability of housing services. This can be seen explicitly from equation (12),

$$\log \frac{S_{t+1}}{C_{t+1}} = sc_{t+1} = (1 - \lambda)sc_t + \mu + (\phi_s - \phi_c) \mathbb{I}_{x>0}x_t - \phi_c \mathbb{I}_{x<0}x_t + e_{sc,t+1}.$$

When $x_t > 0$ the variance of sc_{t+1} is a function of the variance of $(\phi_s - \phi_c)x_t$ and $e_{sc,t}$, whereas when $x < 0$ the variance of sc_{t+1} is a function of $\phi_c x_t$ and $e_{sc,t}$. That is the variance of housing to non-housing consumption has different regimes given the value of the state variable. Given my calibrations (i.e. $\phi_s > 2\phi_c$), I would expect the variance of sc_t to increase in x_t .

To properly see this relationship in the CEX data, I need to foil equation (12) two periods forward given that my model is in discrete time. In addition, similar to equation (15), I need to account for the fact that this data is in expenditure rather than quantity terms using the intra-temporal equilibrium relationship between prices and quantities. Details of this derivation are in APPENDIX

J. My regression specification is

$$\begin{aligned}
 Var_t(\omega_{psu,t+1}) = & \underbrace{b_{int} \left(P_{psu,t}(\mathbb{I}_{x>0} = 1) - P_{psu,t}(\mathbb{I}_{x>0} = 1)^2 \right)}_{\text{Term 1}} \times \underbrace{x_{psu,t}^2}_{\text{Term 2}} + \\
 & \underbrace{b_{var} P_{psu,t}(\mathbb{I}_{x>0} = 1)}_{\text{Term 3}} + \text{Time FE} + \text{PSU FE} + \eta_{psu,t+1}. \quad (16)
 \end{aligned}$$

While equation (16) may seem complicated, it appeals to simple intuition. Term 1 is the variance of a binary random variable and is greatest at $x_t = 0$. Given an unlinked processes, one would assume that sc_t would have the greatest variance around the kink. However, s_{t+1} and c_{t+1} are linked conditionally by x_t ; when $x_t = 0$ only the short-run shocks (i.e. $e_{s,t+1}$ and $e_{c,t+1}$) matter in the expected variance estimate. As these two terms decay and increase at different speeds, their interaction appears in equation (16). Term 1 \times Term 2 is symmetric around $x_t = 0$; for reasonable values of x_t it increases as one moves away from the kink. It also does not load significantly in the regression results I will show. Term 3, on the other hand, reflects the a similar regime switching dynamic as the endowment of housing service growth. Assuming $\phi_s > 2\phi_c$, the value of term 3 will transition from 0 to 1 as x_t goes from $-\infty$ to ∞ . In our regressions, b_{var} will reflect the degree on average to which the variance of the expenditure ratio changes as the economy moves from a very bad to very good economic growth prospect state.

$P_{psu,t}(\mathbb{I}_{x>0} = 1)$ can be estimated via a panel probit. I regress $\mathbb{I}_{x>0}$ next period onto $x_{psu,t}$ today, where the coefficient β_{probit} is pooled across all 22 PSUs. Given the high persistence of $x_{psu,t}$, b_{probit} is well identified (see table (12) for results). I then estimate the left hand side of equation (16) in a three step process. First, I estimate the expected levels of ω_{t+1} given today's x_t from equation (15), subtracting it from the realized expenditure ratio tomorrow to compute an estimate of the residual, $(1 - \frac{1}{\varepsilon}) e_{sc,t+1}$. Second, my measure of *realized* volatility, \widehat{Var}_t^{sc} , is a moving average of these residuals. In order to balance picking up the persistence of expected volatility in the residuals with my relative short panel, I use the three-year (12 quarter) moving average to estimate this variable. Third, my measure of *expected* volatility of the expenditure ratio,

Table 12: Identification Auxiliary Regressions

This table reports results from the auxiliary regressions in section (7.3). The coefficients from these regressions are used to estimate the probability of $x_t > 0$ next period given today's economic state and my measure of expected volatility in equation (16). Model (1) are the results from my probit regression. My measure of realized volatility, \widehat{Var}_t^{sc} , is a three-year (12 quarter) moving average of the residuals from equation (15). Expected volatility is then estimated from a pooled-panel AR(1) regression. Model (2) are the results from that regression.

	(1)	(2)
	β	β
	(t-stat)	(t-stat)
x_t	2.5701*** (5.76)	
\widehat{Var}_t^{sc}		0.5194*** (5.53)
Observations	1,619	1,355
R^2		0.0461

$Var_t(\omega_{psu,t+1})$, is then computed from a predictive regression,

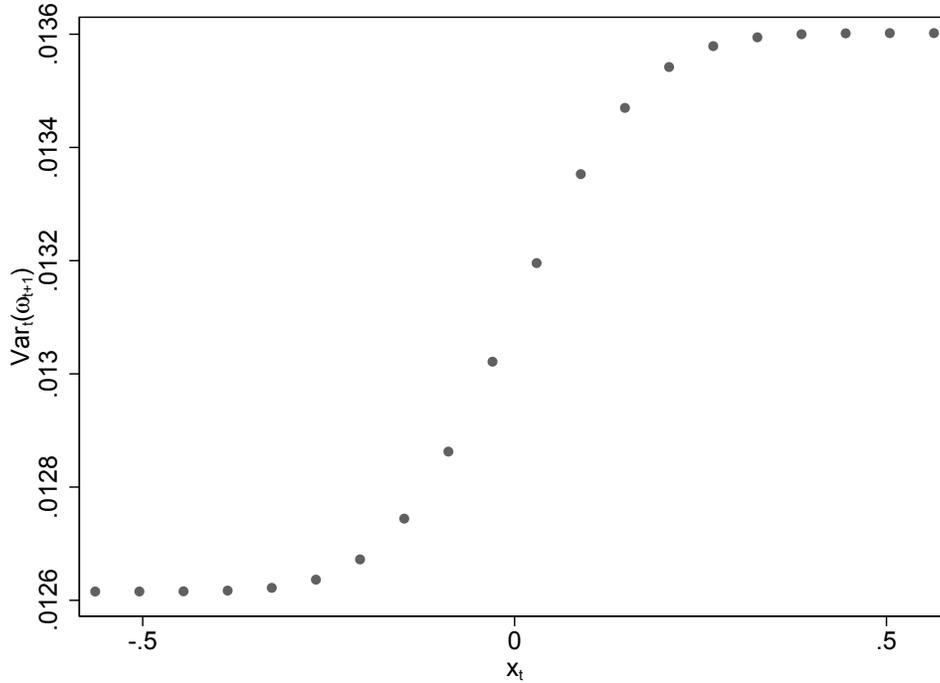
$$\left(1 - \frac{1}{\varepsilon}\right)^2 e_{psu,t+1}^2 = v_{psu,0} + \rho_{var} \left(\widehat{Var}_{psu,t}^{sc} - v_{psu,0}\right) + \eta_{psu,t+1},$$

where $e_{psu,t+1}$ is embedded in the error form equation (15) and $v_{psu,0}$ is a PSU-level fixed effect representing the average variance of relative consumption. The estimated ρ_{var} is presented in table (12). The period t expectations of this auxiliary regression provides us with an estimate of $Var_t(\omega_{t+1})$. Similar to regression (15), I only cluster along the time dimension, but include PSU-level fixed effects. I also standardize all regressors.

If $\phi_s > 2\phi_c$, I should expect b_{var} to be positive and significant. I present the results from this regression on the full-sample of PSUs in column (4) and (5) of table (10). The term loads as expected were there a kink in the Δs_{t+1} process; β_{var} is positive and significant. One would expect b_{int} to load positively given that it corresponds to the model parameters $\phi_s^2 \rho^2$. Although it loads negatively, the term is both statistically and economically insignificant. The non-monotonicity of

Figure 3: Estimated $Var_t(sc_{t+1})$ versus x_t

This figure represents the increase in expected variance of the expenditure ratio of housing services to non-durable consumption versus my measure of long-run prospects, x_t . There are two variables in regression (16), one of which is non-monotonic. The net effect makes interpretation of the regression coefficients presented in table (10) difficult. I thus plot $Var_t(\omega_{t+1})$ given x_t using the estimated coefficients - b_{probit} , b_{int} and b_{reg} - and average (across PSUs) fixed-effects from my regression to show the underlying monotonicity of my result. ω_{t+1} is the relative expenditures of housing to non-housing goods.



terms makes interpretation a little difficult across x_t . Figure (3) plots $Var_t(\omega_{t+1})$ given x_t , using the estimated coefficients - \hat{b}_{probit} , \hat{b}_{int} and \hat{b}_{var} - and average (across PSUs) fixed-effects from my regression. Due to the durability of housing the variance of the expenditure ratio clearly increases in x_t , from roughly σ_{sc}^2 to $b_{reg} + \sigma_{sc}^2$. This monotonic increase requires that Term 3 dominates Term $1 \times$ Term 2 in terms of marginal effect.

The cross section of PSUs seems to provide power in identifying this slow moving, but multi-period phenomenon. This, however, does not change the simple fact that the magnitude of σ_x^2 is so much smaller than that of σ_{sc}^2 . This is exemplified in the low R^2 of my fixed effects regression results and captured in the small relative increase in variance over x_t in figure (3).

8 CONCLUSION

This paper studies how dynamics of local prospects influence local consumption behavior and ultimately drive real estate price dynamics. I first generate a measure that captures heterogeneity in long-run expected growth between different metropolitan areas. I source this heterogeneity from both the cross-section of regional employment composition and industry prospects.

My measure of industry growth prospects is extracted from aggregate industry PD-ratios. My local measure of long-run growth prospects are then estimated using the procedure from Bartik (1991); each aggregate measure of industry prospects is weighted by the employment share within each MSA. I find that when I regress the local excess housing returns onto these real measures of local long-run prospects, a one standard deviation increase in my proxy of local prospects leads to a nearly 120bp increase in annual excess return to housing. Appealing to the Campbell and Shiller (1987) decomposition, I then use novel real estate capitalization rate data from Integra Realty Resource, to show that in “good” times the increase in risk premia dominates the conditional increase in housing services. I find that a one standard deviation increase in local prospects lowers the log price-to-rent by more than 500bps. That is, when prospects are good, the price-to-rent contemporaneously falls.

This implies that bad news about local growth prospects lower the risk to local housing assets, which suggests that housing acts as a hedge against long-run growth prospects. I then develop a consumption based equilibrium model with explicitly durable housing services to better understand this relationship. The durability of housing structures is key to understanding many multi-decade phenomenon in urban and labor economics (see e.g., Glaeser and Gyourko (2005) and Notowidigdo (2013)). I argue that this permanency can extend to the location specific services housing provides. In the model, agents’ non-housing goods and housing service consumption

growth are exposed to a common state variable that mimics my measure of long-run prospects, which are highly persistent AR(1) processes. This persistence amplifies the effects of shocks to local growth prospects on the agent's actions over multiple periods via recursive Epstein and Zin (1989) preferences. The effects of durability on housing's excess returns can then be simulated in my model. Using this simple economic framework, I am able to replicate my empirical findings.

In addition, durability creates endogenous time varying volatility in a key pricing variable, the relative composition of housing services and non-housing good consumption (henceforth, relative consumption). Although small in magnitude the increase in volatility of relative consumption as prospects increase is amplified by the processes persistence. This is the key channel for the increase in risk premia for housing. I identify durability and the characteristics of the time varying volatility in data from the consumer expenditure survey (CEX). During bad economic times, a one standard deviation fall in prospects leads to a 1000bp rise in relative consumption. During good economic times, a one standard deviation rise in prospects leads to a much smaller rise in relative consumption. This suggests that local households raise consumption of housing services in lock step with non-housing consumption during times of good long-run growth prospects, but do not drop them as much during times of bad long-run growth prospects. In addition, the variance of relatively consumption increases by more than 700bps from the low to high growth states.

This paper extends work done in the long-run risk and housing asset pricing literature. Housing services are a major portion of household expenditures. Expenditures are also highly geographically heterogeneous. My empirical strategy makes use of this heterogeneity in understanding the pricing of local assets. My study is different from recent macroeconomic work in housing in that my framework also generalizes beyond just the recent financial crisis. Housing collateral and financial liberalization played an outsized role in the boom-bust of the mid-to-late 2000s. However, fluctuations in multi-decade local growth expectations should also weigh heavily on an agent's mind. In this regard, my study is similar to geography-centric work in urban economics. I add to this literature by focusing on the role that fluctuations in "lower" frequency, long-run dynamics in expectations play in housing returns.

APPENDIX A Gauss-Newton Conversion Algorithm

Some industries do match 1-to-1 between classification systems - e.g., in manufacturing, mining, agriculture and transportation (i.e. “old” industries). Many SIC accounts, however, are partially split between many different NAICS accounts - e.g., services or information technology (i.e. “new” industries). Indicators are given to each correspondence and then rolled-up to the BEA-available primarily 3-digit level. The first line of table (2A) provide summary statistics of how many NAICS industries, N , each SIC account converts. On average each of my 60 SIC industries convert into roughly 4 NAICS accounts. This, however, belies the degree of skewness of this conversion “count” number. While the maximum is 19, even at the 75th percentile a single SIC industry only converts into 5 NAICS accounts.

The BEA SIC and NAICS data have common years from 1977 to 1997. Using the correspondence computed above and these common years, I use a simple, iterative algorithm to estimate SIC-to-NAICS conversion weights for the years NAICS data is not available (details below). These weights are then used to convert SIC value-added and employment data to NAICS for 1947-1976. They are also saved for conversion within other data sets where I have similar limitations. For example, many now defunct or privatized companies are available before 2001 on CRSP and Compustat, but only have associated SIC codes. For my price-to-dividend ratio computation, I aggregate market-cap and dividends along the BEA SIC available accounts for these companies and use the conversion matrices estimated from the Gauss-Newton algorithm on value-added GDP to estimate their contributions to the PD’s of various NAICS accounts.

Given the count data, as I’d expect, the 60 SIC accounts map to roughly 240 (60×4) NAICS accounts. The mean conversion weight is thus $1/4$. As shown in table (2A), the 247 computed weights are skewed towards 1 and 0. Given the industry composition change in employment and output during this period, I also present the analysis using different samples (full 1977-1997, 1977-only, and 1997-only) of weights. Overall the median mapping between SIC and NAICS seems to be much closer to 1-to-1 than implied by the mean of 1-to-4 and is robust over my entire overlap-

ping sample for both employment and GDP. This is unsurprising given that employment compositional change is generally very slow. Both the high degree of correspondence in count and weight statistics thus give me comfort that my converted data is accurate.

Details of Gauss-Newton Algorithm:

1. For the 1997 data an equal weight, $w_{old}^j = \frac{1}{N}$, is given for each NAICS account, j , into which a single SIC account, i , converts for the initial iteration. An estimate of each NAICS account is computed, $\widehat{NAICS}_{1997}^j = SIC_{1997}^i \times w_{old}^j$.
2. Given that I know the actual NAICS value, the error in either GDP or unemployment is then computed for each estimated account, $NAICS_{1997}^j - \widehat{NAICS}_{1997}^j = \varepsilon_{1997}^j$.
3. The old weight is updated to account for the relative size of error. By virtue of my adjustment mechanism, NAICS accounts with larger errors will have relatively larger shifts upwards (positive error) or downwards (negative error). In addition, I maintain unity sum and positive weights, i.e.

$$w_{new}^j = \frac{w_{old}^j \exp(\varepsilon_{1997}^j)}{\sum_{k=1}^N w_{old}^k \exp(\varepsilon_{1997}^k)}.$$

4. Steps 1-3 are repeated using a Gauss-Newton convergence algorithm, using a tolerance parameter, η , and threshold parameter, τ , of 10^{-5} and 10^{-3} , respectively. I stop the iterating when

$$\frac{\|\theta_{j+1} - \theta_j\|}{\|\theta_j\| + \tau} \leq \eta, \quad \text{where, } \|\theta\| = (w_1^2 + w_2^2 + \dots + w_N^2)^{\frac{1}{2}}.$$

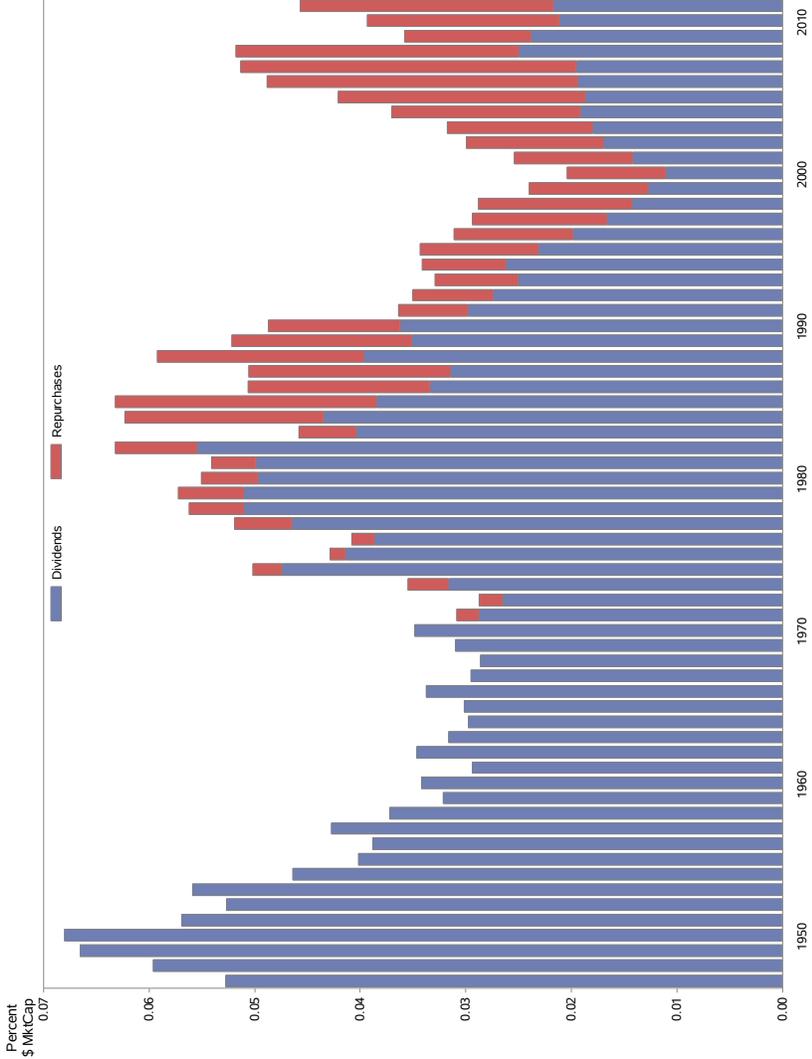
5. Steps 1-4 are repeated backwards for each common year from 1997 to 1977. For each year-iteration (step 4), the initial estimate is the converged upon weights from the Gauss-Newton algorithm using the “proceeding” year’s data. The basic impetus for this approach is that the weights likely shifted as new industries were born (e.g., internet) or changed in

prominence (e.g., health care or banking). I seek to capture these dynamics by computing year-by-year SIC-to-NAICS conversion weights. By using the forward (t) computed as the initial guesses for the previous years (t-1) weights, I also maintain consistency and minimize conversion errors for the 1947 to 1976 SIC data.

6. The 1947 to 1976 SIC data is converted to NAICS using the 1977 computed weights. The BEA also provides NAICS “super” aggregate (e.g., general non-durable manufacturing account) value-added GDP and employment values. I thus scale my converted values such that the sum of aggregate accounts (e.g., all non-durable manufacturing add-values) is equal to the given super-aggregate value. My weights should accurately capture SIC to NAICS conversions near 1977; however, as the years approach 1947, the error between actual (unobserved) and estimated values will naturally increase. The scaling seeks to mitigate these errors. For these “later” years one would also assume that all else equal the greatest errors in measurement will be for the super-aggregate accounts composed of the greatest number of aggregate accounts. The three largest super-aggregate accounts are durables, non-durables, and transportation. Fortunately these are generally “old” industries, which also have the greatest number of 1-1 correspondences.

APPENDIX B Dividend and Repurchase Yield

This figure is an update of the main finding from Boudoukh, Michaely, Richardson and Roberts (2007). For the past 20+ years the value of repurchases have become an important component of the return of capital to shareholders. Today they account for >50 % of the average equity market “yield.” In my case, including this payout is critical in comparing industry prospects due to the significant heterogeneity in repurchases yields (RP_t) across industries. Following Boudoukh, et al. (2007) I define repurchases as expenditures on the purchase of common and preferred stock (Compustat #115) plus any reduction in the value of the net outstanding preferred stock (Compustat #56). Since the universe of Compustat and CRSP stocks do not overlap 100%, I also download the market capitalization of Compustat stocks to compute the cash dividend and repurchase yields separately. My measure of industry and aggregate PD-ratio is then the inverse of the sum of these yields.



APPENDIX C Annual GMM

System (1) generates 10 orthogonality conditions and is thus over-identified. In addition all variables are demeaned (i.e. fixed effects ignored). The moments used are,

1. $E [\eta_t^i] = \frac{1}{T} \sum_{t=1947}^{2011} (pd_t^i - \gamma pd_t^{mkt}) = 0,$
2. $E [pd_t^{mkt} \eta_t^i] = \frac{1}{T} \sum_{t=1947}^{2011} (\eta_t^i) (pd_t^{mkt}) = 0,$
3. $E [e_t^i] = \frac{1}{T} \sum_{t=1948}^{2012} (\Delta y_t^i - \beta pd_{t-1}^{mkt} - \alpha \eta_{t-1}^i) - \delta r_{t-1}^f = 0,$
4. $E [e_t^i pd_{t-1}^{mkt}] = \frac{1}{T} \sum_{t=1948}^{2012} (e_t^i) (pd_{t-1}^{mkt}) = 0,$
5. $E [e_t^i \eta_{t-1}^i] = \frac{1}{T} \sum_{t=1948}^{2012} (e_t^i) (\eta_{t-1}^i) = 0,$
6. $E [e_t^i \eta_{t-1}^i] = \frac{1}{T} \sum_{t=1948}^{2012} (e_t^i) (r_{t-1}^f) = 0,$
7. $E [\epsilon_t^i] = \frac{1}{T} \sum_{t=1947}^{2011} (x_t^i - \rho x_{t-1}^i) = 0,$
8. $E [\epsilon_t^i x_{t-1}^i] = \frac{1}{T} \sum_{t=1947}^{2011} (\epsilon_t^i) (pd_{t-1}^{mkt}) = 0,$
9. $E [\epsilon_t^i x_{t-1}^i] = \frac{1}{T} \sum_{t=1947}^{2011} (\epsilon_t^i) (\eta_{t-1}^i) = 0,$
10. $E [\epsilon_t^i x_{t-1}^i] = \frac{1}{T} \sum_{t=1947}^{2011} (\epsilon_t^i) (r_{t-1}^f) = 0.$

The specification above is for annual data - i.e. $m=11$ when computing PD-ratios in section (4.2). In general, however, companies pay dividends on a quarterly basis. By aggregating dividend data at the lower frequency, I am likely losing information when identifying long-run prospects. Economically, the frequency of data for estimation can also be very important. First, as highlighted by Bansal, Kiku and Yaron (2012), time aggregation is critical for consistently estimating model parameters and accurately mapping annual sampled data to discrete time models where decisions are made at a higher frequencies. Second, given the large cross section and short time series of my dependent variable in the main set of empirical results, short-run and long-run shocks can easily be conflated. Previous literature has shown that business cycle variations are important determinants of local asset prices; this distinction may be important for my analysis.

APPENDIX D Extending Annual GMM to Quarterly

Given that I am now working at different frequencies my notation will slightly change from the annual GMM. In particular, $q \in (1, 2, 3, 4)$ refers to the quarter within a particular year t . The strategy is to maintain an annual GMM structure, which is the lowest frequency data available, but roll-up the quarterly data in a way that maintains the orthogonality of errors - i.e. maintains consistency of parameter and standard error estimates.

For the first equation in system (1), given that $pd_{q,t}^{mkt}$ is exogenous I can simply collapse the quarterly specification of my model within year for my mixed-frequency GMM, i.e.

$$\sum_{q=1}^4 pd_{q,t}^i = \gamma \sum_{q=1}^4 pd_{q,t}^{mkt} + \sum_{q=1}^4 \eta_{q,t}^i. \quad (17)$$

Defining η_t^i as $\sum_{q=1}^4 \eta_{q,t}^i$, I get the first two moment conditions. For the predictive equation, collapsing the moments within year, i.e.

$$\underbrace{\sum_{q=1}^4 \Delta y_{q,t}^i}_{\text{observable}} = \underbrace{\sum_{q=1}^4 x_{q-1,t}^i}_* + \underbrace{\sum_{q=1}^4 e_{q,t}^i}_{**},$$

would yield inconsistent parameter estimates since x_t^i is an AR(1) process - i.e. * and ** are not orthogonal. By expanding out * intra-year I obtain,

$$\begin{aligned} \Delta y_t^i &= (1 + \rho + \rho^2 + \rho^3) x_{4,t-1} + \\ &\underbrace{\rho (\epsilon_{1,t}^i + \epsilon_{2,t}^i + \epsilon_{3,t}^i) + \rho^2 (\epsilon_{1,t}^i + \epsilon_{2,t}^i) + \rho^3 (\epsilon_{1,t}^i) + \sum_{q=1}^4 e_{q,t}^i}_{=e_t^i}. \end{aligned} \quad (18)$$

Aggregating the errors as above, I now satisfy error orthogonality and obtain the next four moment conditions for my GMM. For the autoregressive equation, I similarly need to expand and then

collapse the errors intra-year to obtain the last two moment conditions, i.e.

$$\begin{aligned}\sum_{q=1}^4 x_{q,t}^i &= \rho \sum_{q=1}^4 x_{q-1,t}^i + \sum_{q=1}^4 \varepsilon_{q,t}^i \\ &= \rho (1 + \rho + \rho^2 + \rho^3) x_{4,t-1}^i + \varepsilon_t^i.\end{aligned}\tag{19}$$

The moment conditions used are thus,

1. $\frac{1}{T} \sum_{t=1947}^{2011} \eta_t^i = 0,$
2. $\frac{1}{T} \sum_{t=1947}^{2011} (\sum_{m=1}^4 p d_{m,t}^{mkt}) \eta_t^i = 0,$
3. $\frac{1}{T} \sum_{t=1948}^{2012} e_t^i = 0,$
4. $\frac{1}{T} \sum_{t=1948}^{2012} (e_t^i) p d_{4,t-1}^{mkt} = 0,$
5. $\frac{1}{T} \sum_{t=1948}^{2012} (e_t^i) \eta_{4,t-1}^i = 0,$
6. $\frac{1}{T} \sum_{t=1948}^{2012} (e_t^i) r_{4,t-1}^f = 0,$
7. $\frac{1}{T} \sum_{t=1947}^{2011} \varepsilon_t^i = 0,$
8. $\frac{1}{T} \sum_{t=1947}^{2011} (\varepsilon_t^i) p d_{4,t-1}^{mkt} = 0.$
9. $\frac{1}{T} \sum_{t=1947}^{2011} (\varepsilon_t^i) \eta_{4,t-1}^i = 0.$
10. $\frac{1}{T} \sum_{t=1947}^{2011} (\varepsilon_t^i) r_{4,t-1}^f = 0.$

Table (3B) gives the results of my estimation for the quarterly GMM parameters at the 25th, 50th, and 75th percentiles. The interpretation of the coefficients are the same as with the annual specification. Comparison between the two GMMs requires some explanation, however, due to the change in frequency. First the parameters of the predictive regressions - β , α , and δ - need to be scaled by a factor of 4. The persistence coefficient must also be adjusted such that $\rho_{ann} = \rho_{qtr}^4$.

The Wald test for $\alpha = \beta = \delta = 0$ shows a significant improvement in isolating the industry-level long-run prospects versus the annual specification; the null hypothesis is rejected for more than

55% of the industries at a 5% level. Overall, the parameters between the annual and quarterly specifications seem to reflect the same dynamics as in the annual specification.

APPENDIX E Campbell-Shiller Decomposition

As the return of any asset is its cash flow plus price appreciation, $R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t}$, its price can be iterated forward s.t.

$$\begin{aligned} P_t &= \frac{P_{t+1} + D_{t+1}}{R_{t+1}} \\ &= \frac{D_{t+1}}{R_{t+1}} + \frac{D_{t+2}}{R_{t+1}R_{t+2}} + \dots \\ &= \sum_{j=1}^{\infty} \frac{D_{t+j}}{R_{t,t+j}}. \end{aligned}$$

Taking the log of the initial expression and performing a Campbell-Shiller log-linear approximation around the average price-to-dividend ratio,

$$\begin{aligned} r_{t+1} &= \log(P_{t+1} + D_{t+1}) - \log(P_t) \\ &= d_{t+1} + \log(\exp(pd_{t+1}) + 1) - p_t \\ &\approx d_{t+1} + \log(1 + \exp \bar{pd}) + \frac{\exp \bar{pd}}{1 + \exp \bar{pd}} (pd_{t+1} - \bar{pd}) - p_t \end{aligned}$$

Letting $\rho = \frac{\exp \bar{pd}}{1 + \exp \bar{pd}}$ and $k = \log(1 + \exp \bar{pd}) - \rho \bar{pd}$,

$$p_t - d_t = pd_t = \Delta d_{t+1} + \rho(p_{t+1} - d_{t+1}) + k - r_{t+1}.$$

We can then iterate forward the above result to obtain the decomposition of the price-to-dividend ratio into its expected cash growth and return components,

$$p_t - d_t = \frac{k}{1 - \rho} + \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} - r_{t+1+j}).$$

APPENDIX F Cash Flow Component Extraction

From Campbell-Shiller (1987),

$$p_t - d_t = \frac{k}{1 - \rho} + \sum_{j=0}^{\infty} \rho^j (\Delta d_{t+1+j} - r_{t+1+j}).$$

Both Δd_{t+1+j} and r_{t+1+j} are of quarterly frequency in my primary specification. In order to isolate the cash flow component, I thus need a quarterly estimate of $\sum_{j=0}^{\infty} \rho^j (r_{t+1+j}^{ex} + r_{t+1+j}^f)$. I use the standard ICAPM relationship, $E_t [r_{t+1}^{ex}] = \mu + \gamma Var_t [r_{t+1}]$, to do this.

I run a GARCH(1,1) on daily industry returns from 1947-2011, i.e.

$$h_t = \omega + \alpha \varepsilon_{t-1}^2 + \beta h_{t-1}.$$

Defining the unconditional variance, $E [h_t] = \frac{\omega}{1 - \alpha - \beta}$, as σ^2 , I utilize the GARCH autoregressive structure,

$$\begin{aligned} h_{t+1|t} &= \sigma^2 + (\alpha + \beta) (h_t - \sigma^2) \\ &\vdots \\ h_{t+k|t} &= \sigma^2 + (\alpha + \beta)^k (h_t - \sigma^2), \end{aligned}$$

to estimate of quarterly variance, i.e. the sum of the daily forecasts over the quarter

$$h_{t+1|t}^q = \sum_{l=1}^{66} \sigma^2 + (\alpha + \beta)^{l-1} (h_t^d - \sigma^2).$$

The residual, e_{cf} , extracted from the regression

$$p_t - d_t = \mu + \gamma h_{t+1|t}^q + e_{cf},$$

for each NAICS industry, thus represents my alternative measure of long-run industry prospects. As the conditional and persistent component of consumption drives both the dividend growth and risk-free rate, it is next to impossible to separate the two. My measure will therefore include expected dynamics in both the dividend growth and risk-free rate.

APPENDIX G Solving Model's Equilibrium

From Epstein and Zin (1989) recursive preferences,

$$\frac{U_t^{1-\frac{1}{\psi}} - 1}{1 - \frac{1}{\psi}} = \frac{1}{1 - \frac{1}{\psi}} \left[(1 - \beta) \left(u_t^{1-\frac{1}{\psi}} - 1 \right) + \beta \left(E_t \left[U_{t+1}^{1-\gamma} \right]^{\frac{1-\frac{1}{\psi}}{1-\gamma}} - 1 \right) \right].$$

Applying L'Hospital's Rule and taking $\lim \psi \rightarrow 1$,

$$\ln U_t = (1 - \beta) \ln u_t + \beta \ln E_t \left[U_{t+1}^{1-\gamma} \right]^{\frac{1}{1-\gamma}}.$$

Defining $V_t = \ln U_t - \ln u_t$ and $\theta = \frac{1}{1-\gamma}$, I obtain my value function,

$$V_t = \beta \theta \ln E_t \left[\exp \left(\frac{V_{t+1} + \ln u_{t+1}/u_t}{\theta} \right) \right].$$

From the intra-temporal CES aggregator,

$$\ln \frac{u_{t+1}}{u_t} = \ln \left[\frac{1 - \alpha + \alpha \left(\frac{S_{t+1}}{C_{t+1}} \right)^{1-\frac{1}{\varepsilon}}}{1 - \alpha + \alpha \left(\frac{S_t}{C_t} \right)^{1-\frac{1}{\varepsilon}}} \right]^{1-\frac{1}{\varepsilon}} + \ln \frac{C_{t+1}}{C_t}.$$

Defining $\ln \frac{S_t}{C_t} = sc_t$, the Bellman equation for my value function is thus,

$$V_t = \beta \theta \ln E_t \exp \left[\left(V_{t+1} + \frac{1}{1-\frac{1}{\varepsilon}} \left(\ln \left(1 - \alpha + \alpha \exp(sc_{t+1})^{1-\frac{1}{\varepsilon}} \right) - \ln \left(1 - \alpha + \alpha \exp(sc_t)^{1-\frac{1}{\varepsilon}} \right) + \Delta c_{t+1} \right) \right) / \theta \right].$$

I solve my value function using standard iteration methods along a 5σ grid of x_t and sc_t . Expectations are computed across a 21-point gaussian quadrature, which randomly assigns shocks for e_c , e_s , e_d and ε_x .

The stochastic discount factor is then computed using the non-housing good as the numéraire,

$$M_{t+1} = \left[\frac{\partial \tilde{U}_t}{\partial C_{t+1}} \right] / \left[\frac{\partial \tilde{U}_t}{\partial C_t} \right] \quad \text{s.t.}$$

$$\frac{\partial \tilde{U}_t}{\partial C_t} = \underbrace{\frac{\partial \tilde{U}_t}{\partial u_t}}_{W_{1,t}} \cdot \underbrace{\frac{\partial u_t}{\partial C_t}}_{W_{11,t}} \quad \text{and}$$

$$\frac{\partial \tilde{U}_t}{\partial C_{t+1}} = \underbrace{\frac{\partial \tilde{U}_t}{\partial U_{t+1}}}_{W_{2,t}} \cdot \frac{\partial U_{t+1}}{\partial u_{t+1}} \cdot \frac{\partial u_{t+1}}{\partial C_{t+1}},$$

where I defined $\ln U_t = \tilde{U}_t$. Computing these various partial derivatives I obtain

$$M_{t+1} = \frac{W_{2,t} W_{1,t+1} W_{11,t+1}}{W_{1,t} W_{11,t}} = \left[\frac{\beta \exp(U_{t+1}/\theta)}{E_t \exp(U_{t+1}/\theta)} \right] \left[\frac{u_{t+1}}{u_t} \right]^{\frac{1}{\varepsilon}-1} \left[\frac{C_{t+1}}{C_t} \right]^{-\frac{1}{\varepsilon}}. \quad (20)$$

With algebraic manipulations similar to those done on the value function I obtain

$$\ln M_{t+1} = m_{t+1} = \ln \beta + \frac{V_{t+1}}{\theta} - \left(1 - \frac{1}{\theta} - \frac{1}{\varepsilon} \right) \ln \frac{u_{t+1}}{u_t} - \frac{1}{\varepsilon} \Delta c_{t+1} - \frac{V_t}{\beta \theta}.$$

Thus expectations of the SDF and therefore the risk-free rate ($1/E_t [M_{t+1}]$) for any given state, can be computed as I've characterized the value function.

Furthermore, I use the Euler equation restriction, $E_t [M_{t+1} R_{t+1}] = 1$, to pin-down the price to dividend ratios for my various assets.

$$pc_t = \log E_t \exp \left[m_{t+1} \log (\exp (pc_{t+1}) + 1) + \mu_c + \phi_c x_t + \frac{\lambda}{2} (s_t - c_t) + e_{c,t+1} \right], \quad (21)$$

$$ps_t = \log E_t \exp \left[m_{t+1} \log (\exp (ps_{t+1}) + 1) + \mu_s + I_{(x_t > 0)} \phi_s x_t + \frac{\lambda}{2} (s_t - c_t) + \sigma_s e_{s,t+1} \right], \quad (22)$$

where lower cases signify logs. Given that $r_t^s = \ln (\exp (ps_{t+1}) + 1) - ps_t + \Delta s_{t+1}$, I can now also compute the conditional expected excess returns and excess price appreciation for housing.

APPENDIX H Pricing of Aggregate Markets

My model mimics the spatial equilibrium models in urban economics, where housing acts as a wedge (Rosen (1979) and Roback (1982)). In order to maintain a no-arbitrage condition between cities, dynamics in income growth filter to both housing services and non-housing consumption; although productivity levels can be heterogenous, net utility is identical across cities. In representative agent models this condition is established by existence of a “reservation” city where housing is supplied completely elastically.

In a similar vein, in my model, a shock to “local” long-run prospects positively effect both the consumption of housing services and the non-housing good. In addition, at the aggregate markets, I assume my “reservation” agent is someone who can consume housing services just as a normal, non-housing good - that is they consume housing completely elastically. This marginal investor could, for example, be a high earner whose utility is not tied to a particular location and to whom housing, therefore, does not provide a hedge. Using a simple two MSA example the consumption endowments of the two goods is thus

$$\begin{aligned}\Delta c_{t+1} &= w_{1,t}\Delta c_{1,t+1} + (1 - w_{1,t})\Delta c_{2,t+1} \\ \Delta s_{t+1} &= w_{1,t}\Delta s_{1,t+1} + (1 - w_{1,t})\Delta s_{2,t+1}.\end{aligned}$$

This marginal investor wants zero exposure to housing in expectations and thus chooses $w_{1,t}$ from

$$0 = w_{1,t} (\mathbb{I}_{x_1 > 0} \phi_s x_{1,t}) + (1 - w_{1,t}) (\mathbb{I}_{x_2 > 0} \phi_s x_{2,t}).$$

There are two implications for the marginal investor in the aggregate markets. First there is no exposure to the conditional component of housing. They therefore do not price housing in their consumption risk. In addition, $w_{1,t}$ changes period to period as the two state variables move around. This means they have time varying consumption volatility as well. These two components

lead to the setup of Bansal & Yaron (2004) with both long-run risk and stochastic volatility.

APPENDIX I Simulation Equations

In section (6.4) I use simulations in order to replicate my empirical results from section (5).

What I present below is the formulation for the simulation I run.

$$\ln \frac{S_{t+1}}{C_{t+1}} = s_{t+1} - c_{t+1} = s_t - c_t + \underbrace{\mu_s - \mu_c}_{\mu} + \underbrace{\left(I_{(x_t > 0)} \phi_s - \phi_c \right)}_{\phi_{(x_t > 0)}} x_t + \underbrace{\phi_s e_{s,t+1} - e_{c,t+1}}_{e_{sc,t+1}}.$$

This is an autogressive process; if I foil the expression J times I obtain

$$sc_{t+1} = \mu \sum_{j=0}^J (1-\lambda)^j + \sum_{j=0}^J (1-\lambda)^j e_{sc,t+1-j} + (1-\lambda)^{J+1} sc_{t-J} + \sum_{j=0}^{J-1} \sum_{k=0}^j \phi_{(x_{t-k} > 0)} (1-\lambda)^k \rho^{j-k} \epsilon_{x,t-j} + \left(\sum_{j=0}^J \phi_{(x_{t-j} > 0)} \rho^{J-j} (1-\lambda)^j \right) x_{t-J}.$$

Taking $J \rightarrow \infty$,

$$sc_{t+1} = \mu \sum_{j=0}^{\infty} (1-\lambda)^j + \sum_{j=0}^{\infty} (1-\lambda)^j e_{sc,t+1-j} + \sum_{j=0}^{\infty} \sum_{k=0}^j \phi_{(x_{t-k} > 0)} (1-\lambda)^k \rho^{j-k} \epsilon_{x,t-j}.$$

As I now have sc_t in terms of only orthogonal shocks, I can run a simulation to match the sample of returns and price-to-rent ratios in my empirical specification.

APPENDIX J Estimation of Composition Variance

As x_t is the conditional component of sc_{t+1} , I need to foil out the composition relationship one additional period to see its dependence on shocks to long-run prospects. In addition, I define

$\phi_{(x_t>0)} = \phi_s I_{(x_t>0)} - \phi_c$. Thus,

$$sc_{t+2} = (1 - \lambda)^2 sc_t + (1 + 1 - \lambda)\mu + (1 - \lambda)\phi_{(x_t>0)}x_t + \phi_{(x_{t+1}>0)}x_{t+1} + (1 - \lambda)e_{sc,t+1} + e_{sc,t+2}.$$

In order to compute the variance of sc_{t+2} , I take its time t expectations. One of the RHS expressions becomes $E_t [\phi_{(x_{t+1}>0)}\varepsilon_{x,t+1}] = \phi_s E_t [I_{(x_t>0)}\varepsilon_{x,t+1}] - \phi_c E_t [\varepsilon_{x,t+1}]$. Given that for any binary and continuous random variable D and C , respectively,

$$E[DC] = \begin{cases} E[C], & \text{with prob. } p = P(D = 1) \\ 0, & \text{with prob. } 1 - p, \end{cases} \quad (23)$$

this expectation will be zero as $E_t [\varepsilon_{x,t+1}] = 0$. Therefore,

$$Var_t(sc_{t+2}) = E_t \left[\left((I_{(x_{t+1}>0)} - E_t [I_{(x_{t+1}>0)}]) \phi_s \rho x_t + (I_{(x_{t+1}>0)} \phi_s - \phi_c) \varepsilon_{x,t+1} + \varepsilon_{sc,t+2} + (1 - \lambda) \varepsilon_{sc,t+1} \right)^2 \right].$$

Given the relationship in equation (23), I can expand out the RHS to obtain

$$\begin{aligned} Var_t(sc_{t+2}) &= E_t \left[\left((I_{(x_{t+1}>0)} - E_t [I_{(x_{t+1}>0)}]) \right)^2 \phi_s^2 \rho^2 x_t^2 + (I_{(x_{t+1}>0)} \phi_s - \phi_c)^2 \varepsilon_{x,t+1}^2 + \varepsilon_{sc,t+2}^2 + (1 - \lambda)^2 \varepsilon_{sc,t+1}^2 \right] \\ &= \left(1 + (1 - \lambda)^2 \right) \sigma_{sc}^2 + \phi_c^2 \sigma_x^2 + \left(P_t (I_{(x_{t+1}>0)} = 1) - P_t (I_{(x_{t+1}>0)} = 1) \right)^2 \phi_s^2 \rho^2 x_t^2 + \\ &\quad (\phi_s^2 - 2\phi_c \phi_s) \sigma_x^2 P_t (I_{(x_{t+1}>0)} = 1). \end{aligned}$$

Given the intra-temporal link between prices and quantities

$$Var_t(\omega_{t+2}) = \left(1 - \frac{1}{\varepsilon} \right)^2 Var_t(sc_{t+2}).$$

REFERENCES

- [1] Abraham, J. M. and Hendershott, P. H. Bubbles in Metropolitan Housing Markets. *NBER Working Paper*, 1994.
- [2] Adelino, M., Schoar, A., and Severino, F. Changes in Buyer Composition and the Expansion of Credit During the Boom. *NBER Working Paper*, No. 20848, 2015.
- [3] Adelino, M., Schoar, A., and Severino, F. House prices, collateral, and self-employment. *Journal of Financial Economics*, 117(2):288–306, 2015.
- [4] Albouy, D. What are Cities Worth? Land Rents, Local Productivity, and the Total Value of Amenities. *The Review of Economics and Statistics*, 98(3):477–487, 2016.
- [5] Ang, A. and Bekaert, G. Stock Return Predictability: Is it There? *The Review of Financial Studies*, 20(3):651–707, 2007.
- [6] Autor, D. H. and Duggan, M. G. The Rise in the Disability Rolls and the Decline in Unemployment. *The Quarterly Journal of Economics*, 118(1):157–205, 2003.
- [7] Bansal, R. and Kiku, D. Cointegration and Long-Run Asset Allocation. *Journal of Business & Economic Statistics*, 29(1):161–173, 2011.
- [8] Bansal, R., Kiku, D., and Yaron, A. Risks for the Long Run: Estimation with Time Aggregation. *NBER Working Paper*, No. 18305, 2012.
- [9] Bansal, R. and Yaron, A. Risks for the Long Run. *The Journal of Finance*, 59(4):1481–1509, 2004.
- [10] Bansal, R. and Yaron, A. Risks for the Long Run: Estimation and Inference. *Working Paper*, 2007.
- [11] Bartelsman, E., Haltiwanger, J., and Scarpetta, S. Cross-Country Differences in Productivity: The Role of Allocation and Selection. *American Economic Review*, 103(1):305–334, 2013.
- [12] Bartik, T. J. Boon or Boondoggle? The Debate Over State and Local Economic Development Policies. *Working Paper*, 1991.
- [13] Bekaert, G., Harvey, C. R., and Lundblad, C. Does financial liberalization spur growth? *Journal of Financial Economics*, 77:3–55, 2005.
- [14] Bekaert, G., Harvey, C. R., Lundblad, C., Siegel, S., Bekaert, G., Harvey, C. R., and Lundblad, C. Global Growth Opportunities and Market Integration. *Journal of Finance*, 62(3):1081–1137, 2007.
- [15] Bidder, R. and Dew-becker, I. Long-Run Risk is the Worst-Case Scenario. *Working Paper*, 2016.

- [16] Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. Really Uncertain Business Cycles. *Working Paper*, 2014.
- [17] Board of Governors of The Federal Reserve System,. Q4 Financial Accounts of the United States. Technical report, 2015.
- [18] Boskin, M. J., Dulberger, E. R., Gordon, R. J., Griliches, Z., and Jorgenson, D. W. Consumer Prices, the Consumer Price Index , and the Cost of Living. *The Journal of Economic Perspectives*, 12(1):3–26, 1998.
- [19] Boudoukh, J., Michaely, R., Richardson, M., and Roberts, M. R. On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing. *The Journal of Finance*, 62(2):877–915, 2007.
- [20] Bound, J. and Holzer, H. J. Demand Shifts, Population Adjustments, and Labor Market Outcomes During the 1980s. *Journal of Labor Economics*, 18(1):20–54, 2000.
- [21] Brueckner, J. K. Urban Growth Models with Durable Housing: An Overview. In *Economics of Cities: Theoretical Perspectives*. 2000.
- [22] Campbell, J. Y. and Cocco, J. F. How do house prices affect consumption? Evidence from micro data. *Journal of Monetary Economics*, 54(3):591–621, 2007.
- [23] Campbell, J. Y. and Cochrane, J. H. By Force of Habit : A Consumption Based Explanation of Aggregate Stock Market Behavior. *1Journal of Political Economy*, 107(2):205–251, 1999.
- [24] Campbell, J. Y. and Shiller, R. J. Stock Prices, Earnings, and Expected Dividends. *The Journal of Finance*, 43(3):661–676, 1987.
- [25] Campbell, S. D., Davis, M. A., Gallin, J., and Martin, R. F. What Moves Housing Markets: A Variance Decomposition of the Rent-Price Ratio. *Journal of Urban Economics*, 66(2):90–102, 2009.
- [26] Case, K., Cotter, J., and Gabriel, S. Housing Risk and Return: Evidence From a Housing Asset-pricing Model. *Working Paper*, 2011.
- [27] Chen, B. and Stafford, F. P. Stock Market Participation: Family Responses to Housing Consumption Commitments. *Journal of Money, Credit and Banking*, 48(4):635–659, 2016.
- [28] Chen, H., Michaux, M., and Roussanov, N. Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty. *NBER Working Paper*, No. 19421, 2013.
- [29] Colacito, R. and Croce, M. M. Risks for the Long Run and the Real Exchange Rate. *Journal of Political Economy*, 119(1):153–181, 2011.
- [30] Colacito, R. and Croce, M. M. International Asset Pricing with Recursive Preferences. *Journal of Finance*, 68(6):2651–2686, 2013.
- [31] Colacito, R., Croce, M. M., Ho, S., and Howard, P. BKK the EZ Way: An International Production Economy with Recursive Preferences. *Working Paper*, 2013.

- [32] Constantinides, G. M. and Duffie, D. Asset Pricing with Heterogeneous Consumers. *Journal of Political Economy*, 104(2):219, 1996.
- [33] Constantinides, G. M. and Ghosh, A. Asset Pricing Tests with Long-run Risks in Consumption Growth. *Review of Asset Pricing Studies*, 1(1):96–136, 2011.
- [34] Davidoff, T. Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated with Many Demand Factors. *Critical Finance Review*, forthcoming, 2015.
- [35] Davis, M. A., Fisher, J. D. M., and Veracierto, M. Gross Migration, Housing and Urban Population Dynamics. *Working Paper*, 2014.
- [36] Davis, M. a., Fisher, J. D. M., and Whited, T. M. Macroeconomic Implications of Agglomeration. *Econometrica*, 82(2):731–764, 2014.
- [37] Davis, M. A. and Heathcote, J. Housing and the Business Cycle. *International Economic Review*, 46(3):751–784, 2005.
- [38] Davis, M. A., Lehnert, A., and Martin, R. F. The rent-price ratio for the aggregate stock of owner-occupied housing. *Review of Income and Wealth*, 54(2):279–284, 2008.
- [39] Dew-Becker, I. and Giglio, S. Asset Pricing in the Frequency Domain: Theory and Empirics. *Working Paper*, 2013.
- [40] Dougal, C., Parsons, C. A., and Titman, S. Urban Vibrancy and Corporate Growth. *Journal of Finance*, 70(1):163–210, 2015.
- [41] Duarte, M. and Restuccia, D. The Role of the Structural Transformation in Aggregate Productivity. *Quarterly Journal of Economics*, 125(1):129–173, 2010.
- [42] Engelberg, J., Ozoguz, A., and Wang, S. Know Thy Neighbor: Industry Clusters, Information Spillovers and Market Efficiency. *Working Paper*, 2013.
- [43] Epstein, L. G. and Zin, S. E. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica*, 57(4):937–969, 1989.
- [44] Eraker, B., Shaliastovich, I., and Wang, W. Durable Goods, Inflation Risk, and Equilibrium Asset Prices. *Review of Financial Studies*, 29(1):193–231, 2016.
- [45] Favilukis, J., Kohn, D., Ludvigson, S. C., and Van Nieuwerburgh, S. International Capital Flows and House Prices: Theory and Evidence. In *Housing and the Financial Crisis*, number August, pages 235–299. 2013.
- [46] Ferreira, F. and Gyourko, J. Anatomy of the Beginning of the Housing Boom: U.S. Neighborhoods and Metropolitan Areas, 1993-2009. *NBER Working Paper*, No. 17374, 2011.
- [47] Fillat, J. L. Housing as a Measure for Long-Run Risk in Asset Pricing. *Working Paper*, 2009.

- [48] Flavin, M. and Nakagawa, S. A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence. *The American Economic Review*, 98(1):474–495, 2008.
- [49] Gabaix, X. Zipf ' s Law for Cities : An Explanation. *The Quarterly Journal of Economics*, 114(3):739–767, 1999.
- [50] Gao, Z., Sockin, M., and Xiong, W. Learning about the Neighborhood: The Role of Supply Elasticity for Housing Cycles. *Working Paper*, 2014.
- [51] Garofalo, G. A. and Yamarik, S. Regional Convergence: Evidence from a New State-by-State Capital Stock Series. *Review of Economics and Statistics*, 84(May):316–323, 2002.
- [52] Ghysels, E., Plazzi, A., Valkanov, R., and Torous, W. Forecasting real estate prices. In *Handbook of Economic Forecasting*, pages 509–580. 2013.
- [53] Giglio, S., Maggiori, M., and Stroebel, J. Very Long-Run Discount Rates. *The Quarterly Journal of Economics*, 130(1):1–53, 2015.
- [54] Giglio, S., Maggiori, M., and Stroebel, J. No-Bubble Condition: Model-free Tests in Housing Markets. *Econometrica*, 84(3):1047–1091, 2016.
- [55] Glaeser, E. L. and Gottlieb, J. D. The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States. *Journal of Economic Literature*, 47(4):983–1028, 2009.
- [56] Glaeser, E. L. and Gyourko, J. Urban Decline and Durable Housing. *Journal of Political Economy*, 113(2):345–000, 2005.
- [57] Glaeser, E. L., Gyourko, J., Morales, E., and Nathanson, C. G. Housing dynamics: An urban approach. *Journal of Urban Economics*, 81:45–56, 2014.
- [58] Glaeser, E. L., Gyourko, J., and Saks, R. E. Why Have Housing Prices Gone Up? *The American Economic Review*, 95(2):329–333, 2005.
- [59] Glaeser, E. L., Gyourko, J., and Saks, R. E. Why Have Housing Prices Gone Up? *NBER Working Paper*, No. 11129, 2005.
- [60] Gomes, J. F., Kogan, L., and Yogo, M. Durability of Output and Expected Stock Returns. *Journal of Political Economy*, 117(5):941–986, 2009.
- [61] Gyourko, J., Mayer, C., and Sinai, T. Superstar Cities. *American Economic Journal: Economic Policy*, 5(4):167–199, 2013.
- [62] Gyourko, J. and Molloy, R. Regulation and Housing Supply. *NBER Working Paper*, No. 20536, 2014.
- [63] Han, L. Understanding the puzzling risk-return relationship for housing. *Review of Financial Studies*, 26(4):877–928, 2013.

- [64] Hansen, L. P. and Sargent, T. J. Robust Control and Model Uncertainty. *The American Economic Review*, 91(2):60–66, 2001.
- [65] Harding, J. P., Rosenthal, S. S., and Sirmans, C. F. Depreciation of housing capital , maintenance , and house price inflation : Estimates from a repeat sales model. *Journal of Urban Economics*, 61:193–217, 2007.
- [66] Harvey, C. R. The real term structure and consumption growth. *Journal of Financial Economics*, 22(2):305–333, 1988.
- [67] Hizmo, A. Risk in Housing Markets : An Equilibrium Approach. *Working Paper*, 2015.
- [68] Hsieh, C.-t. and Klenow, P. J. Oxford University Press. *Quarterly Journal of Economics*, 4(November):1403–1448, 2009.
- [69] Hsieh, C.-t. and Moretti, E. Why do Cities Matter? Local Growth and Aggregate Growth. *NBER Working Paper*, No. 21154, 2015.
- [70] Kline, P. M. and Moretti, E. Local Economic Development, Agglomeration Economies, and the Big Push. *NBER Working Paper*, No. 19293, 2013.
- [71] Landvoigt, T., Piazzesi, M., and Schneider, M. The Housing Market(s) of San Diego. *American Economic Review*, 105(4):1–47, 2015.
- [72] Lettau, M. and Ludvigson, S. C. Consumption, Aggregate Wealth, and Expected Stock Returns. *The Journal of Finance*, 56(3):815–849, 2001.
- [73] Lewellen, J. Predicting returns with financial ratios. *Journal of Financial Economics*, 74(2):209–235, 2004.
- [74] Lusitig, H. N. and Van Nieuwerburgh, S. G. Housing Collateral, Consumption Insurance and Risk Premia: An Empirical Perspective. *Journal of Finance*, 60(3):1167–1219, 2005.
- [75] Lusitig, H. N. and Van Nieuwerburgh, S. G. How much does household collateral constrain regional risk sharing? *Review of Economic Dynamics*, 13(2):265–294, 2010.
- [76] Luttmer, E. F. Neighbors as Negatives: Relative Earnings and Well-being. *The Quarterly Journal of Economics*, 120(3):963–1002, 2005.
- [77] Massimiliano Croce, M. Long-run productivity risk: A new hope for production-based asset pricing? *Journal of Monetary Economics*, 66:13–31, 2014.
- [78] Meyer, B. B. D. and Sullivan, J. X. Consumption and Income Inequality and the Great Recession. *The American Economic Review*, 103(3):178–183, 2013.
- [79] Mian, A., Rao, K., and Sufi, A. Household Balance Sheets, Consumption and the Economic Slump. *The Quarterly Journal of Economics*, 128(4):1687–1726, 2013.
- [80] Mian, A. and Sufi, A. What Explains the 2007-2009 Drop in Employment? *Econometrica*, 82(6):2197–2223, 2014.

- [81] Ngai, L. R. and Samaniego, R. M. Accounting for research and productivity growth across industries. *Review of Economic Dynamics*, 14(3):475–495, 2011.
- [82] Nieuwerburgh, S. V. and Weill, P.-O. Why Has House Price Dispersion Gone Up? *The Review of Economics and Statistics*, 77(4):1567–1606, 2010.
- [83] Petersen, M. A. Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies*, 22(1):435–480, 2009.
- [84] Piazzesi, M., Schneider, M., and Tuzel, S. Housing, Consumption and Asset Pricing. *Journal of Financial Economics*, 83(3):531–569, 2007.
- [85] Plazzi, A., Torous, W., and Valkanov, R. Expected Returns and Expected Growth in Rents of Commercial Real. *The Review of Financial Studies*, 23(9):3469–3519, 2010.
- [86] Quigley, J. M. Urban Diversity and Economic Growth. *Journal of Economic Perspectives*, 12(2):127–138, 1998.
- [87] Roback, J. Wages, Rents, and the Quality of Life. *The Journal of Political Economy*, 90(6):1257–1278, 1982.
- [88] Rosen, S. Wage-Based Indexes of Urban Quality of Life. In *Current Issues in Urban Economics*. 1979.
- [89] Rosenthal, S. S. and Strange, W. C. *Evidence on the Nature and Source of Agglomeration Economies*, volume 4. Elsevier Inc., 2004.
- [90] Saiz, A. The Geographic Determinants of Housing Supply. *The Quarterly Journal of Economics*, 125(3):1253–1296, 2010.
- [91] Shilling, J. D. Is There a Risk Premium Puzzle in Real Estate? *Real Estate Economics*, 31(4):501–525, 2003.
- [92] Sinai, T. and Souleles, N. Owner-Occupied Housing as a Hedge against Rent Risk. *The Quarterly Journal of Economics*, 120(2):763–789, 2005.
- [93] Tuzel, S. and Zhang, M. Local Risk, Local Factors, and Asset Prices. *The Journal of Finance*, (Forthcoming), 2016.
- [94] Van Binsbergen, J. H. and Koijen, R. S. J. Predictive regressions: A present-value approach. *Journal of Finance*, 65(4):1439–1471, 2010.
- [95] Villaverde, J. F. and Krueger, D. Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data. *The Review of Economics and Statistics*, 89(3):552–565, 2007.
- [96] Yang, W. Long-run risk in durable consumption. *Journal of Financial Economics*, 102(1):45–61, 2011.