The Economics of Segmented Housing Markets

George Nikolakoudis*
Princeton

May 23, 2024

Abstract

This paper develops a dynamic assignment model in which households purchase indivisible housing units of heterogeneous quality. Theoretically, I characterize how key macroeconomic drivers of housing markets affect the cross-section of prices. Changes in credit conditions mostly affect low-end homes, while changes in future income, transaction costs, and interest rates mostly affect the prices of mid and high-end homes. In contrast to models with a homogeneous housing stock, shocks propagate within the housing market unidirectionally: changes in housing demand "trickle up" in the quality ladder, but do not "trickle down". Empirically, I test the key theoretical prediction of unidirectional shock propagation using plausibly exogenous shocks to different quality segments of the housing market. Consistent with the theory, shocks that affect high-income households only induce changes in the prices of high-end homes, whereas shocks that affect low-income households affect house prices in the entire market.

Keywords: Collateral Constraints, Household Saving, Housing Demand, Housing Supply and Markets, Segmentation, Shock Propagation

*Princeton Department of Economics, Julis Romo Rabinowitz Building, Princeton, NJ 08544. Email: nikolakoudis@princeton.edu.

I am grateful to Mark Aguiar, Narek Alexanian, Markus Brunnermeier, Michael Ehrmann, Pablo Fajgelbaum, Joel P. Flynn, Nobuhiro Kiyotaki, Moritz Lenel, Ernest Liu, Simon Margolin, Gonzalo Paz Pardo, Richard Rogerson, Karthik Sastry, Jiri Slacalek, Oreste Tristani, Gianluca Violante, Atsushi Yamagishi, Wei Xiong, Chansik Yoon, and participants in the Princeton Finance Lunch, Princeton International Trade Lunch, Princeton Macro Lunch, and ECB internal seminar for helpful comments. I am thankful to the European Central Bank for their hospitality when writing the first draft of this paper. First posted version: October 18, 2022.

1 Introduction

Housing markets are segmented: even within a narrow geographical region, there is substantial heterogeneity in relative price changes for different types of houses. Figure 1 illustrates this fact. It plots the annualized capital gains of houses for counties in different metro areas within the United States from 1990-2021, disaggregated by their initial price. In Los Angeles, for example, cheap (low-end) houses depreciated approximately three times as much as expensive (high-end) homes (in percentage terms) in the year 2009. As I will show later, this disparity in relative rates of return is a robust feature for local housing markets across the entire United States. This basic fact implies that we cannot descriptively represent the housing market as a single "market".

Market segmentation can have important macroeconomic consequences. The prices of low-end and high-end homes are differentially exposed, in qualitative ways, to changes in the income distribution, interest rates, and credit. This is because shocks that affect a particular group of households will mostly be reflected in the prices of homes that those households purchase. Moreover, the unequal incidence of housing capital gains that arises from market segmentation can in turn be an important driver of other macroeconomic aggregates, such as consumption expenditures and welfare. Although the urban economics literature has long recognized segmentation as an important feature of housing markets (Sweeney, 1974; Bourassa et al., 1999; Watkins, 2001), macroeconomic models with a housing sector typically model housing as a commodity that can be bought and sold at a common per unit price from a homogeneous stock. Consequently, these models abstract from the role of segmentation in shaping house price movements.

In order to understand the implications of housing market segmentation for the propagation of macroeconomic shocks, I develop an analytically tractable, dynamic assignment model in which housing is segmented by various quality tiers. I show that the presence of housing market segmentation substantially alters the response of house prices to shocks relative to the homogeneous housing benchmark. In the unique equilibrium of the model, shocks propagate within the housing market *unidirectionally* – changes in households' valuations of low-end homes spill over upwards in the house quality ladder, making them important drivers of average house price changes. I characterize the strength of these spillovers analytically and examine how they shape the response of the entire *cross-section* of house prices to macroeconomic shocks, such as changes in credit frictions, interest rates, future income, and transaction costs. The model also gives rise to an amplification channel for the response of house price changes to aggregate consumption expenditure through the heterogeneous incidence of capital gains, as well as novel welfare implications for the role of credit in the macroeconomy: under housing market segmentation, *more* credit can lead to Pareto *inferior* outcomes.

I operationalize this model empirically to measure the importance of housing market segmentation in the data. In the data, shocks that affect high-income households only induce price changes in high-end homes, consistent with the unidirectional propagation of shocks. In contrast, shocks that affect low-income households affect house prices in the entire market through pricing spillovers. These tests empirically reject the predictions of the homogeneous housing benchmark and suggest

Figure 1: Rate of Return Heterogeneity within Counties

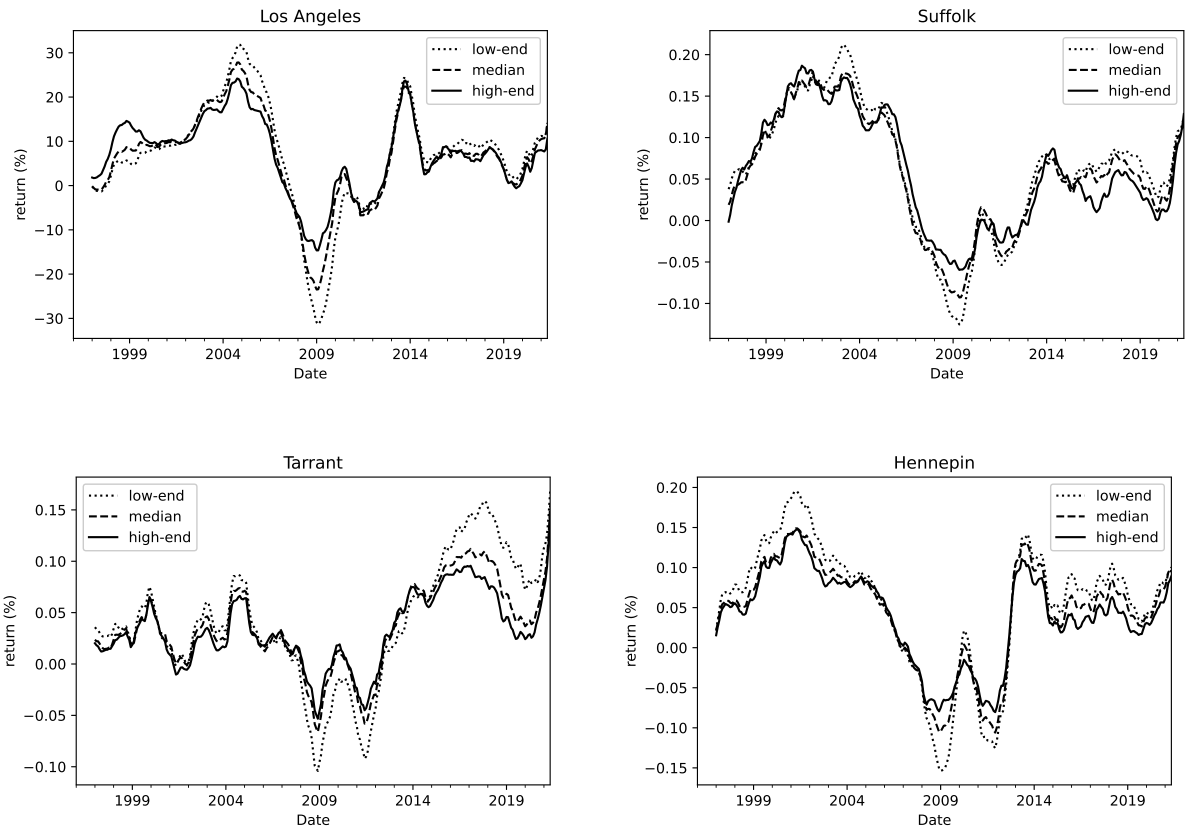

Note: Annualized rates of return for counties in various metro areas (Greater Los Angeles, New York-Newark-Jersey City, Dallas-Fort Worth-Arlington, Minneapolis-Saint Paul). House price data is from the Zillow Research Single Family Homes Index. The low-end (high-end) series plots average annualized returns for ZIP codes in the bottom (top) quartile of the ZIP code price distribution within the given county. The median series plots median house prices within the given county.

that the presence of segmentation is an important feature in order to understand how macroeconomic shocks shape house price movements.

Model. In the model, there exists a continuum of indivisible houses that provide different flows of housing services to households. In equilibrium, households match to houses of varying qualities according to the net present value of their income. Because houses are priced competitively, this gives rise to a continuum of pricing equations, in which each household is a marginal investor for the house that they purchase. These pricing equations relate households' willingness to pay for inframarginal quality units to their marginal rate of substitution between consumption and housing services. Households' marginal rate of substitution, in turn, depends on aggregate economic conditions, such as the level of interest rates or the availability of credit. Heterogeneity in housing capital gains then results from heterogeneous variations in marginal rates of substitution in response to different economic shocks.

Theoretical Results. The existence of housing market segmentation gives rise to distinct asset pricing implications relative to the homogeneous housing benchmark. First, I show that an implication of market segmentation is that "localized" shocks that affect only a subset of households can spill over into neighboring quality tiers by affecting the willingness to pay of neighboring households along the income distribution (as in the static models of Braid, 1981 or Määttänen and Terviö, 2014). This gives rise to a mechanism in which shocks that affect the prices of low-end homes "trickle up" to the valuations of high-end homes in general equilibrium. However, localized shocks that directly affect the prices of high-end homes do not "trickle down" the quality ladder to low-end house prices. Intuitively, the price of a given quality tier is determined by the willingness to pay for additional housing services of households that reside strictly *below* that tier. For this reason, the preferences of wealthy households — that reside in high-end homes — will not affect the prices of lower quality segments. Nevertheless, the preferences of low-income households do matter for the prices of homes occupied by high-income households.

Second, I leverage these results to theoretically characterize the effect of macroeconomic shocks on the cross-sectional distribution of house prices. I formally show that two important macroeconomic forces in the housing market operate on different ends of the wealth and house quality distribution. Changes in credit conditions have the largest effect on the prices of low-end homes, by affecting the willingness to pay of credit constrained, low-income households. In contrast, interest rate changes, changes in future income, as well as transaction costs, have the largest effect on the prices of high-end homes. This is because these shocks are mostly pay-off relevant for unconstrained, high-income households through their effect on the net present value of income and liabilities. As such, the cross-sectional distribution of capital gains is sensitive to the nature of the shock that hits the economy. Moreover, the presence of pricing spillovers implies that credit shocks (and other "low-end" shocks) can have a large impact on prices over the whole quality ladder. Even if few households are credit constrained in equilibrium, the average house price index is sensitive to changes in credit conditions. In all cases, these results differ from the homogeneous housing benchmark (which predicts constant returns for all house qualities) and are robust to a wide range of income and house quality distributions.

Third, I demonstrate that the model can amplify the response of aggregate consumption expenditures to house price changes relative to the homogeneous housing benchmark. Intuitively, this is because market segmentation can create a positive covariance between capital gain incidence (which are wealth effects in the model) and marginal propensities to consume (MPCs). In contrast, in the absence of market segmentation, the housing market gives rise to a form of dampening for *all* aggregate fluctuations. This is because capital gains contrivedly become a decreasing function of liquid assets, which establishes a strictly negative covariance between wealth effects and MPCs. Market segmentation therefore allows the largest capital losses to be incident precisely on those households that are most important in driving aggregate consumption expenditures.

Fourth, the model gives rise to novel predictions for the effect of housing credit on welfare: *more* housing credit can *reduce* the welfare of all households, leading to Pareto inferior outcomes.

This force operates through the general equilibrium transmission channel of the model. In partial equilibrium, higher credit directly increases the welfare of relatively poor households by relaxing their borrowing constraint. In general equilibrium, however, the prices of low-end homes rise to more than offset the direct, positive welfare impact of higher credit — houses become more "unaffordable". Intuitively, this is because low-income households bid up the prices of low-end homes to leave them indifferent between home ownership and their outside option of renting. This increase in prices then spills over upwards to mid- and high-end homes, reducing housing affordability and welfare for all households in the economy.

Empirical Analysis. Finally, I turn to empirically testing the key predictions of the theory. To do so, I exploit cross-sectional variation in the incidence of various shocks across local housing markets in the United States. In the first test, I isolate plausibly exogenous variation in mortgage lending to create a proxy for a localized shock to *low-end* households. I do so by leveraging banks' heterogeneous exposure to the rise in the private label mortgage securitization market (the "PLS market") over 2002-2005 (Mian and Sufi, 2022). I show that this credit shock is almost entirely incident on low-end segments of the housing market. Moreover, this shock creates monotonic capital gains across the quality ladder: low-end homes appreciate more than high-end homes, as predicted by the theory. However, I also provide evidence for the existence of "trickle up" forces, which cause an economically significant increase in the prices of homes at the top end of the quality distribution.

In the second test, I provide empirical support for the absence of "trickle down" effects by isolating plausibly exogenous variation in stock market wealth changes (Chodorow-Reich et al., 2021). I show that stock market wealth is primarily held in high-end segments of the house quality distribution. These high-end shocks cause a statistically and economically significant increase on the prices of high-end homes, but do not affect lower quality tiers. Moreover, the results are robust to local general equilibrium spillovers, modeled as in Guren et al. (2021) or Nakamura and Steinsson (2018). Consequently, the data suggests that there is strong *directionality* in the propagation of shocks within the housing market. Taken together, both the theory and data highlight the importance of housing market segmentation in understanding the propagation of macroeconomic shocks to house prices, and *vice versa*.

Relation to Existing Literature. This paper lies at the intersection of several strands of literature. First, this paper relates to a large empirical and theoretical literature on the effect of macroeconomic shocks on house price movements (Cox and Ludvigson, 2021; Favilukis et al., 2017; Garriga and Hedlund, 2020; Greenwald, 2018; Greenwald and Guren, 2021; Kaplan et al., 2020; Kiyotaki et al., 2011; Mian and Sufi, 2022; Piazzesi and Schneider, 2016). This literature primarily characterizes how shocks affect average house prices *across* different housing markets. In contrast, I analyze the propagation of shocks to house prices in a framework that features segmentation *within* a housing market. I demonstrate that within-market segmentation gives rise to distinct, testable asset pricing and macroeconomic predictions which I provide evidence for in the data.[1]

[1] Greenwald and Guren (2021) analyze a distinct form of segmentation between rental and owner-occupied housing. However, they do not model segmentation *within* owner-occupied housing, which is the focus of this paper.

Second, this paper contributes to the literature on within-city house price dynamics. Ortalo-Magne and Rady (2006) construct a life-cycle model of the housing market in which households live for four periods and show that the prices of flats display more volatility than houses. My model features a simpler life-cycle specification (in which households only live for two periods), but permits a more transparent characterization of the cross-section of capital gains in response to different shocks. In studying how macroeconomic shocks affect house prices, I primarily focus on how *demand-side* changes affect house prices. Related works by Guerrieri et al. (2013) and Nathanson (2020) demonstrate how *supply-side* movements affect house prices through construction and gentrification.

Braid (1981) and Määttänen and Terviö (2014) show that income shocks spill over upwards to house prices in a static assignment model that features no-trade equilibria. I show that uni-directional shock propagation is also a feature of my dynamic economy which features housing trade across generations. However, the presence of dynamics qualitatively alters the nature of these spillovers: these spillovers "dissipate" as one moves up the house quality ladder, and are modulated by a composite parameter that reflects future discounting (and are otherwise constant in a static model). This allows for richer capital gain heterogeneity in my model that informs the general equilibrium response of house prices to aggregate shocks. Moreover, because of their static nature, the frameworks of Braid or Määttänen and Terviö cannot analyze how credit frictions, interest rates, or perceptions regarding *future* house values affect capital gains and propagate onto the rest of the economy, which is the focus of this paper. Landvoigt et al. (2015) also use an assignment model approach to show that looser lending standards can quantitatively generate higher cross-sectional capital gains at the low-end of the housing quality distribution in a calibration to San Diego County. I complement their quantitative work by theoretically characterizing how a wide range of economic shocks affects the cross-sectional distribution of capital gains in a general macroeconomic environment, and subsequently test for these predictions in the data. Hacamo (2021) uses an assignment model to empirically investigate how interest rates affect the distribution of house prices. In contrast, I offer evidence on market segmentation by testing for spillovers using plausibly exogenous shocks to different quality segments of the housing market.

Finally, this work relates to papers by Mian et al. (2013), Mian and Sufi (2014), Berger et al. (2018), and Guren et al. (2021), which analyze the effect of house price changes on consumption expenditures. My work offers a complementary mechanism for the propagation of house price changes to the macroeconomy through heterogeneity in capital gain incidence. Moreover, this literature primarily examines the consumption response to house price changes in partial equilibrium. In contrast, I emphasize the importance of general equilibrium forces in shaping the transmission mechanism of housing wealth effects through the cross-sectional incidence of capital gains.

Outline. The rest of the paper proceeds as follows. Section 2 documents additional motivating facts regarding housing market segmentation. Section 3 sets up the model. Sections 4 and 5 present the main theoretical results. Section 6 tests the key predictions of the theory. Section 7 concludes.

2 Facts About Housing Market Segmentation

I first provide quantitative evidence for the premise of this paper — housing market segmentation. To this end, I analyze the rates of return of different types of houses within various local housing markets across the United States. My definition of a local housing market is that of a county: this is consistent with prior studies on the within-market variation in housing prices (Guerrieri et al., 2013; Landvoigt et al., 2015), as well as empirical work that quantifies the effect of net worth changes on local employment (Mian and Sufi, 2014). By defining local housing markets in this way, I implicitly assume that the set of reasonable housing choices for households lies within the county they reside in.[2]

2.1 Data

My measure of house prices comes from the Zillow ZIP code level price indices. The Zillow index is a hedonically adjusted price index. It uses detailed information on properties, collected from public records, to create a seasonally adjusted measure of the typical home value in a given region.[3] I restrict the analysis to the years 2000-2020. Many ZIP codes are missing prior to the year 2000, which makes it difficult to characterize the cross-sectional evolution of house prices preceding this time period. In order to construct weighted averages of house price growth, I compute the total number of houses by county and ZIP code from the American Community Survey (ACS). I use the ACS five-year survey whose corresponding year is closest to the year in which I disaggregate houses into different quality tiers.

2.2 Methodology

I treat each county as a common housing market that is segmented into various quality tiers. I follow Landvoigt et al. (2015) in assuming that market prices approximately reflect a one-dimensional index of housing quality. The reasoning for this is that any changes in the payoff relevant characteristics of the house, such as land, location, and structure, should be incorporated into house valuations in competitive equilibrium. Moreover, in the theoretical model presented in the next section, I show that price is a sufficient statistic for a given housing type's quality ranking.

Consequently, I sort ZIP codes in a particular county into different quality quantiles according to their initial price in the sample. When looking at changes in prices experienced by houses of a given quality over a particular time frame, the year in which houses are sorted into quality tiers is somewhat arbitrary. However, the relative ranking of median house prices across ZIP codes within counties is very persistent. As such, all the empirical results are insensitive to the year in which ZIP codes are sorted into different quality quartiles.

[2] The results are quantitatively similar when using MSAs as the market definition. There are more sophisticated methods to identify housing markets using household search behavior as in Piazzesi et al. (2020). Doing so requires data on search behavior for houses, which is challenging to obtain and beyond the scope of this paper.

[3] My measure comes from the ZHVI Single-Family Homes Time Series index in `https://www.zillow.com/research/data/`. I use data on single family homes to maintain consistency with my theoretical model in Section 4.

Figure 2: Heterogeneity in Within-Market Housing Returns

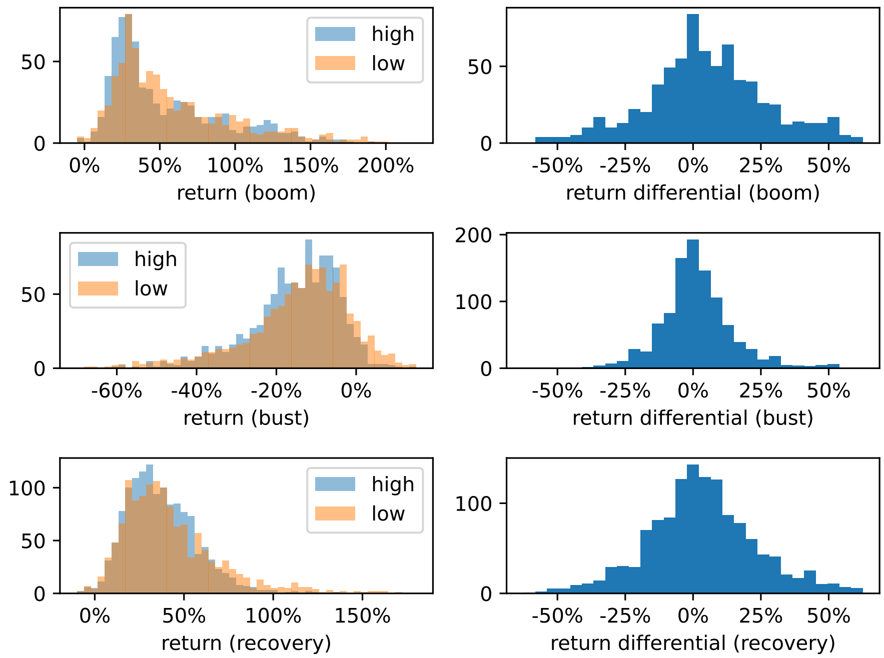

Note: Relative rates of return for all counties during the boom (first row), bust (middle row), and recovery (third row). The first column plots the histograms of cumulative returns for the bottom and top quartile of house qualities for the corresponding time period. The second column plots the return difference of the bottom quartile relative to the top quartile. Counties with an average house price change of less than 2% (in absolute terms) have been dropped from the sample.

2.3 Cross-Sectional Patterns

The national house price cycle can be trisected into the 2000-2006 boom, the 2006-2012 bust, and its subsequent recovery. Hence, it is natural to analyze the cross-sectional distribution of house prices within each time frame separately.

The left column of Figure 2 plots the average growth rates of house prices for the top and bottom quality quartiles for all counties and each time period under consideration. The right column of Figure 2 plots the differences between the bottom and top quality quartiles *within* each county. Two observations are apparent. First, there is significant variation in the *across* county growth rate of house prices, even in the same relative quality tier. Second, there is significant variation in the *intra*-county relative growth rates of various quality tiers.

Moreover, in the first and third rows of Figure 2, where the national housing market was trending upwards (boom and recovery), the distribution of low quality growth rates places more mass on its right tail relative to its high quality counterpart. Hence, in periods of national house price appreciation, low quality housing tends to appreciate more on average relative to high-end housing. This observation is consistent with the results in Guerrieri et al. (2013), who demonstrate that houses with lower initial prices experience larger capital gains during nationwide housing increases.

Table 1: Return Heterogeneity for Different Periods

	return differential > 10%	
	percentage of counties	average return differential
boom (2000-06)	47%	25%
bust (2006-12)	23%	19%
recovery (2012-19)	43%	24%

Note: This table depicts summary statistics for return differences across counties in the United States. Return differential is defined as the average absolute return between the lowest and highest quality quartiles within a given county over the corresponding period.

For counties that experienced a county-wide average house price growth of more than 10% during the boom (recovery) low-end housing appreciated by 4.6% (4.9%) more than high-end housing, on average. However, the reverse observation is true during the bust: high-end homes *depreciated* by 1.4% more on average.

Although these numbers might appear small, a closer inspection of the data reveals substantial heterogeneity in the relative rates of returns for different quality tiers. For example, across the 862 counties in the sample during the boom, 47% experienced an absolute difference in the returns of low-end relative to high-end homes of more than 10%. Of these counties, 66% saw low-end homes appreciate more, and the *average* difference in returns across quality tiers was 25%. These numbers are non-negligible. During the bust, of the 1195 counties in the sample, about a quarter (23%) experienced a similar disconnect in their relative rates of return. Of these, the majority of counties (59%) saw high-end homes depreciate more than low-end homes, and the average absolute difference in returns across quality tiers was 19%. Table 1 summarizes this within market return heterogeneity.

These observations suggest that one cannot descriptively characterize the housing market as a single "market". Consequently, understanding the determinants of house price changes requires understanding the *joint* price evolution of different segments within the housing market. This analysis merits a theoretical basis.

3 A Model of Segmented Markets

To study the key drivers of heterogeneity in capital gains, as well as its role in amplifying macroeconomic shocks, I develop a tractable macroeconomic model with a housing sector. The model embeds two key features of housing markets that will be crucial in generating heterogeneous capital gains across different types of houses, (i) *segmentation* across houses of varying quality and (ii) *indivisibility* of housing units.

Agents and Preferences. Time is discrete and indexed by $t \in \mathbb{N}$. The economy is populated by overlapping generations of households. Each period, a mass $m > 1$ of households is born and all households live for two periods. Household age is indexed by $j \in \{0, 1\}$ and will be referred to interchangeably as "young" and "old".

Households have perfect foresight and there is no aggregate uncertainty. They derive utility over consumption streams of a numeraire good and housing. I assume that households purchase housing when they are young, and sell on to the next generation's young when they are old. Payoffs are specified as follows:

$$\mathcal{U}\left(\{c_{t+j,j}\}_{j=0}^{1}, h_{t,1}\right) = \log(c_{t,0}) + \beta \log(c_{t+1,1}) + \chi \log(h_{t,1}) \tag{1}$$

where $c_{t,j}$ and $h_{t,j}$ denote the consumption of the numeraire good and housing quality purchased at time t and age j, respectively. The parameter $\chi > 0$ is a utility weight on housing services and $0 < \beta < 1$ is the household's discount factor. Finally, I assume that households incur a proportional transaction cost $0 < \delta < 1$ from selling a house.

Housing Market. The housing market consists of a collection of indivisible housing units. Houses are heterogeneous and are indexed in terms of their quality $h \in \mathbb{R}^+$. This one-dimensional index can be thought of as an appropriately weighted aggregator of the various amenities households care about in a house (including location, number of bathrooms, quality of neighboring schools, etc.). House quality in any given period is distributed according to a strictly increasing cumulative distribution function $H(\cdot)$, where the support of $H(\cdot)$ is $[\underline{h}, \overline{h}] \subset \mathbb{R}^+$. Each house can accommodate at most one household, who must be an owner-occupier.

Households can purchase at most one house of a given quality. A house of quality h at time t trades at a price $q_t(h)$ in a competitive market. Note that this gives rise to a price *function* – a price for each house quality tier. Households have the option to not become home-owners, in which case they receive an exogenous outside flow of housing services of quality $0 < b < \underline{h}$ at zero cost. This assumption ensures that household utility is well-defined if households choose not to enter into the housing market. I normalize the total quantity of housing units available for purchase to unity. Since $m > 1$, at least some households will opt for this outside option in equilibrium.

Credit Frictions. Households can participate in credit markets in their youth by either lending or borrowing at a gross interest rate of $R > 1$. However, the presence of credit frictions implies that households are subject to a loan-to-value constraint

$$d_{t,1} \leq (1 - \theta) q_{t+1}(h_{t,1}) \tag{2}$$

where $d_{t,1}$ is the household's choice of debt at time t, and θ is a proportional down-payment required by the household.[4] Note that this constraint implies that non-home owning households cannot borrow at all with respect to their second-period income.

[4] This can be microfounded as in Iacoviello (2005). Borrowers can repudiate on their debt obligations while lenders can only repossess the borrower's assets by paying a proportional transaction cost $\theta q_{t+1}(h_{t,j})$. Therefore, lenders will never lend more than $(1 - \theta) q_{t+1}(h_{t,j})$ in a subgame perfect equilibrium. With this microfoundation, the fact the down-payment is paid in proportion to *future* house prices is needed to avoid default in equilibrium following unanticipated shocks. However, this particular specification is inessential for the main results.

In order to isolate the direct effect of various changes in macroeconomic conditions on the cross-section of house prices, I take R to be exogenous. However, in the analysis that follows, I also characterize how changes in the real rate affect house prices in order to inform the full general equilibrium effect.

Endowments. Households are born without any initial wealth, but are heterogeneous in their first period income. In particular, each household receives income y at age 0, which is distributed according to a strictly increasing cumulative distribution function $F(\cdot)$ with support $[\underline{y}, \overline{y}] \subset \mathbb{R}^+$. In the second period, households receive a common income stream $y_1 > 0$. This single margin of heterogeneity is a tractable way to generate sorting between households and house quality type, but is not necessary to obtain the main results. In Online Appendix B.1, I demonstrate how the model can be extended to incorporate household heterogeneity in old income as well as income risk.

3.1 Equilibrium

We are now in a position to state the household problem. A household with income y that is born at time t takes the pricing functions $\{q_t(h)\}_t^{t+1}$ as given, and solves

$$\sup_{\{c_{t+j,j}\}_{j=0}^1, h_{t,1}, d_{t,1}} \mathcal{U}\left(\{c_{t+j,j}\}_{j=0}^1, h_{t,1}\right) \tag{3}$$

$$\text{s.t. } c_{t,0} + q_t(h_{t,1}) = y + d_{t,1} \tag{4}$$

$$c_{t+1,1} = y_1 - Rd_{t,1} + (1-\delta)q_{t+1}(h_{t,1}) \tag{5}$$

$$d_{t,1} \leq (1-\theta)q_{t+1}(h_{t,1}) \tag{6}$$

where (4)-(5) are the household's first and second-period budget constraint, respectively, and (6) is the loan-to-value constraint described above.

In the following analysis, I focus on the economy's response to unanticipated changes in various parameters from a stationary competitive equilibrium, defined as follows:

Definition 1. *Given a distribution of income earnings $F(\cdot)$, a distribution of house qualities $H(\cdot)$, and a gross interest rate R, a stationary equilibrium is a set of time-invariant individual decision functions $\{c_j^*\}_{j=0}^1$: $[\underline{y}, \overline{y}] \to \mathbb{R}^+$, d_1^*: $[\underline{y}, \overline{y}] \to \mathbb{R}$ and h_1^*: $[\underline{y}, \overline{y}] \to \{b\} \times [\underline{h}, \overline{h}]$, as well as a pricing function q^*: $[\underline{h}, \overline{h}] \to \mathbb{R}^+$, such that*

1. *The individual decision functions solve the household problem (3) - (6)*
2. *The market for each housing quality h clears*

$$H(h) = m \int_{\underline{y}}^{\overline{y}} 1[\underline{h} \leq h_1^*(y) \leq h] \mathrm{d}F(y), \quad \forall\, h \in [\underline{h}, \overline{h}] \tag{7}$$

3. *The housing pricing function is time-invariant $q_t(\cdot) = q^*(\cdot)$, $\forall\ t \in \mathbb{N}$*

Equation (7) ensures that the number of household that choose a housing quality less than h is equal to the total number of available houses below that quality. The key difference in this definition from the standard case with homogeneous housing is that there exists a *continuum* of market clearing conditions. This extends the standard market clearing condition in which the per unit price of housing clears through the *total* amount of housing demanded. In this case, the equivalent market clearing condition is given by the single equation:

$$\bar{H} = \int_{\underline{h}}^{\bar{h}} h \mathrm{d}H(h) = m \int_{\underline{y}}^{\overline{y}} h_1^*(y) \mathrm{d}F(y) \tag{8}$$

where the homogeneous stock quantity $\bar{H}$ is given by the sum of all housing services available in this economy.[5] We will often be interested in contrasting the predictions of our segmented economy with this *homogeneous housing benchmark*.

3.2 Further Extensions and Comments

Housing Supply. In analyzing the response of house price changes to aggregate shocks, I intentionally abstract from housing *supply* movements to isolate *demand*-side factors. This can be viewed as a useful approximation over short-term horizons, when housing supply is inelastic. Moreover, declining housing markets are generally associated with prices that are below the minimum profitable production cost, which makes the supply schedule across all housing qualities inelastic (Glaeser and Gyourko, 2018). Online Appendix B.2 presents an extension of this model with free entry into a construction sector in which houses are demand-determined following contractions.

Rental Markets. The absence of rental markets is inessential to the main analysis. The model can be extended to allow for a fixed rental rate of housing, in which case the assumption $b < \underline{h}$ would reflect the observed segmentation between rental and owner-occupied markets (Garriga and Hedlund, 2020; Greenwald and Guren, 2021). Online Appendix B.3 presents an extension of the model in which households can purchase rental units from a rental stock.

Bequests. Online Appendix B.4 extends the baseline model to incorporate bequests. The introduction of bequests leaves the main theoretical results unchanged. However, bequests can introduce additional variation in marginal propensities to consume, thereby affecting the propagation of house price movements to aggregate consumption expenditures. This mechanism is studied in greater detail in Section 5.2.

Trading-Up. Note that houses are priced by *young* households because they enjoy housing services later in life. Hence, the model abstracts from the "trading-up" of houses over the life-cycle, but naturally incorporates the well-documented importance of first-buyers in shaping aggregate housing fluctuations (Ortalo-Magne and Rady, 2006). Moreover, although this is a simplified model of housing that abstracts from long-term mortgages (given that it is only a two-period model), it captures various realistic features of the housing market. First, housing is a good that can be used

[5] Note that this definition of a homogeneous housing stock accommodates the restriction that a minimum level of housing services $\underline{h}$ must be purchased (as in *e.g.* Kaplan et al., 2020), while ensuring that the price per housing unit is constant.

to transfer resources across states. Second, credit frictions impede the household from borrowing up to the full value of the house. Third, house price changes during the second period of life constitute pure *wealth effects*. This is consistent with a large empirical and theoretical literature documenting that the effect of house price changes on consumption is well approximated by marginal propensities to consume (Mian et al., 2013; Berger et al., 2018).

4 Equilibrium Characterization

I now characterize the main features of the stationary equilibrium. I show that it admits a unique, differentiable price function characterized by two implicit ordinary differential equations. Next, I show that the equilibrium exhibits *positive assortative matching* (PAM) between household income and house quality. Finally, I consider how the price function responds to "localized" shocks that only affect a subset of home-purchasing households. I show that an implication of PAM is that these shocks can propagate upwards across house quality tiers (a "trickle up" effect), but not downwards. Finally, I leverage these results to analyze how house prices respond to aggregate shocks — unanticipated changes to the loan-to-value constraint, interest rates, transaction costs, and future income.

4.1 Spillovers and Positive Assortative Matching

First, I show that the pricing function is strictly increasing and everywhere differentiable for all house qualities. This implies that we may use the households' first-order condition to derive a pricing equation for each house quality tier, which links the distribution of households' willingness to pay for housing services to house prices across all quality tiers.

Proposition 1. *$q^*(\cdot)$ is strictly increasing and differentiable $\forall\ h \in (\underline{h}, \overline{h})$. For an unconstrained household with income y, the pricing function satisfies:*

$$\left(1 - \frac{1-\delta}{R}\right) \frac{\mathrm{d}}{\mathrm{d}h} q^*(h_1^*(y)) = \underbrace{\chi \frac{c_0^*(y)}{h_1^*(y)}}_{MRS} \tag{9}$$

where $c_0^*(y) = \frac{1}{1+\beta}\left(y + \frac{y_1}{R} - \frac{R+\delta-1}{R} q^*(h_1^*(y))\right)$. For a constrained household:

$$\left(1 - \frac{1-\delta}{R}\right) \frac{\mathrm{d}}{\mathrm{d}h} q^*(h_1^*(y)) = \chi \frac{1}{h} \left(\frac{x}{\tilde{c}_0^*(y)} + \frac{(1-x)\beta R}{\tilde{c}_1^*(y)}\right)^{-1} \tag{10}$$

where $x \equiv \frac{\theta R}{R+\delta-1}$, $\tilde{c}_0^*(y) = y - \theta q^*(h_1^*(y))$ and $\tilde{c}_1^*(y) = y_1 - \left[R(1-\theta) - (1-\delta)\right] q^*(h_1^*(y))$.

Proof. See Appendix A.1. □

Equation (9) demonstrates that the slope of the pricing function at a given house quality tier is dictated by the *marginal rate of substitution* (MRS) between consumption and housing services of the household that purchases that housing type in equilibrium. The willingness of households to sacrifice consumption goods for additional housing services therefore determines the *rate of change* of prices across the quality spectrum. This is in contrast to the homogeneous housing stock case, in which the MRS of all households prices a common *per-unit* price of housing.

Equation (10) illustrates the impact of credit constraints on the pricing function. Relative to the unconstrained case, we see that what determines the rate of change of house prices is a *distorted* marginal rate of substitution: consumption enters the MRS through a linear combination of marginal utilities across young and old age. The constant $x \equiv \theta R/(R+\delta-1)$ parameterizes the distribution of household liabilities over time. As the down payment θ increases, x increases, thereby making *current* marginal utility relatively more important in pricing houses. Clearly, if $\tilde{c}_1^*(y) = \beta R \tilde{c}_0^*(y)$, the household lies on its Euler equation and Equations (9) and (10) are identical.

Notably, the presence of market segmentation and indivisibilities imply that each household becomes a "marginal" investor for the house quality that it optimally chooses to purchase in equilibrium. In other words, changes in the MRS for a particular household will affect the rate of change of prices only for that household's neighboring quality tiers. Shocks that heterogeneously affect the MRS of households along the income distribution will therefore naturally give rise to heterogeneous capital gains along the house quality distribution.

However, there are two difficulties in obtaining a clean characterization of capital gain heterogeneity in response to different kinds of shocks. First, it is difficult to establish how exactly a shock will affect a household's MRS. For example, a higher weight on housing utility χ will uniformly increase the MRS for all households, which will raise house prices. In general equilibrium, this increase in house prices will reduce the amount of resources available for consumption, having a counteracting effect on the MRS through the budget constraint. The challenge is to disentangle how these general equilibrium effects shape the pricing function.

Second, it is difficult to establish how changes in a household's MRS will affect house prices in the absence of additional information regarding the housing policy function $h_1^*(y)$. This is needed in order to relate changes in the willingness to pay for housing services along the *income* distribution to changes in prices along the house *quality* distribution. The next proposition demonstrates that every equilibrium features an assignment of households to houses in which house quality is *strictly* increasing in income.

Proposition 2 (PAM)**.** *The following two statements are true in any stationary equilibrium.*

1. *There exists a $y_p \in (\underline{y}, \overline{y})$ such that households with $y > y_p$ are homeowners.*
2. *The policy function $h_1^*(y)$ is strictly increasing in y for all $y > y_p$.*

Proof. See Appendix A.2. □

This result, which follows from the weak supermodularity in consumption and housing in the households' payoffs, allows us to characterize the allocation of houses to income types. To see this, note that monotonicity in h_1^* implies that the market clearing condition given by Equation (7) must now equate the mass of all home-owners below a given quality with the corresponding supply:

$$m\left(F(y(h)) - F(y_p)\right) = H(h), \quad \forall\, h > \underline{h} \tag{11}$$

or, equivalently:

$$y(h) = F^{-1}\left(\frac{1}{m}\left(H(h)\right) + F(y_p)\right), \quad \forall\, h > \underline{h} \tag{12}$$

which gives rise to the *inverse* matching function $y(\cdot) \equiv \left(h_1^*\right)^{-1}(\cdot)$ in closed form. As such, we can index household's policy functions through their *equilibrium* choice of housing. This observation will be very useful in characterizing how the pricing function will respond to shocks that affect different quality segments. Before doing so, I settle the question of equilibrium existence and uniqueness.

Proposition 3. *Assume $x < 1$. Then, a unique stationary equilibrium exists.*

Proof. See Appendix A.3. □

The assumption $x < 1$ ensures that the consumption of constrained households at old age is decreasing in the price of the home they had purchased when young (once one factors in mortgage payments). This suffices to uniquely pin down the price of the lowest house quality $q^*(\underline{h})$. I maintain this assumption for the remaining analysis.

4.2 Price Response to Localized Shocks

This section shows that there is a sense in which shocks propagate unidirectionally in the housing market. It extends a "no trickle down" property that is present in *static* assignment models featuring no-trade equilibria (*e.g.* Sweeney, 1974; Braid, 1981; Määttänen and Terviö, 2014; Higgins, 2023) to a dynamic environment with credit constraints and capital gains across households. In Section 5, I build on these results to understand how the cross-section of house prices responds to changes in inherently dynamic variables, such as credit frictions, transaction costs, and future income.

I begin by considering an unanticipated, permanent change to the income distribution $F(\cdot)$ to $\tilde{F}(\cdot)$. I denote all equilibrium variables following the shock with a tilde.

Theorem 1 (No Trickle Down)**.** *Suppose $F(\cdot)$ and $\tilde{F}(\cdot)$ satisfy $F(y) = \tilde{F}(y)\ \forall\, y < y'$, where $y' \in [\underline{y}, \overline{y}]$. Then, $\tilde{q}^*(h) = q^*(h)\ \forall\, h \leq h_1^*(y')$.*

Proof. See Appendix A.4. □

This theorem states that shocks that affect richer households do not affect house prices below their income group. Intuitively, this is because price changes only flow in the same direction in which *incentives* bind. Increasing the income of a given household will increase the demand for housing services *above* its original housing tier, but will leave the prices below that tier unaffected. Thus, shocks do not "trickle down". House prices at low quality tiers are invariant to distributional changes in the income of households that match to *better* houses.

However, the "reverse" reasoning does not apply. Changing the income of poorer households — while leaving the income of richer households unaffected — will generally change the *entire* distribution of house prices. The following lemma illustrates this statement for a particular kind of shock: one which causes the income distribution for poor households to first-order stochastically dominate the initial distribution (*i.e.* low-income households become richer).

Lemma 1 (Trickle Up)**.** *Suppose $\tilde{F}(y) \leq F(y)$ for all $y \in [\underline{y}, y']$ where $y' \in [y_p, \overline{y}]$ and where the inequality is strict for some open interval. Suppose further $\tilde{F}(y) = F(y)$ for all $y \in (y', \overline{y}]$. Then, $q^*(h) \leq \tilde{q}^*(h)$ for all $h \in [\underline{h}, \overline{h}]$. Moreover, $q^*(h) < \tilde{q}^*(h)$ for all $h \in (h_1^*(y'), \overline{h}]$.*

Proof. See Appendix A.5. □

Hence, increasing the income of households below a particular income quantile will increase the prices of *all* homes. At its core, this is a non-trivial implication of PAM within the model's competitive equilibrium. Intuitively, incentives for housing choice bind *upwards*. Absent any price changes, increasing the income of a particular household will induce it to demand better quality homes. However — due to PAM — richer households must still reside in the homes that now a less wealthy household is demanding. The only way these richer households can continue to reside in their original quality tiers is if they bid more for housing. One therefore gets pricing spill-overs that occur from the bottom of the house quality distribution to the top.

The presence of market segmentation implies that changes in marginal rates of substitution can "trickle up" along the quality ladder, but shocks can never "trickle down". In this sense, shocks to poorer households are most important in shaping the cross-section of housing returns because they affect the house prices of all qualities *above* them through spill-overs in quality tiers. This result stands in stark contrast to the predictions from models with a homogeneous housing stock, in which changes in the marginal rate of substitution of any given household will affect the house prices of all households in the economy. Moreover, this result is robust to the functional forms of the income and housing distribution, or the value of credit constraints and interest rates. The generalizability of this result begs for a data-driven test and motivates my empirical analysis in Section 6.

5 Segmentation and Aggregate Shocks

In this section, I consider how different *aggregate* shocks affect the cross-sectional distribution of capital gains. Because these shocks typically affect the marginal rates of substitution of all households in the economy, we cannot use the results in the previous section to characterize the distribution of capital gains. However, with segmented housing markets, the incidence of these shocks will be unequally distributed across households. Consequently, price changes will be more pronounced for quality tiers bought by households who experience a larger change in their marginal rate of substitution. The distribution of capital gains thus crucially depends on the *type* of shock under consideration.

5.1 The Price Response to Aggregate Shocks

Characterizing the distribution of capital gains for arbitrary functional forms is challenging because the pricing function is generally non-linear across quality tiers. This can induce non-monotonicities in constraint status, where wealthy households are constrained because the pricing function increases at a faster rate than their income. In order to characterize the effect of aggregate shocks on the pricing function in closed form, I maintain the following assumption:

Assumption 1. *$\beta R y_p > y_1$ and χ is not too large.*[6]

Assumption 1 is a sufficient condition that ensures that the borrowing constraint for wealthy households is not violated by bounding the marginal utility of housing. Intuitively, sufficiently small χ bounds households' MRS to ensure that house prices and debt do not grow at a faster rate than income. This ensures that constraint status is monotone in income. The next proposition characterizes the pricing function under Assumption 1.

Proposition 4. *There exists a threshold $y_c \in [\underline{y}, \overline{y}]$ such that households with $y \geq y_c$ are not credit constrained. Moreover, $q^*(h)$ satisfies the following equation for all $h \geq h_1^*(y_c)$:*

$$q^*(h) = \underbrace{\frac{\chi}{\omega(1+\beta)} \int_{h_c}^{h} \left(\frac{\tilde{h}^{\frac{\chi}{1+\beta}-1}}{h^{\frac{\chi}{1+\beta}}} \right) y(\tilde{h}) d\tilde{h}}_{(1):\ \textit{slope}} + \underbrace{\frac{y_1}{R\omega} \left(1 - \left(\frac{h_c}{h} \right)^{\frac{\chi}{1+\beta}} \right)}_{(2):\ \textit{future income}} + \underbrace{\left(\frac{h_c}{h} \right)^{\frac{\chi}{1+\beta}} q^*(h_c)}_{(3):\ \textit{spillovers}} \tag{13}$$

where $\omega \equiv \left(1 - \frac{1-\delta}{R}\right)$ and $h_c \equiv h_1^*(y_c)$.

Proof. See Appendix A.6. ☐

This result establishes three features of the effect of changes in economic fundamentals on the distribution of capital gains. First, changes in *future* income y_1 have a relatively higher impact on the prices of *high*-end homes through term (2). This is because changes in future income are mostly payoff relevant to wealthier households not on their credit constraint. The marginal rate of substitution of these households increases, and these spillovers accumulate upwards across the quality spectrum. Consequently, capital gains in response to future income changes are *increasing* in house quality.

Second, changes in credit conditions will mostly affect the prices of *low*-end homes. Note that θ does not directly appear in Equation (13) since credit is not pay-off relevant to unconstrained households. However, credit still affects high-end home prices through changes in $q^*(h_c)$. Hence, the effects of credit are larger for low house qualities, and dissipate as one moves upwards the quality ladder. Note that these spill-overs dissipate at a rate that is proportional to $h^{-\frac{\chi}{1+\beta}}$. Intuitively, this is linked to the slackness in the housing market — when χ is high (or β is low), small changes in consumption have a large effect on the household's marginal rate of substitution. This makes it more difficult for households to accommodate pricing pressures from below.

[6] Appendix A.7 contains exact parametric restrictions in terms of model primitives.

Third, high-end homes are particularly sensitive to changes in interest rates or transaction costs. This is embodied through ω, which captures the proportional change in the net present value of liabilities of a household when purchasing a more expensive home. Whenever term (1) is increasing in house quality, changes in interest rates or transaction costs will have a *proportionately* larger effect on high-end homes. The next result formalizes this discussion: it characterizes the effects of aggregate shocks on the cross-section of house prices for a large class of wealth and housing distributions.

Theorem 2 (Absolute Capital Gains)**.** *Assume that the slope of the inverse matching function is uniformly bounded from below: $y'(h) \geq \epsilon$ for all $h \in (\underline{h}, \overline{h})$ and some $\epsilon > 0$. Then, the pricing function responds as follows to permanent, unanticipated changes in parameters at time t:*

1. *Higher quality tiers experience smaller capital losses when credit tightens — $\frac{\partial q_t(h)}{\partial \theta \partial h} > 0$ for all $h > h(y')$, where $y' > 0$.*
2. *Higher quality tiers experience larger capital losses when interest rates rise — $\frac{\partial q_t(h)}{\partial R \partial h} < 0$ for $h > h(y')$, where $y' > 0$.*
3. *Higher quality tiers experience larger capital losses when transaction costs rise — $\frac{\partial q_t(h)}{\partial \delta \partial h} < 0$ for $h > h(y')$, where $y' > 0$.*
4. *Higher quality tiers experience larger capital gains when future income rises — $\frac{\partial q_t(h)}{\partial y_1 \partial h} > 0$.*

Proof. See Appendix A.7. □

This result highlights that the incidence of capital gains along the income distribution depends on the type of shock under consideration. The assumption that the slope of the inverse matching function is uniformly bounded is a technical restriction that ensures that the house price-to-income ratio remains strictly positive asymptotically. The top-left panel in Figure 3 plots the change in the pricing function due to a 10% increase in the net interest rate. The shaded areas indicate the first and third quartile of the housing quality distribution in the parameterization. An interest rate change has little effect on the prices of low-end homes that constrained households match to. This is because changes in the interest rate affects the net present value of liabilities, which is something that is mostly pay-off relevant to unconstrained households. These losses accumulate as one moves up the quality ladder due to spillovers. As such, capital losses are monotonic in quality. A similar reasoning applies to changes in transaction costs and future income changes.

In contrast, a ten percent permanent increase in the loan-to-value ratio (bottom left panel) results in larger capital losses amongst low-end homes. Moreover, these capital losses are monotonically decreasing for $h > h_c$, where the credit threshold h_c is indicated by the black dots in Figure 3. This is because changes in the loan-to-value constraint are *only* directly payoff relevant to constrained households. An increase in credit frictions magnifies the discrepancy between the marginal valuation of consumption across periods for constrained households. As a result, low-end homes experience greater capital losses. However, because of spillovers, mid-end and high-end homes also experience large capital losses. In light of Proposition 4, these spillovers dissipate at a rate of $\chi/(1 + \beta)$ and therefore tend to zero as house quality becomes sufficiently large.

Figure 3: Absolute Price Differences across Quality Tiers

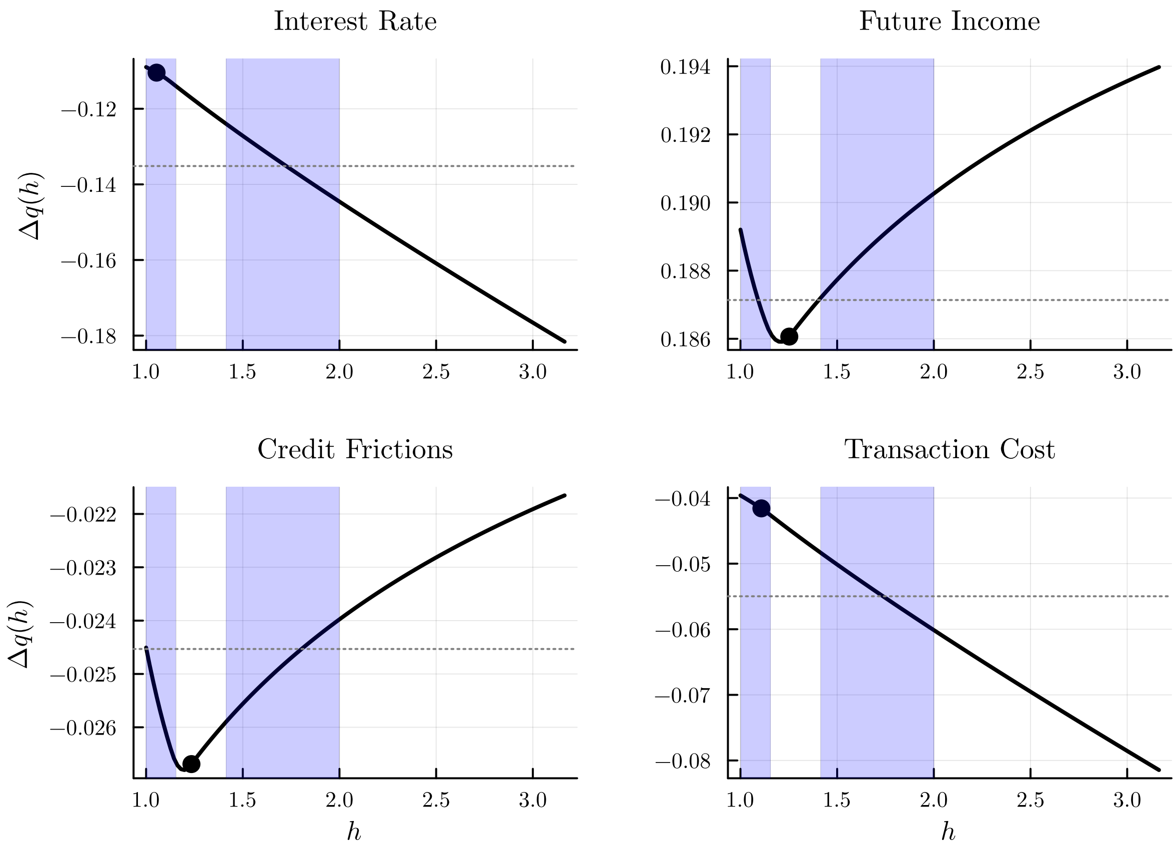

Note: Absolute price differences in response to various shocks. Each subplot depicts the capital gains of various quality tiers in response to a 10% increase in the parameters in the title of each figure. Parameters are $(\beta, R, \theta, y_1, \underline{h}, b, \underline{y}, \delta, m) = (0.8, 1.1, 0.2, 0.8, 1.0, 0.2, 0.7, 1.5, 0.1)$. The distribution for housing and income are Pareto with a common shape parameter of two. The black dots mark the credit threshold h_c after the shock. The shaded blue regions correspond to the quartiles of the housing quality distribution: the first shaded region includes houses in the bottom 25% of the quality distribution, and the second shaded region includes houses in the 50%-75% ranks. The dotted black line plots the average capital gain for all quality tiers.

The preceding analysis characterized how the pricing function responded to shocks in *absolute* terms. Crucially, these predictions are robust to arbitrary distributions for income and house qualities, so long they satisfy the uniform boundedness condition on the matching function. In Section 5.2, I show that absolute price changes are the relevant statistic in order to understand the effects of house price changes on consumption expenditures. However, the economic analyst may also be interested in the effect of shocks on the cross-section of *percentage returns*. The challenge in characterizing percentage returns is that the pricing equation in Proposition 4 is generally not linear in house qualities. However, under an additional parametric assumption on the distribution of income and house qualities, it is possible to characterize the effect of aggregate shocks on percentage returns as well.

Assumption 2. *Income and housing qualities are Pareto distributed with a common shape parameter $\alpha > 0$:*

$$F(y) = 1 - \left(\frac{\underline{y}}{y}\right)^{\alpha} \quad \text{and} \quad G(h) = 1 - \left(\frac{\underline{h}}{h}\right)^{\alpha}$$

This assumption ensures that the logarithm of the pricing function becomes *asymptotically* log-linear in house quality, which permits a closed-form characterization of the effect of shocks on percentage house price changes.

Proposition 5 (Percentage Returns)**.** *The pricing function responds as follows to permanent, unanticipated changes in parameters at time t:*

1. *Higher quality tiers experience smaller percentage losses when credit tightens — $\frac{\partial \log q_t(h)}{\partial \theta \partial h} > 0$ for $h > h_c$. Moreover, $\lim_{h\to\infty} \frac{\partial \log q_t(h)}{\partial \theta} = 0$.*

Moreover, if Assumption 2 is satisfied:

2. *In response to an interest rate increase: $\lim_{h\to\infty} \frac{\partial \log q_t(h)}{\partial R} = -\frac{1-\delta}{R(R+\delta-1)}$.*
3. *In response to a change in transaction costs: $\lim_{h\to\infty} \frac{\partial \log q_t(h)}{\partial \delta} = -\frac{1}{R+\delta-1}$.*
4. *Higher quality tiers experience lower percentage gains when future income increases: $\frac{\partial \log q_t(h)}{\partial y_1 \partial h} < 0$ for $h > h(y')$, where $y' > 0$. Moreover, $\lim_{h\to\infty} \frac{\partial \log q_t(h)}{\partial y_1} = 0$.*

Proof. See Appendix A.8 □

Figure 4 plots the change in the pricing function in percentage terms in response to a 10% increase in each of the variables in the subpanels. In contrast to Theorem 2, the only robust predictions one can obtain for *percentage* returns in the cross-section is for changes in credit and future income: tighter credit and higher future income mostly affect the returns of low-end homes. This is because credit is not directly payoff relevant to rich households, while future income does not *proportionately* affect the valuation of a house. For this reason, percentage returns dissipate as one moves up the housing quality ladder for these two shocks.

In the parameterization of Figure 4, interest rates mostly affect the returns of *low-end* homes. However, interest rates have two confounding effects on house prices: (i) they reduce the opportunity cost of purchasing a house by increasing mortgage payments, and (ii) they decrease the net present value of income. The first effect changes the *percentage returns* of high-end homes relatively more, while the second effect is relatively more important for *percentage returns* of low-end homes (in light of the preceding discussion). The effect of interest rates on percentage returns across the housing distribution is therefore generally ambiguous and depends on the precise form of the house quality and income distribution. In contrast, transaction costs leave the net present value of income unchanged. Consequently, an increase in transaction costs causes high-end homes to depreciate relatively more in Figure 4.

Importantly, all of these effects are absent in a model with a homogeneous housing stock, in which percentage returns are constant for all house quality tiers.

Figure 4: Percentage Price Changes across Quality Tiers

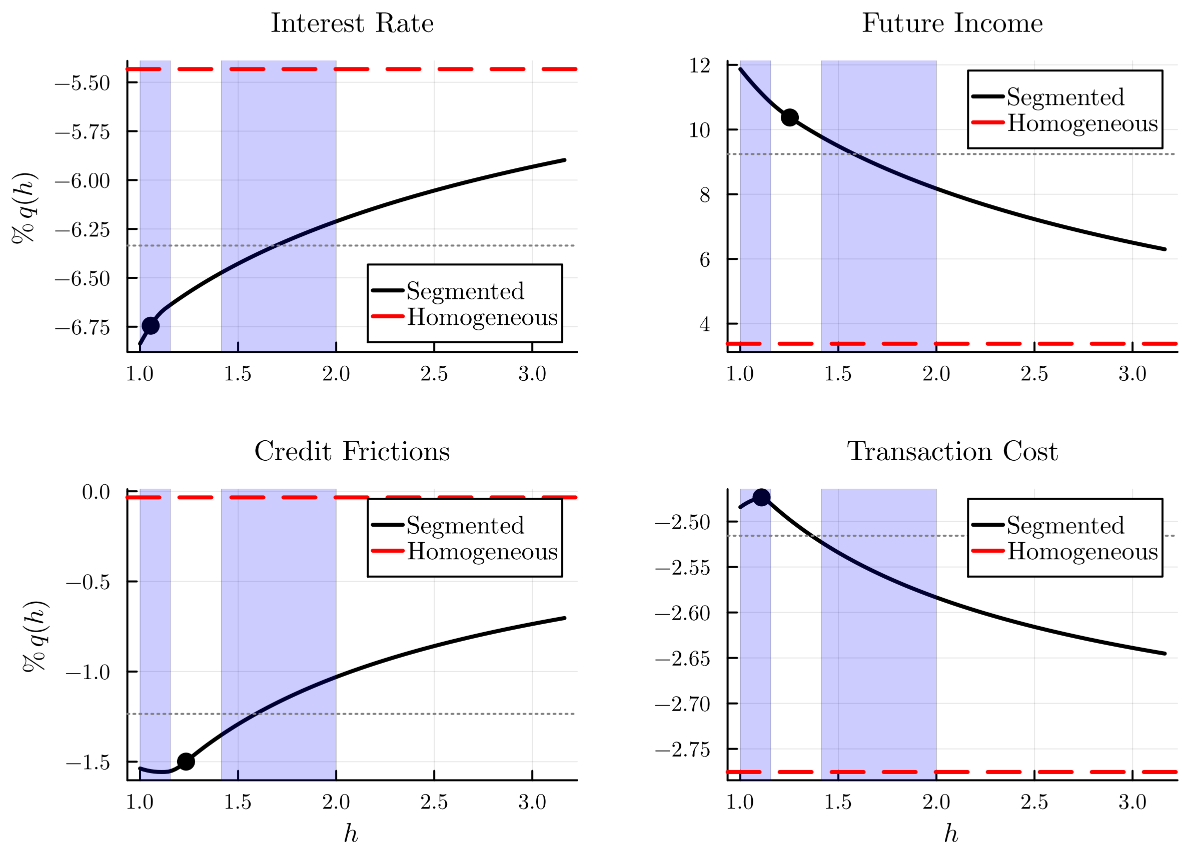

Note: Capital gains (in percent) in response to various shocks. Each subplot depicts the capital gains of various quality tiers in response to a 10% increase in the parameters in the title of each figure. See Figure 3 for parameterization. The black dots mark the credit threshold h_c after the shock. The shaded blue regions correspond to the quartiles of the housing quality distribution: the first shaded region includes houses in the bottom 25% of the quality distribution, and the second shaded region includes houses in the 50%-75% ranks. The dotted black line plots the average capital gain change across all quality tiers. The dashed red line plots the capital gains that would arise in the homogeneous housing benchmark. All parameter values between the two models are the same, except for θ which is calibrated so that the number of credit constrained households are the same in both the segmented and homogeneous housing models.

Lemma 2. *Consider a permanent, unanticipated increase in parameter $x \in \{\theta, R, \delta, -y_1\}$ at time t in the homogeneous housing benchmark:*

1. *Absolute price declines are larger for higher house qualities:* $\frac{\partial q_t(h)}{\partial x} < 0$.
2. *Percentage losses are constant for all house qualities:* $\frac{\partial^2 \log q_t(h)}{\partial x \partial h} = 0$.

Proof. See Appendix A.9 □

This is formally shown in Lemma 2, and graphically depicted by the red dotted line in Figure 4. The counterfactual homogeneous stock economy the segmented housing returns are compared to is one in which all parameter values are held constant, but in which θ is calibrated to ensure that both economies have the same number of credit constrained households. Returns are constant in the homogeneous housing economy and differ substantially from its segmented housing counterpart.

Amplification of Credit Shocks. Observe that the effect of credit on house prices converges to zero as their quality increases, but prices still drop substantially for mid-tier homes through spillovers in quality tiers. Hence, market segmentation gives rise to a mechanism in which changes in credit can have a large effect on *average* house prices, even if few households are credit constrained. In Figure 4, for example, only 30% of households are credit constrained, and yet there is an economically meaningful drop in average house prices (depicted by the dotted black line). This is in contrast to models with a homogeneous stock of housing, where changes in credit can only have a large effect on the per unit price of housing if many households are credit constrained (Greenwald and Guren, 2021). In the counterfactual homogeneous stock economy of Figure 4, for example, the drop in average house prices is negligible (red dashed line). This is because unconstrained, richer households hold most of the housing stock and are therefore more important in determining the per-unit price of housing. As the marginal rate of substitution of these households is independent of the level of credit, per-unit prices are insensitive to credit conditions. In this sense, market segmentation can amplify the effect of credit conditions on house prices.

Drivers of the Boom and Bust. The above analysis shows that different shocks have different predictions for the cross-sectional distribution of house prices, a result that is obscured when only looking at aggregate returns. Consequently, one can identify the key shocks that drive housing market returns by analyzing the distributional changes at the micro level. For example, a recent literature analyzes whether the boom-bust cycle of the Great Recession was driven by credit conditions or beliefs about house values (Cox and Ludvigson, 2021; Kaplan et al., 2020). If changes in credit conditions were the dominant force in the boom-bust cycle, we should expect low-end homes to have more volatile returns relative to mid-tier and high-end homes and *vice versa* for changes in beliefs (as captured by changes in transaction costs, which affect the perceived future resale value of a house). Moreover, Figure 4 demonstrates that the *quantitative* predictions one obtains for the effect of these shocks on *average* house prices are sensitive to the assumption of housing stock homogeneity.

Further Closed Form Characterizations. It is difficult to obtain a closed-form characterization of the entire pricing function without further specifying the distribution of house qualities and income or imposing parametric restrictions on the loan-to-value constraint. Online Appendix C.1 derives the pricing function for all house quality tiers in closed-form under the additional restriction that income and house qualities are Pareto distributed and credit constraints are maximally tight ($\theta = 1$). In this setting, the house prices of constrained households are independent of the real interest rate and increase proportionally in response to a loosening of credit. This is consistent with the survey evidence of Fuster and Zafar (2021), who find that poor households' willingness to pay for housing services is insensitive to changes in the mortgage rate, but is sensitive to credit conditions.

5.2 Market Segmentation and Consumption

This section illustrates that the heterogeneous incidence of capital gains can be an important driver of other macroeconomic aggregates, such as aggregate consumption expenditures. Intuitively, market segmentation allows larger capital losses to be experienced precisely by those households that are most important in driving aggregate consumption. In particular, the effect of house price changes on consumption will depend on how capital gains covary with a "distorted" marginal propensity to consume (MPC): a weighted difference in the MPC of young households, keeping their housing choice fixed, and the MPC of old households. Since the type of shock is a crucial determinant of the incidence of capital gains along the income distribution, the transmission of house price changes to aggregate consumption expenditures depends on the nature of the shock that hits the economy.

Concretely, let $\widetilde{MPC}_0(h)$ denote the additional consumption of a young household that arises from an incremental increase in their income, holding their housing policy choice fixed at h. Moreover, let $MPC_1(h)$ denote the marginal propensity to consume of an old household that has purchased a house quality of h. The next proposition characterizes the effect of house price changes on aggregate consumption expenditures.

Proposition 6 (Heterogeneity and Consumption)**.** *Suppose the economy is in a stationary equilibrium at time t and consider an unanticipated, permanent change to $s \in \{\theta, R, \delta, y_1\}$. Then, the effect of house price changes on aggregate consumption dC_t is:*

$$dC_t = \underbrace{\mathbb{E}[\overline{MPC}(h)]\mathbb{E}[dq_t(h)]}_{\textit{homogeneous response}} + \underbrace{cov(\overline{MPC}(h), dq_t(h))}_{\textit{matching "multiplier"}} \tag{14}$$

where $\overline{MPC}(h) = (1-\delta)MPC_1(h) - \widetilde{MPC}_0(h)$ and the expectation is over house qualities.

Proof. See Appendix A.10. □

The key insight in proving this result is to notice that changes in house prices constitute a pure *negative* income effect for young, home-purchasing households, and a *positive* wealth effect for old, home-owning households. Hence, the aggregate effect on consumption expenditures depends on the relative MPCs of home-owners (adjusted for transaction costs) to young purchasers, holding their housing policy function fixed. Intuitively, the presence of housing indivisibility and positive assortative matching implies that any given change in house prices will not affect the relative quality tier that a household will target in equilibrium. Hence, we can abstract from *substitution* effects that arise due to house price changes in the presence of market segmentation. This is in contrast to models with a homogeneous housing stock, in which changes in house prices create substitution, income, and collateral effects for household consumption, making such a simple characterization generally intractable (Berger et al., 2018).

This proposition highlights the role of redistribution in driving aggregate consumption. Conditional on the *average* fall in house prices in an economy, the *covariance* of MPCs with capital gains is an additional force that shapes the transmission of house price changes to aggregate consumption expenditures. Modeling house prices in general equilibrium is therefore important to discipline the incidence of capital gains. This is once again in contrast to models with a homogeneous housing stock, in which *any* kind of shock generally creates a negative covariance of capital gains with MPCs.[7] Although prior work has emphasized the role of redistribution in shaping the transmission mechanism of monetary policy (Auclert, 2019; Bilbiie, 2018), the preceding discussion demonstrates that housing market segmentation is a *necessary* ingredient in order for redistribution to matter for the transmission of house price changes to the rest of the economy.

In Online Appendix B.4, I extend the model to allow for more realistic MPC heterogeneity through a *warm-glow* bequest motive, as modeled in De Nardi (2004). This extended model captures the empirically relevant feature that MPCs are declining in wealth (Kaplan and Violante, 2014), while leaving the main theoretical results unchanged. Whenever MPCs are declining in wealth, the homogeneous response and matching multiplier term in Proposition 6 work in the same direction in response to a credit supply shock, because credit supply shocks result in larger capital losses precisely for those households that have high MPCs. In contrast, other shocks, such as changes in interest rates or transaction costs, induce capital losses that are mostly incident upon low MPC households. Consequently, market segmentation can give rise to a "matching multiplier" (Patterson, 2023), an amplification channel that is a product of the model's equilibrium sorting.

5.3 Market Segmentation and Welfare

This subsection shows that aggregate shocks have qualitatively different welfare implications for households in the presence of market segmentation compared to the homogeneous housing benchmark. Concretely, changes in credit have two counteracting effects on household welfare. First, more credit directly increases the welfare of constrained households by allowing them to smooth consumption intertemporally. In general equilibrium, however, the associated house price increase from this credit expansion reduces welfare. The following proposition shows that this second, general equilibrium effect is always dominant under housing market segmentation.

Proposition 7 (Credit and Welfare)**.** *Suppose θ is sufficiently low. Consider an unanticipated, permanent increase in θ at time t. Household welfare improves for all generations born at time $t \geq 0$:*

$$\frac{\mathrm{d}\mathcal{U}_t(y)}{\mathrm{d}\theta} \geq 0 \tag{15}$$

for all $y \in [\underline{y}, \bar{y}]$, where $\mathcal{U}_t(y)$ is the lifetime utility of household with income y born at time t.

Proof. See Appendix A.11. □

To understand this result, observe that the marginal homeowner is indifferent between purchasing the house of the lowest quality and their outside option. As increases in the loan-to-value

[7] By Lemma 2, changes in the per-unit house price q generate constant returns to housing across the wealth distribution. If poorer households have lower MPCs, all shocks must generate a negative covariance in response to a per-unit increase in the price of housing: $cov(MPC(h), \mathrm{d}q(h)) = cov(MPC(h), h\mathrm{d}q) < 0$.

Figure 5: Welfare Changes in Response to a Credit Reduction

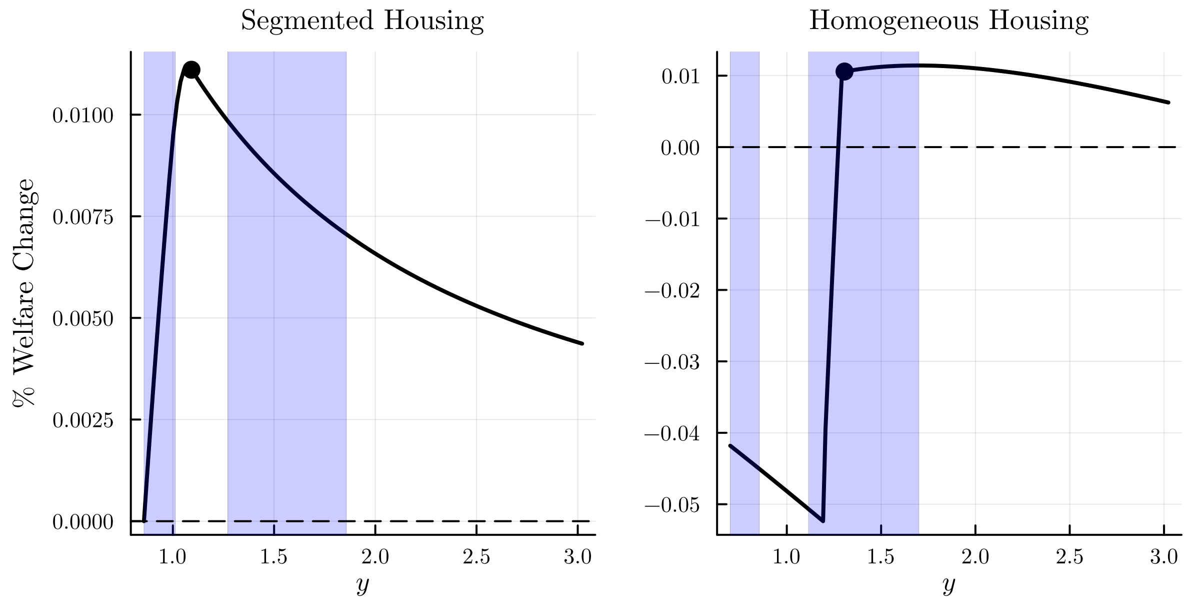

Note: Percentage welfare changes in response to a credit supply reduction (increase in θ) for the segmented housing market economy (left panel) and homogeneous housing stock benchmark (right panel). See Figure 3 for parameterization. The black dots mark the credit threshold after the shock. The shaded blue regions correspond to the first and third quartiles of the household income distribution.

constraint leaves the households' outside option unaffected, this implies that the house price of the marginal homeowner must decrease to leave their welfare exactly unchanged.[8] In light of Theorem 2, this price reduction then spills over into neighboring quality tiers, making houses more affordable for all young households and future generations, thereby raising their lifetime utility. Of course, the old generation at time t suffers welfare losses in response to this unanticipated price reduction. A reduction in credit can therefore be viewed as a redistributive transfer from the current old generation to current young households (and future generations).

This is to be contrasted with the homogeneous housing stock economy, in which a reduction in the loan-to-value constraint reduces the welfare of credit constrained households. With a homogeneous housing stock, it is *all* home-owning households that determine the per-unit price of housing. As such, prices do not decrease sufficiently to offset the direct welfare impact of lower credit for poor households. Figure 5 illustrates these results graphically. The left panel plots the welfare changes arising from a contraction in credit (an increase in θ) in the segmented housing benchmark, while the right panel plots the welfare changes that arise in the homogeneous housing economy. A reduction in credit with homogeneous housing substantially reduces the welfare of credit constrained households (left of the black dot), and increases the welfare of unconstrained households by making housing cheaper. In contrast, welfare improves for *all* households in the

[8] This result would no longer hold if the availability of credit affected the households' outside option. However, credit does not affect rental prices in common macroeconomic models (e.g. Kaplan and Violante, 2014; Greenwald and Guren, 2021). Indeed, this result is still true when rental prices are endogenized as in Appendix B.3.

economy featuring market segmentation. This discussion highlights that the welfare implications of aggregate shocks hinge critically on the assumed segmentation structure of the housing market. Although Proposition 7 formally holds for a sufficiently loose loan-to-value constraint, simulations suggest that this result is quantitatively robust to all values for θ (including maximally tight credit conditions, $\theta = 1$).

6 Testing the Theory

Section 4 demonstrated that a robust prediction of market segmentation is unidirectional shock propagation. In this section, I provide empirical evidence in favor of the model's mechanism by testing this prediction. First, I show that shocks to high-end quality tiers do not *trickle down* to lower quality tiers by exploiting plausibly exogenous variation to stock market wealth changes through a shift-share design. I show that households residing in low-end homes hold little stock market wealth, thereby making stock market wealth changes a valid shock to high-end tiers. Second, I show that shocks to low-end tiers *trickle up* to high-end tiers by exploiting plausibly exogenous variation in mortgage growth originating from an expansion of the private label securitization market (the "PLS" market). I show that this mortgage growth was mostly incident on low-end houses, making it a valid shock to test the model's trickle-up mechanism.

6.1 Test I: Trickle Down Effects of a Stock Market Wealth Shock

The test for the absence of a trickle down effect exploits the regional heterogeneity in stock market wealth using a Bartik-style shift-share instrument, as in Chodorow-Reich et al. (2021). Intuitively, we may obtain exogenous variation in stock market wealth by taking advantage of the fact that changes in aggregate stock prices (the shifter) should affect housing markets with different levels of stock market wealth (the shares) differentially.

Data. The goal is to analyze the effect of changes in stock market wealth in a given housing market on the cross-sectional distribution of house prices across quality tiers. In order to do so, I collect data (at a monthly frequency) on house prices, stock market wealth, and stock market returns. I continue using Zillow ZIP code level price indices for single family homes as my measure of house prices and define a housing market as a county as in Section 2.

I follow the model in constructing quality tiers by sorting ZIP codes according to their relative price ranking within a county. Moreover, I obtain per-capita stock market wealth at the ZIP code level using the IRS Statistics of Income (SOI) data over the period 2005-2019 (the shares), and use returns on the S&P 500 as a measure of stock market returns (the shifter). Additional details regarding data collection can be found in Online Appendix D.1.

Methodology. I denote my main regressor as $Shock_{c,t} \equiv W_{c,t-1} \times R_{c,t-1}$, where $W_{c,t-1}$ is stock market wealth per capita in period $t-1$ in county c and $R_{c,t-1}$ is the return on the S&P 500

Table 2: Average Shock Incidence Across Quality Quartiles

Quartile	1	2	3	4
Average $|Shock_{z,t}|$ (per capita)	106.2	166.2	261.3	716.6

Note: Average absolute value of the stock market wealth shock across quality quartiles over 2005-2019. Quartile 1 is the lowest quality quartile and Quartile 4 is the highest quality quartile.

between $t-1$ and t. Motivated by the theoretical results presented in Section 4, I assume th†e following specification:

$$\triangle_h p_{z,c,t} = \sum_{q=1}^{4} \left[\beta_{q,h} \times \mathbb{1}\{z \in q\} \times Shock_{c,t}\right] + \Gamma_h' X_{c,z,t} + \varepsilon_{z,c,t,h} \tag{16}$$

where $\triangle_h p_{z,c,t}$ is the absolute change in prices for a particular ZIP code z in county c between $t+h$ and $t-1$, $\mathbb{1}\{\cdot\}$ is an indicator for whether a ZIP code belongs to a housing quartile q, $X_{c,z,t}$ is a vector of ZIP-county-time specific controls, and $\varepsilon_{z,c,t,h}$ is an idiosyncratic error component. The coefficients $\{\beta_{q,h}\}_{q=1}^{4}$ capture the heterogeneous response of house prices across quality quartiles to changes in total stock market wealth per capita for different time horizons h.[9]

The key assumption for identification in this framework is that high and low wealth counties are not heterogeneously affected by other aggregate variables that co-move with stock market returns, conditional on the controls. This condition mirrors the parallel trends assumption in a continuous difference-in-difference design with multiple treatments (Goldsmith-Pinkham et al., 2020). Chodorow-Reich et al. (2021) argue extensively that this assumption is satisfied in their specification that studies the response of labor market variables to the stock market wealth shock. Below, I provide further evidence for exogeneity of the shock by illustrating the absence of pre-trends prior to the treatment.

Results. I first begin by providing evidence that the shock is sufficiently targeted to high-end ZIP codes to make it a valid high-end shock. Table 2 shows the average value of the shock over 2005-2019 for different quality quartiles: households that reside in ZIP codes within the highest quality quartile hold the highest amount of stock market wealth. Moreover, they hold approximately seven times as much stock market wealth, on average, than households residing in ZIP codes in the bottom quality quartile. The theory predicts that the effect of the shock on low-end homes should be small, since shocks cannot trickle down in the presence of segmentation.

Table 3 reports the effect of the shock on all quality quartiles over a 12 month horizon, for various specifications of Equation (16). In my preferred specification (final column), I non-parametrically control for time-invariant aggregate shocks that covary with stock market returns *and* the share

[9] $Shock_{c,t}$ is interpreted as the dollar change in stock market wealth per capita in a given county at time t. For this reason (and mirroring Chodorow-Reich et al., 2021), the dependent variable is in dollar terms (absolute differences).

Table 3: Wealth Shocks and Price Impacts on the Quality Distribution

	(1)	(2)	(3)	(4)
	Outcome: Absolute Price Changes ($\triangle_{12} p_{z,c,t}$)			
Quartile 1 ($\beta_{1,12}$)	0.453	0.214	0.357	0.054
	(0.284)	(0.280)	(0.266)	(0.166)
Quartile 2 ($\beta_{2,12}$)	0.441	0.339	0.401	0.206
	(0.237)	(0.230)	(0.228)	(0.162)
Quartile 3 ($\beta_{3,12}$)	0.661*	0.591*	0.590*	0.377*
	(0.276)	(0.255)	(0.253)	(0.170)
Quartile 4 ($\beta_{4,12}$)	1.672**	1.614**	1.456**	1.274**
	(0.488)	(0.469)	(0.474)	(0.432)
Lagged Returns	✓	✓	✓	✓
Time FE	✓	✓	✓	
County FE		✓	✓	✓
Quartile FE			✓	✓
State×Time FE				✓
N	2,397,225	2,397,225	2,397,225	2,397,225

Note: The table reports alternative specifications for the baseline regression (16). The dependent variable is the six-month change in the house price level for a given ZIP code. Robust standard errors double-clustered by county and time in parentheses. $*$ denotes significance at the 5% level, and $**$ at the 1% level.

of wealth within a given county by including county fixed effects. I also control for unobservable variables that might heterogeneously affect different quartiles depending on their initial level of stock market wealth by including quartile fixed effects. Furthermore, I include time fixed effects and clean up any residual correlation in the stock market wealth shock by including 12 lags of $Shock_{c,t}$. I also exploit only within state variation by including state×time fixed effects. I double cluster all standard errors at the county and time level. Clustering by time soaks up any remaining serial correlation on stock market returns, while clustering by county allows households in the same county to experience common shocks.

The impact of a one dollar stock wealth increase per capita for a given county increases the prices of high-end homes by 1.27 dollars over a six-month horizon.[10] Moreover, the effect of the shock is statistically indistinguishable from zero for the bottom two quartiles, and is small (but in line with the direct effect of the shock) for the third quality quartile. In all cases, the standard errors are small. Hence, these observations are entirely in line with Theorem 1: shocks to high-end homes do not trickle down — all price changes are contained within the highest quality quartile.

In order to show that the results are robust to different time horizons, Figure 6 plots the local

[10]This does not imply an MPC that is greater than one for non-durable consumption because — as shown in Table 2 — the incidence of the treatment is unequally distributed across quality quartiles.

Figure 6: Price Responses from a Stock Market Wealth Shock

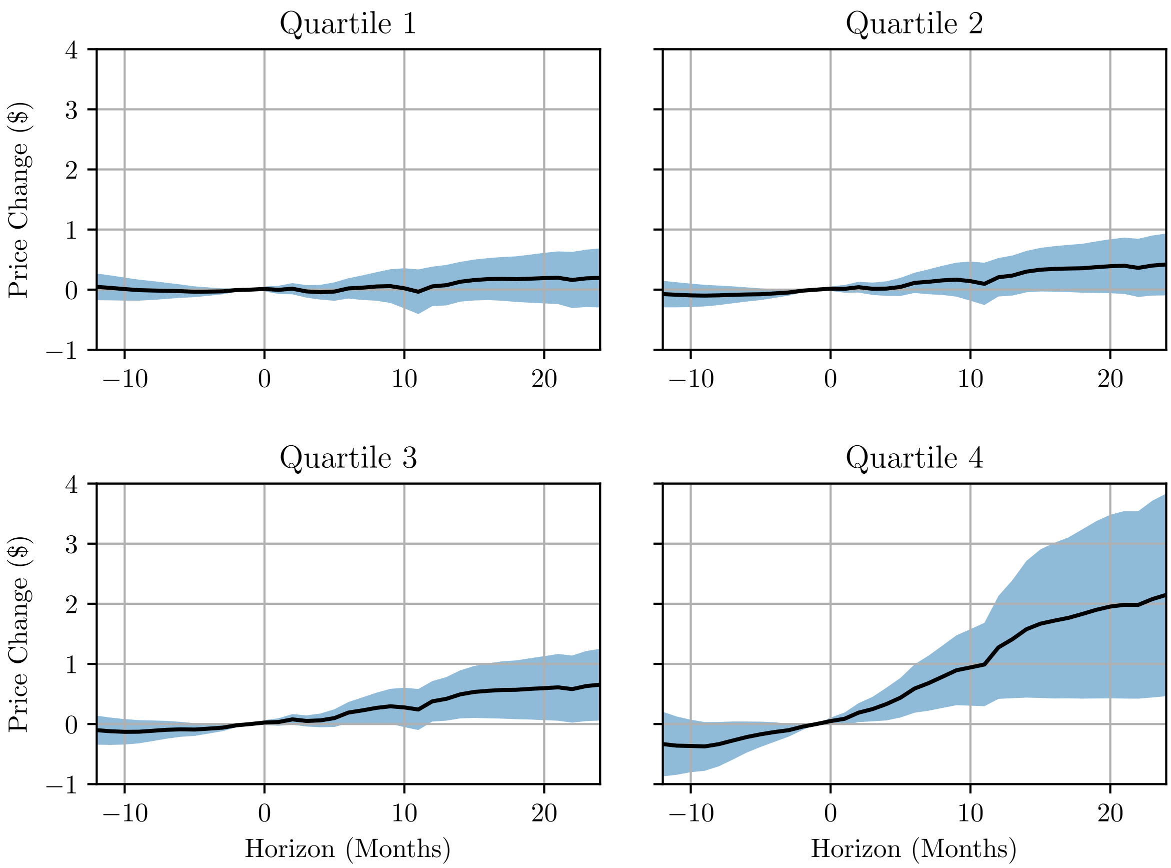

Note: Local projections to column 4 of Equation (16). The shaded area depicts 95% confidence intervals, based on standard errors clustered two-way by county and time. The shock occurs at a horizon of −1.

projections of Equation (16). The shock occurs at period −1 and is equivalent to a one dollar per capita increase in the stock market wealth for a given county. Figure 6 provides further evidence for exogeneity by showing the absence of pre-trends. The only pronounced increase in house prices occurs in the highest quality quartile, exactly as predicted by the theory. Note that the increase appears persistent — this is consistent with the interpretation that the stock market is perceived to be close to a random walk, mirroring the results of Chodorow-Reich et al. (2021) for labor market variables. The price effects for lower quality quartiles are small and statistically indistinguishable from zero at long time horizons.

Additional General Equilibrium Spill-Overs. Note that the local projections validate the predictions of Theorem 1 despite any *county*-level general equilibrium effects likely to occur through the stock market wealth shock (e.g. through aggregate demand externalities), which have *not* been controlled for (Nakamura and Steinsson, 2018). Consequently, the results suggest that the model's mechanism is robust to the presence of additional county-level general equilibrium spillovers, insofar as these spillovers positively affect the *county*-level demand for housing.

6.2 Test 2: Trickle-Up Effects of a Sub-Prime Credit Shock

I next test for trickle-up effects from low- to high-end quality tiers. I do so by isolating a subprime credit shock that was mostly incident on low-end ZIP codes. In particular, I exploit the differential city-level exposure to the sudden, pronounced increase in the PLS market in 2003 (Justiniano et al., 2019). I follow Mian and Sufi (2022) in hypothesizing that banks that relied more heavily on non-core-deposit financing were able to expand their mortgage lending more aggressively during this time period. Mian and Sufi (2022) argue that exposure to high non-core-liability (NCL) lenders provides a plausibly exogenous source of variation in mortgage credit supply across housing markets, and show that there were no differential pre-trends in credit in more exposed markets. I also document that counties that were more exposed to high NCL lenders experienced greater mortgage growth, but that this increase was mostly concentrated in *low*-end ZIP codes. Hence, this natural experiment provides an ideal setting to test for spillovers across quality tiers.

Data. Testing the trickle-up mechanism requires data on house prices, mortgage origination growth, and the liabilities of issuing institutions. For house price data, I continue using the publicly available ZIP code level Zillow series. I collect data on the flow of new mortgage loans originated in the years 2000-2010 through the "Home Mortgage Disclosure Act" (HMDA) data set. I obtain the non-core liabilities of each lender by using Call Report data and merge the two datasets using a key provided by the Federal Reserve Board. Online Appendix D.2 provides further details on the data construction.

Methodology. I follow Mian and Sufi (2022) in constructing a county-level exposure to high-NCL lenders. Concretely, I construct a county-level average of the 2002 NCL ratios of mortgage lenders, where the average is weighted by a lender's amount of mortgage originations in 2002:

$$NCLShare_{c,2002} = \sum_{b} \omega_{c,b,2002} \times NCL_{b,2002} \tag{17}$$

where

$$\omega_{c,b,2002} = \frac{Originations_{c,b,2002}}{\sum_b Originations_{c,b,2002}} \tag{18}$$

and where $NCL_{b,2002}$ denotes the non-core liabilities of bank b in 2002, and $Originations_{c,b,2002}$ denotes the amount of mortgage originations from bank b to county c in 2002.

I then regress the percentage change in house prices directly on the 2002 NCL ratio to analyze the impact of this subprime loan expansion on house prices.[11] I consider the following specification, which closely mirrors the one used for the stock-market wealth shock:

$$\triangle_{2005,2002} Price_{z,c} = \sum_{q=1}^{4} [\beta_q \times \mathbb{1}\{z \in q\} \times NCLShare_{c,2002}] + \Gamma' X_{z,c} + \varepsilon_{z,c} \tag{19}$$

[11] This approach is followed by Greenwald and Guren (2021) and Mian and Sufi (2022), who also regress outcomes directly on the instrument. This is sufficient to provide evidence for the trickle-up mechanism, although it implies that we cannot interpret the coefficients in terms of units of credit. Online Appendix E shows that the trickle-up mechanism is robust to a 2SLS procedure, although the coefficients are estimated with greater noise.

Table 4: Response of Variables to NCL Exposure

	% Change 2002-2005		
	(1) Originations	(2) Price	(3) Income
Quartile 1	7.303*	2.911**	-0.073**
	(3.014)	(0.974)	(0.028)
Quartile 2	5.868**	3.349**	-0.066
	(2.002)	(0.941)	(0.037)
Quartile 3	2.391	3.023**	-0.093*
	(1.484)	(0.878)	(0.043)
Quartile 4	0.378	2.298*	0.035
	(1.681)	(0.789)	(0.023)
Quartile FE	✓	✓	✓
N	9453	9135	9135

Note: The table reports the coefficients and standard errors from Equation (20) for the percentage growth in originations over 2002-2005 (column 1), prices (column 2) and income (column 3) for quality quartiles. Robust standard errors clustered at the county-level in parentheses. $*$ denotes significance at the 5% level and $**$ at the 1% level.

where $\triangle_{2005,2002}Price_{z,c}$ is the log-change in house prices in ZIP code z and county c over 2002-2005 and $\varepsilon_{z,c}$ reflects unmodeled determinants of house price growth. The identifying assumption is that counties did not experience shocks post-treatment that were correlated with their level of the NCL share as of 2002. Mian and Sufi (2022) argue that this orthogonality assumption is satisfied. I provide further evidence for this below by showing a lack of pre-trends.

Results. I first show that the additional supply of credit that is induced by being in a high NCL county is mostly incident on low-end ZIP codes within a county. This is necessary to ensure that price variation that is attributed to the trickle-up effect that occurs in general equilibrium is not confounded with the direct effect of providing more credit to high-end ZIP codes. To this end, I run the following regression:

$$\triangle_{2005,2002}Originations_{z,c} = \sum_{q=1}^{4} [\gamma_q \times \mathbb{1}\{z \in q\} \times NCLShare_{c,2002}] + \Gamma' X_{z,c} + \varepsilon_{z,c} \tag{20}$$

where $\triangle_{2005,2002}Originations_{z,c}$ denotes the log-change in mortgage originations. Column 1 of Table 4 depicts the coefficients γ_q for each quality quartile. These coefficients capture the expected increase in mortgage originations that occur in quality tier q when the NCL exposure of a given housing market increases from zero to one.

Table 4 shows that the treatment due to changes in the NCL share is almost entirely localized in low-end ZIP codes. An increase in the NCL share from zero to one results in a seven-fold increase in the amount of mortgages that are originated in the lowest quality quartile, but an imprecisely estimated 0.3-fold increase for the highest quality quartile. This is consistent with the notion that the vast majority of privately securitized loans were originated to buyers in low-end, and therefore low-income and low-credit score neighborhoods.[12]

Having shown that the treatment is targeted to low-end ZIP codes, I next turn my attention to its effect on house prices across the quality spectrum. Column 2 of Table 4 depicts the coefficients β_q from this regression. Changing the county-wide NCL ratio from zero to one increases the prices of homes in the bottom two quartiles of ZIP codes by an additional 290% and 335%, respectively. Moreover, there are significant trickle-up effects: a unit increase in a county's NCL ratio causes a 230% increase in house prices for the highest quality quartile. This is despite the relatively few (albeit imprecisely estimated) originations to these high-end homes. These results are consistent with Theorem 2 in that $\beta_1 > \beta_4 > 0$. Low quality homes do appreciate more in response to a credit expansion relative to high-end homes, but high-end homes also appreciate through spillovers that dissipate across the quality ladder.

The main threat to a causal interpretation of these findings is that counties with a high NCL ratio are also impacted by other covariates that affect house prices and co-move with local mortgage growth. In order to provide further evidence for the exogeneity of the NCL-based credit shock, I employ an event-study approach and run the following regression separately for each housing quality quartile:

$$\triangle_{t,2002} Prices_{z,c,t} = \xi_c + \psi_t + \sum_{h \neq 2002} \left[\beta_{q,h} \times \mathbb{1}\{t = h\} \times NCLShare_{c,2002}\right] + \varepsilon_{z,c,t} \tag{21}$$

where $\triangle_{t,2002} Prices_{z,c,t}$ denotes the percentage change in house prices that occurred between t and 2002 in ZIP code z and county c and ξ_c and ψ_t denote county and year fixed effects, respectively. The coefficients $\{\beta_{q,h}\}$ give the relative growth in prices that is attributed to being in a high NCL county in 2002 (which is the omitted year) for quartile q.

Figure 7 plots the estimated coefficients $\{\beta_{q,h}\}$ for each quality. The estimated coefficients show that there was no pre-trend and a sharp rise in house prices for high NCL areas starting in 2003. The rise is most pronounced for houses in the lowest quality quartile, but we also see an economically significant increase in the prices of high-end homes. At its peak, the coefficient for ZIP codes in the highest quality quartile was around sixty percent (4.6) the magnitude of the corresponding coefficient for low-end homes (7.1). This analysis suggests that trickle-up forces that occur in general equilibrium can have a meaningful effect on house prices across the quality distribution. As predicted by the theory, shocks that affect low-income households are crucial in shaping the response of house prices across the *entire* market. However, shocks that affect high-

[12] Mian and Sufi (2022) provide direct evidence that high NCL exposure is correlated with lower income, low credit scores, and a younger population using data from TransUnion.

Figure 7: Price Responses From a Credit Supply Shock

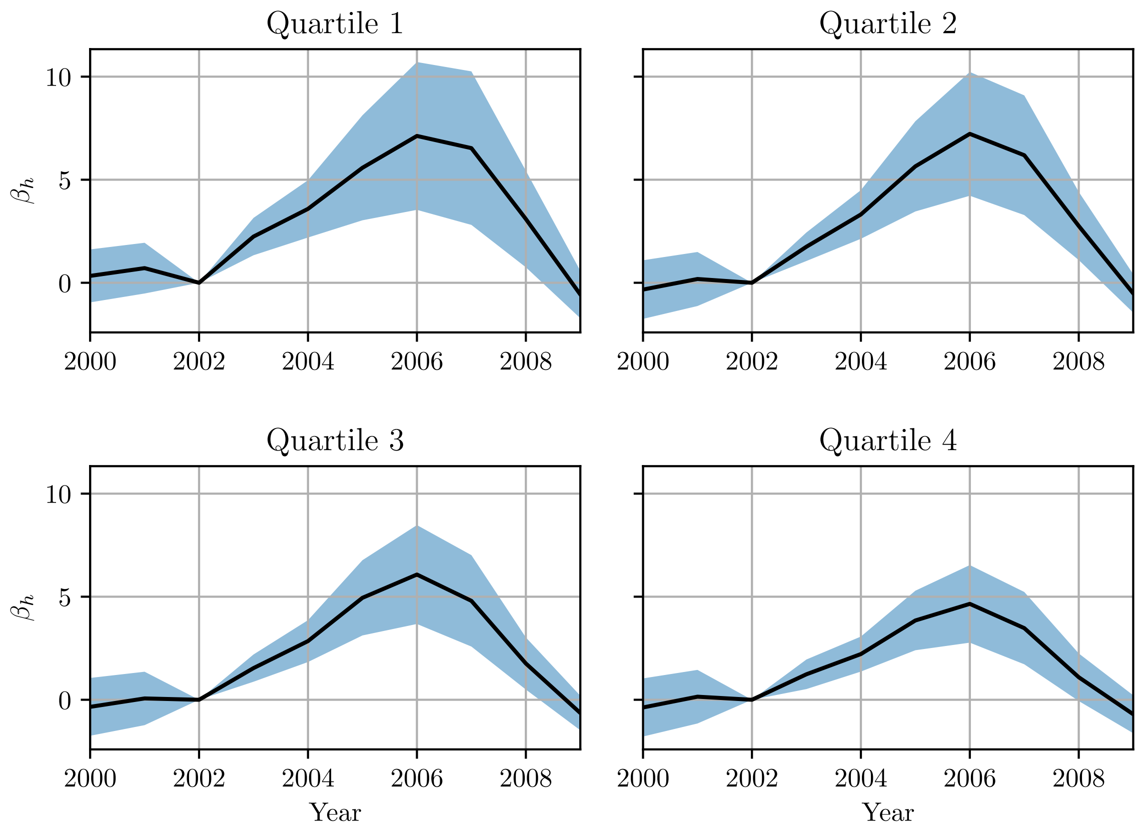

Note: The figure reports the coeffients β_h for each quality tier from Equation (21). The figure shows 95% confidence bands, based on standard errors clustered at the county-level.

income households do not trickle down to lower quality tiers. These two results in tandem highlight the role of segmentation in shaping house price dynamics.

Additional General Equilibrium Spill-Overs. It is important to note that these regressions do not control for other general equilibrium forces that might further increase the price of high-end homes. For example, high NCL counties may experience greater income growth from credit exposure through changes in aggregate demand. This would increase house prices for all quality quartiles, thereby biasing the coefficients in (19) upwards. This "contamination" of the partial equilibrium effect with county-level general equilibrium effects is a general feature of empirical designs that use variation across regions for identification (Guren et al., 2021).[13]

In order to address these concerns, I use IRS ZIP code level SOI data to see if high NCL areas also experienced higher income growth relative to NCL areas. Table 4 reports the coefficients from

[13] As discussed previously, these general equilibrium spillovers do *not* pose identification concerns for the previous trickle down test to the extent that these spillovers raise house prices.

regression (19) using income as the dependent variable. The effect of NCL exposure on income growth is low for all quality quartiles and statistically indistinguishable from zero for the highest quality quartile. In all cases, the standard errors are small. These results suggest that it is the direct, partial equilibrium, effect of increased credit that is responsible for house price growth as opposed to indirect effects that arise through Keynesian channels. Online Appendix E shows that these results are also robust to longer time horizons by employing a similar event-study regression as in Equation (21).

7 Concluding Remarks

In this paper, I develop a dynamic assignment model in which housing is segmented by various quality tiers. I theoretically characterize how the entire cross-section of house prices responds to aggregate shocks and show that two important forces in the housing market operate on different ends of the wealth and housing quality distribution: changes in credit conditions mostly affect low-end homes, while changes in interest rates, transaction costs, and future income mostly affect high-end homes. In turn, this unequal incidence of capital gains can be an important driver of household welfare and other macroeconomic aggregates by disproportionately affecting those households that are most important in driving aggregate consumption expenditures.

The data lends support to the model's robust prediction of unidirectional shock propagation. Trickle-up effects due to low-end shocks (as measured by subprime mortgage growth) have large, but dissipating effects on house prices in higher quality tiers. In contrast, high-end shocks (as measured by stock market wealth changes) only affect house prices in high-end tiers. This supports the model's prediction of no trickle down forces operating in the housing market. These results are jointly consistent with the theory and highlight the role of housing market segmentation in shaping house price dynamics in the data. Crucially, both of these tests reject models with a homogeneous housing sector, which imposes uniform rates of return for all households.

Although the homogeneous housing stock framework has dominated macroeconomic modeling of the housing market, this paper argues that treating housing as a heterogeneous, private-value asset is indispensable in obtaining a complete view of housing within the macroeconomy. Although this paper has investigated how *demand-side* factors shape pricing and welfare in the presence of segmentation, a theoretical investigation of these issues in the context of optimal *supply-side* policies is bound to be a promising avenue for future research. Nathanson (2020) and Abramson (2021) have taken first steps in exploring this issue quantitatively.

References

Abramson, B. (2021). The welfare effects of eviction and homelessness policies.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–2367.

Berger, D., Guerrieri, V., Lorenzoni, G., and Vavra, J. (2018). House prices and consumer spending. *The Review of Economic Studies*, 85(3):1502–1542.

Bilbiie, F. O. (2018). Monetary policy and heterogeneity: An analytical framework.

Bourassa, S. C., Hamelink, F., Hoesli, M., and MacGregor, B. D. (1999). Defining housing submarkets. *Journal of Housing Economics*, 8(2):160–183.

Braid, R. M. (1981). The short-run comparative statics of a rental housing market. *Journal of Urban Economics*, 10(3):286–310.

Chodorow-Reich, G., Nenov, P. T., and Simsek, A. (2021). Stock market wealth and the real economy: A local labor market approach. *American Economic Review*, 111(5):1613–57.

Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., and Van Rooij, M. (2021). Heterogeneous wealth effects. *European Economic Review*, 137:103805.

Cox, J. and Ludvigson, S. C. (2021). Drivers of the great housing boom-bust: Credit conditions, beliefs, or both? *Real Estate Economics*, 49(3):843–875.

De Nardi, M. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies*, 71(3):743–768.

Favilukis, J., Ludvigson, S. C., and Van Nieuwerburgh, S. (2017). The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy*, 125(1):140–223.

Fuster, A. and Zafar, B. (2021). The sensitivity of housing demand to financing conditions: evidence from a survey. *American Economic Journal: Economic Policy*, 13(1):231–65.

Garriga, C. and Hedlund, A. (2020). Mortgage debt, consumption, and illiquid housing markets in the great recession. *American Economic Review*, 110(6):1603–1634.

Glaeser, E. and Gyourko, J. (2018). The economic implications of housing supply. *Journal of economic perspectives*, 32(1):3–30.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.

Greenwald, D. (2018). The mortgage credit channel of macroeconomic transmission.

Greenwald, D. L. and Guren, A. (2021). Do credit conditions move house prices? Technical report, National Bureau of Economic Research.

Guerrieri, V., Hartley, D., and Hurst, E. (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics*, 100:45–60.

Guren, A. M., McKay, A., Nakamura, E., and Steinsson, J. (2021). Housing wealth effects: The long view. *The Review of Economic Studies*, 88(2):669–707.

Hacamo, I. (2021). Interest rates and the distribution of housing wealth. *Available at SSRN 2687277*.

Higgins, B. E. (2023). Racial segmentation in the us housing market. Technical report, Technical report, Stanford University. URL https://web. stanford. edu ....

Iacoviello, M. (2005). House prices, borrowing constraints, and monetary policy in the business cycle. *American economic review*, 95(3):739–764.

Justiniano, A., Primiceri, G. E., and Tambalotti, A. (2019). Credit supply and the housing boom. *Journal of Political Economy*, 127(3):1317–1350.

Kaplan, G., Mitman, K., and Violante, G. L. (2020). The housing boom and bust: Model meets evidence. *Journal of Political Economy*, 128(9):3285–3345.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Kiyotaki, N., Michaelides, A., and Nikolov, K. (2011). Winners and losers in housing markets. *Journal of Money, Credit and Banking*, 43(2-3):255–296.

Landvoigt, T., Piazzesi, M., and Schneider, M. (2015). The housing market (s) of san diego. *American Economic Review*, 105(4):1371–1407.

Määttänen, N. and Terviö, M. (2014). Income distribution and housing prices: An assignment model approach. *Journal of Economic Theory*, 151:381–410.

Mian, A., Rao, K., and Sufi, A. (2013). Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics*, 128(4):1687–1726.

Mian, A. and Sufi, A. (2014). What explains the 2007–2009 drop in employment? *Econometrica*, 82(6):2197–2223.

Mian, A. and Sufi, A. (2022). Credit supply and housing speculation. *The Review of Financial Studies*, 35(2):680–719.

Nakamura, E. and Steinsson, J. (2018). Identification in macroeconomics. *Journal of Economic Perspectives*, 32(3):59–86.

Nathanson, C. G. (2020). *Trickle-down housing economics*. Northwestern University.

Ortalo-Magne, F. and Rady, S. (2006). Housing market dynamics: On the contribution of income shocks and credit constraints. *The Review of Economic Studies*, 73(2):459–485.

Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review*, 113(4):982–1012.

Piazzesi, M. and Schneider, M. (2016). Housing and macroeconomics. *Handbook of macroeconomics*, 2:1547–1640.

Piazzesi, M., Schneider, M., and Stroebel, J. (2020). Segmented housing search. *American Economic Review*, 110(3):720–59.

Shapley, L. and Shubik, M. (1977). Trade using one commodity as a means of payment. *Journal of political economy*, 85(5):937–968.

Sweeney, J. L. (1974). A commodity hierarchy model of the rental housing market. *Journal of Urban Economics*, 1(3):288–323.

Watkins, C. A. (2001). The definition and identification of housing submarkets. *Environment and Planning A*, 33(12):2235–2253.

Appendices

A Mathematical Appendix

A.1 Proof of Proposition 1

We first prove the first part of the proposition.

Lemma 3. *$q^*(\cdot)$ is strictly increasing and differentiable for all $h \in (\underline{h}, \overline{h})$.*

Proof. The fact that $q^*(\cdot)$ is strictly increasing follows from market clearing and the fact that payoffs are strictly increasing in housing quality. By Lebesgue's Differentiation Theorem, monotonicity of $q^*(\cdot)$ then implies that $q^*(\cdot)$ is differentiable almost everywhere on $(\underline{h}, \overline{h})$. Where the derivative exists, the first-order condition satisfies

$$(q^*)'(h_1^*(y)) = f(q^*(h_1^*(y)), y) \tag{22}$$

for some continuous function $f(\cdot, \cdot)$ and $y \in (\underline{y}, \overline{y})$.

Now, in order to prove differentiability everywhere, I first argue that $q^*(\cdot)$ must be continuous. Suppose it is not and consider a point of discontinuity $h' \in (\underline{h}, \overline{h})$. Because the utility derived from housing is continuous in h, any household would prefer adjusting their housing choice in either direction as opposed to a discrete jump to $q(h)$. Moreover, this adjustment is feasible because $H(\cdot)$ is strictly increasing and the support of $H(\cdot)$ is a connected set. This proves continuity.

Next, consider the decision functions $c_j^*(\cdot)$ for $j \in \{0, 1\}$ and $h_1^*(\cdot)$. Since $q^*(\cdot)$ is continuous, the conditions of the Maximum Theorem are satisfied and the decision functions are also continuous with respect to y. Hence, $f(q^*(h_1^*(y)), y)$ is a continuous function, being the composition of continuous functions. Moreover, the fact that $F(\cdot)$ is strictly increasing implies that the function is well-defined on the interval $[\underline{y}, \overline{y}]$. The continuity of $f(q^*(h_1^*(\cdot)), \cdot)$ on the open interval $(\underline{y}, \overline{y})$ then implies that $q'(\cdot)$ is differentiable on $(\underline{h}, \overline{h})$ by the Fundamental Theorem of Calculus.

□

We now prove the second part of the proposition.

Proof. I suppress the dependence of the decision functions on y and asterisks for notational simplicity. Using the Karush-Kuhn-Tucker conditions for the household problem, we obtain the following equation for constrained households.

$$q'(h_1)\left(\frac{(1-\theta)\beta R - (1-\delta)\beta}{c_1} + \frac{\theta}{c_0}\right) = \chi \frac{1}{h_1} \tag{23}$$

Multiplying and dividing by a factor of $\frac{R+\delta-1}{R}$, and letting $x \equiv \frac{\theta R}{R+\delta-1}$ we obtain

$$q'(h_1)\left(1 - \frac{1-\delta}{R}\right) = \chi \frac{1}{h_1}\left(\frac{x}{c_0} + \frac{(1-x)\beta R}{c_1}\right)^{-1} \tag{24}$$

Finally, we substitute for c_0 and c_1 using (4) - (6). This yields (10) directly.
Households that are not constrained lie on their Euler equation

$$c_1 = \beta R c_0 \tag{25}$$

The pricing equation for unconstrained households is therefore

$$q'(h_1)\frac{1-\frac{1-\delta}{R}}{c_0} = \chi\frac{1}{h_1} \tag{26}$$

Since (6) is not binding, we substitute d_1 out of (4) and (5) and use (25) to solve for c_0. This yields

$$c_0 = \frac{1}{1+\beta}\left(y + \frac{y_1}{R} - \frac{R+\delta-1}{R}q(h_1)\right) \tag{27}$$

□

A.2 Proof of Proposition 2

Proof. I set up the problem to make use of Topkis' Theorem. The challenge is to express the household's payoff only in terms of y and h and characterize its second derivative. To this end, we substitute for optimal consumption in terms of y and h. For an unconstrained household, this is given by

$$c_0(y,h) = \frac{1}{1+\beta}\left(y + \frac{y_1}{R} - \frac{R+\delta-1}{R}q(h)\right) \tag{28}$$

$$c_1(y,h) = \frac{\beta R}{1+\beta}\left(y + \frac{y_1}{R} - \frac{R+\delta-1}{R}q(h)\right) \tag{29}$$

Similarly, we have for a constrained household

$$\tilde{c}_0(y,h) = y - \theta q(h) \tag{30}$$

$$\tilde{c}_1(y,h) = y_1 - (R(1-\theta) - (1-\delta))q(h) \tag{31}$$

The consumption functions will therefore depend on the level of debt d_1 (because this dictates whether households will be constrained or not), the housing quality choice h, and the level of first-period income y. Let $V_c(h,y)$ and $V_u(h,y)$ be the payoff of a household with first-period income y that chooses a house of quality h for a constrained and unconstrained household, respectively. We must show that these functions are supermodular in h,y.

Step 1. I first show that $V_c(h,y)$ and $V_u(h,y)$ are twice differentiable in h,y. We have

$$V_c(h,y) = \log(c_0(y,h)) + \beta\log(c_1(y,h)) + \chi\log(h) \tag{32}$$

and similarly for a constrained household. By Proposition 1, we know that $q'(h)$ is everywhere differentiable on $(\underline{h}, \overline{h})$. Twice differentiability of $V_c(\cdot,\cdot)$ and $V_u(\cdot,\cdot)$ on $(\underline{h}, \overline{h}) \times (\underline{y}, \overline{y})$ then follows.

Step 2. I now show that the cross-derivatives are strictly positive. The twice differentiability of $V_c(h, y)$ and $V_u(h, y)$ implies that we can characterize supermodularity through the second-partial derivative. For the constrained household:

$$\frac{\partial^2 V_c(h, y)}{\partial h \partial y} = q'(h)\left(\frac{\theta}{c_0^2}\right) > 0 \tag{33}$$

Similarly, it is straightforward to show for an unconstrained household:

$$\frac{\partial^2 V_u(h, y)}{\partial h \partial y} = q'(h)\frac{R + \delta - 1}{(1+\beta)Rc_0^2} > 0 \tag{34}$$

Hence, the value function is supermodular in h, y. By Topkis's theorem $h_1^*(y)$ must be strictly increasing for both constrained and unconstrained households. Finally, a household that switches from being constrained to unconstrained (or vice versa) would not want to decrease their housing quality due to the continuity of the house quality decision function. Finally, at the boundary $\underline{h}$, $q(\underline{h})$ must be such so that the marginal household is indifferent to purchasing the lowest quality home or obtaining the outside option. By the supermodularity of the payoff function, it follows that there must exist a $y_p \in [\underline{h}, \overline{h}]$ such that all households with $y > y_p$ purchase housing.

□

A.3 Proof of Proposition 3

Proof. The *allocation* of households to house qualities is unique by Proposition 2 and the argument outlined in the main text. Equilibrium uniqueness and existence then follows by showing that there is a unique pricing function that satisfies (9) and (10).

First, I argue that the initial condition $q(\underline{h})$ is well-defined and unique. It is the price that makes the marginal household with income y_p indifferent to owning or obtaining their outside option. There are four cases to consider: (i) the household is credit constrained when purchasing the house, and unconstrained otherwise, (ii) unconstrained when purchasing the house and constrained otherwise, (iii) constrained in both cases, or (iv) unconstrained in both cases. For case (i), $q(\underline{h})$ must solve the following equation:

$$\frac{(\beta R)^\beta}{(1+\beta)^(1+\beta)}\left(y_p + \frac{y_1}{R}\right)^{1+\beta} b^\chi = (y_p - \theta q(\underline{h}))\left(y_1 - ((1-\theta)R - (1-\delta))q(\underline{h})\right)^\beta \underline{h}^\chi \tag{35}$$

Note that the left hand-side is the payoff to an unconstrained household that opts for the outside option (after exponentiating), whereas the right-hand side is the payoff to a constrained household that purchases the lowest quality house (after exponentiating). Both terms on the RHS are positive due to the non-negativity of consumption in both periods. Similarly, for case (ii), $q(\underline{h})$ must solve:

$$(y_p)(y_1)^{\beta}b^{\chi} = \frac{\beta R}{(1+\beta)^2}\left(y_p + \frac{y_1}{R} - \frac{R+\delta-1}{R}q(\underline{h})\right)^{1+\beta}\underline{h}^{\chi} \tag{36}$$

For case (iii), $q(\underline{h})$ must solve

$$(y_p)(y_1)^{\beta}b^{\chi} = (y_p - \theta q(\underline{h}))\left(y_1 - ((1-\theta)R - (1-\delta))q(\underline{h})\right)^{\beta}\underline{h}^{\chi} \tag{37}$$

For case (iv), $q(\underline{h})$ must solve

$$\frac{(\beta R)^{\beta}}{(1+\beta)^{1+\beta}}\left(y_p + \frac{y_1}{R}\right)^{1+\beta}b^{\chi} = \frac{\beta R}{(1+\beta)^2}\left(y_p + \frac{y_1}{R} - \frac{R+\delta-1}{R}q(\underline{h})\right)^{1+\beta}\underline{h}^{\chi} \tag{38}$$

Furthermore, the income of the marginal home-owning household, y_p, is the unique solution to

$$m(1 - F(y_p)) = 1 \tag{39}$$

Note that the marginal household that does not purchase a house is unconstrained if and only if

$$y_p > \frac{y_1}{\beta R} \tag{40}$$

Suppose this condition is satisfied, so that the relevant cases are (i) and (iv). The left-hand side of Equations (35) and (38) are independent of $q(\underline{h})$. Further, the assumption that $x \equiv \theta R/(R+\delta-1) \in (0,1)$ implies that $(1-\theta)R - (1-\delta) > 0$. Hence, the RHS of both equations is strictly decreasing in $q(\underline{h})$.

I now argue that a unique solution exists. Note that the RHS of (35) is equal to the RHS of (38) at the constraint threshold. When $d_1(y(\underline{h})) \leq (1-\theta)q(\underline{h})$, the home-owning household is unconstrained and $q(\underline{h})$ given by Equation (38). When $d_1(y(\underline{h})) > (1-\theta)q(\underline{h})$, the solution is given by Equation (35). We may therefore construct a continuous function that is strictly decreasing and is given by the RHS of either (35) or the RHS of (38) depending on the households' constraint status. This function is greater than the LHS of either (35) or (38) for $q(\underline{h}) = 0$ (in which the household is unconstrained by assumption) and tends to zero as $q(\underline{h})$ increases. Hence, a unique, positive solution exists by the intermediate value theorem. A similar argument shows the initial condition is unique and well-defined when $y_p \leq y_1/(\beta R)$.

I now show that the solution to the ODEs (10) and (9) is unique. We may substitute the inverse matching function to obtain the following equation for an unconstrained household:

$$\left(1 - \frac{1-\delta}{R}\right)\frac{\mathrm{d}}{\mathrm{d}h}q^*(h) = \chi\frac{c_0^*(h)}{h} \tag{41}$$

where $c_0^*(y) = \frac{1}{1+\beta}\left(y(h) + \frac{y_1}{R} - \frac{R+\delta-1}{R} q^*(h)\right)$. For a constrained household, we obtain:

$$\left(1 - \frac{1-\delta}{R}\right) \frac{\mathrm{d}}{\mathrm{d}h} q^*(h) = \chi \frac{1}{h} \left(\frac{x}{\tilde{c}_0^*(h)} + \frac{(1-x)\beta R}{\tilde{c}_1^*(h)}\right)^{-1} \tag{42}$$

where $\tilde{c}_0^*(h) = y(h) - \theta q^*(h)$ and $\tilde{c}_1^*(h) = y_1 - [R(1-\theta) - (1-\delta)]\, q^*(h)$.

We now verify the conditions for the Picard-Lindelöf theorem. Continuity in h follows from the continuity of the inverse matching function. Lipschitz continuity follows from the fact that the RHS of the ODEs are continuously differentiable on the compact interval $h \in [\underline{h}, \overline{h}]$. Therefore, the derivative is bounded and the RHS is Lipschitz continuous in h.[14]

We construct the pricing function iteratively, using either (9) or (10) depending on the constraint status of the household. Since the initial constraint status of the household is well defined and the construction on each interval is unique, the pricing function is unique on the interval $[\underline{h}, \overline{h}]$. □

A.4 Proof of Theorem 1

Proof. The matching function is identical in both economies for all $y < y'$ from Equation (12). From Proposition 3, the price function is unique on $[y, y']$ for a given matching function. Hence, the two price functions must be identical on that interval. □

A.5 Proof of Lemma 1

Proof. The inverse matching function in the two economies is:

$$y(h) = F^{-1}\left(\frac{1}{m}\left(H(h) + F(y_p)\right)\right), \quad \forall\, h > \underline{h} \tag{43}$$

$$\tilde{y}(h) = \tilde{F}^{-1}\left(\frac{1}{m}\left(H(h) + \tilde{F}(y_p)\right)\right), \quad \forall\, h > \underline{h} \tag{44}$$

Note that $\tilde{F}(y_p) = F(y_p) = \frac{m-1}{m}$ by market clearing. Since $\tilde{F}(\cdot) \le F(\cdot)$, we have $\tilde{F}^{-1}(y) > F^{-1}(y)$ and $\tilde{y}(h) > y(h)$. By (9) and (10), we observe that both pricing functions satisfy the following ODEs:

$$\tilde{q}(h) = \tilde{f}(h, q(h)) \tag{45}$$

$$q(h) = f(h, q(h)) \tag{46}$$

where $\tilde{f}(h, q(h)) > f(h, q(h))$. For the same initial condition $q(\underline{h})$, we have $\tilde{q}(h) \ge q(h)$ by Petrovitsch's Theorem. However, the initial conditions must satisfy $\tilde{q}(\underline{h}) \ge q(h)$. But this can only increase $\tilde{q}(h)$ further at every $h \in [\underline{h}, \overline{h}]$. Towards a contradiction, suppose that $\tilde{q}(h^*, q(\underline{h})) >$

[14] This follows from the mean value theorem: $f(q_2) - f(q_1) = f'(q_3)(q_2 - q_1)$ for some $q_3 \in (q_2, q_1)$ and a function $f(\cdot)$. But $f'(q)$ is continuous on a compact set, and therefore bounded. Lipschitz continuity follows.

$\tilde{q}(h^*, \tilde{q}(\underline{h}))$ for some $h^* \in [\underline{h}, \overline{h}]$, where $q(h, a)$ is the unique solution to the ODE

$$\tilde{q}'(h) = \tilde{f}(h, \tilde{q}(h)), \quad q(\underline{h}) = a \tag{47}$$

Since $\tilde{q}(h, q(\underline{h}))$ and $\tilde{q}(h, \tilde{q}(\underline{h}))$ are continuous functions of h, and $\tilde{q}(\underline{h}, q(\underline{h})) < \tilde{q}(\underline{h}, \tilde{q}(\underline{h}))$, and $\tilde{q}(h^*, q(\underline{h})) > \tilde{q}(h^*, \tilde{q}(\underline{h}))$, there must exist an $h^{**} \in (\underline{h}, h^*)$ such that the two functions are equal by the Intermediate Value Theorem. But this contradicts equilibrium uniqueness of the pricing function for the initial condition $q(h^{**})$. □

A.6 Proof of Proposition 4

Proof. First, I derive Equation (13). From Equation (41), we have:

$$\left(1 - \frac{1-\delta}{R}\right) \frac{\mathrm{d}}{\mathrm{d}h} q^*(h) = \frac{\frac{\chi}{1+\beta}\left(y(h) + \frac{y_1}{R} - \frac{R+\delta-1}{R} q^*(h)\right)}{h} \tag{48}$$

Recall that $\omega \equiv 1 - (1-\delta)/R$. Rearranging and multiplying by $h^{\frac{\chi}{1+\beta}-1}$ yields:

$$\omega \left(\frac{\chi}{1+\beta} h^{\frac{\chi}{1+\beta}-1} q^*(h) + h^{\frac{\chi}{1+\beta}} \frac{\mathrm{d}}{\mathrm{d}h} q^*(h) \right) = h^{\frac{\chi}{1+\beta}-1} \frac{\chi}{1+\beta} \left(y(h) + \frac{y_1}{R} \right) \tag{49}$$

Note that we may factor the left-hand side as follows:

$$\omega \frac{\mathrm{d}}{\mathrm{d}h} \left(h^{\frac{\chi}{1+\beta}} q^*(h) \right) = \frac{\chi}{1+\beta} \left(h^{\frac{\chi}{1+\beta}-1} y(h) + h^{\frac{\chi}{1+\beta}-1} \frac{y_1}{R} \right) \tag{50}$$

Integrating the above from h_c to h then yields Equation (13).

Next, I show that constraint status is monotone. To this end, I characterize when households are constrained as a function of house prices. Note that the first-period consumption for an unconstrained household is

$$c_0^*(h) = \frac{1}{1+\beta} \left(y(h) + \frac{y_1}{R} - \frac{R+\delta-1}{R} q^*(h) \right) \tag{51}$$

Moreover, debt satisfies

$$d_1^* = c_0^*(h) + q^*(h) - y(h) \tag{52}$$

A household is therefore unconstrained if and only if

$$c_0^*(h) + q^*(h) - y(h) \leq (1-\theta) q^*(h) \tag{53}$$

Which is equivalent to:

$$\left(\theta - \frac{R+\delta-1}{(1+\beta)R} \right) q^*(h) \leq \frac{\beta}{(1+\beta)} \left(y(h) - \frac{y_1}{\beta R} \right) \tag{54}$$

Under Assumption 1, the RHS is strictly positive, since $y(h)$ is strictly increasing and $y_p > y_1/(\beta R)$.

If $\left(\theta - \frac{R+\delta-1}{(1+\beta)R}\right) \leq 0$, then all households are not credit constrained, irrespective of the price of housing.[15] If $\left(\theta - \frac{R+\delta-1}{(1+\beta)R}\right) > 0$, then households are not credit constrained if and only if:

$$q^*(h) \leq \frac{\beta}{R(1+\beta)\theta - R + \delta - 1} y(h) - \frac{y_1}{R(1+\beta)\theta - R + \delta - 1} \tag{55}$$

A sufficient condition for constraint status monotonicity is that $q^*(h)$ does not grow as fast as the RHS of Equation (55) for $h > h_c$. Using the Leibniz formula, the derivative of (13) is:

$$\begin{aligned} q^{*'}(h) =& \frac{\chi}{\omega(1+\beta)} \frac{y(h)}{h} - \int_{h_c}^{h} \frac{\chi^2}{\omega(1+\beta)^2} \frac{\tilde{h}^{\frac{\chi}{1+\beta}-1}}{h^{\frac{\chi}{1+\beta}+1}} y(\tilde{h}) \mathrm{d}\tilde{h} \\ &+ \frac{\chi y_1}{(1+\beta)R\omega} h_c^{\frac{\chi}{1+\beta}} h^{-\frac{\chi}{1+\beta}-1} - \frac{\chi}{1+\beta} h_c^{\frac{\chi}{1+\beta}} h^{-\frac{\chi}{1+\beta}-1} q^*(h_c) \end{aligned} \tag{56}$$

A sufficient condition is that the RHS of Equation (56) is smaller than the RHS of (55) for all $h > h_c$, which is satisfied if χ is small enough. This proves that constraint status is monotone. □

A.7 Proof of Theorem 2

Proof. First, I note that the proofs for interest rates and transaction costs require the uniform boundedness assumption on the inverse matching function, while the proofs for credit frictions and future income changes do not.

Credit Frictions. First, I prove that $\frac{\partial q(h_c)}{\partial \theta} < 0$. To this end, I show that $q(\underline{h})$ is decreasing in θ. θ only affects $q(\underline{h})$ if the household is constrained. Hence, we must show that the RHS of Equation (37) is decreasing. Note that this is equivalent to:

$$-\frac{\theta}{c_0^*(y_p)} + \frac{\theta\beta R}{c_1^*(y_p)} < 0 \tag{57}$$

For a constrained household, $c_0^* > \beta R c_1^*(y_p)$. Hence, the initial condition is decreasing.

Next, I show that the ODE (10) is decreasing in θ. Differentiating the RHS of (10) yields:

$$\begin{aligned} &-\chi \frac{1}{h} \left(\frac{x}{c_0^*(h)} + \frac{(1-x)\beta R}{c_1^*(h)} \right)^{-2} \\ &\times \left(\frac{x}{c_0^{*2}} - \frac{(1-x)\beta R^2}{c_1^{*2}(h)} + \frac{R}{(R+\delta-1)c_0^*(h)} - \frac{\beta R^2}{(R+\delta-1)c_1^*(h)} \right) \end{aligned} \tag{58}$$

The last two terms in the second parentheses are positive by the Euler inequality. We can bound the sum of the first two terms as follows:

$$\frac{x}{c_0^{*2}} - \frac{(1-x)\beta R^2}{c_1^{*2}(h)} > \frac{(\beta R)^2 x - (1-x)\beta R^2}{c_1^{*2}} \tag{59}$$

[15] Intuitively, this is because θ is relatively small (the credit constraint is slack) and first period consumption decreases as house prices increase.

where we have used the Euler inequality. The RHS is positive if and only if $x \geq \frac{1}{1+\beta}$, which is a necessary condition for a household to be credit constrained by Equation (55) (which characterizes the constraint threshold).

By Petrovitsch's Theorem, it follows that $\frac{\partial q(h)}{\partial \theta} \leq 0$ for all $h > \underline{h}$, with strict inequality if some households are credit constrained.

Interest Rates. From Equation (13), we may write

$$\begin{aligned}\frac{\partial^2 q^*(h)}{\partial R \partial h} =& \frac{\partial^2}{\partial R \partial h}\left[\frac{\chi}{\omega(1+\beta)}\int_{h_c}^{h}\left(\frac{\tilde{h}^{\frac{\chi}{1+\beta}-1}}{h^{\frac{\chi}{1+\beta}}}\right)y(\tilde{h})\mathrm{d}\tilde{h}\right] \\ &+ \frac{\partial^2}{\partial R \partial h}\left[\frac{y_1}{R\omega}\left(1-\left(\frac{h_c}{h}\right)^{\frac{\chi}{1+\beta}}\right)+\left(\frac{h_c}{h}\right)^{\frac{\chi}{1+\beta}}q^*(h_c)\right]\end{aligned} \tag{60}$$

Note that the derivatives in the second term can be made arbitrarily small for large h. To the extent that the first term does not vanish, the sign of $\frac{\partial^2 q^*(h)}{\partial R \partial h}$ will therefore be determined by the sign of the first term in the Equation above for large h. Using the Leibniz rule, the first term can be expressed as:

$$\frac{\partial^2}{\partial R \partial h}\left[\frac{\chi}{\omega(1+\beta)}\int_{h_c}^{h}\left(\frac{\tilde{h}^{\frac{\chi}{1+\beta}-1}}{h^{\frac{\chi}{1+\beta}}}\right)y(\tilde{h})\mathrm{d}\tilde{h}\right] = -\frac{\partial \omega}{\partial R}\frac{1}{\omega^2}\frac{\chi}{1+\beta}\left[\frac{y(h)}{h}-\frac{\chi}{1+\beta}\frac{1}{h}\int_{h_c}^{h}\frac{\tilde{h}^{\frac{\chi}{1+\beta}-1}}{h^{\frac{\chi}{1+\beta}-1}}y(\tilde{h})\mathrm{d}\tilde{h}\right]$$

Simplifying the RHS using integration by parts then yields

$$\frac{\partial^2}{\partial R \partial h}\left[\frac{\chi}{\omega(1+\beta)}\int_{h_c}^{h}\left(\frac{\tilde{h}^{\frac{\chi}{1+\beta}-1}}{h^{\frac{\chi}{1+\beta}}}\right)y(\tilde{h})\mathrm{d}\tilde{h}\right] = -\frac{\partial \omega}{\partial R}\frac{1}{\omega^2}\frac{\chi}{1+\beta}\left[\frac{h_c^{\frac{\chi}{1+\beta}}y(h_c)}{h^{\frac{\chi}{1+\beta}+1}}+\frac{1}{h^{\frac{\chi}{1+\beta}+1}}\int_{h_c}^{h}y'(\tilde{h})\tilde{h}^{\frac{\chi}{1+\beta}}\mathrm{d}\tilde{h}\right]$$

Under the assumption that $y'(\tilde{h}) \geq \varepsilon$ uniformly for all h, the term in the integral is uniformly bounded above zero for all h. The result then follows by observing that $\partial\omega/\partial R > 0$.

Transaction Costs. The proof is identical to that for interest rates, replacing $\partial\omega/\partial R$ with $\partial\omega/\partial\delta$.

Future Income. From Equation (13), we have:

$$\frac{\partial^2 q(h)}{\partial h \partial y_1} = \left(\frac{h_c}{h}\right)^{\frac{\chi}{1+\beta}}\frac{1}{h}\left(\frac{1}{R+\delta-1}-\frac{\partial q(h_c)}{\partial y_1}\right) \tag{61}$$

I claim that $\frac{\partial q(h)}{\partial y_1} < \frac{1}{R+\delta-1}$ for all $h < h_c$. Suppose otherwise. Let h_a be the first occurrence of h for which $\frac{\partial q(h)}{\partial y_1} \geq \frac{1}{R+\delta-1}$. If the marginal household is not credit constrained, then it is straightforward to see that $\frac{\partial q(\underline{h})}{\partial y_1} < \frac{1}{R+\delta-1}$ from Equation (38). If the marginal household is constrained, we can implicitly differentiate Equation (35) to obtain:

$$\left(-\frac{1+\frac{\beta}{R}}{y_p+\frac{y_1}{R}}+\frac{\beta}{c_1^*}\right)\mathrm{d}y_1=\left(\frac{\theta}{c_0^*}+\frac{(1-\theta)R-1+\delta}{c_1^*}\right)\mathrm{d}q(\underline{h}) \tag{62}$$

If the LHS is negative, then it trivially follows that $\frac{\partial q(\underline{h})}{\partial y_1}<\frac{1}{R+\delta-1}$. Otherwise, we can construct the following bound on the response of the $\mathrm{d}q^*(\underline{h})$:

$$\mathrm{d}q(\underline{h})\leq\frac{\frac{\beta}{c_1^*}}{\frac{\theta}{c_0^*}+\frac{(1-\theta)R-1+\delta}{c_1^*}}\mathrm{d}y_1 \tag{63}$$

$$\leq\frac{\beta}{\theta\beta R+(1-\theta)R-1+\delta} \tag{64}$$

where the second line uses the Euler inequality. Hence, a sufficient condition is:

$$\frac{\beta}{R\theta(\beta-1)+R-1+\delta}\leq\frac{1}{R+\delta-1} \tag{65}$$

$$\Longrightarrow\ (\beta-1)(R+\delta-1)\leq R\theta(\beta-1) \tag{66}$$

Which is always true given our assumption that

$$x\equiv\frac{R\theta}{R+\delta-1}\leq 1 \tag{67}$$

Hence, $\frac{\partial q(\underline{h})}{\partial y_1}<\frac{1}{R+\delta-1}$. By the continuity of this derivative, the mean value theorem thus implies that $h_a>\underline{h}$. Differentiating (10) with respect to y_1, we obtain the following equation

$$\begin{aligned}\left(1-\frac{1-\delta}{R}\right)\frac{\partial q'(h)}{\partial y_1}=&\chi\frac{1}{h}\left(\frac{x}{c_0^*(h)}+\frac{(1-x)\beta R}{c_1^*(h)}\right)^{-2}\left(\frac{x}{c_0^{*2}}\left(-\theta\frac{\partial q(h)}{\partial y_1}\right)\right.\\&\left.+\frac{(1-x)\beta R}{c_1^{*2}}(1-((1-\theta)R-(1-\delta))\frac{\partial q(h)}{\partial y_1}\right)\end{aligned}$$

Using the fact that $c_1^*>\beta Rc_0^*$ (since agents are credit constrained), we obtain

$$\begin{aligned}\left(1-\frac{1-\delta}{R}\right)\frac{\partial q'(h)}{\partial y_1}<&\frac{\chi\beta R}{hc_1^{*2}}\left(\frac{x}{c_0^*(h)}+\frac{(1-x)\beta R}{c_1^*(h)}\right)^{-2}\left(\beta Rx\left(-\theta\frac{\partial q(h)}{\partial y_1}\right)\right.\\&\left.+(1-x)(1-((1-\theta)R-(1-\delta))\frac{\partial q(h)}{\partial y_1}\right)\end{aligned}$$

We can simplify the terms in the second parentheses to obtain:

$$x(R(1-\theta)-(1-\delta)-\beta R\theta)\frac{\partial q(h)}{\partial y_1}+(1-\delta-R(1-\theta))\frac{\partial q(h)}{\partial y_1}+1-x \tag{68}$$

Note that the first term is negative under the assumption that households are credit constrained from Equation (55). Moreover, $\frac{\partial q(h)}{\partial y_1}$ is positive. Hence, we can construct an upper bound to the term given by:

$$(1-\delta-R(1-\theta))\frac{\partial q(h)}{\partial y_1}+1-x \tag{69}$$

But this implies that a sufficient condition for $\frac{\partial q'(h)}{\partial y_1}<0$ is that $\frac{\partial q(h)}{\partial y_1}\geq\frac{1}{R+\delta-1}$. But since $\frac{\partial q(\underline{h})}{\partial y_1}<\frac{1}{R+\delta-1}$, continuity implies implies that no such presumed h_a can exist, which is a contradiction. Hence, $\frac{\partial q(h)}{\partial y_1}<\frac{1}{R-1+\delta}$ for all $h<h_c$. This proves that $\frac{\partial^2 q(h)}{\partial h\partial y_1}>0$.

□

A.8 Proof of Proposition 5

Proof. Substituting the Pareto functional forms into the market clearing equation given by Equation 7 yields a linear inverse matching function:

$$y(h)=kh \quad \text{where} \quad k=\frac{y_p}{\underline{h}} \tag{70}$$

Substituting for this inverse matching function into Proposition 4 yields

$$\begin{aligned} q(h)=&\frac{k\chi}{\omega(\chi+1+\beta)}\left(h-\left(\frac{(h_c)^{\frac{\chi+1+\beta}{1+\beta}}}{h^{\frac{\chi}{1+\beta}}}\right)\right)\\ &+\left(\frac{y_1}{R}\right)\left(\frac{1}{\omega}\right)\left(1-\left(\frac{h_c}{h}\right)^{\frac{\chi}{1+\beta}}\right)+\left(\frac{h_c}{h}\right)^{\frac{\chi}{1+\beta}}q(h_c)\end{aligned} \tag{71}$$

where the equation holds for $h>h_c$ and recall that $\omega\equiv\frac{R+\delta-1}{R}$

Credit Frictions. The proof follows from Theorem 2 and the observation that the pricing function is strictly increasing in house quality.

Interest Rates. To prove that $\lim_{h\to\infty}\frac{\partial\log q(h)}{\partial R}=-\frac{1-\delta}{R(R+\delta-1)}$, note that

$$\lim_{h\to\infty}\frac{\partial\log q(h)}{\partial R}=\frac{\partial}{\partial R}\log\left(\frac{k\chi}{\omega(\chi+1+\beta)}\right)=-\frac{(1-\delta)}{R(R+\delta-1)} \tag{72}$$

Transaction Costs. We have

$$\lim_{h\to\infty}\frac{\partial\log q(h)}{\partial\delta}=\frac{\partial}{\partial\delta}\log\left(\frac{k\chi}{\omega(\chi+1+\beta}\right)=-\frac{1}{R+\delta-1} \tag{73}$$

Future Income. To prove that there exists a y' such that such that $\frac{\partial\log q(h)}{\partial y_1}<0$ for all $y>y'$, note that:

$$\frac{\partial^2\log q(h)}{\partial y_1\partial h}=\frac{qq_{yh}-q_hq_y}{q^2} \tag{74}$$

The numerator of this ratio contains only one constant equal to

$$-\frac{k\chi}{\omega^2(\chi+1+\beta)R} \tag{75}$$

while all other terms in the numerator are at least of order $O\left(h^{-\frac{-\chi}{1+\beta}}\right)$. Hence, $\frac{\partial \log q(h)}{\partial y_1} < 0$ for all $y > y'$. □

A.9 Proof of Lemma 2

The pricing function in the homogeneous housing benchmark is given by $q(h) = \bar{q}h$ for some constant $\bar{q} > 0$. The fact that percentage losses are constant for all house qualities in the homogeneous housing benchmark follows trivially. The fact that absolute price declines are larger for higher house qualities follows from the observation that an increase in $x \in \{\theta, R, \delta, -y_1\}$ reduces household demand for housing for all $y \in [\underline{y}, \overline{y}]$. Consequently, $\partial\bar{q}/\partial x < 0$.

A.10 Proof of Proposition 6

Proof. House price changes for old households are pure wealth effects. Hence, their effect on consumption is $(1-\delta)MPC(h)$. Assuming the absence of substitution effects, the effect of house price changes on consumption for young households is $-\widetilde{MPC}(h)$.

It remains to show that there are no substitution effects for young households. To this end, note that house price changes do not change relative income rankings. Hence, by Lemma 1, the matching function before and after the credit shock is unchanged for young households. But this implies that there can be no substitution effects. The result then follows by integrating over house quality distribution and decomposing the expectation function. □

A.11 Proof of Proposition 7

Proof. First, note that an increase in credit does not alter the relative income rankings of households. Hence, by positive assortative matching (Proposition 2), all households purchase the same house quality before and after the credit shock. Next, by Theorem 2, $\frac{\mathrm{d}}{\mathrm{d}\theta}q^*(h) < 0$ for all $[\underline{h}, \overline{h}]$. Since credit does not directly affect the welfare of unconstrained households, this necessarily implies that the welfare of unconstrained households weakly increases because their homes become cheaper. It remains to show that the welfare of constrained households increases.

By Proposition 1, $\mathcal{U}(y)$ is strictly increasing and differentiable for all $(y_p, \bar{y})$. We may therefore write:

$$\mathcal{U}(y) = \int_{y_p}^{y} \mathcal{U}'(\tilde{y})\mathrm{d}\tilde{y} + \mathcal{U}(y_p) \tag{76}$$

where $y \in [y_p, y_c]$ and where y_c is the credit constraint income cutoff. Further, using the envelope theorem, we can write the above expression as

$$\mathcal{U}(y) = \int_{y_p}^{y} \frac{1}{c_0^*(\tilde{y})} \mathrm{d}\tilde{y} + \mathcal{U}(y_p) \tag{77}$$

Note that the outside option of the marginal home-owning household is given by:

$$\frac{(\beta R)^{\beta}}{(1+\beta)^{1+\beta}} \left(y_p + \frac{y_1}{R} \right)^{1+\beta} b^{\chi} \tag{78}$$

if the household is credit constrained and

$$y_p y_1^{\beta} b^{\chi} \tag{79}$$

otherwise. In both cases, the outside option of the household is invariant to θ. Therefore, $\frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{U}(y) = 0$ for all $y \in [\underline{y}, y_p]$. We may therefore write

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{U}(y) = \int_{y_p}^{y} \frac{\mathrm{d}}{\mathrm{d}\theta} \frac{1}{c_0^*(\tilde{y})} \mathrm{d}\tilde{y} \tag{80}$$

For constrained households, we have that

$$\frac{\mathrm{d}}{\mathrm{d}\theta} \frac{1}{c_0^*(y)} = \frac{q^*(h_1^*(y)) + \theta \frac{\mathrm{d}}{\mathrm{d}\theta} q^*(h_1^*(y))}{c_0^*(y)^2} \tag{81}$$

In order to prove that $\frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{U}(y) > 0$ for $y \in (y_p, y_c]$, it suffices to prove that Equation 81 is strictly positive. Hence, it remains to show that price reductions cannot be too negative, *i.e.*

$$\frac{\mathrm{d}}{\mathrm{d}\theta} q^*(h_1^*(y)) > -\frac{q^*(h_1^*(y))}{\theta} \tag{82}$$

But observe that $\frac{\mathrm{d}}{\mathrm{d}\theta} q^*(h_1^*(y))$ must remain uniformly bounded for all $y \in [y_p, y_c]$ and $\theta \in [0, 1]$ by the Lipschitz continuity of Equation 10 (as consumption in both periods is bounded away from zero). But this implies that this inequality is always satisfied for $\theta \to 0$. Hence,

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{U}(y) \geq 0 \tag{83}$$

for all $y \in [\underline{y}, y_c]$, where the inequality is not strict to allow the possibility that there are no households that are credit constrained. The proof follows. □

ONLINE APPENDICES

B Model Extensions

B.1 Income Heterogeneity

This subsection extends the model to the case in which income at old age is heterogeneous. I show that the trickle-up result also holds in this more general environment (Theorem 1), and the effects of different shocks on the cross-sectional distribution of house prices characterized in the main text (Theorem 2) continues to hold under the assumption of credit constraint monotonicity and linear inverse matching functions. I consider the following form of old income heterogeneity:

Assumption 3. *$y_1 = g(y) + \varepsilon_y$, where*

1. *$g(\cdot)$ is a weakly increasing function with positive range*
2. *ε_y is a random variable with positive and bounded support*

Assumption 3 ensures that greater income in youth is associated with weakly higher income, but the distribution of idiosyncratic risk can be quite arbitrary. This is sufficient to ensure that utility is strictly increasing in first-period income.

Proposition 8. *Suppose Assumption 3 is satisfied and the two following parametric conditions hold: (i) $(1-\delta) \leq (1-\theta)R$, and (ii) $(R-1)(1-\delta+\beta R) < \delta R(1+\beta)$. Then,*

1. *There exists a $y_p \in (\underline{y}, \overline{y})$ such that households with $y > y_p$ are homewoners.*
2. *The policy function $h_1^*(y)$ is strictly increasing in y for all $y > y_p$.*

Proof. The proof is an application of Topkis' Theorem. The twice differentiability of the value function follows (with only minor modifications) from the proof of Proposition 2. It remains to show that the cross-derivative of the value function is strictly positive in h and y. For a constrained household, we have

$$\frac{\partial V}{\partial y} = \frac{1}{c_0} + \beta g'(y)\mathbb{E}\frac{1}{c_1} \tag{84}$$

Differentiating with respect to h, we have:

$$\frac{\partial^2 V}{\partial y \partial h} = \frac{\theta}{c_0^2}q'(h) + \beta g'(y)\left((1-\delta) - (1-\theta)R\right)\mathbb{E}\frac{q'(h)}{-c_1^2} \tag{85}$$

Under the assumptions of the proposition, each term is strictly positive. I next turn to unconstrained households. The difficulty here is that we do not have a closed form expression for consumption. However, we may use the Euler equality:

$$\frac{1}{c_0} = \beta R \mathbb{E}\left[\frac{1}{c_1}\right] \tag{86}$$

We may differentiate with respect to h and y to obtain the following two equations

$$\frac{1}{c_0^2}\left(1+\frac{\partial d}{\partial y}\right)=\beta R\mathbb{E}\left(\frac{1}{c_1^2}\right)\left(g'(y)-R\frac{\partial d}{\partial y}\right) \tag{87}$$

$$\frac{1}{c_0^2}\left(-q'(h)+\frac{\partial d}{\partial h}\right)=\beta R\mathbb{E}\left(\frac{1}{c_1^2}\right)\left((1-\delta)q'(h)-R\frac{\partial d}{\partial h}\right) \tag{88}$$

where d is the agent's choice of debt. Differentiating the value function twice, we obtain:

$$\begin{aligned}\frac{\partial^2 V}{\partial y\partial h}=-\frac{1}{c_0^2}\left(-q'(h)+\frac{\partial d}{\partial h}\right)\left(1+\frac{\partial d}{\partial y}\right)+\frac{1}{c_0}\frac{\partial^2 d}{\partial y\partial h}+\beta\mathbb{E}\left(\frac{1}{c_1}\right)\left(-R\frac{\partial^2 d}{\partial y\partial h}\right)\\ -\beta R\mathbb{E}\left(\frac{1}{c_1^2}\right)\left(g'(y)-R\frac{\partial d}{\partial y}\right)\left((1-\delta)q'(h)-R\frac{\partial d}{\partial h}\right)\end{aligned} \tag{89}$$

Using (86) and (87), we obtain

$$\frac{\partial^2 V}{\partial y\partial h}=-\frac{1}{c_0^2}\left(1+\frac{\partial d}{\partial y}\right)\left(-\delta q'(h)+\frac{\partial d}{\partial h}(R-1)\right) \tag{90}$$

I now provide bounds for $\frac{\partial d}{\partial y}$ and $\frac{\partial d}{\partial h}$ such that the cross-derivative is strictly positive. Using Jensen's inequality, (86), and (87), we obtain

$$1+\frac{\partial d}{\partial y}>g'(y)-R\frac{\partial d}{\partial y} \tag{91}$$

Note that this implies $\frac{\partial d}{\partial y}>-1$ if $g'(y)>0$. Hence, the first parentheses of the RHS of Equation (90) is positive. Moreover, using (86) and using Jensen's inequality, we obtain

$$1<(\beta R)^2\mathbb{E}\left(\left(\frac{c_0}{c_1}\right)^2\right) \tag{92}$$

Combining this expression with (88), we obtain

$$\beta R\left(-q'(h)+\frac{\partial d}{\partial h}\right)<(1-\delta)q'(h)-R\frac{\partial d}{\partial h} \tag{93}$$

It is then easily verified that the cross-derivative is strictly positive if the conditions on R outlined in the proposition are satisfied. Global supermodularity then follows from continuity of the policy function. Hence, $h_1^*(y)$ cannot decrease as one crosses the constraint threshold.

□

The presence of positive assortative matching implies that a variant of Theorem 1 continues to hold. In particular consider two economies with different incomes CDFs $F(\cdot)$ and $\tilde{F}(\cdot)$ that are otherwise identical. Furthermore, let $q(\cdot)$ and $\tilde{q}(\cdot)$ be the respective price functions in these two economies. We have:

Proposition 9. *Suppose $F(\cdot)$ and $\tilde{F}(\cdot)$ satisfy $F(y)=\tilde{F}(y)\ \forall\, y<\tilde{y}$, where $\tilde{y}\in[\underline{y},\overline{y}]$. Suppose* *further that $\tilde{q}(\underline{h}) = q(\underline{h})$ (i.e. the prices of the lowest home qualities are identical in these two economies). Then, $\tilde{q}(h) = q(h)\ \forall\, h \le h(\tilde{y})$.*

Proof. The Lipschitz continuity of marginal utility implies uniqueness given an initial condition. The result then follows from the proof of Theorem 1. □

We may also characterize the effects of aggregate shocks on the distribution of house prices under the assumptions of (i) credit constraint monotonicity, (ii) linear inverse matching functions, (iii) $g(y) = ky$ for some $k \in \mathbb{R}^+$, and (iv) $\varepsilon_y = 0$ (no idiosyncratic risk). In this case, the results of Theorem 2 continue to hold (with the caveat that the comparative statics with respect to future income are now meaningless), which follows with minor modifications from the proof of the original proposition.

B.2 Supply-Side Movements

This section extends the main model to incorporate a construction sector. I show that all the results in the main text continue to hold in the presence of *contractionary* shocks (defined as shocks that reduce the average value of housing), as the price of housing falls below the cost of producing a new housing unit.

Developers. I assume that there exist a large number of potential developers that can choose to enter the housing market at a cost of c. Once a developer enters the housing market, they draw a random house quality from the distribution of quality quartiles $G(\cdot)$, which they may then sell to the household at its competitive price. I take $G(\cdot)$ to be exogenous, and interpret the random draw of housing as idiosyncratic risk within the development process. As such, the presence of developers can change the *mass* of the total housing stock, but will not change the relative densities between different quality tiers. Developers are risk-neutral. Free entry for developers therefore implies

$$\mathbb{E}\left[q_t(h)\right] \ge c \tag{94}$$

I assume that c is sufficiently low so that a positive stock of housing exists in equilibrium. Let s be the total mass of housing in equilibrium. In order to maintain market clearing for all quality quartiles, I further assume that c is sufficiently high so that $s < m$.[16] This implies that the home ownership rate is strictly less than one.

Analysis. I first present a simple result that relates developer costs to house prices.

Proposition 10. *A reduction in c strictly decreases house prices along all quality tiers.*

Proof. The inverse matching function in this economy is given by

$$y(h) = F^{-1}\left(\frac{s}{m}G(h) + F(y(\underline{h}))\right) \tag{95}$$

[16] Having $s > m$ would lead to vacant houses, which would necessitate a new equilibrium definition.

where $y(\underline{h}) = F^{-1}\left(1 - \frac{s}{m}\right)$. Differentiating with respect to s, we obtain

$$\frac{\partial y(h)}{\partial s} = \frac{\frac{1}{m}(G(h) - 1)}{f\left(F^{-1}\left(\frac{s}{m}G(h) + F(y(\underline{h}))\right)\right)} < 0 \tag{96}$$

It is also straightforward to show that the price which makes a household indifferent between the outside option b and the lowest house quality $\underline{h}$ is strictly decreasing in y. The initial price function was given by an equation of the form

$$q'(h) = r(q, y(h)) \tag{97}$$

whereas the new pricing function is given by

$$\tilde{q}'(h) = \tilde{r}(q, y(h)) \tag{98}$$

with $\tilde{r}(q, \cdot) < r(q, \cdot)$ and with a lower initial condition. Hence, the new price function must be uniformly lower follows from a direct application of Petrovitsch's Theorem. □

The above result also sheds light on the effect of migration on house prices. An increase in m raises house prices for all house quality tiers. Consider now a shock that decreases average house prices permanently.[17] Because the housing stock does not depreciate over time, the total supply of housing is such so that the new average price is below the minimum profitable cost. As such, contractionary shocks in this setting are isomorphic to an economy with inelastic housing supply, a point also illustrated in Glaeser and Gyourko (2018) in a model with a homogeneous housing stock.

B.3 Rental Markets

This section endogenizes the household's outside option by incorporating a homogeneous rental market into the main model. I derive conditions under which positive assortative matching holds in this new environment, and show that the main results are robust to this richer framework.

I assume that renting households can now purchase rental units from a homogeneous housing stock at a per unit price of p_r. The pay-off of a renting household that consumes c_0, c_1, and purchases r rental units is given by

$$\log(c_0) + \log(c_1) + \chi\gamma \log(r) \tag{99}$$

where $0 < \gamma < 1$ is a utility penalty of renting, as opposed to owning, a home. Moreover, I assume that the total quantity of the rental stock is fixed at $\overline{R} > 0$. In the remainder of the analysis, I assume that $\underline{y}$ is sufficiently large relative to y_1 so that renting households are not credit constrained. This considerably simplifies some of the algebra in the remainder of the analysis, but is inessential to the main results. I now derive a condition under which positive assortative matching holds in

[17] Note that an increase in θ, or a decrease in R, δ, or y_1 lower house prices for all quality tiers in the main text.

this environment. To this end, define the average income of renting households

$$\overline{Y} \equiv \int_{\underline{y}}^{y_p} y dF(y)$$

where y_p is the income for the marginal home-owning household (I will show that the average income takes this form momentarily). We have the following lemma:

Proposition 11. *Suppose the following condition is satisfied:*

$$\underline{h} > \left(\frac{y_p + \frac{y_1}{R}}{m\overline{Y} + \frac{y_1}{R}} \overline{R} \right)^{\gamma} \tag{100}$$

where $y_p = F^{-1}\left(1 - \frac{1}{m}\right)$*. Then, there exists an equilibrium such that (i) all households with* $y > y_p$ *become homeowners, and (ii) the housing policy function* $h_1^*(y)$ *is strictly increasing in* y *for all* $y > y_p$*.*

Proof. The proof proceeds by conjecturing that all households with $y \leq y_p$ are renters, and then verifying the conditions for positive assortative matching. The household's first-order condition gives rise to the policy function:

$$h = \frac{\chi\gamma \left(\frac{1}{1+\beta} \left(y + \frac{y_1}{R} - p_r r \right) \right)}{p_r} \tag{101}$$

where it is here that we have used the fact that renters are not credit constrained. Aggregating and imposing market clearing for the rental market, we can obtain the price for rents in closed form:

$$p_r = \frac{\frac{\chi\gamma}{1+\beta} \left(m\overline{Y} + \frac{y_1}{R} \right)}{\left(1 + \frac{\chi\gamma}{1+\beta} \right) \overline{R}} \tag{102}$$

We may then substitute the price back into the household's rental policy function (101) to obtain

$$r = \frac{y + \frac{y_1}{R}}{m\overline{Y} + \frac{y_1}{R}} \overline{R} \tag{103}$$

Clearly, rental choice is increasing in income. Under the assumption of the proposition, the utility from the lowest house quality tier exceeds the utility derived from optimally choosing one's rental housing in the rental market. The price of the lowest quality tier is such so as to make the household with income y_p indifferent to renting and owning. Since the pay-off function is supermodular in house quality and income, all households with income $y < y_p$ strictly prefer renting to paying y_p's indifference price. This proves (i). The fact that the housing policy function is strictly increasing in income then follows from Proposition 2. □

Note that the average price of rents is determined *independently* of house prices – it is the price of the lowest house quality quartile that adjusts to make the marginal household indifferent to renting and home-owning. This observation, in conjunction with positive assortative matching, implies that the main results are robust to the inclusion of a rental sector. The only additional consideration here is that certain aggregate shocks will now induce "variation" in the outside option by changing the price of rents. In particular, we obtain the following alternative to Theorem 2.

Theorem 3 (Capital Gains with a Rental Sector)**.** *Suppose Assumption 1 is satisfied. Then, there exists an $h_c > 0$ such that households with $y > y(h_c)$ are not credit constrained (and all other home-purchasing households are credit constrained). Moreover, the pricing function responds as follows to various permanent, unanticipated change to various parameters:*

1. *Loan-to-value constraint (θ) — $\frac{\partial q(h)}{\partial \theta \partial h} > 0$ and $\frac{\partial \log q(h)}{\partial \theta \partial h} > 0$ for $h > h(y_c)$. Moreover,* $\lim_{h\to\infty} \frac{\partial \log q(h)}{\partial \theta} = 0$
2. *Interest rate (R) — There exists a threshold y' such that $\frac{\partial q(h)}{\partial R \partial h} < 0$ for $h > h(y')$. Moreover,* $\lim_{h\to\infty} \frac{\partial \log q(h)}{\partial R} = -\frac{1-\delta}{R(R+\delta-1)}$
3. *Transaction cost (δ)— There exists a threshold y' such that $\frac{\partial q(h)}{\partial \delta \partial h} < 0$. Moreover,* $\lim_{h\to\infty} \frac{\partial \log q(h)}{\partial \delta} = -\frac{1}{R+\delta-1}$

Proof. Follows directly from the proof of Theorem 2. ☐

B.4 Extension to Bequests

In order to introduce MPC heterogeneity in a tractable way, I extend the model to allow for a *warm-glow* bequest motive as in De Nardi (2004). For simplicity, I assume that bequests leave the economy.[18] I assume that payoffs take the form:

$$\begin{aligned}\mathcal{U}\left(\{c_{t+j,j}\}_{j=0}^{1}, v_{t+1,1}, h_{t,1}\right) =& \log(c_{t,0}) + \beta \log(c_{t+1,1}) \\ &+ \beta \log(v_{t+1,1} + \phi) + \chi \log(h_{t,1})\end{aligned} \tag{104}$$

where $v_{t+1,1}$ denotes the choice of bequests at old age and $\phi > 0$ parameterizes the extent to which bequests are a *luxury* good. To see how this creates heterogeneity in the marginal propensity to consume out of house price changes, consider a sudden, proportional decrease in house prices at time t by a factor $d\lambda$. The consumption response of old households to this wealth effect is

$$dc_{t,1}(h) = \begin{cases} -(1-\delta)q_t(h)d\lambda & \text{if} \quad v_{t,1}(h) = 0 \\ -\frac{1}{2}(1-\delta)q_t(h)d\lambda & \text{if} \quad v_{t,1}(h) > 0 \end{cases} \tag{105}$$

Note that $v_{t,1} > 0$ if and only if a household's income is greater than some threshold y_b. This capture the empirically relevant feature that MPCs are declining in wealth (Kaplan and Violante, 2014).[19]

[18] Allowing bequests to be passed on the next generation's young is inessential to the main results.

[19] The effect of house price changes on the consumption expenditures of non-home-owning households is zero. Home-owning households with low MPCs can therefore be viewed as wealthy hand-to-mouth (Kaplan and Violante, 2014) due to their low *liquid* wealth.

The next corollary illustrates that the homogeneous response and matching multiplier term in Proposition 6 work in the same direction in response to a credit supply shock, thereby giving rise to an amplification channel for the response of aggregate consumption expenditures to house price changes.

Proposition 12. *Suppose* $1-\delta > \frac{2}{1+2\beta}$ *and Assumption 1 holds. Consider a permanent reduction in* θ *from a stationary equilibrium and suppose* ϕ *is large enough, so that* $h_1^*(y_b) > h_1^*(y_c)$*. Then,*

$$cov(\overline{MPC}(h), dq(h)) > 0, \quad \text{for} \quad h > h_c \tag{106}$$

Proof. See Online Appendix C.2. □

The assumption $1-\delta > \frac{2}{1+2\beta}$ is needed to ensure the marginal propensity to consume out of housing wealth is larger for old households than for young households, so that unanticipated house price drops result in an initial reduction in consumption.[20]

C Additional Theoretical Results

C.1 Special Case: Closed Form Characterization of Pricing Function

In this section, I provide a simple example which obtains the entire pricing function in closed form and illustrate the response of house prices to various shocks. I assume the income distribution and the house quality distribution is Pareto with a common shape parameter $\alpha > 0$:

$$F(y) = 1 - \left(\frac{\underline{y}}{y}\right)^{\alpha}, \quad G(h) = 1 - \left(\frac{\underline{h}}{h}\right)^{\alpha} \tag{107}$$

and cutoffs $\underline{y}$ and $\underline{h}$, respectively. Assuming $m = 1$ (so that all households are home-owners) and using (11), we obtain the following inverse matching function in terms of house qualities:

$$y(h) = \frac{\underline{y}}{\underline{h}} h \tag{108}$$

for all $h \geq \underline{h}$, which is linear in housing quality. In order to obtain the constrained pricing function in closed form, I also assume that borrowing constraints are maximally tight $\theta \approx 1$, that housing depreciates fully $\delta = 1$ (or, equivalently, that transaction costs are equal to the market value of the house) and that households are perfectly patient $\beta = 1$. Finally, I take $\underline{h} \to 0$ and $\underline{y} \to 0$ such that $\frac{\underline{y}}{\underline{h}} \to 1$,[21] This yields the following pricing ODE for constrained households:

$$q'(h) = \frac{\chi}{\theta}\frac{h - \theta q(h)}{h} \tag{109}$$

[20] Christelis et al. (2021) provide some evidence that housing wealth effects are larger for older households.

[21] Formally, we can consider a sequence of economies in which cut-offs of the income and housing distributions converge to zero at the same rate. Taking the limits of the cut-offs to zero is necessary in order to prevent indeterminacy in equilibrium. Otherwise, multiple equilibria would arise due to the presence of discrete choice. This is common in assignment models with indivisible assets due to the generic non-uniqueness of the core (Shapley and Shubik, 1977).

This is a separable ordinary differential equation, the solution of which is

$$q(h) = \frac{\chi}{(1+\chi)\theta} h \tag{110}$$

since the price of the house of lowest quality must be zero given that the poorest household has zero income and cannot borrow.

Hence, we can readily observe the impact of credit constraints along all households that will be credit constrained. In particular, in response to a temporary tightening in the borrowing constraint, the price change will be

$$\frac{d \ln q(h)}{d\theta} = -\frac{1}{\theta} < 0 \tag{111}$$

which is *independent* of the house quality h. Therefore, the model predicts that as long as households are credit constrained (and credit constraints are sufficiently tight), house price appreciation or transaction costs in response to changes in credit conditions will be uniform across all such house qualities. What can be said about house qualities for which the matched households are not credit constrained? Letting h_c, y_c denote the lowest housing and income level above which households are not constrained,[22] one can show that the pricing ODE for an unconstrained household satisfies

$$q'(h) = \chi \frac{h + \frac{y_1}{R} - q(h)}{2h} \tag{112}$$

the unique solution of which is

$$q(h) = \frac{\chi}{2+\chi} h - \phi \left(\frac{1}{h}\right)^{\frac{\chi}{2}} \left(\frac{\theta y_1}{R}\right)^{\frac{\chi}{2}+1} + \frac{y_1}{R} \qquad \text{for} \qquad h > h_c \tag{113}$$

where $\phi \equiv \frac{\chi(1+\chi)^{\frac{\chi}{2}+1}}{2+\chi}$

In the limit as $h \to \infty$, pricing for high quality houses becomes linear, but is non-linear for lower-tier houses. Moreover, we can transparently see how different shocks affect the cross-sectional distribution of capital gains. A tightening in the borrowing constraint will induce constant capital gains for all housing qualities up to h_c, and decreasing capital gains thereafter, exactly as predicted by Theorem 2.

Changes in y_1 will only affect house prices for households that can afford to tap into their future earnings. Indeed, an increase in y_1 leaves house prices for constrained households entirely unaffected, but raises house prices for higher quality tiers. Similarly, a decrease in the interest rate increases the present value of income, which boosts house prices for all qualities greater than h_c. However, it does not change the prices of low-end homes because the owners of these homes are credit constrained: marginal changes in the future income of these households does not affect their willingness to pay in their first period.

[22] Under the preceding assumptions, this is given by $y_c = \theta \frac{1+\chi}{R} y_1$.

C.2 Proof of Proposition 12

Proof. The introduction of bequests do not change the positive assortative matching of the equilibrium, which leaves the inverse matching function unchanged. Recall that $h_b \equiv h(y_b)$ is the lowest house quality for which households start leaving bequests. For ϕ large enough so that $h_b > h_c$, it also does not change the initial conditions of the pricing function. Hence, it remains to show that credit constraint status monotonicity holds under a linear matching function with bequests. Following the arguments of Proposition 4, we can show that the bequests change the pricing function as follows in a stationary equilibrium:

$$q(h) = \frac{\chi}{\omega(\chi + 1 + 2\beta)} \left(h - \left(\frac{(h_b)^{\frac{\chi+1+2\beta}{1+2\beta}}}{h^{\frac{\chi}{1+2\beta}}} \right) \right) + \left(\frac{y_1 + \phi}{R} \right) \left(\frac{1}{\omega} \right) \left(1 - \left(\frac{h_b}{h} \right)^{\frac{\chi}{1+2\beta}} \right) + \left(\frac{h_b}{h} \right)^{\frac{\chi}{1+2\beta}} q(h_b) \tag{114}$$

for all $h > h_b$, where again h_b is the threshold at which households start leaving bequests

$$h_b = \inf\{h : v_1^*(h) > \phi\} \tag{115}$$

and $v_1^*(h)$ is the stationary bequest function.

It is straightforward to see that the response of the pricing function to various shocks are the same as those outlined in Theorem 2 by following the arguments in Appendix A.7. Next, note that the $\widetilde{MPC}$ of a young household that is not credit constrained and does not leave bequests is given by $\frac{1}{1+\beta}$. The MPC of an old household that does leaves bequests is unity, and any house price changes are attenuated in the budget constraint by a factor of $(1-\delta)$. Similarly, the $\widetilde{MPC}$ for young households that leaves bequests is $\frac{1}{1+2\beta}$ and is $\frac{1}{2}$ for old households that leave bequests. The condition

$$1 - \delta > \frac{2}{1 + 2\beta} \tag{116}$$

therefore ensures that $\overline{MPC}(h) > 0$ for all $h \in (h_c, \overline{h}]$.

The condition that ϕ is large enough ensures that households start leaving bequests above the credit constraint threshold, which results in non-trivial variation in marginal propensities to consume above h_c. By Theorem 2, capital losses are a strictly decreasing function of housing quality for $h > h_c$. Moreover, MPCs and $\widetilde{MPC}$s are a non-decreasing function of housing quality. Hence, the conditions of Theorem 1 and Theorem 2 in Behboodian (2006) are satisfied and we must have $cov(\overline{MPC}_t(h), dq_t(h)) > 0$ for $h > h_c$ in response to a reduction in θ. □

D Data Appendix

D.1 Additional Details on Stock Market Wealth Shock Construction

Comprehensive stock market wealth data at the ZIP code level does not exist. I therefore follow the approach of Chodorow-Reich et al. (2021) in imputing stock market wealth at the ZIP code level. In particular, I use publicly available IRS Statistics of Income (SOI) data, which contains ZIP code aggregates of annual dividend income reported on individual tax returns over the period 2005-2019. I then convert dividend income to stock market wealth by multiplying dividends by the price to earnings ratio of the S&P 500 in the relevant year. I am then able to compute variations in stock market wealth across ZIP codes using this measure in conjunction with the monthly total return on the S&P 500 (which I use as the shifter). SOI data is available at a yearly frequency. Consequently, I interpolate total stock market wealth at the monthly level in order to conduct my analysis at the monthly frequency. In practice, this introduces little measurement error because the time series for total stock market wealth is persistent. IRS SOI data also reports the number of tax returns by ZIP code, which I use in order to construct stock market wealth measures on a per capita basis.

D.2 Additional Details on Credit Shock Construction

I collect data on the flow of new mortgage loans originated in the years 2000-2010 through the "Home Mortgage Disclosure Act" (HMDA) data set. HMDA records various characteristics of each mortgage and applicant at the loan application level. I collect information on the applicant's total loan amount, purpose of borrowing (refinancing/home purchase/home improvement), income, and whether or not (and to which kind of agency) the loan was sold to in the secondary market within that given year. Each institution in the HMDA dataset is given a unique HMDA identifier. I obtain the non-core liabilities of each lender by linking these financial institutions to Call Report data using a key provided by the Federal Reserve Board.[23]

I continue to define low quality tiers as ZIP codes that fall within the lowest price quartile within a county for a given base year. I use 2002 as the base year as it precedes the expansion in the PLS market, but the choice of a base year is immaterial given that the ranking of ZIP codes across years is persistent. I drop counties for which we have fewer than four ZIP code-observations, leading to a final sample of 9135 ZIP codes and 136 counties. Finally, I only consider mortgages originated for the purpose of purchasing a house.

[23]The key is available on Neil Bhutta's website: https://sites.google.com/site/neilbhutta/data?authuser=0

E Additional Tables and Figures

Table 5

	(1)
	$\triangle_{2005,2002} Price$
β_1	1.800
	(1.260)
β_2	0.760**
	(0.182)
β_3	1.380*
	(0.650)
β_4	1.672*
	(0.844)
Quartile FE	✓
N	9135

Note: The table reports the coefficients from Equation (19), using the NCL share as an instrument for the total change in credit growth in the corresponding county from 2002-2005. The dependent variable is the percentage change in the house price in a given ZIP code. Robust standard errors are in parentheses. $*$ denotes significance at the 5% level, and $**$ at the 1% level.

Figure 8

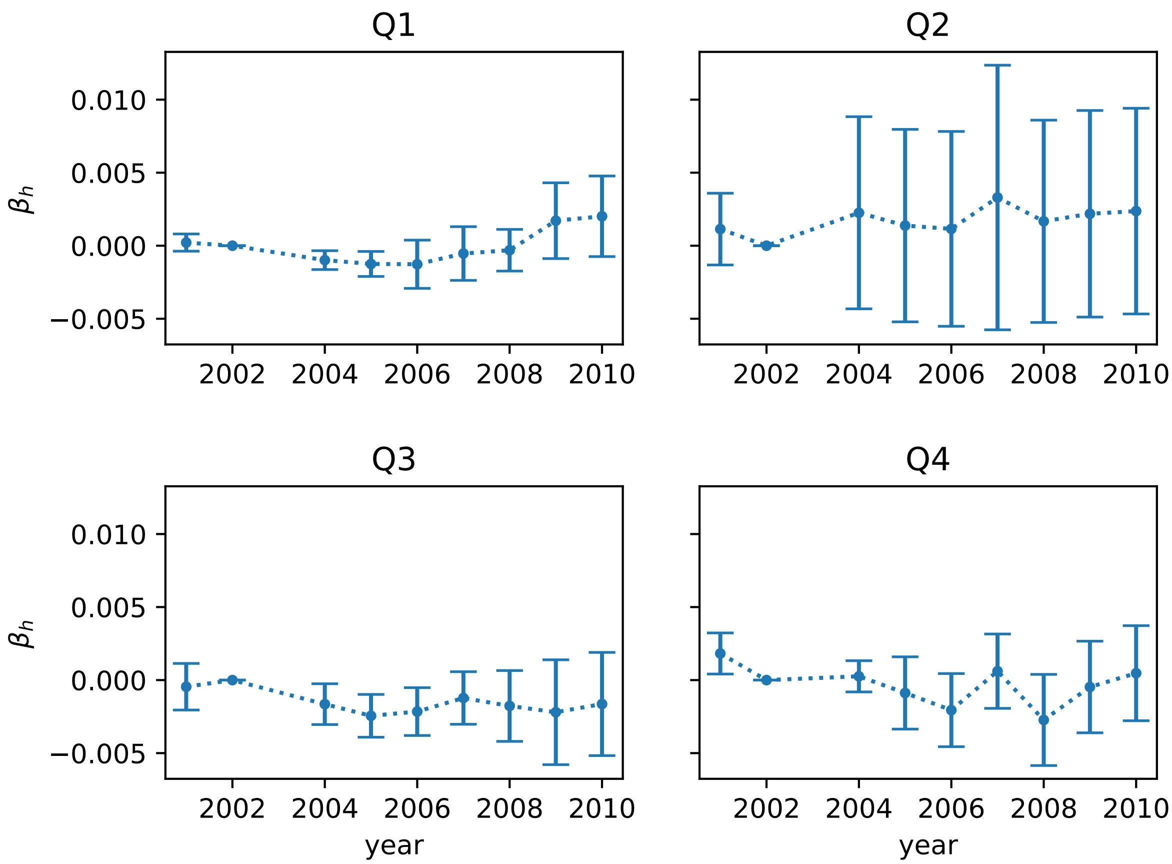

Note: The figure reports the coefficients β_h for each quality tier from Equation (21) using the percentage change in income as the dependent level. 2002 is the base year and observations are at the ZIP code level. The figure shows 95% confidence bands, based on standard errors clustered at the county-level. The years 2000 and 2003 have been omitted due to the lack of ZIP code level income data.