ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Sakong, Jung

Working Paper

Cyclical transactions and wealth inequality

Working Paper, No. WP 2022-05

Provided in Cooperation with:
Federal Reserve Bank of Chicago

Suggested Citation: Sakong, Jung (2022) : Cyclical transactions and wealth inequality, Working Paper, No. WP 2022-05, Federal Reserve Bank of Chicago, Chicago, IL, https://doi.org/10.21033/wp-2022-05

This Version is available at:
https://hdl.handle.net/10419/251015

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

Federal Reserve Bank of Chicago

Cyclical Transactions and Wealth Inequality

Jung Sakong

January 26, 2022

WP 2022-05

https://doi.org/10.21033/wp-2022-05

Working papers are not edited, and all opinions and errors are the responsibility of the author(s). The views expressed do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

Cyclical Transactions and Wealth Inequality*

Jung Sakong†

January 26, 2022

Abstract

Wealth is distributed more unevenly than income, and one contributing factor might be that richer households earn higher portfolio returns. I uncover one channel that causes portfolio returns to be increasing in wealth: Poorer households consistently buy risky assets in booms—when expected returns are low—and sell after a bust—when expected returns are high. Although time-varying expected returns are a robust empirical fact, theories are ambiguous on whether poorer or richer households engage in such cyclical trading patterns. I estimate the trading patterns for households across wealth levels, in the US housing market for 1988-2013. I interact housing ownership patterns from deeds records with household-level wealth, which I infer from merging owners' surnames with their name-based income in the 1940 full Census. The estimated dispersion in expected returns from this "buy-high-sell-low" channel is large: The interquartile-range difference is 60 basis points per year. The channel predicts that geographies with historically higher volatility will feature more wealth inequality than income inequality: I verify this implication in the data. These results suggest that a government policy intended to boost poorer households' wealth via homeownership can backfire if it ignores the status of house prices.

JEL classification: D15, D31, E32, G11, R21
Keywords: Wealth inequality, business cycles, real estate asset, timing of transaction, wealth return

*I am extremely grateful to my committee chair Amir Sufi and members Marianne Bertrand, Raghuram Rajan and Luigi Zingales for their continuous guidance and support. I also thank Greg Buchak, Peter Chen, Lancelot Henry de Frahan, Lars Hansen, John Heaton, Erik Hurst, Adam Jorring, Paymon Khorrami, Sehwa Kim, Helen Koshy, Simone Lenzu, Stefan Nagel, Krisztina Orban, Canice Prendergast, John Shim, Dan Sullivan, Ana-Maria Tenekedjieva, Douglas Xu, Alex Zentefis, Xiao Zhang, Eric Zwick, as well as seminar participants at the Corporate Finance Reading Group, Economic Dynamics Working Group, Finance Brownbag and Stigler Lunch Workshop, for helpful suggestions. All errors are my own. The views expressed in this paper are those of the author and do not reflect those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

†Federal Reserve Bank of Chicago (email: jung.sakong@chi.frb.org)

1 Introduction

Wealth is distributed more unevenly than income, even below the top 1%, which is the part of the wealth distribution where the literature has focused most.[1] One reason might be that the rate of return on wealth increases in wealth. If that is the case, poorer households could earn a lower return in two ways: (1) They participate less in risky assets that yield higher returns, or (2) they consistently participate at the "wrong" times–when prices are high and expected returns are low. Many papers have focused on the first channel. The second channel has received less attention.

In this paper, I use the US housing market to study this second channel. Constructing a new dataset, I estimate the trading patterns of households across wealth levels. Lower-wealth households do indeed consistently purchase housing when prices are high, and they sell when prices are low. I find that this "buy-high-sell-low" channel has a significant impact on wealth accumulation: the interquartile range of annual returns across wealth levels is 60 basis points.

Housing, especially ownership of a primary residence, is often seen as a vehicle for accumulating wealth by middle- and lower-wealth households.[2] Campbell (2006) shows that housing is the asset class with the highest share of total assets between the 30th and 96th percentiles of the total-asset distribution. Housing may help wealth accumulation for multiple reasons. One is that present-biased individuals may benefit by tying up wealth in an illquid asset like a house.[3] Partly to encourage wealth accumulation by the middle class, government policies have also encouraged and incentivized homeownership at least since the 1930s.[4] My findings caution government policies that encourage buying a home, however. If such policies disproportionately incentivize home purchases when prices are high, they can backfire by impeding wealth accumulation and worsening wealth inequality.

Before describing the empirical exercise, I should first clarify what I mean by poorer households "buying high and selling low." Given any data series, there will always be households who trade at the "wrong" times ex post. In order to have a lasting impact on wealth accumulation, poorer households must *consistently* buy when expected returns are low and sell when expected returns are high. If expected returns were constant, poorer households might be unlucky in some periods, but this outcome

[1] For example, compared to the bottom 50% of the income distribution, the next 49% make 4.7 times as much in income, but own 6.5 times as much in net worth, based on the 2013 Survey of Consumer Finance.

[2] Charles and Hurst (2002); Wainer and Zabel (2019)

[3] A long literature on present bias and its implication for savings took off starting with Laibson (1997). In an earlier work with co-authors, I found in field experiments that experimentally increasing the illiquidity of a savings account attracted more savings from subjects (Beshears et al. (2015)).

[4] Carliner (1998)

would balance out in other periods when they are lucky.

When expected returns are time-varying and predictable, however, households who consistently buy high and sell low will earn lower expected returns in a way that can be anticipated. Whether any household will regularly buy high and sell low is theoretically ambiguous, and some standard examples give opposite predictions. For instance, if mortgage availability increases when prices are high, *poorer* households might be more likely to buy because at other times they are rationed out of the credit market. On the other hand, if prices rise in economic booms because investors perceive overly-optimistic returns, *richer* households might be more likely buy when prices are high because they have better capacity to take advantage of the higher expected returns. This theoretical ambiguity justifies constructing a dataset and estimating who "buys high and sells low."

To precisely measure who engages in what kind of trading behavior, a dataset that contains both identifying information and observed actual quantities traded is needed. This is because even within a broad asset class such as housing or stocks, there are actual assets that differ in how their prices behave. Therefore, even if I find that poorer households' housing wealth rises more, I cannot conclude that they bought more housing units, because they may just own houses whose prices rise more. Luckily for housing, all trades are publicly observable from deeds records. Private information beyond just names and residential addresses is missing though. For this reason, the wealth of home buyers and sellers needs to be imputed.

My empirical solution is to use the house ownership data and attribute wealth levels to surnames.[5] Surnames are passed down through generations. Wealth levels can be estimated by surname using the 1940 full-count Census, which was the first Census to ask about income and is the last Census that is publicly available in full detail, because the Census Bureau only releases a full Census after 72 years.[6] In my concurrent work with a co-author, we find that the income averaged at the surname level from the 1940 Census is a strong predictor of those surnames' average-wealth levels today, constructed from individual-ownership-level data (Henry de Frahan and Sakong (2020)).

Sorting surnames into percentiles using their historical income from the 1940 Census, I find that

[5]Using surname-level variation in wealth works in the US, because there are about 160,000 surnames with 100 or more individuals. In China, by contrast, 100 most common surnames account for 85% of the population and hence using surname-level variation would not be informative. In another paper, I use this latter fact to identify Chinese buyers in the US housing market (**?**).

[6]"This '72-Year Rule' 92 Stat. 915; Public Law 95-416; October 5, 1978) restricts access to decennial census records to all but the individual named on the record or their legal heir." https://www.census.gov/history/www/genealogy/decennial_census_records/the_72_year_rule_1.html?CID=CBSM+history

poorer households buy more housing (in quantity units) than rich households when prices increase. In other words, lower income households have a higher sensitivity, or "beta", in their choices of housing quantity to price. The negative slope in beta along the wealth distribution is shown in Figures 2a and 2b.[7]

The overall negative relationship between these betas and the wealth level is driven by differences between racial groups: Non-White population groups exhibit highly pro-cyclical ownership of housing. By constrast, after controlling for the racial share at the surname level, the betas are slightly increasing in wealth level. Two interpretations are possible: (1) Belonging to a racial-minority group may be an independent predictor of low wealth, or (2) racial minorities may be particularly vulnerable to cyclical downturns.[8]

Going back to the overall negative relationship between the betas and my proxies for wealth levels (i.e., surname-level historical incomes from 1940), I wish to know how much dispersion in return on housing is generated by the timing of trades? To convert the estimated betas into interpretable differences in returns along the wealth distribution, I make two sets of transformations: First, I map the wealth-proxies to the present-day percentiles in the wealth distribution, and second, I map the betas to returns on housing.

To map each percentile of the 1940 income by surnames to the corresponding place in the present-day wealth distribution, I take two steps: (1) Using surname-level data on average primary residence value in 2012-2013, I map each 1940-income-percentile to its future housing consumption; and (2) To map housing consumption to the corresponding place in the wealth distribution, I estimate the relationship between these two variables in the 2013 Survey of Consumer Finances (SCF). Combining these two steps, I convert the the surname-level 1940 income to the present-day wealth percentiles.

To map the estimated, housing-quantity-to-price betas to returns on housing, I use the formulas I derived linking these two quantities along with estimates of expected-return variations taken from Cochrane (2011).

After conducting these transformations, I find that returns on housing go up 60 basis points per

[7] I construct and use two samples: One that maximizes the number of counties covered and another that maximizes the number of years. Sample selection is discussed in detail in Section 3.

[8] In an earlier work, I found empirical evidence that economic downturns in a local geographical area cause racial prejudice in that area to rise (Sakong (2018)). In another earlier work, I used close electoral victory of Black politicians as an instrument for local areas' racial prejudice against Black people, I found that such increase in racial prejudice caused Black people's employment to fall and mortgages to be denied more (**?**). Combining the two results, it is possible that business cycles disproportionately affect racial-minority groups through counter-cyclical racial prejudice. Also see Bayer et al. (2016); Bayer et al. (2017).

year between the interquartile range of the wealth distribution (Figure 4c).[9]

I connect the estimated return differentials to the level of wealth inequality using a wealth accumulation equation and a back-of-the-envelope calculation. Simple manipulations of this equation reveal that two key factors largely determine how return differentials translate to wealth inequality above and beyond differences in income. First, even if some households earn lower returns on wealth, their wealth share does not vanish because labor income replenishes wealth; hence, the labor-income-to-wealth share modulates the impact of differential returns on wealth. Second, this stabilizing effect of labor income is itself softened by expenditures out of current income; hence, the consumption-expenditure-to-current-income ratio matters. Using estimates of these two quantities from widely used survey data, I calculate that the estimated 60-basis-point return differential explains roughly 20% of the observed wealth inequality between the interquartile range in the US above the part attributable to income inequality.

Beyond explaining part of wealth inequality in the aggregate, the "buy-high-sell-low" channel has a cross-sectional prediction: Geographies with larger time-variation in expected returns in the housing market should have greater wealth inequality, over and above income inequality. This is because in those areas, even the same beta-differences will generate a greater dispersion in wealth returns between rich and poor households. And the greater dispersion in wealth returns persists in the geographical area, because households typically own housing assets near where they live even for investment homes and because families are reluctant to move once settled. I test and confirm this cross-sectional implication of the channel.

To test this cross-sectional implication, I first sort US counties by historical business-cycle cyclicality, which itself predicts how much expected housing returns would vary (Cochrane (2011)). Using a new set of imputed inequality measures and controlling for labor income inequality, I indeed find that current wealth inequality is greater in those areas with higher historical cyclicality.

Executing this cross-sectional test faces additional data issues: Data on wealth are rare, and there are no existing measures for wealth-inequality levels across US geographies. I impute wealth-inequality levels by metropolitan areas, by combining multiple administrative data sources on the assets and debts held by US zip codes.[10] The imputed, between-zip-code wealth-inequality measures correlate strongly

[9] By comparison, Fagereng et al. (2016) look at total returns on financial wealth in Norway and find about a 1% return differential between the interquartile range per year.

[10] Mian et al. (2013); Saez and Zucman (2016)

with historical cyclicality (Figure 6a), consistent with the "buy high and sell low" mechanism of this paper.

The rest of the paper proceeds as follows. Section 2 presents the theoretical framework connecting housing transactions to wealth inequality. Section 3 presents the data, empirical methodology, and estimation results. Section 4 presents the geographical cross-sectional results. Section 5 concludes. The appendix contains theoretical derivations, additional empirical tests and data details.

1.1 Contribution to the Literature

This paper sits at the intersection between literatures on wealth inequality, household portfolio choice, and business cycles. Relative to the wealth-inequality literature, I study a dynamic mechanism and micro-found the return heterogeneity with micro-data evidence. Relative to the portfolio-choice literature, I focus on the rebalancing. This focus is more important in highly incomplete market settings, where a full set of Arrow-Debreu securities cannot be used to fully replicate all dynamic trading.

The literature on wealth inequality decomposes contributions to wealth inequality into three categories: income inequality, differences in savings rates, and differences in returns generated on wealth portfolios. This paper is in third category. Even within the literature on the heterogeneity in returns on wealth, there are four broad subcategories: (1) By far the largest literature examines differences in the average risky-asset-market participation, mainly in housing, stocks and pensions;[11] (2) A largely structural literature uses heterogeneous-agent frameworks to quantify how much return differential there must be in order to generate the observed wealth inequality, given observed data on income inequality and savings rates, often without taking a stance on what would generate such return differential[12][13]; (3) Another literature targets the upper tail of the wealth distribution by exploring a particular mechanism of entrepreneurial return; and (4) A recent empirical literature uses rich micro-data on wealth holdings to calculate the actual returns earned on wealth by households across wealth levels, thus far exclusively with financial assets.[14] My paper differs in a few dimensions: (1) I estimate

[11] Most theoretical mechanisms that have been proposed to explain this phenomenon are static in nature: Low-wealth households do not participate in high-return activity (e.g., Campbell et al. (2018); Boulware and Kuttner (2020) for racial gap).

[12] Wealth distribution is wider than income distribution (Bisin and Benhabib (2017)). To explain why, the literature on wealth inequality finds an important role for return heterogeneity and return loading on wealth (Benhabib et al. (2015a); Gabaix et al. (2016); De Nardi and Fella (2017)).

[13] Theoretical papers have used the heterogeneous-agent model to understand wealth inequality (Bisin and Benhabib (2017); De Nardi and Fella (2017); Guvenen (2011); Heathcote et al. (2009); Quadrini and Ríos-Rull (2015)).

[14] Recent empirical work is starting to find evidence that wealth returns are increasing in income and wealth (Bach et al.

returns to the housing portfolio; (2) I take a stance on what generates the differential expected returns, and impute the expected-return differential from trading behavior; and (3) The mechanism I propose is dynamic in that it deals with the timing of trades.

The literature on household portfolio choice has a huge sub-literature on housing-market participation. Relative to that literature, my paper focuses on the timing of trades, or "portfolio rebalancing" or "active changes" to use language from the closest paper, Calvet et al. (2009). What enables me to study this under-explored dynamic mechanism is the construction of a panel on housing ownership by wealth levels.

The literature on business cycles and asset pricing has begun to incorporate the interaction between time series volatility and the cross-sectional distribution. More of the literature has rightly focused on how cross-sectional inequality affects the behavior of aggregate quantities such as asset prices or gross domestic product (GDP). By contrast, I focus on how time series volatility can affect the cross-sectional distribution, highlighting a potential two-way feedback loop.

Methodologically, this paper contributes to the greater housing literature in general. Being able to identify ownership of housing by poorer and richer households over the annual frequency can be useful for questions beyond the impact on housing portfolio returns and wealth: For example, in a work in progress, I study the cyclical behavior of residential segregation by income and race, using the panel data I construct in this paper.

2 What to Measure and What It Means

I first begin by spelling out the behavior that corresponds to consistently "buying high and selling low" and that will be estimated in the data. A contrast with an alternative measurement may be most helpful: For some sample period, I can observe ownership spells for which I observe both the purchase and the sale; with these I can estimate a realized return on that set of trades. This realized return differs from the average-return differential that leads to wealth inequality for three reasons.

First, conceptually I use variations in expected returns as opposed to realized returns. For example, from 2007 to 2008, average US housing stock experienced a realized return of -8% (or -20% relative to the time series average). But standing in 2007, based on just the rent-to-price ratio at that point and

(2016) in Sweden; Fagereng et al. (2016) in Norway; Garbinti et al. (2017) in France; Kuhn and Rios-Rull (2016) and Wolff (2017) in the US).

what we knew from the return predictability literature, the expected return was 8% (or -4% relative to the average). Because I study persistent "buy high, sell low" behaviors, I use the -4% expected return as opposed to the -20% realized return. That is, given any finite time period, there will always be people who buy and sell at the ex-post "wrong" time; but over the long run, differences in the unpredictable realized returns would cancel each other out. Given a long enough sample, the realized returns will average out to an average-return difference. In this paper I have at most 25 years of data, so I estimate changes in expected returns directly.

Second and related to the first reason, I estimate how housing-ownership levels in quantity change with house-price changes as well as how expected returns change with house-price changes, and then multiply them. This is because I can compute expected returns only at some market level. More specifically, I impute expected returns as a linear function of the rent-to-price ratio, which can only be observed at some local level. Therefore, the formulas derived in this section specify how to convert the observed housing quantities and prices to units of portfolio returns.

Third, I compute expected returns on the whole stock of housing owned, and not just on those that are traded. The alternative measurement using actual trades would omit households who just hold onto the housing and thus generate a 0% realized return. The return differential conditional on trading would be higher, but for overall wealth inequality, I would have to average that with households that do not trade.

For these reasons, I derive formulas that convert the sensitivity of housing holdings in quantity to some market-level house prices to wealth returns, and ultimately to the contribution to wealth inequality levels, in the first sub-section. The second sub-section highlights that in theory, richer or poorer households could exhibit higher sensitivity of housing holdings in quantity to changes in house prices. The last sub-section highlights how the conceptual object maps to the data, and what complications arise when using data that do not meet the requirements.

2.1 How Would the Timing of Ownership Affect Expected Returns

I take a top-down approach, to start from a wealth-accumulation equation, to zero in on the term that needs to be estimated in the data. The following assumption specifies how observed wealth is accumulated.

Assumption 1 (Measured-wealth accumulation)**.** *Household i's measured-wealth accumulation is*

given by

$$\frac{dW_{it}}{W_{it}} = \left(\frac{Y_{it}}{W_{it}} - \frac{C_{it}}{W_{it}}\right) dt + \underbrace{\sum_k \theta_{it}^k dR_t^k}_{\equiv \overline{dR}_{it}}$$

for measured wealth W_{it} of household i in time t. Asset classes are denoted by k with investment share θ_{it}^k and return dR_t^k, labor income flow Y_{it} and expenditure C_{it} inclusive of both paid rents and the imputed user cost of housing (i.e. the opportunity cost for owner-occupants). In particular, assets k are defined such that all households get the same return, namely, $dR_{it}^k = dR_t^k \forall i$.

First note that the income variable Y_t here is only labor income, because the cash flow component of capital incomes is included in the total returns dR_t^k. Non-durable expenditure C_{it} is defined to include the user cost of housing for owner-occupants, as with rent paid for renters. For owner-occupants, the hypothetical rents (i.e., cash flow) from dR_t^k of housing and the hypothetical user cost included in C_{it} cancel each other out. The wealth W_t is all non-human wealth, inclusive of financial assets and real assets (i.e., the wealth concept in Piketty (2015)). It is important to note that W_{it} is the measured wealth; the equation above is an accounting equation.

The assumption is straightforward for investment assets and housing owned as investment homes. Further explanations are necessary for owner-occupied housing, for which economists have explored whether housing wealth is actual wealth. Changes in the current price of housing for owner-occupants appear on both the income side (i.e., in dR_t^k for housing) and the expenditure side (i.e., the user-cost-of-housing component of C_{it}). There is no net effect on observed wealth. Changes in the future price of housing that lead to changes in house prices show up only on the income side, in the form of house-price appreciation in dR_t^k. In economic terms, the household is not richer, in that the increase in the house price on the wealth side corresponds to the net present value of higher expected costs of housing going forward.[15] But in accounting terms, the house-price appreciation is nevertheless an increase in measured wealth. In terms of accounting wealth, the equation is true for both owner-investors and owner-occupants.

The accounting wealth W_{it} need not correspond to the economic concept of wealth. In particular, endogenous policy functions C_{it} and θ_{it}^k are functions of economic concepts of wealth. For example, increases in W_{it} driven by an increase in the price of the primary residence for owner-occupants need

[15] Sinai and Souleles (2005)

not lead to an increase in C_{it}. How C_{it} behaves is important when translating an average return to a level of wealth inequality later. For most of the paper, I deal with the accounting concept of wealth. The empirical literature on wealth inequality uses the accounting, observed measure of wealth as well.

The key assumption here is that the assets k are defined such that all households earn the same return on them (more precisely, they have to have the same expected return). Surveys often bundle together assets (e.g., housing), but heterogeneous asset returns are likely, especially in an asset class such as housing. How violation of this assumption can affect estimated average returns is discussed in more detail below.

Lemma 2 (Return decomposition)**.** *The average return on wealth can be decomposed:*

$$E\left[\overline{dR}_{it}-\overline{dR}_t\right]=\sum_k\left\{E\left(\theta_{it}^k\right)-E\left(\theta_t^k\right)\right\}E\left(dR_t^k\right)$$
$$+\sum_k\left\{\underbrace{cov\left(\theta_{it}^k,E_t dR_t^k\right)-cov\left(\theta_t^k,E_t dR_t^k\right)}_{"timing"}\right\}$$

Note that the covariance is with respect to the expected return, $E_t dR_t^k$. Importantly, the second term is non-zero only if expected returns are time-varying.

The lemma follows immediately from the fact that the expectation of the product is the sum of the product of expectations and the covariance, that is, $E\left[\theta_t E_t\left(dR_t\right)\right]=cov\left(\theta_t,E_t\left(dR_t\right)\right)+E\left(\theta_t\right)E\left(E_t\left(dR_t\right)\right)$ (i.e., first take the conditional expectation, and then take the unconditional one).

The first term on the right-hand side (the product of expectations) is the focus of an enormous literature in economics, finance, and sociology on risky asset market participation. These papers typically multiply an average differential in participation $E\left(\theta_{it}^k\right)-E\left(\theta_t^k\right)$ by the expected return $E\left(dR_t^k\right)$ to obtain the the return differential.

The second term (the covariance) is the focus of this paper. If conditional asset returns were not time-varying, that is, if $E_t dR_t^k=\overline{\mu}^k$, the covariance terms would drop out. With time-varying expected returns, agents who increase exposure when expected returns are high accumulate wealth faster.

With time-varying expected returns, the magnitude of the variations in the expected return partly determines the contribution of the timing of trades on average returns. While realized returns clearly vary over time, whether expected returns vary over time and in predictable ways has been one of

the central questions in the asset pricing literature. In particular, Cochrane (2011) estimates expected returns using the rent-price ratio by regressing the realized return,

$$\log \mu_t^k \equiv \log E_t \left(R_{t+1}^k \right) = a^k + b^k \log \frac{D_t^k}{P_t^k}$$

An extensive literature documents return predictability in the housing market, with different papers using different house-price indices.[16]

The contribution of the trade timing to returns is in terms of wealth shares θ_{it}^k, which are themselves products of prices and quantities in units. The co-movement with expected returns consists of two parts: passive and active change, mirroring the decomposition of risky portfolio share in Calvet et al. (2009).[17] In particular, in response to price changes in asset k, the asset share θ_{it}^k moves in the same direction in the absence of active adjustment, and the expected return $E_t dR_t^k$ likely moves in the opposite direction. The active change can offset or amplify the change in asset share θ_{it}^k. I work with formulas linking each piece to price changes $d \log P_t^k$ to derive an easily computable formula.

Lemma 3 (Covariance approximation)**.** *Given that expected returns are a function of the rent-to-price ratio, co-movement between kth asset share and the expected return can be decomposed as*

$$cov\left(\theta_{it}^k, \mu_t^k\right) \approx E\left(\theta_{it}^k\right) E\left(\mu_t^k\right) \left[\underbrace{cov\left(\log P_t^k - \log W_{it}, \log \mu_t^k\right)}_{passive} + \underbrace{cov\left(\log Q_{it}^k, \log \mu_t^k\right)}_{active} \right]$$

See Appendix for the derivation.

The formula shows that passive change induces a negative relationship between θ and μ. However, with a big enough contrarian active change, the covariance can even be positive.

Proposition 4 (Return differential from active trades)**.** *The timing of active changes widens wealth inequality if and only if*

$$acov\left(\theta_{it}^k, \mu_t^k\right) \approx -b^k E\left(\mu_t^k\right) var\left(\log P_t^k\right) E\left(\theta_{it}^k\right) \frac{cov\left(\log Q_{it}^k, \log P_t^k\right)}{var\left(\log P_t^k\right)}$$

[16] Relatively fewer papers study return predictability in local housing markets with proper returns. The key data limitation is measuring local rent levels at a frequency higher than the decade frequency, for the cash-flow component. In more recent years, both the Department of Housing and Urban Development and Zillow have begun to release local rent indices. Several papers study return predictability with log house price instead (e.g., see papers in the survey in Ghysels et al. (2013)). Other papers that document return predictability in the housing market are Piazzesi and Schneider (2016) (using Zillow Research for local house prices), Glaeser and Nathanson (2015) (using Federal Housing Finance Agency data for the price index and Department of Housing and Urban Development data for rent), Campbell et al. (2009), Davis and Van Nieuwerburgh (2015), and Cochrane (2011) for aggregate return.

[17] Table 4 in Calvet et al. (2009) discusses heterogeneity by household financial characteristics.

is increasing in wealth level and $b^k \neq 0$ (return is predictable in asset k).

See Appendix for the derivation.

The rest of the paper estimates the empirical elasticity, $\frac{cov\left(\log Q_{it}^k, \log P_t^k\right)}{var\left(\log P_t^k\right)}$.

2.2 Theoretical Ambiguity on Who Would Own When

Before addressing the empirical challenges in estimating what type of households would hold more housing when prices are high, I first discuss how even standard theories predict that it could be either poor or rich households. The theoretical ambiguity arising from commonly accepted forces further justifies dealing with the empirical challenges. In this sub-section, I discuss the possibilities in words; for a more rigorous discussion based on a standard consumption-savings problem with portfolio choices, refer to the Appendix.

On one hand, suppose cyclical booms are accompanied by independent increases in the supply of household credit. Such pro-cyclical expansion of credit may be driven by market forces or by government policies that assist homeownership especially in boom years when budgets are plentiful.[18] Increased access to mortgages and other household credit would disproportionately affect poorer households who were likely to be closer to borrowing or collateral constraints. Therefore, pro-cyclical credit supply would lead to poorer households buying and owning more housing in booms than in busts, relative to richer households.

On the other hand, suppose cyclical booms are accompanied by widespread expectations of higher asset returns, for example, because households mistakenly believe high past asset returns will continue on.[19] Again, if poorer households are closer to borrowing or collateral constraints, richer households are better positioned to change portfolio holdings to take advantage of the perceived higher expected returns in booms. Therefore, extrapolative expected returns in the housing market would lead to richer households buying and owning more housing in booms than in busts, relative to poorer households.[20]

Even widely studied theories on what happens over business cycles have opposing predictions for what types of households would own assets more in booms versus busts, and it is possible that both forces are simultaneously at play. Therefore, who accumulates wealth faster through the timing of

[18] Rajan (2011)

[19] Barberis et al. (2015)

[20] This implication can be seen in Kaplan et al. (2017), where a positive shock to expected house-price appreciation leads to a lower homeownership rate.

trade is an open, empirical question.

2.3 Empirical Challenge in Estimating the Timing of Ownership

Whether rich or poor households trade against expected returns is an empirical question that cannot be conclusively answered with any accessible, off-the-shelf dataset. This section briefly discusses what issues arise if I attempt the estimation in commonly used datasets.

Estimating the timing-of-ownership, $cov\left(\theta_{it}^{k}, E_t dR_t^k\right)$, is difficult in the Survey of Consumer Finances (SCF), for example, where we observe only the product of price and quantity, because of the assumption, $dR_{it}^k = dR_t^k \forall i$. That is, $\overline{dR}_i = \sum_k \theta_i^k dR^k$ is true only if returns are common across individuals. Data at such a level of disaggregation are rare: Wealth-tax data (e.g., Calvet et al. (2009)) and housing microdata are two examples.

With coarser asset class K (suppressing time subscripts where obvious),

$$\overline{dR}_i = \sum_K \theta_i^K dR_i^K$$

where dR^{iK} retains the i superscript because the asset composition within the asset class K affects the expected returns. The two terms can be decomposed into:

$$\theta_i^K = \sum_{k \in K} \theta_i^k$$
$$dR_i^K = \sum_{k \in K} \frac{\theta_i^k}{\theta_i^K} dR^k$$

θ_i^K can be calculated in the SCF, with only wealth share, but dR_i^K cannot be. Effectively, what I do when I assume dR^K for average return to asset class is ignore the variation in $\left(dR_i^K - dR^K\right)$. This omission would compress the dispersion of return heterogeneity, in the same sense that assuming common $\overline{dR}$ instead of $\overline{dR}_i$ does. Using only the coarse asset classes, the actual individual-specific expected return terms are

$$E\left[\theta_{it}^K dR_{it}^K\right] = E\left[\theta_{it}^K\right] E\left[dR_{it}^K\right] + cov\left(\theta_{it}^K, dR^K\right) + cov\left(\theta_{it}^K, dR_{it}^K - dR_t^K\right)$$

Both the return-on-asset-class-K earned by household i, $E\left[dR_{it}^K\right]$, and the co-movement of that household-i-specific return and share, $cov\left(\theta_{it}^K, dR_{it}^K - dR_t^K\right)$, cannot be estimated from the SCF.

If we use average return instead of individual-specific returns for coarse asset classes (let dR^K denote the average return on asset class K), the last term, $cov\left(\theta_{it}^K, dR_{it}^K - dR_t^K\right)$, can differ due to skill

or different portfolio risk profiles. Note from above that passive change makes $cov(\theta, dR)$ negative mechanically. The discrepancy $cov\left(\theta_{it}^K, dR_{it}^K - dR_t^K\right)$ will be more negative for poor households because of passive change, if they live in areas with more volatile house prices.

Most importantly, asset pricing tells us that if covariance is driven by price dynamics, a tight negative link exists between $cov\left(\theta_{it}^K, dR^K\right)$ (as well as $cov\left(\theta_{it}^K, dR_{it}^K - dR_t^K\right)$) (i.e., household i takes on more risk quantity) and $E\left[dR_{it}^K\right]$ (i.e., higher risk premium for taking on higher risky).

In the housing micro-data from deed records in which I get to observe the quantity, I can bypass the issue discussed in this sub-section by dealing with quantity changes directly and multiplying them using changes in expected returns. Furthermore, by going from aggregate housing returns to location-specific returns, I get closer from K to k.

For a discussion of why this measurement is not possible in existing datasets such as survey data and mortgage origination data, see the Appendix.

3 Estimating the Timing of Ownership by Wealth Levels

In order to find out who holds more housing in booms versus busts and to estimate the consequent wealth-return gradient against wealth levels, in this section I construct the dataset, plot raw-data patterns, estimate the housing-quantity-to-price betas by wealth proxies and finally convert those betas into wealth returns.

3.1 Compiling the Dataset

For the main dataset, I merge a panel of housing ownership (by property and year) to the owners' wealth levels, using surnames of the owners. Constructing the dataset is a non-trivial exercise, but it allows me to observe unit-quantity-holdings of housing by wealth levels, which is essential to estimate the "buy high, sell low" channel. On the asset side, the housing data from CoreLogic (more details below) are comprehensive, disaggregated and reliable. What was missing had been the identity of the owners, buyers and sellers, for whom we see the names and their mailing addresses, but do not have readily usable covariates such as income, wealth or race. The essence of the dataset-construction is in extracting such covariates from the names and the mailing addresses.

The final dataset is dominated by survey data on the precision of the owner characteristics, but is

superior on the asset-side details. For studying dynamic portfolio choice, for which one major challenge is in dealing with assets of different price dynamics, this trade-off is useful.[21]

After describing the CoreLogic data and samples in the first subsection, I explain in detail how I infer household characteristics from names. I then describe the data sources with demographic information on names, followed by validation exercises.

3.1.1 CoreLogic Data and Samples

I first describe the CoreLogic data on housing and how I select the main samples for the analyses in this paper. CoreLogic is a private data-provider that acquired DataQuick, which compiles public records on housing assessments and transaction deeds from various jurisdictions into a unified data set. I use two components of the CoreLogic (formerly DataQuick) data: The assessor file and the transaction deed records. I describe each component in turn.

The assessor file collects a single cross section in 2012-2013 of property assessment, for assessing the amount of property tax. Because the purpose of the assessment is to assign a value to the property, there are a lot of details on the property (e.g., number of bedrooms, total square footage, number of floors, whether it has a view) and its value (e.g., assessment value, the value it would get if sold on the market, the latest actual transaction value). Each property is also associated with the owner's name, the type of owner (individual or institutional), whether it is a primary residence, and the mailing address of the owner. Other versions of the CoreLogic data contain multiple years of cross sections, but I have access to only a single year. The assessor file contains roughly 104 million records, from jurisdictions covering roughly 94% of the US population.

The transaction-deed records are the official records that are signed and mailed when some transaction takes place involving a real estate property. Transactions that could lead to a deed record are sales of an existing property, sales of newly constructed property, mortgage originations, etc.[22] I focus on deeds that result from an ownership transfer, whether of an existing or a new property. For each transaction, the deed data contain the date of transfer, the value of transaction, and the names of both buyers and sellers, among a few other details.

These CoreLogic data come from multiple jurisdictions. Jurisdictions are based on counties,

[21] This dataset is also useful in questions that require more detail on the asset side. Another example is residential segregation. In an ongoing project, I use the same dataset to study the high-frequency dynamics of residential segregation by wealth and race.

[22] Recording a deed is not required in every place, but it is almost always done.

and most jurisdictions are unique within a county.[23] In particular for the transaction-deed records, jurisdictions are added to the CoreLogic database over time. I know when a jurisdiction enters the database, after which I have the full set of transactions that took place with the properties in that jurisdiction.

I want to construct a balanced sample of properties to track who owns them over time. Over time, more jurisdictions are added to the CoreLogic database. At the same time, I do not want the results to be driven by the selection of jurisdictions over time; I want to form samples of jurisdictions for which I can construct consistent panel. There is a trade-off: For a longer time series, I am forced to use fewer jurisdictions, whereas to use more jurisdictions, I am forced to use a shorter time series. This trade-off is shown in Figure A.1c: For each year in the x-axis, the figure plots the number of counties that would be included consistently between that year and 2013 (in blue, dashed line; left y-axis) and the share of total US population covered in that sample (in hollow circles and red, solid line; right y-axis). Figure A.1c shows that many jurisdictions were added in 1996-1998 and then in 2004.

To maximize the amount of data used, I pick two samples. The first sample spans 1998-2013, and is selected to cover the largest fraction of the population as possible, while giving a full picture of at least one boom-bust episode. The second sample spans 1988-2013, and is selected to retain the longest time period. The counties included in each sample are graphed in Figures A.1a (for the 1998-2013 sample) and A.1b (for the 1988-2013 sample). The first 1998-2013 sample spans 36 states and more than 60% of the US population; the second 1988-2013 sample spans 11 states and 21% of the population (Figure A.1c).[24] Even the broader, 1998-2013 sample contains counties that are more likely to be urban and are not representative of the US as a whole (Figure A.1d shows that the house price boom-bust was larger).

For each sample, I construct a balanced sample of ownership, at the property and year level. For each sample, I start from the annual cross section of ownership in the 2012-2013 assessor file. Then, I work backwards and change owners when there is a transaction, using the transaction-deeds data. For some properties, there are multiple ownership changes within a year. In constructing the annual panel, I keep the owner on December 31 of each year.[25]

I omit properties that do not exist in 2012-2013. Properties that were constructed in the middle

[23]The exceptions occur in six states: CT (21), MA (25), ME (37), NH (19), RI (8), and VT (18), with the average number of jurisdictions per county in parentheses.

[24]The second sample contains 148 counties (674 jurisdictions) in AZ, CA, CT, MA, NC, NJ, NV, OR, RI, TN, and WA.

[25]Higher-frequency panels can also be constructed (up to the daily frequency), and may be more useful for studying cycles.

of the sample period appear throughout, but are not assigned owners until they are constructed and transferred to individuals. How I deal with new constructions is important, especially given that the cyclicality of constructions is stronger in some areas than others. This issue is intimately tied to the empirical specification, and will be discussed in more detail.

In the main analysis, I only use the information that can be extracted from surnames. There are two additional sources of information. First, for investment homes, where the owners live is informative about how rich they may be, and residential address can be inferred from the mailing address (recorded to receive the signed deed). Second, rarer full names can be matched to individuals for whom public data are available. In particular, I have lists of financial brokers, medical doctors, corporate executives, hedge-fund managers and politicians, by name. Individuals in these high-paying occupations are likely richer than the population average. I use both of these auxiliary approaches in the online appendix.

3.1.2 Information on Surnames

In the main analysis of this paper, I use surnames in the housing records to infer the owners' wealth levels. The information on the surnames comes from two sources.

The first is the 1940 full-count Census, as processed in my earlier work with a co-author (Henry de Frahan and Sakong (2020)). The 1940 Census is the latest full Census that has been released to the public following the 72-Year Rule.[26] It is also the first Census that explicitly asked for households' income.[27] In addition to the full names and household income, the full-count Census also contain variables on rents paid for renters and value of homes for homeowners, years of education, race, and where they live.

The main variable used to assign wealth levels to surnames is the household-level wage income. To be precise, we average the wage incomes of households whose head has the surname Smith, and assign that as a proxy for the average wealth of all Smiths in today's data. In our earlier work, we document that this variable is a strong predictor of various proxies for wealth in today's data. For example, it strongly predicts the average primary-residence value of Smiths who are homeowners in the CoreLogic assessor data for 2012-2013 (Figure 1a and also in Henry de Frahan and Sakong

[26] "This '72-Year Rule' 92 Stat. 915; Public Law 95-416; October 5, 1978) restricts access to decennial census records to all but the individual named on the record or their legal heir." https://www.census.gov/history/www/genealogy/decennial_census_records/the_72_year_rule_1.html?CID=CBSM+history

[27] Previous waves of the Census do contain a variable called "occupational prestige score," which basically assigned the average income of the occupation to households by the household head's occupation.

(2020)). The historical wage income is predictive of today's wealth through both the human capital (i.e., grandsons of high-income grandfathers are more likely to earn higher income today) and the non-human capital (i.e., the higher-income grandfathers left more wealth for their descendants).

The raw estimations in this paper are at the surname-level (i.e., Mackenzies consistently behave differently from the Smalls in the past 30 years). Interpreting the surname-level estimation for family-level relationships requires additional assumptions. The full econometric framework to interpret the surname-level estimation is discussed in our earlier work (Henry de Frahan and Sakong (2020)).

The second source of information on last names is the recent Census tabulations of surnames in 2000 and 2010. Recent waves of the Decennial Census include tabulations of surnames, for which there are 100 or more individuals with that surname, along with those surnames' composition by major racial groups (Word et al. (2008)). In 2000, the criteria of having 100 or more people left 151,671 surnames and 242 million people covered by the surname data, relative to the total population of 282 million, implying an 85.8% coverage. In 2010, the same criteria left 162,254 surnames and 295 million people, implying a 95.6% coverage. Two sets of variables are used from these data: the counts by each surname and the shares that of the major racial/ethnic groups (Asian, Black, Hispanic and White).

The surname-level population counts are used for two purposes: (1) They are used as denominators in computing the per-capita-housing holdings (i.e., I divide the number of properties held by Smiths in the CoreLogic ownership panel by the total number of Smiths in the US from the Census); (2) They are used to weigh the surname-level data. In the main analysis, I only use the surnames for which I can observe the total counts in both 2000 and 2010. For the years other than 2000 and 2010, I linearly interpolate and extrapolate in logs to get the population count (i.e., I assume a constant population-growth rate by surname).

3.1.3 The Constructed Dataset

Here I briefly describe the structure of the constructed dataset. The structure is common across the 1998-2013 and the 1988-2013 samples.

Each balanced panel has observations at the property and year levels. Each property-year is associated with a set of surnames, one for each owner in that year (often the property-years have just one owner). For property-years with multiple owners, still the weights for that property-year add up to one. For each surname, there is an associated surname-level average-household-wage income from

the 1940 Census and the corresponding percentile value (I will refer to this variable as the "1940 income percentile"). For property-years before the construction of the property, there is no associated owner. Later, I collapse such a panel to "1940 income percentile" and year levels, to facilitate visual examination and estimation.

Each surname is also associated with a total population count for the given year as well as racial shares for major racial groups.

The main quantity unit is the number of properties, that is, each observation in the property-year-level panel has the same weight. In robustness analysis, I use two other quantity units: number of bedrooms and square-footage. These quantity measures are available in the 2012-2013 assessor file only. Therefore, each property is assigned the same number of bedroom sand the same square-footage throughout the years.

3.1.4 Validation against Census Data

My methodological contribution is to use information on surnames to attribute wealth levels to owners on housing records. While intuitive, the newly constructed dataset needs to be validated. Here, I provide one additional validation against 2000 Census. More validation exercises can be found in my earlier paper with a co-author (Henry de Frahan and Sakong (2020)).

Here is the main idea: I take the 2000 cross section from the constructed panel for the 1998-2013 sample, covering roughly 60% of the US population. For each zip code in the sample, I compute the average 1940 income across all owners in that zip code, separately for owner-occupied housing and investment housing. Note again that the average 1940 income was assigned to each owner in the property-level housing data using the surnames. I take the zip-code-level averages of average 1940 income computed solely from surnames, and compare them to zip-code-level income measures from the 2000 Census. Note that while I use the historical incomes from 1940 as a proxy for today's wealth, I use the zip-code-level income from the 2000 Census, because there is no comprehensive wealth measure even in the Census.

I run three sets of validation regressions, with the results reported in Table 1.

The first set of tests are for averaged incomes on the CoreLogic side for owner-occupants only (the first two columns of Table 1). I verify that the average 1940 household-wage income of homeowners in a zip code (from CoreLogic) is correlated with the median household income from the Census. Column

(2) adds county fixed effects and the correlation is still strong.

The second set of tests are for averaged incomes on the CoreLogic side for investment homes only (columns (3) and (4) of Table 1). For investment housing, there is a zip code associated with the location of the property and a zip code associated with where the owner lives (the zip codes can be the same too). I regress the Census 2000 income of the zip codes where the owners live on two variables: the averaged 1940 income from surnames (from the CoreLogic data) and the Census 2000 income of the zip codes where the properties are located. Column (4) adds county fixed effects. Here is an illustration of what I verify: If a Mackenzie from a high-income neighborhood own a property in a low-income neighborhood, I expect the wealth level attributed to Mackenzies to be associated with the high income level of the where the Mackenzie lives.

The last sets are for all housing on the CoreLogic side (columns (5) through (7) of Table 1). The right-hand side variables in these regressions are the two averaged incomes from the CoreLogic panel, one for owner-occupied properties and another for investment properties. The prediction is that the zip codes' incomes from the 2000 Census should be more informative about the permanent income of the owners who live in the area more than those that do not (i.e., investment-owners). The left-hand side variables are from the 2000 Census: zip-code-level income for columns (5) and (6), and zip-code-level median house price for column (7). Columns (5) and (6) show the area's household income is more highly correlated with the average 1940 income of owner-occupants, because those owners are the ones whose incomes are reported to the Census. For home values in column (7), the two types of owners have similar magnitudes of correlation.

Across the three sets of validation, the correlations are as expected and strong, adding confidence to the use of surnames and their associated historical income from the 1940 Census in assigning wealth or permanent income proxies to the housing owners in the CoreLogic data.

3.2 Raw-Housing-Ownership Patterns by Wealth Proxy

In this sub-section, I show raw-data patterns by transforming the CoreLogic housing panel constructed above, one step at a time. This sub-section serves two purposes: (1) show the differences in housing ownership between rich and poor households are evident even from raw plots, and (2) show exactly what variations in the data are used for the estimation of "betas" in the next sub-section (i.e., the elasticity of housing quantity to house price, by wealth levels).

Starting with the CoreLogic panel at the property and year levels for the 1998-2013 sample covering a larger set of counties, I sum over the "1940 income percentile" assigned using the owners' surnames, to create an annual time series of total number of housing properties owned, for each of the 100 "1940 income percentile" groups.[28] For each percentile group, I divide the total number of properties owned by owners with surnames in that percentile group, by the total number of individuals in that percentile group, to compute the per-capita ownership by number of properties. I will refer to this per-capita number of housing properties owned as q_{it} for percentile group i in year t. For alternative quantity measures using number of bedrooms or square footage, see the online appendix.[29]

Raw time series: Figure 1b plots the average q_{it} for selected decile groups (i.e. the second decile includes percentiles 11-20), divided by the corresponding level in 1998.[30] Higher deciles correspond to surnames with higher historical incomes in the 1940 Census. In the plot, all holdings are increasing over time because of new constructions.[31] Beyond the broad-based increase in holdings, the rate of increase is higher for the lower decile groups during the boom years up to 2007 (marked with the red vertical line), then either decreases or plateaus afterward.

That poorer households bought more pre-2007 and less afterwards is evident from the raw patterns. To calculate the differences in elasticity of quantity to price, however, I need to transform the data further.

Cyclical variation in housing ownership: Figure 1c plots the residuals ε_{it}^{q} from

$$\log q_{it} = \alpha_i + \alpha_t + \gamma_i t + \varepsilon_{it}^{q}$$

for "1940 income percentile" group i in year t. I describe each transformation in turn. The percentile-group-fixed effect α_i allows us to focus on changes rather than differential levels across groups.[32]

The year-fixed-effect α_t allows us to exclude variations driven by new constructions; for the purposes of understanding return *differences* between wealth groups, what matters is the differences in when they own. Where new constructions occur is an important issue in and of itself, but it is

[28] For the raw patterns for the longer 1988-2013 sample for fewer counties, see the online appendix. The actual estimation-results are included in the next sub-section for both samples.

[29] Both raw-data patterns and estimation results are qualitatively the same when I use either the number of bedrooms per capita or square footage per capita, instead of the number of properties per capita.

[30] I plot the relative quantity, because higher-percentile groups have higher level of q_{it} throughout the sample period. See the online appendix for $q_{i,1998}$. It looks just like the average-primary-residence values by percentile groups in Figure 1a.

[31] Also, given that the set of properties in my sample are those in the 2012-2013 assessor file, I omit properties that had been in place but were not owned by anyone by 2012-2013. This omission would lead to an over-estimation of the growth rate in the number of properties held by all groups.

[32] It is similar to what I did when I plotted the raw time series relative to some base year.

orthogonal to observing who owns more when. First, even if new houses are more likely to be built in poorer neighborhoods in booms, who owns those houses may be the new owner-occupants of that area who switch from renting or rich owners who buy the new houses and rent them out. Second, even if the poorer residents of those neighborhoods own the newly constructed units in booms, for portfolio returns, it is true that poorer households are acquiring risky assets when expected returns are lower. Because the outcome of interest in this paper is return *differential* and wealth inequality, *how* households acquire housing is not central (although that question is extremely important in and of itself).

The 1940-income-percentile-group-specific linear trend γ_i removes any long-term trends that differ between the percentile groups. The goal is to remove the effects of long-term changes in population, inequality and homeownership, which do not vary at the same frequency that expected returns vary at. I discuss the rationale in more detail below, after I present the de-trended price time series.

After these transformations, the boom-bust patterns by wealth groups are more evident. Figure 1c shows exactly the variation that will be used in the estimation in the next sub-section.

Cyclical variation in house price: Figure 1d plots the residuals $\tilde{\varepsilon}_t^p$ from

$$\log P_t = \tilde{\gamma} t + \tilde{\varepsilon}_t^p$$

where P_t is the national house-price index from CoreLogic.[33] Because house prices are growing over time, I remove the linear trend in logs, akin to the transformation for quantities.

I remove a log-linear trend from both the housing quantities and the house price series. There are several reasons, but they all spring from having a finite sample period, for which realized returns can differ from expected returns even after averaging. I use house-price level as the proxy for expected returns (Cochrane (2011)). House-price levels are non-stationary, and ideally I could use the rent-to-price ratio, but the observed rent series does not correspond to the house prices the same way dividends correspond to stock prices. In this context, taking out a linear trend is akin to assuming that the rent series increases at a constant rate. This is effectively how Cochrane (2011) deals with the rent series.[34] Similarly for quantities, suppose one group increases its holding of housing at a constant rate due to secular changes in inequality. Over time, that group would earn no extra return due to the timing of

[33] CoreLogic constructs the house-price indices from the transactions micro-data that I use in this paper. The data documentation states, "The CoreLogic HPI measures changes in housing market prices from 1976 to present. The HPI is a repeat sales, value weighted, econometric Home Price Index Model. Base year is 2000 set at 100."

[34] He takes the rent value every ten years and interpolates for the years in-between.

trades. A log-linear trend removes this variation from the housing-quantity series. With the shorter 1998-2013 sample, using a linear trend to remove long-term trend may be problematic, for the precise estimation. The 1988-2013 sample is longer and this problem is less egregious, but the short time-series dimension of the panel is an over-arching issue in this paper.

The estimated "betas" in the next sub-section are basically the ordinary-least-squares (OLS) regression-coefficients of ε_{it}^{q} (the quantity residual) on $\tilde{\varepsilon}_{it}^{p}$ (the price residual).

3.3 Estimation of Betas

The raw plots show quantity time series for poor and rich households differ. The estimation turns the transparent, visual relationship into one number that summarizes the co-movement between quantity and price, by wealth levels. I basically regress the quantity-residuals plotted in Figure 1c on the price-residuals plotted in Figure 1d, percentile-group by percentile-group.

For the shorter but wider sample covering 1998-2013, Figure 2a plots the estimated betas (that is, the housing-quantity-to-price elasticities) β_i for each percentile group i, from

$$\log q_{it} = \beta_i \log P_t + \alpha_i + \alpha_t + \gamma_i t + \xi_{it}$$

where q_{it} is the per-capita number of housing properties held by "1940 income percentile" group i in year t, and P_t is the national house-price index. This estimation framework uses the same variation as Hoopes et al. (2016).[35] The estimated betas average to zero by construction.

Except for a few percentiles on the extremes, the estimated betas are decreasing in the "1940 income percentile," which proxies for wealth. That is, poorer households hold more housing more pro-cyclically. The pattern is especially pronounced for surnames in the bottom 20% of the 1940 average income distribution: For households with those surnames, a 10% increase in prices is associated with more than 1% higher ownership of housing relative to the population average.

[35] To see that I use the same variation as Hoopes et al. (2016), my specification is $\log q_{it} = \beta_i \log P_t + \alpha_i + \alpha_t + \gamma_i t + \xi_{it}$. I take first difference to obtain, $\Delta \log q_{it} = \beta_i \Delta \log P_t + \Delta \alpha_t + \gamma_i + \Delta \varepsilon_{it}$. Hoopes et al. (2016) run a two-stage estimation. In the first stage, they estimate percentile means of log changes in quantity (i.e., $\Delta \log q_{it}$) after regressing out time fixed effects (this plays the role of $\Delta \alpha_t$); in the second stage, they regress the percentile means estimated in the first stage against a group-specific constant (i.e., playing the role of γ_i) and changes in the aggregate state (i.e., playing the role of $\Delta \log P_t$). The group-specific linear trends, and equivalently group-specific constant in changes in Hoopes et al. (2016), are more justified in the daily-frequency setting of Hoopes et al. (2016). In the annual-frequency setting in this paper, the short sample length makes the distinction between cyclical variation and long-term variation difficult. One solution is to obtain a longer time series, which I plan to do in future research.

The relationship is similar for the longer but narrower sample covering 1988-2013 (Figure 2b). We want to know whether the differences in cyclical ownership are specific to the recent boom-bust episode or true more generally. To that end, I run the same estimation for the 1988-2013 sample, but only for the 1988-2002 sub-period.[36]

Figure 2c shows that the betas are similar for the bottom 20% of the "1940 income percentile" distribution, but for the top 80%, the betas are *increasing*, unlike for the full-period estimation.

These estimation results highlight two issues. First, the results are not robust for the 1988-2002 sub-period, for the top 80%. Second, the stark non-linearity is problematic: Because I have already aggregated up to surnames and then to the corresponding "1940 income percentile" groups, even non-linearity in the individual relationship would largely be smoothed out. The issues suggest that the surname-level aggregation may be picking up variation apart from the differences in their incomes in 1940.

Figure 3a plots the White racial share in 2010 of the individuals with surnames in each "1940 income percentile" group. Surnames are informative about the racial composition (Henry de Frahan and Sakong (2020)) and surnames that are more White included higher-income households in 1940. Most notably, the White-share-to-"1940 income percentile" relationship shows the same non-linearity as in the estimated betas, in the bottom 20%. I explore the variations in betas between-race and within-race in the next sub-section.

3.4 Betas Between- and Within-Race

To see why this is a decomposition,

$$y_i = \sum_{r \in R} \alpha_i^r + \beta x_i + \varepsilon_i$$
$$\bar{y}_\ell = \sum_{r \in R} \gamma^r S_\ell^r + \beta \bar{x}_\ell + \bar{\varepsilon}_\ell$$

The estimated betas (i.e., housing-quantity-to-price elasticity) exhibited a non-linear pattern (Figure 2a) , which had a similar shape as the shape of each "1940 income percentile" group in the White racial group (Figure 3a). Therefore, in this sub-section, I decompose the variation in the estimated betas to the variation between-race and within-race.

[36]The end-point 2002 was chosen arbitrarily, to include the most years before the 2000s housing boom starts.

For a quick summary, I find that the negative beta-wealth relationship is driven by the between-race variation: Surnames that include more non-White individuals in 2010 have households who hold more housing in booms (these are also surnames with lower average income in 1940). Conditional on the non-White share, the relationship between the estimated betas and the "1940 income percentiles" is in fact positive.

Before describing the details and presenting results, there is an important caveat in interpreting the results in this section. The observed race indicator proxies for multiple conceptual factors, including race per se (e.g., racial prejudice of Becker (1971)) as well as wealth and permanent income (e.g., White people may own more wealth than Black people conditional on wage income, even in 1940). Therefore, it is misleading to interpret the results in this sub-section as independent contributions of race and wealth. Rather, I interpret them as variation in the estimated betas between racial groups (e.g., difference between White and non-White population groups) and within the racial groups (e.g., difference conditional on racial-group differences), treating races just as measurement. Moreover, the composite relationship between betas and "1940 income percentile" (Figures 2a and 2b) is still the overall relationship between the betas and wealth levels as predicted by surname-level historical income.

Keeping the caveat in mind, I plot residuals from linear regressions of the 100 estimated betas against the share of each percentile group that is White or against the actual percentile values, residualized by the other variable. That is, Figure 3b plots the residuals $\tilde{\xi}_i$ from $\hat{\beta}_i = \tilde{\delta}_1 \left(\text{1940 income percentile}\right)_i + \tilde{v}_i$ for "1940 income percentile" group i, against the share of group i that is White in 2010, for the shorter-but-broader-1998-2013 sample. Figures 3d and 3f plot the equivalent values for the longer-but-narrower-1988-2013 sample; Figure 3d is for the full period, and Figure 3f is for the sub-period 1988-2002.

By contrast, I plot the residuals from regressing the estimated betas on the White shares, against the "1940 income percentile" in Figure 3c, for the shorter but broader 1998-2013 sample. The equivalent plots for the longer-but-narrower-1988-2013 sample are presented in Figures 3e (for the full period) and 3g (for 1988-2002).

That is, the fitted slopes in the figures on the left (Figures 3b, 3d and 3f) are δ_0, and the slopes in the figures on the right (Figures 3c, 3e and 3g) are δ_1, from

$$\hat{\beta}_i = \delta_0 \left(\text{share White}\right)_i + \delta_1 \left(\text{1940 income percentile}\right)_i + v_i$$

for the "1940 income percentile" groups i in the respective samples.

Put together, these plots show that across the samples and sub-sample, the pro-cyclicality of housing ownership is decreasing in the White share but increasing in the "1940 income percentile" conditional on the racial shares. That is, the overall negative relationship between the estimated betas (the housing-quantity-on-price elasticities) and the wealth proxies is driven by the differences between racial groups.

Consistent with the decomposition, the gap between wealth inequality and income inequality is the widest between racial/ethnic groups (see Figures A.3a and A.3b for the wealth gaps and income gaps from the Survey of Consumer Finances, for Black and Hispanic households respectively). An enormous literature studies the Black-White gap in wealth (e.g., **??**); see the online appendix for a discussion of the racial wealth inequality.

As discussed above, the race indicator can be interpreted as race per se or as another proxy for wealth. Distinguishing the two possibilities is beyond this paper, but I discuss one under-explored possibility by which racial minorities may be disproportionately affected by business cycles. In an earlier work, I found some empirical evidence that cyclical downturns cause racial prejudice of a metropolitan area to rise (Sakong (2018)). In another earlier work, I use close electoral victory of Black politicians in local elections as an instrumental variable that *increase* those areas' racial prejudice, and found that an increase in the local racial prejudice causes Black people to lose more jobs and face more mortgage denials than White people (**?**). Putting these two together, business cycles may affect racial minorities disproportionately, because racial prejudice itself is counter-cyclical. This possibility will be explored more in future research.

For the rest of this paper, as discussed above, the overall relationship between the estimated betas and the wealth levels as proxied for using the "1940 income percentile" is still valid. The quantification exercises in the next two sub-sections use the overall relationship between the estimated betas and the "1940 income percentile" groups. For the corresponding quantification exercises for the between-race and within-race variations, see the online appendix.

3.5 Conversion from Betas to Return Differentials

Thus far, I have estimated how housing-ownership quantity co-varies with the national house-price index differentially by the "1940 income percentile." Both the betas and the wealth proxies do not

correspond to anything that can be interpreted with respect to portfolio returns or wealth inequality. Therefore, in this sub-section, I convert the estimated beta-"1940 income percentile" relationship to a meaningful relationship between the consequent return differential against today's wealth groups. The goal of this sub-section is to arrive at the following summary result: Between the interquartile range of today's wealth distribution, richer households earn an annual 60-basis-point higher return on housing than poorer households, because of the differences in the timing of ownership.

I proceed in two steps. First, I translate the right-hand-side variable, "1940 income percentile," to today's wealth percentile. This step takes two sub-steps: (1) I map the "1940 income percentile" to the average home value among homeowners today; and (2) I translate the average home value to the corresponding place in today's wealth distribution, using a linear relationship estimated from the 2013 Survey of Consumer Finances (SCF). Second, I translate the left-hand-side variable, the "betas," to a return differential, using the first-order approximation in Proposition 4.

For the quantification in this sub-section, I use the estimate from the shorter-but-wider 1998-2013 sample. The estimates from the other sample are similar.

3.5.1 From 1940 Income to Today's Wealth

I first translate the "1940 income percentile" to the predicted place in the 2013 wealth distribution. The basic idea has two components: (1) For each surname, I can observe the average primary-residence value in the 2012-2013 assessor file (Henry de Frahan and Sakong (2020)); and (2) Because richer households live in more expensive homes, I infer how rich the surnames must be based on the value of their first homes (Engel (1857)). I use the value of first homes rather than total housing wealth, because rich households will live in more expensive homes even if they choose to invest in other assets for their investment portfolios (Campbell (2006)).[37]

First, I translate the "1940 income percentile" groups to their average home values in 2012-2013, to estimate how the estimated betas co-vary with today's home values. This estimation takes the following two-stage form: For "1940 income percentile" group i, I estimate

$$\beta_i = \gamma x_i + \varepsilon_i$$

$$x_i = Z_i \Gamma$$

[37]While only monotonicity is required, it helps for the second transformation that housing consumption also has an income elasticity close to one (Henry de Frahan and Sakong (2020)).

where $x_i \equiv E\left[\text{log home value} \mid \text{own}\right]_i$ is the average home value among homeowners in the 2012-2013 CoreLogic assessor file (Figure 1a), and Z_i is a vector of dummies for each of the "1940 income percentile" groups. This "second-stage" relationship between β_i and x_i is plotted in Figure 4a.

Second, I infer what wealth levels today must have led to the observed home values, x_i. For this relationship, I use the 2013 Survey of Consumer Finances (SCF), which is the most reliable source of micro-data on wealth (Pfeffer et al. (2016)). Figure 4b groups households by the net-worth percentile they belong to and then plots the average log-first-home-value, only among the homeowners in that group. The bin scatter (by percentiles) shows that for most of the distribution in the middle, the log-home-value is linearly increasing in the wealth percentiles.

The slope of the home-value-on-wealth-percentile relationship in Figure 4b is estimated from

$$\begin{aligned} x_j \equiv E\left[\text{log home value} \mid \text{own}\right]_j &= a + b\left(\text{wealth percentile}_j\right) \\ &\equiv f\left(\text{wealth percentile}_j\right) \end{aligned}$$

for household j. I use the estimated *linear* relationship to turn the average home-values by "1940 income percentile" groups from the CoreLogic data, into their corresponding wealth percentiles, i.e., $\left(\text{wealth percentile}\right)_i = f^{-1}\left(x_i\right)$.[38] This variable will be on the x-axis in the final return-differential-on-wealth-percentile relationship.

3.5.2 From Estimated Betas to Return Differentials

The estimated betas measure the elasticity of the quantity of housing held by one group of surnames in response to the national house-price index. Proposition 4 gives a simple, linear approximation of how this beta translates to a return differential arising from the timing of trades. As a reminder, the linear relationship is given by $-var\left(\mu_t^k\right)\left(\tilde{b}^k \frac{\overline{D}^k}{\overline{P}^k}\right)^{-1}\hat{\beta}_i$, where $\hat{\beta}_i$ are the estimated betas for the "1940 income percentile" groups i. The coefficient on $\hat{\beta}_i$ is composed of expected-return properties of housing: $var\left(\mu_t^k\right)$ measures the variance of expected returns, $\tilde{b}^k$ measures how expected returns co-vary with the rent-to-price ratio $\frac{D_t^k}{P_t^k}$, and $\frac{\overline{D}^k}{\overline{P}^k}$ is the time-averaged rent-to-price ratio. Given those housing-market-level values, the estimated betas $\hat{\beta}_i$ can be translated to a return differential linearly.

[38] Because I estimate the function $f(\cdot)$ at the household-level in the 2013 Survey of Consumer Finances (SCF), but apply the inverted function to surname-level data, I make the assumption is that the non-wealth determinants of homeownership and home values relate to wealth in the same way at the individual level as at the surname level. For a formalization of what assumptions are required to translate surname-level relationships to household-level relationships, see Henry de Frahan and Sakong (2020).

In this section, I take the characteristics of the expected returns of the national housing market from Cochrane (2011), who showed that return predictability is comparable between the aggregate stock market and aggregate housing market. I also verify similar numbers by using the national house-price index from CoreLogic myself. I assume the following numbers: $var\left(\mu_t^k\right) \approx (0.0546)^2$ and $\tilde{b}^k \approx 3.8$ from Cochrane (2011) for the aggregate stock market, and the average rent-to-price ratio of $\frac{\overline{D}^k}{\overline{P}^k} \approx \frac{1}{16}$.

Note that for transparency, I used the characteristics of the aggregate housing market (i.e., the asset k is the aggregate housing stock). In reality, expected-return dynamics may vary by local housing markets and may be correlated with the beta-against-wealth-level relationships. I explore this heterogeneity by local housing markets k, in the next section on the geographical cross-section.

3.5.3 Interpreting the Return Differential by Wealth Percentiles

Figure 4c plots the imputed return differential against the imputed net-worth percentile. The linear fit shows that a 10% increase in the net-worth percentile translates to a roughly 12-basis-points-higher annual return on wealth.

The estimate is quite large. For the sake of comparison, I "extrapolate" onto the interquartile range, with a 60-basis-point return differential.[39] Note that this extrapolation is more justified than appears: Because I have aggregated population to surname-level data, the linear relationship is the first-order approximation of the true relationship in the population (see a similar explanation for zip-code-level data in Mian and Sufi (2016)).

There is only a limited number of estimates of how portfolio returns vary by wealth level.[40] The most comprehensive estimate is from Fagereng et al. (2016), who study financial wealth in Norway. They find a total financial-portfolio-return-differential of roughly 1% in the interquartile range.

If we think the overall return differential for the US housing portfolio is similar (in the absence of any real estimate), then my estimation argues that 60% of the overall return differential may be due to the timing of trades given the same assets. Whereas most of the literature on wealth-return differentials focuses on the asset-side heterogeneity, this estimate argues that the timing of trades matters, even though it is much less explored.

[39] A related calculation is how much realized higher returns were for the richer households, given the realized returns in the recent boom and bust (1998-2013). That calculation uses the same formula, but instead of using the variance in expected return, I would use the variance of realized returns, which was roughly 0.01 post-1999. This estimate translates to an average annual return differential of 2% from the timing of trades.

[40] Some examples include Piketty (2014) (using university endowments), Saez and Zucman (2016) (using US foundations), and Piketty and Zucman (2015) (using *Forbes* global wealth rankings over 1987-2013).

3.6 From Return Differentials to Wealth Inequality

In this sub-section, I derive a back-of-the-envelope calculation to translate the estimated return differential to wealth inequality. This last calculation is non-trivial. While wealth inequality exhaustively decomposes into contributions from income inequality, savings-rate differences, and portfolio-return differences, the contribution of return-differentials to wealth-inequality levels depends on the other two factors. For example, in the neoclassical benchmark in which human capital is tradable (i.e., idiosyncratic labor-income shocks are fully insured), any agent who can earn a systematically higher return will own all the wealth in the long-run, however small the return differential (Levy (2003)).

In a more realistic model where human capital is not tradable, labor income flows keep the wealth distribution from diverging even with systematically different returns; given a fixed level of labor income, as my wealth gets larger, the labor income acts as a higher proportional inflow into my wealth portfolio (Gabaix et al. (2016) call it "stabilization"). How much labor income can "stabilize" wealth portfolios in turn depends on how much households spend out of the labor-income inflow. The back-of-the-envelope calculation makes approximate assumptions on these forces.

As a reminder, the accumulation of measured, non-human wealth W_{it} was given by: $\frac{dW_{it}}{W_{it}} = \left(\frac{Y_{it}}{W_{it}} - \frac{C_{it}}{W_{it}}\right) dt + \overline{dR}_{it}$, where Y_{it} is labor income, C_{it} is expenditure inclusive of user cost of housing, and $\overline{dR}_{it}$ is wealth return.

For the back-of-the-envelope calculation, I assume idiosyncratic labor incomes are not fully insured and shut down more complex savings-rate differences with the following assumption.

Assumption 5 (Approximate expenditure policy)**.** *Assume the following approximation for the expenditure C_{it}:*

$$C_{it} = c_y Y_{it} + c_w W_{it}$$

for the same constants c_y and c_w for all household types.

This expenditure policy encompasses some benchmarks,[41] but it is not theoretically sound. For example, the approximation predicts that house-changes for homeowners would lead to proportional changes in expenditures, but the wealth effect of house-price changes is debated (e.g., see Berger et al.

[41] For example, a hand-to-mouth agent would have $c_w = 0$. Heterogeneous-agent models with two states (income and asset) would imply $C_t = c_y(Y_t, W_t) Y_t + c_w(Y_t, W_t) W_t$. Here the arguments of the average propensity to consume out of income and wealth are suppressed.

(2017)). The policy also implies a marginal propensity to consume (MPC) that does not vary with wealth level, which is at odds with the data. For these reasons, the approximate policy is meant as an approximation and simplification.

To be precise, what matters for the accumulation of measured wealth is the average propensity to consume (APC).[42]

Lemma 6 (Measured wealth share)**.** *Given all Assumptions and that aggregate labor income (Y) and measured wealth (W) are co-integrated, in the long-run stationary distribution even with aggregate shocks,*

$$E\left[\frac{Y_{it}}{W_{it}}\right] - E\left[\frac{Y_t}{W_t}\right] = -\frac{E\left[\overline{dR}_{it} - \overline{dR}_t\right]}{1-c_y}$$

Wealth shares inherit income shares if returns are identical.

The lemma follows immediately from $E\left[d\log W_{it} - d\log W_t\right] = 0$ for agent types i relative to the population average in the long-run stationary wealth distribution.[43] This formula translates the portfolio return differential, $E\left[\overline{dR}_{it} - \overline{dR}_t\right]$, to the difference between wealth inequality to income inequality: $E\left[\frac{Y_{it}}{W_{it}}\right] - E\left[\frac{Y_t}{W_t}\right]$.

The actual equation mapping the return differential to wealth inequality is difficult to manipulate due to the expectations. To get an order of magnitude and for insight into the formula, I treat all quantities as constants for a back-of-the-envelope calculation.

Corollary 7 (Back-of-the-envelope calculation)**.** *The relationship between the wealth share and income share of household type i is approximate by*

$$\frac{W_i}{W} \approx \left\{ 1 + \underbrace{\left(E\overline{dR}_i - E\overline{dR}\right)}_{\textit{return differential}} \overbrace{\frac{W}{Y}\frac{1}{1-c_y}}^{\textit{income stabilization}} \right\} \frac{Y_i}{Y}$$

See the appendix for the derivation.

The relationship between wealth share and income share of an agent type i is given by the return differential, offset by the average income flow rate, tempered by $(1-c_y)$ (i.e., savings rate out of

[42]While MPC is decreasing in income and liquid wealth, the APC gradient is theoretically ambiguous, even among models with downward-sloping MPC. Models with precautionary savings by the poor would exhibit increasing APC (Aiyagari (1994)). Models in which the rich prefer to save (i.e., wealth is a luxury) would exhibit decreasing APC (Carroll (2000)). The interaction is insignificant in the Panel Study of Income Dynamics (PSID).

[43]The short-term drift $E_t\left[d\log W_t\right]$ need not equate, in the presence of aggregate shocks.

current income). Again, the APC out of current labor income (c_y) works to scale the return differential up and down in determining wealth-to-income inequality. To get the order of magnitude, I estimate a rough consumption policy from the PSID, which has information for consumption, income, and wealth. The estimated coefficient is $c_y \approx 0.25$. See the online appendix for the estimation.

Using the measured-wealth-to-labor-income ratio of $\frac{W}{Y} \approx 10$ from the SCF and $c_y \approx 0.25$, the interquartile return differential per year of 60 basis points translates to a wealth share that is 8% higher than the income share. Comparing this to the wealth-to-income elasticities in the SCF (Figure A.5a), the timing of trades explains roughly a fifth of the residual wealth inequality above and beyond income inequality.

The approximation exercises in this section are complementary to fully fledged models of steady-state wealth distributions.[44]

4 Comparing across Geographies

Given poor and rich households' different ownership elasticity to price, the magnitude of the return differential depends on the variability of expected returns in the housing market. Since households are more likely to own housing assets near where they live and residence is sticky over time, we would expect the trade-timing mechanism to lead to greater return differential and greater wealth inequality in housing markets with more volatility. In this section, we test this cross-sectional implication of this paper's main mechanism.

In the first sub-section, I sort local housing markets by US counties by one predictor of variation in expected returns: how much the local economy fluctuates along national business cycles. The predictive power is true in the data, and justified by: (1) Returns in housing as an asset class have strong geographical components (Piazzesi and Schneider (2016)), and (2) Business cycles are strong predictors of expected returns (Cochrane (2011)).

In the second sub-section, I test whether the historical cyclicality of the area did accompany both more volatile prices and higher elasticities of quantity to price for poor households.

In the last sub-section, I test for the long-run, ultimate implication of whether wealth inequality is indeed higher. The second test can be considered the ultimate outcome, but I have to first overcome an

[44] Early examples include Champernowne (1953), Vaughan (1979), Laitner (1979), Stiglitz (1969) (using the neoclassical growth model), and Hopenhayn and Prescott (1992). Recent models that relate returns and wealth inequality include Benhabib et al. (2011), Benhabib et al. (2015b), and Nirei et al. (2009).

empirical challenge: A local measure of wealth inequality does not exist.

4.1 Defining Housing Asset Sub-classes

I first sort geographies on the cyclicality of their expected returns and consequently the effect of the timing of trades on wealth inequality. The sorting variable defined in this sub-section will be the independent variable in the next two sub-sections, where I show that it predicts (1) higher gradient of the housing-quantity-on-price elasticity (i.e., the "beta") and (2) higher wealth-inequality levels relative to income inequality.

Housing markets are potentially segmented between distant geographies (Piazzesi and Schneider (2016)). I divide housing assets into sub-classes by counties in which those properties are located. In this paper, I focus on the cyclicality of the county's local economy. The local business-cycle cyclicality is computed by regressing county-level log income change on the aggregate log change using data from the Bureau of Economics Analysis for 1969-2015. That is, I take π_c from

$$\Delta \log Y_{ct} = \pi_c \Delta \log \overline{Y}_t + \eta_{ct}$$

where Y_{ct} is the per-capita income in county c in year t, and $\overline{Y}_t$ is the national per-capita income in year t. This regression is estimated using per-capita income data from the Bureau of Economic Analysis, for 1969-2015. The coefficient π_c captures the cyclicality of the local economy in county c. The distribution of cyclicality is seen in Figure A.4c.

An alternative way to sort locations by the expected-return volatility is to use proxies for the supply elasticity in the housing market. With more inelastic supply of housing, prices will be more volatile, leading to more variations in the expected returns. One measure of the housing-supply elasticity is from Saiz (2010). Using the Saiz (2010) housing-supply elasticity instead of the π_c leads to qualitatively the same results.

4.2 Beta-gradient by Local-Market Cyclicality

Before moving onto the level of wealth inequality, I check whether the beta-gradients (i.e., the relationship between the housing-quantity-to-price elasticity and wealth levels) are higher or lower in areas where expected returns are more volatile. The direction is theoretically ambiguous. The higher expected-return variation could have induced poorer households to engage in less selling in busts via a

strong price effect. Alternatively, bigger credit deterioration in those areas with bigger price drops may have more adversely affected poor households and caused them to sell more via a strong income effect.

To see how the elasticity gradient relates to the cyclicality in the local market, I run the following two regressions, for each county c:

$$\log q_{ict} = \tilde{\delta}_c \left(\log P_{ct} \times 1940 \text{ income percentile}_i\right) + \alpha_{ic} + \alpha_{ct} + \gamma_{ic} t + \xi_{ict}$$

where q_{ict} is the number of real estate properties located in county k, held by individuals with surnames in the "1940 income percentile" group i in year t, and P_{ct} is the CoreLogic house-price index in county c in year t. This regression is estimated using the 1998-2013 CoreLogic sample covering 60% of the US population. The estimate $\tilde{\delta}_c$ captures the extent to which poorer households hold properties pro-cyclically in county c.

Figure 5 plots the county-level beta-gradient $\tilde{\delta}_c$ against business-cycle loading π_c. The plot shows the elasticity gradient was more negative in areas with bigger cycles. That is, for real estate properties in geographical areas with historically higher cyclicality, poorer households exhibited more pro-cyclical holdings.

4.3 Wealth-inequality Level versus Cyclicality

The return-gradient driven by timing of trades is larger in areas with more cyclical economies because the expected returns vary more and the beta-gradient is steeper (i.e., poor households' housing-quantity ownership responds more sensitively to house prices). Consequently, in areas with higher historical cyclicality, we should see a bigger wealth gap than income gap, compared to households who invest in less volatile housing markets.

Conveniently for testing, housing investment exhibits a strong home bias. Owner-occupants live in that same house. Even for investors, because housing is heterogeneous and requires local information to invest, investors too would exhibit strong home bias in terms of geography. These factors imply that in cities where business cycles and house prices are more volatile, we would see greater wealth inequality above and beyond income inequality. This sub-section documents the correlation between business-cycle cyclicality and wealth inequality, above and beyond income inequality. This correlation is not meant to rule out other mechanisms, but it is a necessary implication of the mechanism discussed in this section.

4.3.1 Measuring CBSA-level Wealth Inequality

I first address an empirical challenge: Local-level measures of wealth inequality do not exist. To overcome this challenge, I first form zip-code-level balance sheets, and then form between-zip-code inequality measures for metropolitan areas, and argue that they are informative for household-level inequality measures.

Following Mian et al. (2013) and Saez and Zucman (2016), the balance sheet is given by

$$NW = F^i + F^{ni} + H - D$$

for net worth NW. F^i is income-generating financial assets. F^{ni} is non-income-generating financial assets: life insurance and pension funds, currency and non-interest deposits (~1% of total wealth today), and offshore wealth held through foreign institutions (~4% of net financial wealth) (Saez and Zucman (2016))[45]. H is housing (both owner-occupied and investment housing). D is liability.

I first describe components for which I have direct measures at the zip-code level. Data on housing ownership come from the assessor file of the CoreLogic data, described in detail in the second section. I assign housing to zip codes, using the mailing address of the owner. As described earlier, I have the single cross-sectional assessor data for 2012-2013, so I form the rest of the balance-sheet measures for 2012.

Data on household liability come from Equifax. Equifax is a consumer-credit-reporting agency, which collects data on consumer-credit histories to assign credit scores. I have access to zip-code-level aggregate amounts of various household-debt instruments. I have access to Equifax annual panel up to 2011, so I use the zip-code-level debt amounts for that year.

Zip-code-level financial-asset holdings require imputation, because no data on financial-wealth holdings by zip codes exist. The basic idea is to obtain cash flows by asset categories (e.g., dividend for equity and interest for bond), and to capitalize them into the stock, assuming households earn the same yield within a given asset class, following Mian et al. (2013) and Saez and Zucman (2016).[46][47] I

[45] Non-taxable fixed income claims (state/local government bonds) are tax-exempt but reported on individual tax returns since 1987 (Saez and Zucman (2016)). Wealth held by individuals through trusts flows directly to dividends, realized capital gain, interest, and to Schedule E fiduciary income (rents/royalties).

[46] Saez and Zucman (2016) use capital income: "Capital income includes dividends, taxable interest, rents, estate and trust income, the profits of S-corporations, sole proprietorships and partnerships; we also present a series including realized capital gains... For the post-1962 period, we impute wealth at the individual level by assuming that within a given asset class, everybody has the same capitalization factor."

[47] The capitalization technique has a long history, going back to King (1927), Stewart (1939), Atkinson and Harrison

obtain zip-code-level total dividends, total interests and total private-business profits from the Internal Revenue Services (IRS) Statistics of Income (SOI). I take the capitalization factors from Saez and Zucman (2016) Appendix Table A11: "Capitalization factors by asset class." In 2004, for example, the capitalization factor is 51.4 for taxable interest and 43.6 for dividends.[48]

Lastly, I also get zip-code-level labor income from the IRS SOI ("wage and salary"). This variable is used to form labor-income-inequality measures, used as controls so that I can focus on residual wealth inequality above and beyond income inequality.

After I form the zip-code-level NW, I form two measures of inequality for each metropolitan area (or core-based statistical areas (CBSA)).

The first inequality measure is the coefficient of variation (CV), defined as the standard deviation divided by the mean. The squared CV is subgroup-decomposable, so the total CV in wealth at the household level can be decomposed into between-zip-code variance (observed) and within-zip-code variance. A key assumption in using the between-zip-code CV is that the between- and within-zip-code variances are proportional.

Figure A.4a plots the imputed CV in net worth by CBSAs, and Figure A.4b plots the imputed CV for labor income. Note the two maps have overlap but also have differences.

I validate this between-zip-code measure using total income, for which I can form individual-level inequality measures at the CBSA level using the American Community Survey accessed via Integrated Public Use Microdata Series (IPUMS). Note the IPUMS data are a sample and thus have sampling error. Figure 6c plots the zip-code-level total-income CV from the IRS against the IPUMS household-level total-income CV. A significant positive relationship exists across CBSAs.

The second inequality measure is elasticity of wealth to labor income, or the wealth-to-labor-income ratio. The mechanism posits that the same amount of income would translate to more wealth if returns are increasing in wealth. The wealth-to-labor-income ratio can be calculated at the zip-code level. The zip-code-level wealth-to-labor-income elasticity of 1.4 for net worth (Table 3b column 1) is close to the household-level relationship in Figure A.5a, computed from the Survey of Consumer Finances.

(1978), and Greenwood (1983)

[48]Capital gains are ignored for now, although they may be useful for inferring equity holdings.

4.3.2 Wealth-inequality Level Results

Figure 6a plots the between-zip-code CV in net worth for CBSAs against the average π_c (income cyclicality) in each CBSA, controlling for the CV in wage income. The strong positive relationship is shown in Table 3a, in column (2), using the specification,

$$\text{CV}_m = \phi \pi_c + \gamma \text{wage CV}_m + \Gamma X_c + \varepsilon_c$$

where CV_m is the coefficient of variation in assets or net worth in CBSA m, π_c is the cyclicality measure computed in the first sub-section, and X_c is a vector of controls.

The significant positive relationship is robust to varying the empirical specification. In both tables, column (3) shows the coefficient on the business cycle goes up if controlling for size of the CBSA or the average price level. Column (4) uses the average equity and bond holdings over 2003-2012 (formed using each year's IRS SOI dividend and interest income and each year's capitalization factors from Saez and Zucman (2016)) instead of the value for 2012. Column (5) adds state fixed effects to focus on more local differences. Across these specifications, higher cyclicality in the past half-century predicts higher wealth inequality if the area had higher loading on the aggregate cycle.

An alternative test is to see if wealth-to-labor-income ratios are increasing in the local economy's cyclicality. Table 3b reports results from:

$$\begin{aligned} \log W_z = & \psi \left[\log(\text{wage})_z \times \pi_c\right] + \Gamma_1 \left[\log(\text{wage})_z \times X_c\right] \\ & + \Gamma_0 X_{cz} + \delta \pi_c + \gamma \log(\text{wage})_z + \varepsilon_z \end{aligned}$$

for net worth W_z in zip code z in county c, where X_{cz} includes log population size and log house-price level.

The average elasticity in column (1) is the same as the household-level elasticity in Figures A.5a and A.5b.[49] Column (2) shows the elasticity is higher (i.e., wealth-to-labor-income ratio is higher) in counties with higher cyclicality. This correlation is robust to controlling for city size and average house price (column 3), using average capital income over 2003-2012 (column 4), and including state fixed effects (column 5).

[49] Using the 2013 Survey of Consumers Finances (SCF), I plot log net worth against log income (Figure A.5a) and log total asset against log income (Figure A.5b). Taking logs, these plots restrict the sample to positive amounts of wealth. The elasticities are roughly constant and significantly greater than 1, so that higher income translates to disproportionately higher wealth. The coefficient is roughly 1.4 for net worth and 1.5 for total asset.

Interpreting the level evidence requires several caveats: (1) The inequality measures are computed using data at the zip-code level as opposed to individual-level data (one main concern is residential segregation by income, differentially between CBSAs); (2) I use the capital-income flow to impute financial-wealth-stock based on constant capitalization factors; (3) The relationship is not causal, from cyclicality to wealth inequality; (4) Even if I can establish causality, mechanisms other than the trade-timing mechanism of this paper can generated the causal relationship. Yet the cross-sectional implication on wealth inequality is borne out in the correlations.

5 Conclusion

Why is wealth distributed so unevenly even among the bottom 99%, and even more so than income is? This paper gives one partial answer: Poorer households own more housing during booms when house prices are high and expected returns are low, and vice versa in busts. The return-differential generated from this channel is large: 60 basis points per year between the interquartile range of the wealth distribution.

To arrive at this estimate, I construct and use a panel dataset on quantity of housing held by wealth levels. I assign wealth levels to owners in the housing-deed records from CoreLogic, by matching them by surnames to the average incomes of those surnames in the 1940 full-count Census. I derive approximations that translate the quantity-ownership patterns to return differentials.

This trade-timing mechanism behind wealth differentials arises because expected returns on housing are time-varying and predictable. It further implies that time-series volatility would widen wealth inequality: I verify this implication across US metropolitan areas.

This paper also makes broader points: Wealth is about accumulation, so dynamic mechanisms are important and asset-price movements (not just average returns) are important. On the methodological side, the nexus between deed records and full count Census can be extended back much further in history. This nexus can be a lens through which to study cycles (and other topics in economics) going back hundreds of years.

References

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684.

Atkinson, A. B. and Harrison, A. J. (1978). *Distribution of personal wealth in Britain*. Cambridge Univ Pr.

Bach, L., Calvet, L. E., and Sodini, P. (2016). Rich pickings? risk, return, and skill in the portfolios of the wealthy.

Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2015). X-capm: An extrapolative capital asset pricing model. *Journal of financial economics*, 115(1):1–24.

Bayer, P., Casey, M., Ferreira, F., and McMillan, R. (2017). Racial and ethnic price differentials in the housing market. *Journal of Urban Economics*, 102:91–105.

Bayer, P., Ferreira, F., and Ross, S. L. (2016). The vulnerability of minority homeowners in the housing boom and bust. *American Economic Journal: Economic Policy*, 8(1):1–27.

Becker, G. S. ([1957] 1971). *The Economics of Discrimination*. University of Chicago press.

Benhabib, J., Bisin, A., and Luo, M. (2015a). Wealth distribution and social mobility in the us: A quantitative approach. Technical report, National Bureau of Economic Research.

Benhabib, J., Bisin, A., and Zhu, S. (2011). The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica*, 79(1):123–157.

Benhabib, J., Bisin, A., and Zhu, S. (2015b). The wealth distribution in bewley economies with capital income risk. *Journal of Economic Theory*, 159:489–515.

Beraja, M., Hurst, E., and Ospina, J. (2016). The aggregate implications of regional business cycles. Technical report, National Bureau of Economic Research.

Berger, D., Guerrieri, V., Lorenzoni, G., and Vavra, J. (2017). House prices and consumer spending. *The Review of Economic Studies*, 85(3):1502–1542.

Beshears, J., Choi, J. J., Harris, C., Laibson, D., Madrian, B. C., and Sakong, J. (2015). Self control and commitment: can decreasing the liquidity of a savings account increase deposits?

Bisin, A. and Benhabib, J. (2017). Skewed wealth distributions: Theory and empirics. *Journal of Economic Literature*.

Boulware, K. D. and Kuttner, K. N. (2020). Wealth stratification and portfolio choice. In *AEA Papers and Proceedings*, volume 110, pages 411–15.

Calvet, L. E., Campbell, J. Y., and Sodini, P. (2009). Fight or flight? portfolio rebalancing by individual investors. *The Quarterly journal of economics*, 124(1):301–348.

Campbell, J. Y. (2006). Household finance. *The journal of finance*, 61(4):1553–1604.

Campbell, J. Y., Ramadorai, T., and Ranish, B. (2018). Do the rich get richer in the stock market? evidence from india.

Campbell, J. Y., Viceira, L. M., Viceira, L. M., et al. (2002). *Strategic asset allocation: portfolio choice for long-term investors*. Clarendon Lectures in Economic.

Campbell, S. D., Davis, M. A., Gallin, J., and Martin, R. F. (2009). What moves housing markets: A variance decomposition of the rent–price ratio. *Journal of Urban Economics*, 66(2):90–102.

Carliner, M. S. (1998). Development of federal homeownership policy. *Housing Policy Debate*, 9(2):299–321.

Carroll, C. D. (2000). Why do the rich save so much? *Does Atlas Shrug?: The Economic Consequences of Taxing the Rich*, page 465.

Champernowne, D. G. (1953). A model of income distribution. *The Economic Journal*, 63(250):318–351.

Charles, K. K. and Hurst, E. (2002). The transition to home ownership and the black-white wealth gap. *The Review of Economics and Statistics*, 84(2):281–297.

Cochrane, J. H. (2011). Presidential address: Discount rates. *The Journal of finance*, 66(4):1047–1108.

Davis, M. A. and Van Nieuwerburgh, S. (2015). Housing, finance, and the macroeconomy. In *Handbook of regional and urban economics*, volume 5, pages 753–811. Elsevier.

De Nardi, M. and Fella, G. (2017). Saving and wealth inequality. *Review of Economic Dynamics*, 26:280–300.

Engel, E. (1857). Die produktions-und konsumptionsverhältnisse des königreichs sachsen. *Zeitschrift des Statistischen Bureaus des Königlich Sächsischen Ministeriums des Innern*, 8:1–54.

Fagereng, A., Guiso, L., Malacrino, D., and Pistaferri, L. (2016). Heterogeneity and persistence in returns to wealth. Technical report, National Bureau of Economic Research.

Gabaix, X., Lasry, J.-M., Lions, P.-L., and Moll, B. (2016). The dynamics of inequality. *Econometrica*, 84(6):2071–2111.

Garbinti, B., Goupille-Lebret, J., and Piketty, T. (2017). Accounting for wealth inequality dynamics: Methods, estimates and simulations for france (1800-2014).

Ghysels, E., Plazzi, A., Valkanov, R., and Torous, W. (2013). Forecasting real estate prices. In *Handbook of economic forecasting*, volume 2, pages 509–580. Elsevier.

Glaeser, E. L. and Nathanson, C. G. (2015). Housing bubbles. In *Handbook of regional and urban economics*, volume 5, pages 701–751. Elsevier.

Greenwood, D. (1983). An estimation of us family wealth and its distribution from microdata, 1973. *Review of Income and Wealth*, 29(1):23–44.

Guvenen, F. (2011). Macroeconomics with hetereogeneity: a practical guide. *Economic Quarterly*, (3Q):255–326.

Guvenen, F., Ozkan, S., and Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660.

Heathcote, J., Storesletten, K., and Violante, G. L. (2009). Quantitative macroeconomics with heterogeneous households. *Annu. Rev. Econ.*, 1(1):319–354.

Henry de Frahan, L. and Sakong, J. (2020). Intergenerational elasticity of consumption and income.

Hoopes, J., Langetieg, P., Nagel, S., Reck, D., Slemrod, J., and Stuart, B. (2016). Who sold during the crash of 2008-9? evidence from tax-return data on daily sales of stock. Technical report, National Bureau of Economic Research.

Hopenhayn, H. A. and Prescott, E. C. (1992). Stochastic monotonicity and stationary distributions for dynamic economies. *Econometrica (1986-1998)*, 60(6):1387.

Kaplan, G., Mitman, K., and Violante, G. L. (2017). The housing boom and bust: Model meets evidence. Technical report, National Bureau of Economic Research.

King, W. I. (1927). Wealth distribution in the continental united states at the close of 1921. *Journal of the American Statistical Association*, 22(158):135–153.

Kuhn, M. and Rios-Rull, J.-V. (2016). 2013 update on the us earnings, income, and wealth distributional facts: A view from macroeconomics. *Quarterly Review, Federal Reserve Bank of Minneapolis, April*, pages 1–75.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478.

Laitner, J. (1979). Household bequest behaviour and the national distribution of wealth. *The Review of Economic Studies*, 46(3):467–483.

Levy, M. (2003). Are rich people smarter? *Journal of Economic Theory*, 110(1):42–64.

Longstaff, F. A. (2001). Optimal portfolio choice and the valuation of illiquid securities. *The Review of Financial Studies*, 14(2):407–431.

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. (2016). Who bears the cost of recessions? the role of house prices and household debt. In *Handbook of Macroeconomics*, volume 2, pages 255–296. Elsevier.

Mian, A., Sufi, A., and Verner, E. (2017a). Household debt and business cycles worldwide. *The Quarterly Journal of Economics*, 132(4):1755–1817.

Mian, A., Sufi, A., and Verner, E. (2017b). How do credit supply shocks affect the real economy? evidence from the united states in the 1980s. Technical report, National Bureau of Economic Research.

Nirei, M. et al. (2009). Pareto distributions in economic growth models. *Institute of Innovation Research Working Paper*, (09-05).

Pfeffer, F. T., Schoeni, R. F., Kennickell, A., and Andreski, P. (2016). Measuring wealth and wealth inequality: Comparing two us surveys. *Journal of economic and social measurement*, 41(2):103–120.

Piazzesi, M. and Schneider, M. (2016). Housing and macroeconomics. In *Handbook of Macroeconomics*, volume 2, pages 1547–1640. Elsevier.

Piketty, T. (2014). *Capital in the Twenty-First Century*. Harvard University Press.

Piketty, T. (2015). About capital in the twenty-first century. *American Economic Review*, 105(5):48–53.

Piketty, T. and Zucman, G. (2015). Wealth and inheritance in the long run. In *Handbook of income distribution*, volume 2, pages 1303–1368. Elsevier.

Quadrini, V. and Ríos-Rull, J.-V. (2015). Inequality in macroeconomics. In *Handbook of Income Distribution*, volume 2, pages 1229–1302. Elsevier.

Rajan, R. G. (2011). *Fault lines: How hidden fractures still threaten the world economy*. princeton University press.

Saez, E. and Zucman, G. (2016). Wealth inequality in the united states since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, 131(2):519–578.

Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3):1253–1296.

Sakong, J. (2018). Impact of economic condition on racial prejudice.

Sinai, T. and Souleles, N. S. (2005). Owner-occupied housing as a hedge against rent risk. *The Quarterly Journal of Economics*, 120(2):763–789.

Stewart, C. (1939). Income capitalization as a method of estimating the distribution of wealth by size groups. In *Studies in Income and Wealth, Volume 3*, pages 95–146. NBER.

Stiglitz, J. E. (1969). Distribution of income and wealth among individuals. *Econometrica: Journal of the Econometric Society*, pages 382–397.

Vaughan, R. (1979). Class behaviour and the distribution of wealth. *The Review of Economic Studies*, 46(3):447–465.

Viceira, L. M. (2001). Optimal portfolio choice for long-horizon investors with nontradable labor income. *The Journal of Finance*, 56(2):433–470.

Wainer, A. and Zabel, J. (2019). Homeownership and wealth accumulation for low-income households. *Journal of Housing Economics*, page 101624.

Werning, I. (2015). Incomplete markets and aggregate demand. Technical report, National Bureau of Economic Research.

Wolff, E. N. (2017). Household wealth trends in the united states, 1962 to 2016: Has middle class wealth recovered? Technical report, National Bureau of Economic Research.

Word, D. L., Coleman, C. D., Nunziata, R., and Kominski, R. (2008). Demographic aspects of surnames from census 2000. *Unpublished manuscript, Retrieved from http://citeseerx. ist. psu. edu/viewdoc/download*.

FIGURE 1: Raw-data pattern, sorting by proxies

These figures show raw-data patterns, sorting owners by surnames (and associated 1940 income). Panel (a) plots the average value of primary residence conditional on owning in 2012-2013, for CoreLogic's assessor record, which covers almost the entire US population in a single year. Panel (b) plots the per capita holdings of any real estate asset (i.e., count), for selected decile groups, relative to the 1998 levels. Panel (c) plots the residuals ε_{it} for the same set of selected deciles from the regression:

$$\log(q_{it}) = \alpha_i + \alpha_t + \gamma_i t + \varepsilon_{it}$$

where the regression is weighted by the number of individuals in each decile group, and q_{it} is the holdings of all real estate by number of property by members of the decile group in a given year. For comparison, panel (d) plots the same residuals for CoreLogic national house-price index, i.e., ε_t from

$$\log(P_t) = \gamma_0 t + \varepsilon_t$$

where P_t is the house-price index. The vertical red line indicates 2007.

(A) By surname: primary-residence value

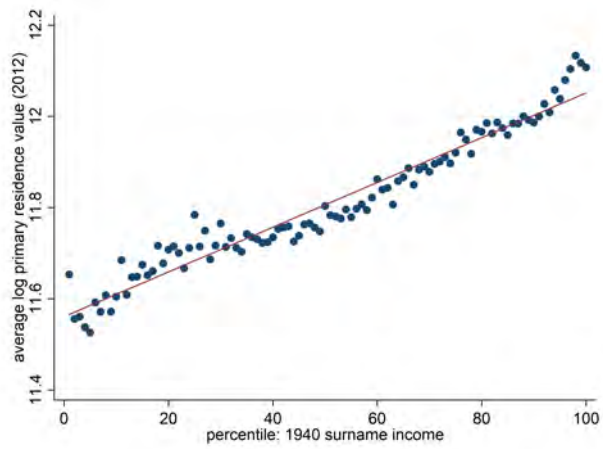

(B) By surname: Per capita holding vs. 1998

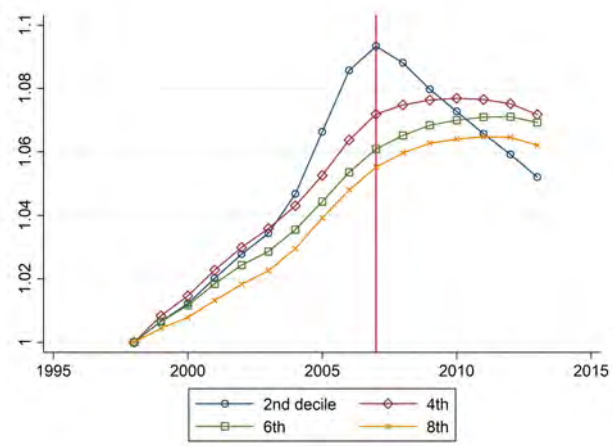

(C) By surname: Detrended log residual

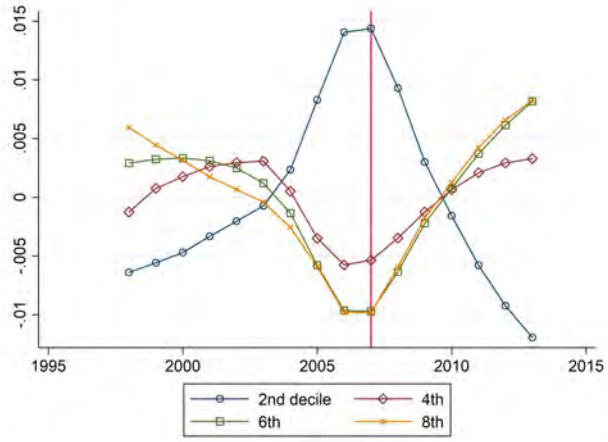

(D) National house price: Detrended log residual

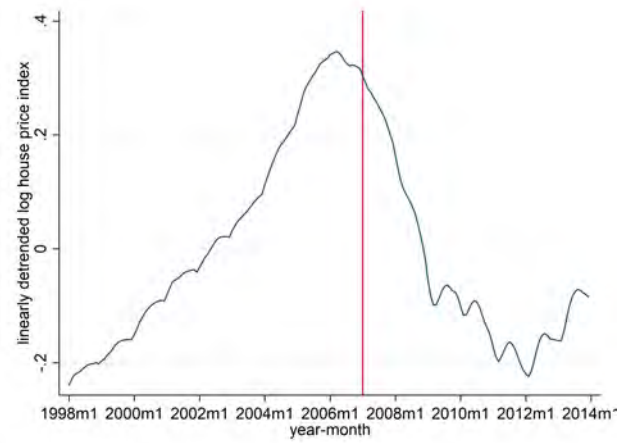

FIGURE 2: Estimated Quantity Elasticity versus Wealth Level

These figures plot the elasticity β_i from

$$\log(q_{it}) = \beta_i \log(P_t) + \alpha_i + \alpha_t + \gamma_i t + \xi_{it}$$

where q_{it} is the total number of properties held per capita for the percentile group i, where the percentiles are sorted using the associated surnames' average household-wage income from the 1940 full Census, and P_t is CoreLogic national house-price index. Panel (a) is estimated using the 1998-2013 CoreLogic sample covering roughly 60% of the US population. Panels (b) and (c) are estimated using the 1988-2013 CoreLogic sample covering roughly 25% of the US population. Panel (b) uses the entire 1988-2013 period for estimation; panel (c) uses only 1988-2002 to exclude the subprime boom and bust.

(A) For 1998-2013 sample

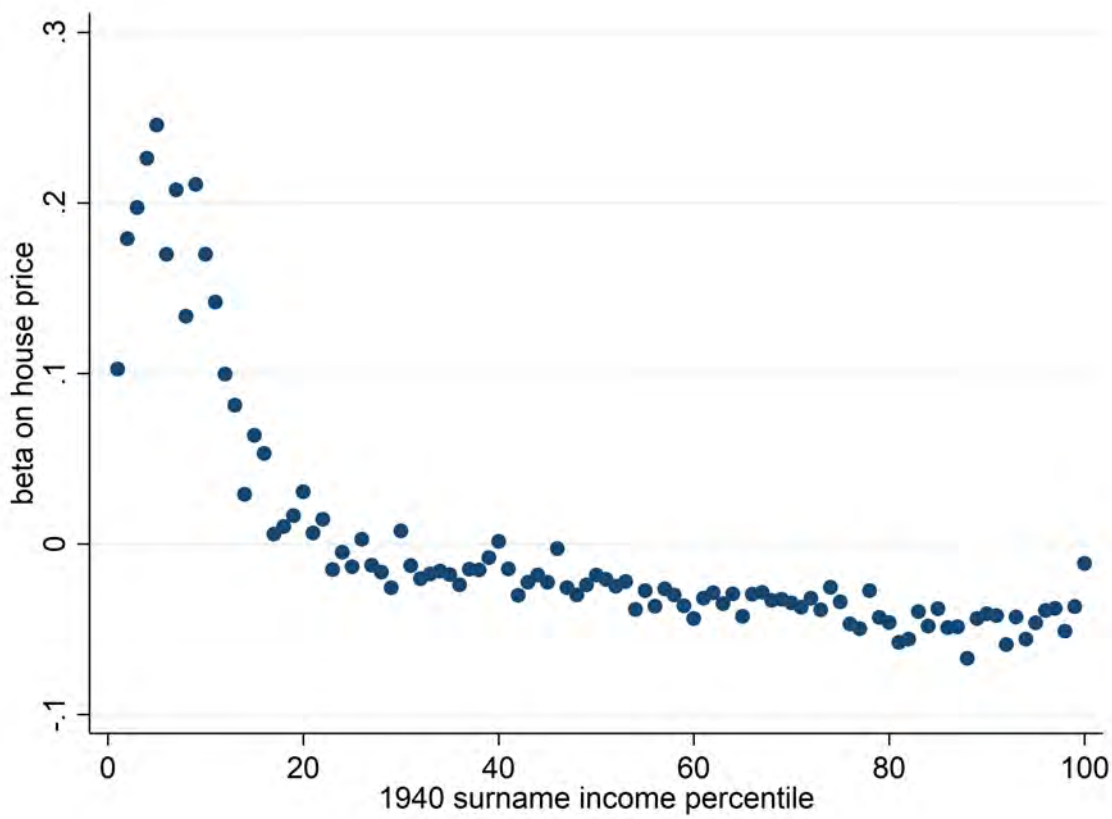

FIGURE 2: Estimated Quantity Elasticity versus Wealth Level (continued)

(B) For 1988-2013 sample: Full period

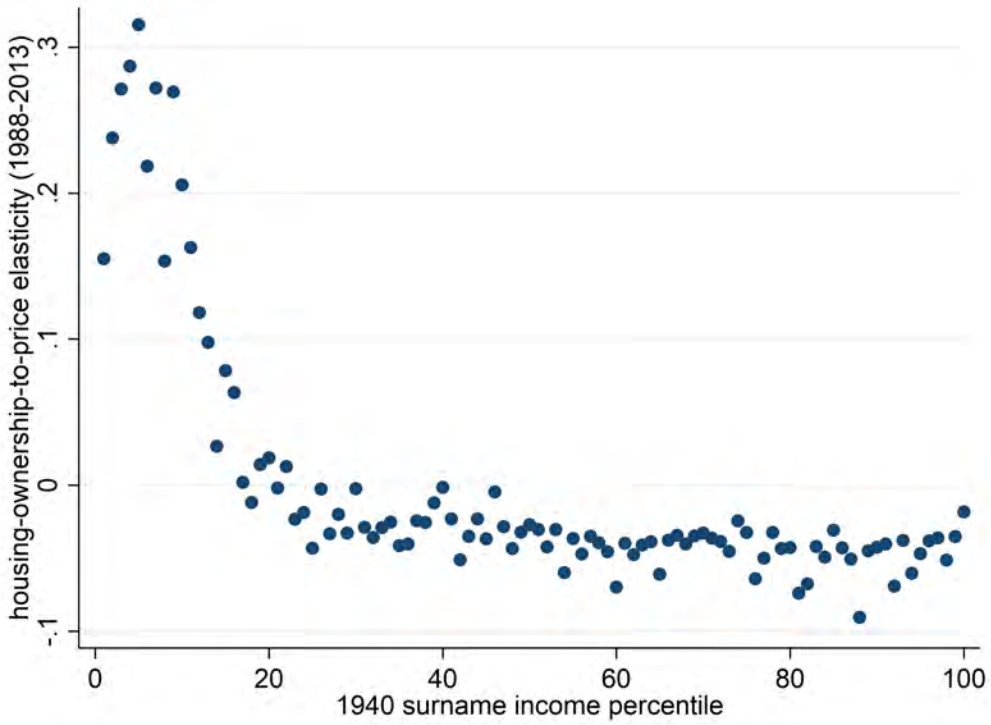

(C) For 1988-2013 sample: Only 1988-2002

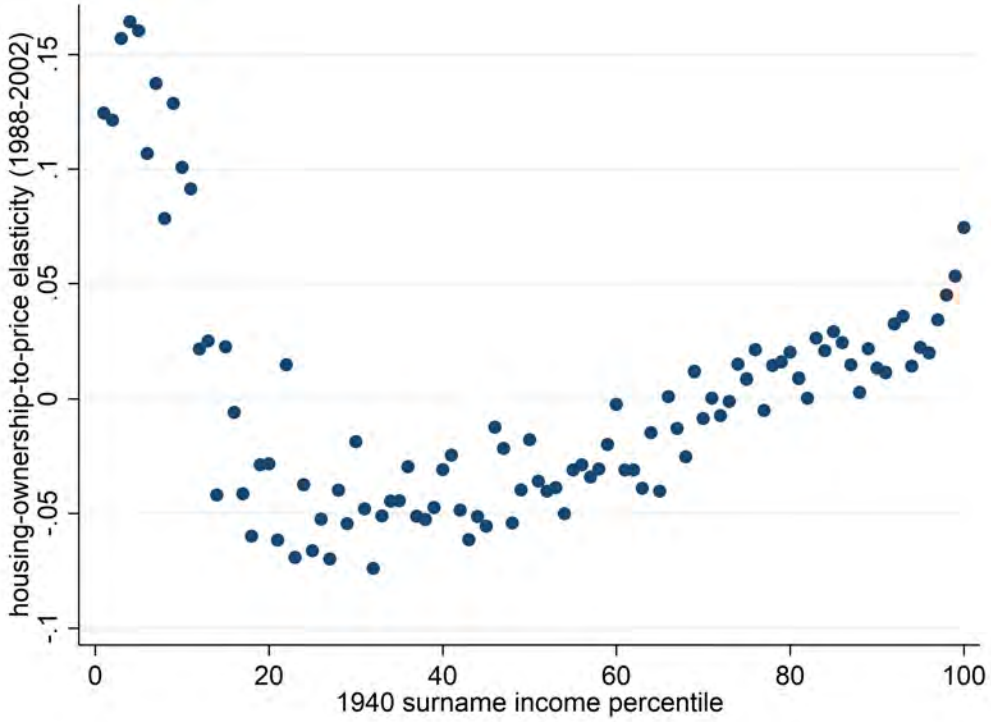

FIGURE 3: "Beta" between- and within-racial share

Panel (a) plots the average share of each "1940 income percentile group" that is racially White. To create the rest of the figures, I first estimated the elasticity β_i from

$$\log(q_{it}) = \beta_i \log(P_t) + \alpha_i + \alpha_t + \gamma_i t + \xi_{it}$$

where q_{it} is the total number of properties held per capita for the percentile group i, where the percentiles are sorted using the associated surnames' average household-wage income from the 1940 full Census, and P_t is CoreLogic national house-price index. Then, for different samples, I plot the β_i against the average White share, controlling linearly for the numerical "1940 income percentile" values, in panels (b), (d) and (f); I plot the β_i against the numerical "1940 income percentile" values, controlling for the average White share, in panels (c), (e) and (g). Panels (b) and (c) are estimated using the 1998-2013 CoreLogic sample covering roughly 60% of the US population. Panels (d), (e), (f) and (g) are estimated using the 1988-2013 CoreLogic sample covering roughly 25% of the US population. Panels (d) and (e) use the entire 1988-2013 period for estimation; panels (f) and (g) use only 1988-2002 to exclude the subprime boom and bust.

(A) White share (2010) by 1940 surname income

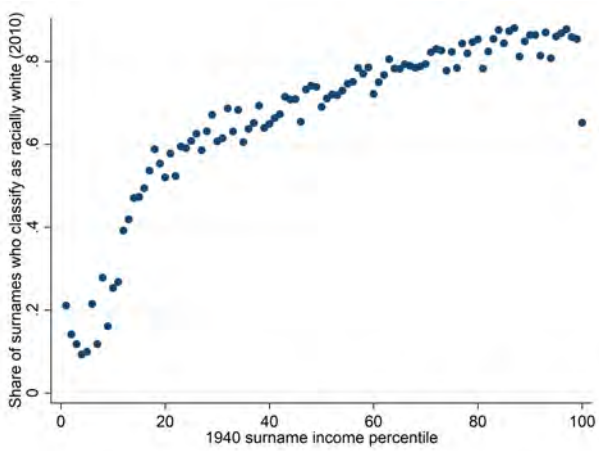

(B) 1998-2013 sample: White share (residual)

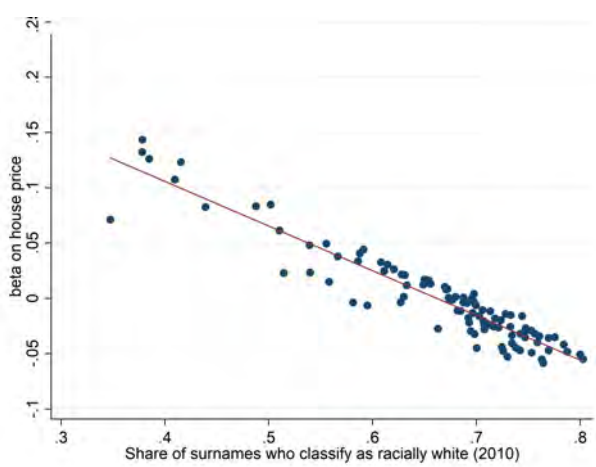

(C) 1998-2013 sample: 1940 income (residual)

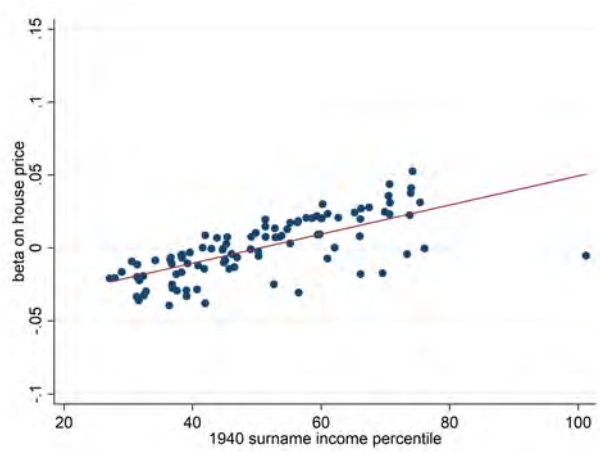

FIGURE 3: "Beta" between- and within-racial share (continued)

(D) 1988-2013 sample (full): White share (residual)

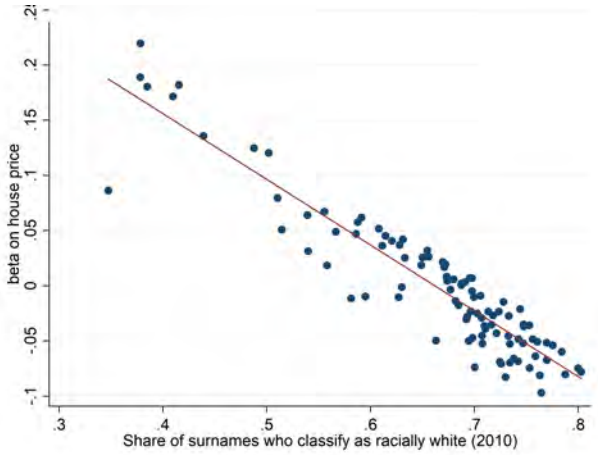

(E) 1988-2013 sample (full): 1940 income (residual)

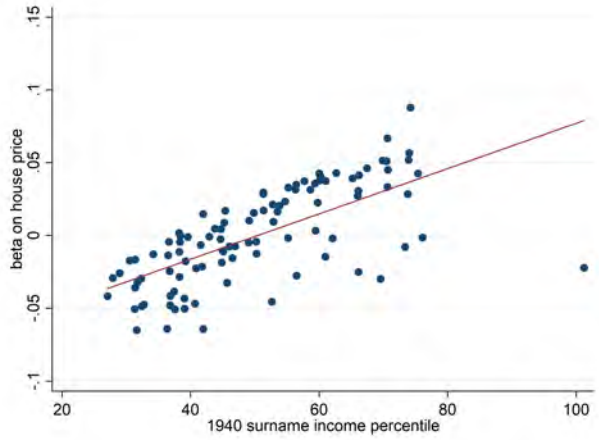

(F) 1988-2013 sample (1988-2002): White share (residual)

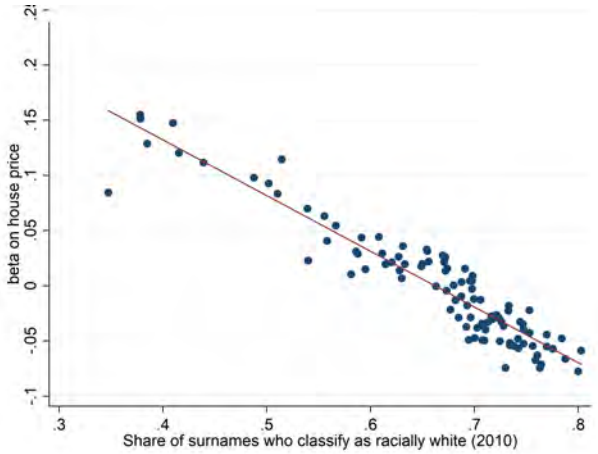

(G) 1988-2013 sample (1988-2002): 1940 income (residual)

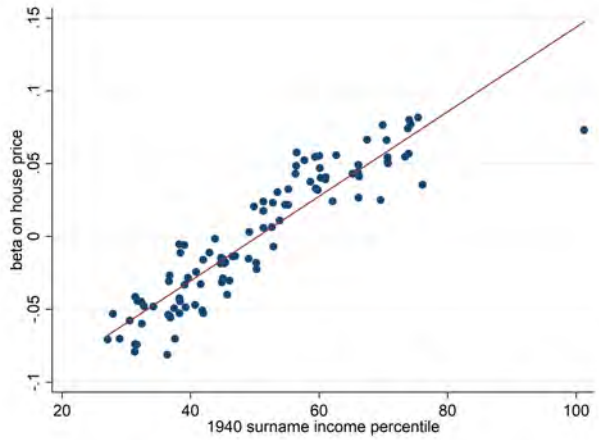

FIGURE 4: Conversion to return differential

These figures convert the estimated housing quantity elasticity to house price by surnames' associated 1940 income percentiles, to return differential by the corresponding wealth percentiles today. Panel (a) plots $\frac{d\log q_{it}}{d\log P_t}$ estimated from the holdings panel for each income-percentile group (by income in 1940 Census), against those percentile groups' average primary-residence value in 2012, conditional on owning. Each dot represents a percentile group. The plotted relationship can be viewed as a "second stage" of quantity elasticity against wealth level as proxied for using home value, with surnames as the instruments.

Panel (b) is estimated using the 2013 Survey of Consumer Finances (SCF), in order to map primary-residence value to the corresponding place in the wealth distribution, using an Engel curve argument. Panel (b) plots average log home value conditional on owning primary residence, against the net worth percentile. Estimation yields:

$$\begin{aligned} E\left[\log \text{ home value} \mid \text{own}\right] &= 0.026 \text{ net worth percentile } + 10.408 \\ &\equiv f\left(\text{net worth percentile}\right) \end{aligned}$$

Panel (c) plots imputed return differential (relative to population average) against imputed net worth percentile today. Each dot represents a percentile group defined by surnames' associated 1940 income, as with panel (a). primary-residence value in 2012-2013 is converted to net worth percentile using $f^{-1}\left(E\left[\log \text{ home value} \mid \text{own}\right]\right)$. Elasticity $\frac{d\log q_{it}}{d\log P_t}$ is converted to return differential using

$$\text{return differential}_i = -var\left(E_t dR_t\right)\left(\tilde{b}\frac{\overline{D}}{\overline{P}}\right)\overline{\theta}\left(\frac{d\log q_{it}}{d\log P_t}\right)$$

with the following coefficients from Cochrane (2011): $var\left(E_t dR_t\right) \approx (0.0546)^2$, $\tilde{b} \approx 3.8$. I further use $\frac{\overline{D}}{\overline{P}} \approx \frac{1}{16}$ and $\overline{\theta} \approx 1$. All regressions are estimated using the 1998-2013 CoreLogic sample covering roughly 60% of the US population.

(A) Beta vs. primary-residence value

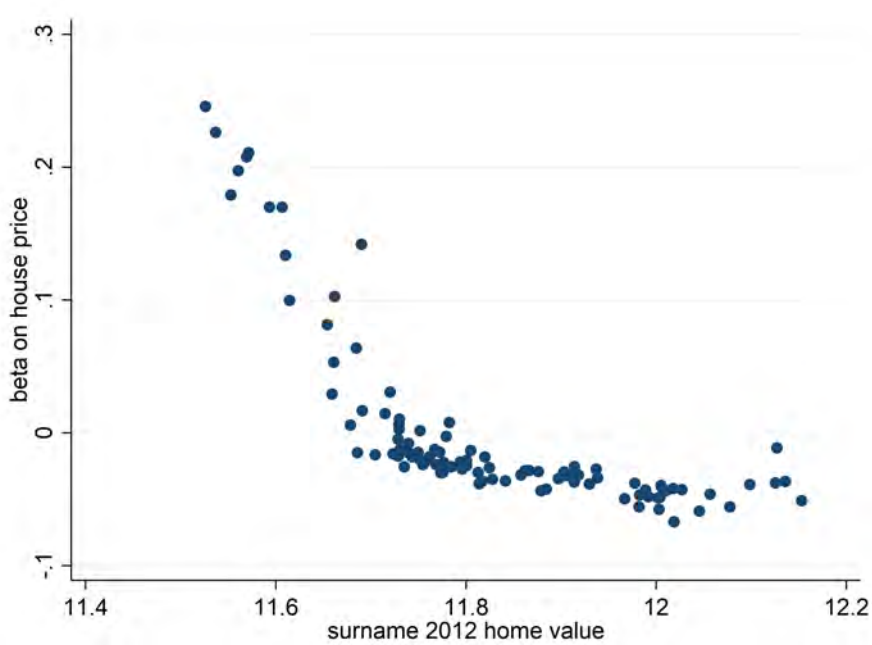

FIGURE 4: Conversion to return differential (continued)

(B) SCF: primary-residence value vs. wealth percentile

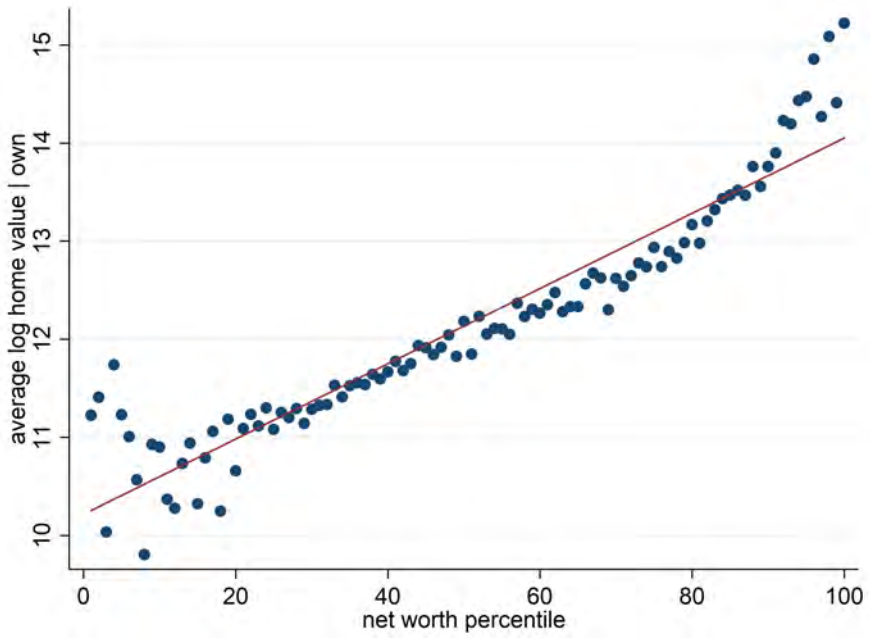

(C) Return differential vs. wealth percentile

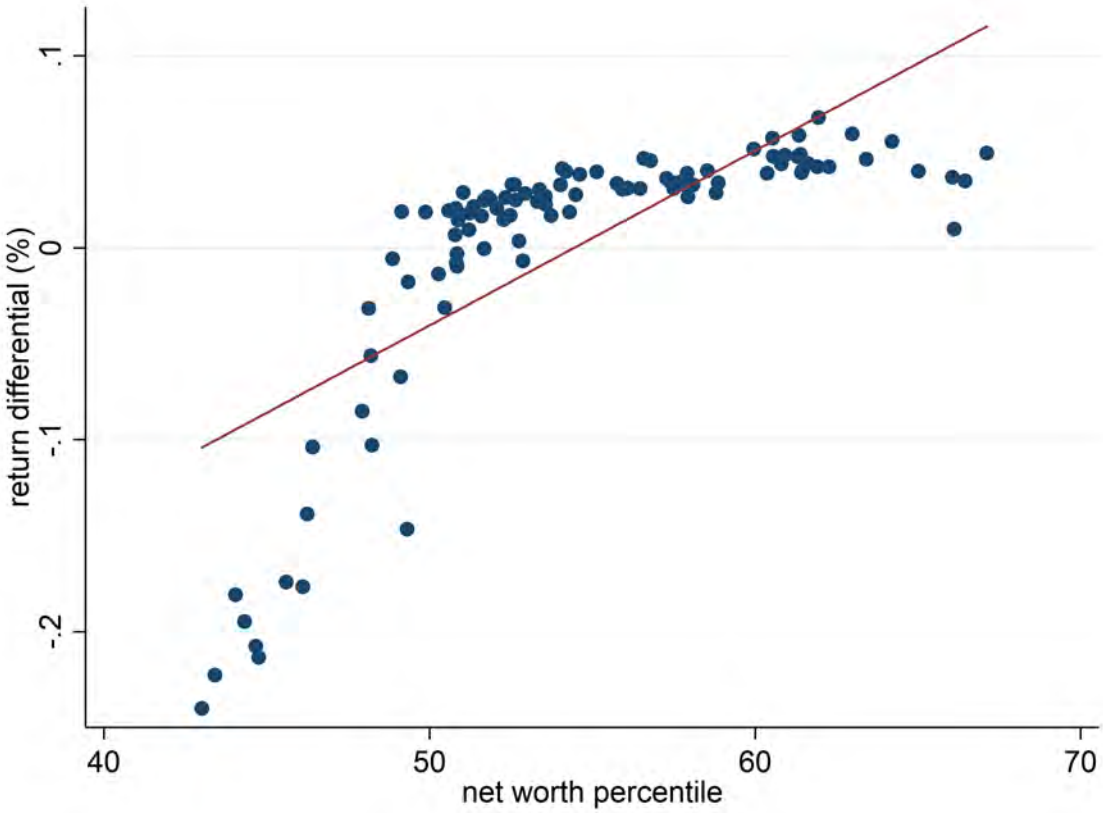

FIGURE 5: Elasticity gradient by local cyclicality

This figure plots a bin scatter at the county-level. For each county k, it plots the quantity-to-price elasticity gradient $\tilde{\beta}_k$ against business-cycle loading δ_k. $\tilde{\beta}_k$ are estimated from:

$$\log q_{ikt} = \tilde{\beta}_k \left(\log P_{kt} \times 1940 \text{ income percentile}_i\right) + \alpha_{ik} + \alpha_{kt} + \gamma_{ik} t + \xi_{ikt}$$

where q_{ikt} is the number of real estate properties located in county k, held by individuals of surname i in year t, and P_{kt} is the house-price index in county k in year t. The estimate $\tilde{\beta}_k$ captures the extent to which poorer households hold properties procyclically in county k. δ_k are estimated from:

$$\Delta \log y_{kt} = \delta_k \Delta \log Y_t + \nu_{kt}$$

where y_{kt} is the per-capita income in county k in year t, and Y_t is the national per-capita income in year t. δ_k captures the cyclicality of the local economy in county k. All regressions are estimated using the 1998-2013 CoreLogic sample covering roughly 60% of the US population.

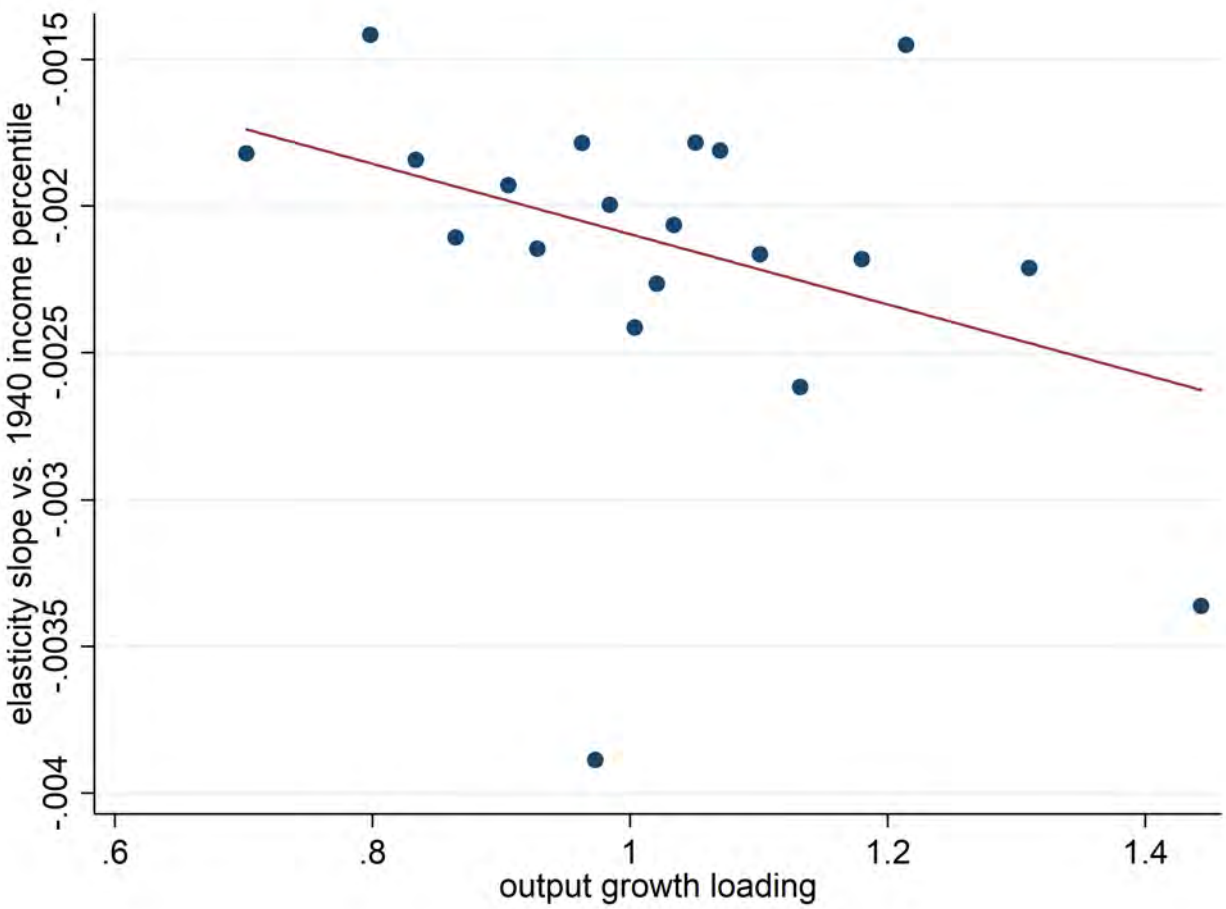

FIGURE 6: Wealth inequality level: Coefficient of variation from zip code data

Panel (a) plots Core-based Statistical Area (CBSA)-level coefficient of variation of asset in 2012 against the area's "income loading," controlling for the CBSA-level coefficient of variation of wage income. Income loadings have been calculated at the county-level by regressing changes in county-level log per-capita income on changes in aggregate log per-capita income, using data from the Bureau of Economic Analysis 1969-2015. Coefficients of variations in asset, wealth and wage have been calculated for CBSAs using zip code-level variation. Wage comes directly from "Salaries and wages" in the IRS Statistics of Income. Asset and net worth are imputed using capital income from the IRS Statistics of Income, capital income capitalization factors from Saez and Zucman (2016), housing ownership from CoreLogic assessor records, and zip code-level debt stocks from Equifax.
Panel (b) plots the CBSA-level coefficient variation of net worth against the area's income loading, again controlling for the CBSA-level coefficient of variation of wage income.
In panel (c), CBSA's are grouped by their coefficient of variation of zip code-level adjusted gross income from the IRS Statistics of Income, computed as above for 2012. It plots the household-level coefficient of variation of household income from the 2008-2012 5-year American Community Survey, accessed via IPUMS. The two measures are both meant to measure household-level inequality in income: the x-axis variable uses zip code-level data to compute; the y-axis variable is based on sample of household-level data.

(A) Net worth coefficient of variation (CBSA)

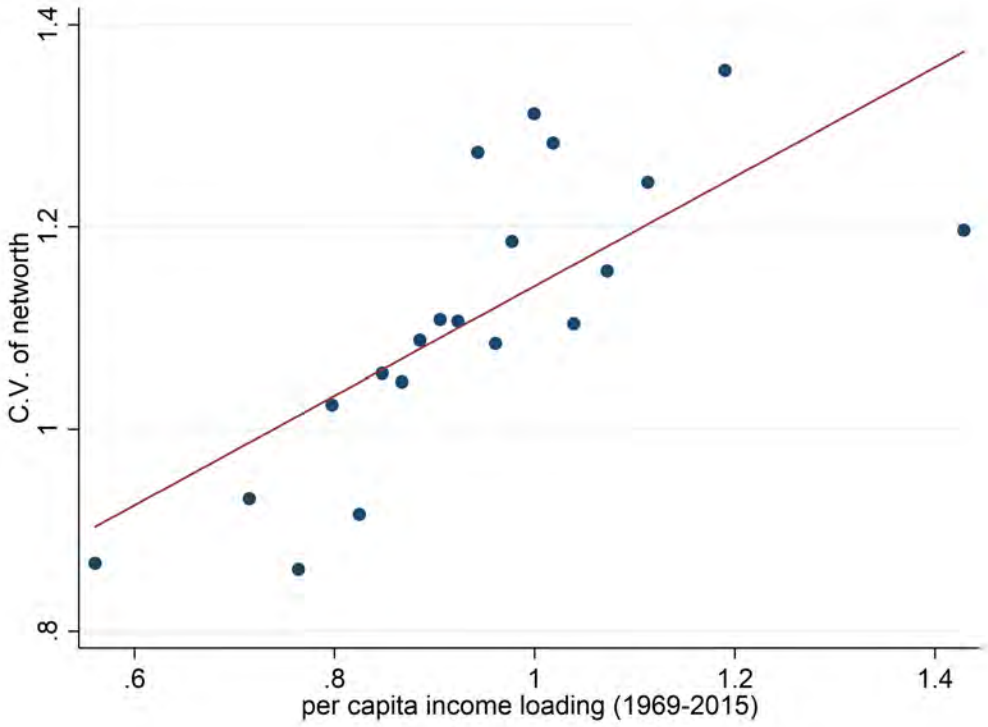

FIGURE 6: Wealth inequality level: Coefficient of variation from zip code data (continued)

(B) Asset coefficient of variation (CBSA)

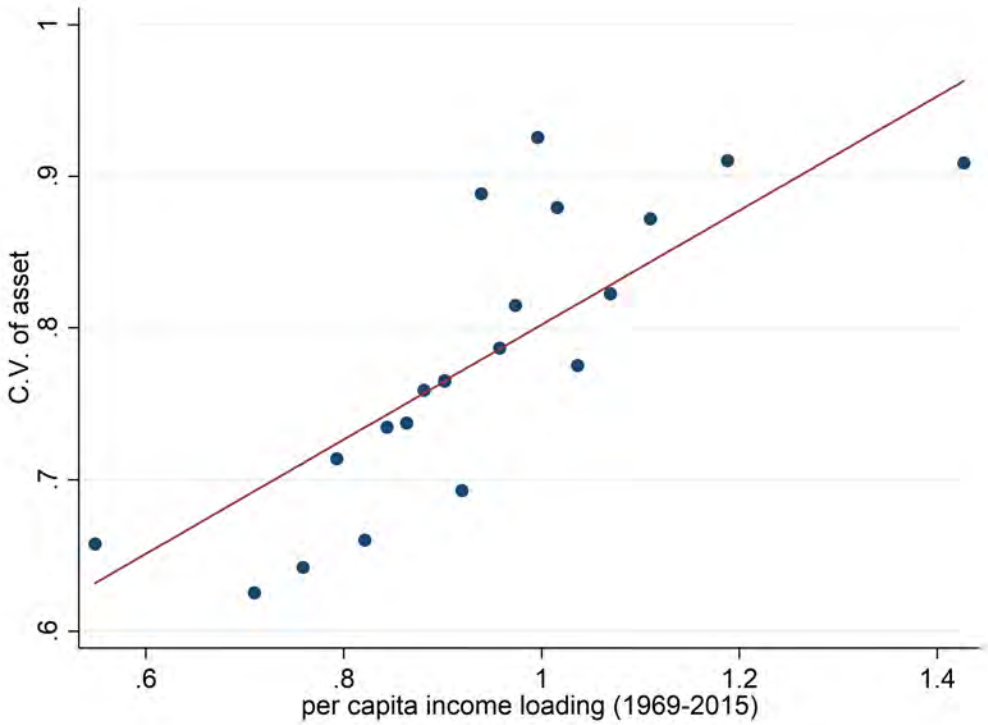

(C) Household income coefficient of variation: IPUMS vs. IRS

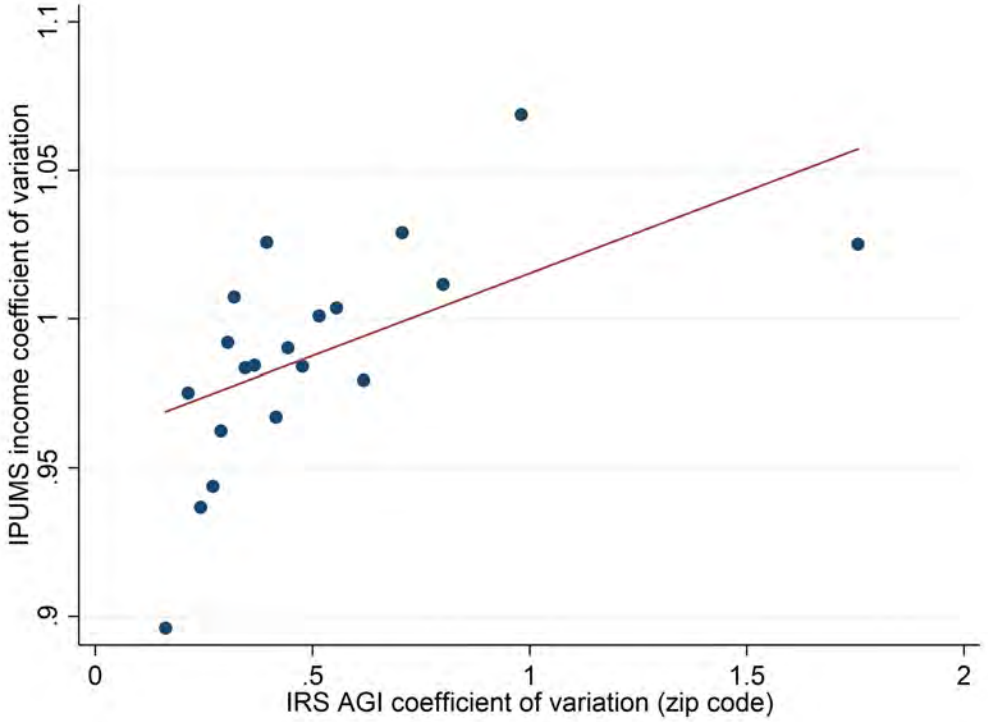

TABLE 1: Validation of surname-based historical income against Census 2000

These regressions validate the use of surnames to predict housing owners' wealth. From the 1998-2013 sample of CoreLogic data with owners matched via surnames to their average 1940 income, I take the single-year data for 2000. I then average the average-1940-incomes by the zip code where the property is located and the zip code where the owner lives. Each regression relates the zip-code-level averages from the CoreLogic data to the corresponding zip-code-level data from the 2000 Census. Columns (1) and (2) are for owner-occupied housing only: they regress the log-median-household-income from Census 2000 against CoreLogic's 1940 income. Column (2) includes county fixed effects. The data for columns (1) and (2) are at the zip-code-level. Columns (3) and (4) are for non-owner-occupied housing only: they regress the log median household income of the zip code where the owners live (i.e., the mailing address zip code) against CoreLogic's 1940 income, controlling for the median income of the zip code where the property is located. Column (4) includes county fixed effects for both the property site and the owners' residential area. The data for columns (3) and (4) are at the zip code of property site $\times$ zip code of owner's residence level, and standard errors are clustered by the zip code of owner's residence. Columns (5), (6) and (7) include both owner-occupied and non-owner-occupied housing, and for each tenure status there is a separate variable for CoreLogic's 1940 income. Columns (5) and (6) regress the log median household income of the zip code of the property site from Census 2000 against CoreLogic's 1940 income, with separate variables for owner-occupants and investor-occupants. Column (6) includes county fixed effects. The data for columns (5), (6) and (7) are at the property-site-zip-code level. Column (7) runs the same specification as column (6), but with the average log home value from Census 2000 in the property-site-zip-code as the dependant variable.

	Census income		owner residence zip income		property-site income		home value
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1940 log wage	1.633***	1.824***	0.169***	0.071***			
	(0.036)	(0.061)	(0.005)	(0.002)			
site area log income			0.192***	0.071***			
			(0.003)	(0.002)			
1940 log wage (not own)					0.586***	0.882***	1.587***
					(0.112)	(0.120)	(0.180)
1940 log wage (owner occupants)					1.162***	1.236***	1.423***
					(0.110)	(0.113)	(0.163)
Constant	-0.535*		7.588***		-1.324***		
	(0.247)		(0.058)		(0.230)		
Adjusted R^2	0.216	0.481	0.046	0.460	0.226	0.490	0.696
county FE		O		O		O	O
# of clusters			28271	28099			
Observations	9878	9727	1796362	1793999	9808	9666	9654

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

TABLE 2: Race decomposition

(A) Return (imputed from "beta"; unit in percent)

	1998-2013		1988-2013		1988-2002	
	(1)	(2)	(3)	(4)	(5)	(6)
wealth percentile (2013)	0.014***	-0.003***	0.020***	-0.001*	0.005***	-0.014***
	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Asian share		-0.003		0.018		0.146*
		(0.034)		(0.044)		(0.063)
Black share		-0.170***		-0.153***		-0.246***
		(0.009)		(0.012)		(0.025)
Hispanic share		-0.419***		-0.544***		-0.477***
		(0.006)		(0.008)		(0.016)
Adjusted R^2	0.152	0.554	0.090	0.299	0.001	0.015
Observations	119420	119420	118667	118667	117835	117835

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

(B) Average primary residence value ("first stage")

	2013 wealth percentile		1940 income
	(1)	(2)	(3)
1940 income	15.480***	10.207***	
	(0.193)	(0.152)	
Asian share		13.350***	-0.155***
		(0.756)	(0.031)
Black share		-10.236***	-0.865***
		(0.260)	(0.011)
Hispanic share		-6.865***	-0.606***
		(0.142)	(0.014)
Adjusted R^2	0.504	0.662	0.462
Observations	119432	119432	119531

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

TABLE 3: Level of wealth inequality

For county c in CBSA m,

$$\text{CV}_m = \beta \text{income loading}_c + \gamma \text{wage CV}_m + \Gamma X_c + \varepsilon_c$$

where X_c includes log population size and log house price level. Income loadings have been calculated at the county-level by regressing changes in county-level log per-capita income on changes in aggregate log per-capita income, for 1969-2015. Coefficients of variations in asset, wealth and wage have been calculated for CBSAs using zip code-level variation. Wage comes directly from "Salaries and wages" in the IRS Statistics of Income. Asset and net worth are imputed using capital income from the IRS Statistics of Income, capital income capitalization factors from Saez and Zucman (2016), housing ownership from CoreLogic assessor records, and zip code-level debt stocks from Equifax. All standard errors are clustered at the CBSA-level.

(A) CBSA coefficient of variation

	networth c.v. (2012)			with avg cap income 2003-2012	
	(1)	(2)	(3)	(4)	(5)
beta: income per cap	0.747***	0.541***	0.804***	0.556**	0.651***
	(0.151)	(0.110)	(0.177)	(0.197)	(0.188)
wage c.v.		1.489**	1.317**	1.181	1.149*
		(0.452)	(0.473)	(0.608)	(0.523)
log population			0.070	0.282	0.265*
			(0.044)	(0.145)	(0.108)
log house price			-0.330*	0.569	0.523
			(0.135)	(0.592)	(0.594)
Constant	0.406***	0.078	0.780	-5.395	
	(0.113)	(0.149)	(0.734)	(4.230)	
Adjusted R^2	0.032	0.312	0.321	0.213	0.306
state FE					O
# of CBSA	881	881	649	649	649
Observations	1707	1707	1092	1092	1091

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

TABLE 3: Level of wealth inequality (continued)

For zip code z in county c,

$$\log(Y)_z = \beta\left[\log(\text{wage})_z \times \text{income loading}_c\right] + \Gamma_1\left[\log(\text{wage})_z \times X_c\right] + \Gamma_0 X_{cz} + \delta\text{income loading}_c + \gamma\log(\text{wage})_z + \varepsilon_z$$

where X_{cz} includes log population size and log house price level. Income loadings have been calculated at the county-level by regressing changes in county-level log per-capita income on changes in aggregate log per-capita income, for 1969-2015. Wage, asset and net worth vary at the zip code level. Wage comes directly from "Salaries and wages" in the IRS Statistics of Income. Asset and net worth are imputed using capital income from the IRS Statistics of Income, capital income capitalization factors from Saez and Zucman (2016), housing ownership from CoreLogic assessor records, and zip code-level debt stocks from Equifax. All standard errors are clustered at the county-level.

(B) Zip code wealth-wage elasticity

	log networth per capita ('12)			with avg cap income '03-'12		
	(1)	(2)	(3)	(4)	(5)	(6)
income beta × wage		0.809***	0.769**	0.689**	0.735***	0.958***
		(0.195)	(0.244)	(0.223)	(0.204)	(0.197)
population × wage			0.056*	0.057**	0.065**	0.066**
			(0.028)	(0.022)	(0.023)	(0.023)
house price × wage			0.093	0.228	0.151	0.565
			(0.221)	(0.207)	(0.220)	(0.316)
income beta		-8.413***	-7.922**	-7.074**	-7.457***	
		(2.041)	(2.604)	(2.381)	(2.172)	
log population			-0.600*	-0.674**	-0.763**	-0.739**
			(0.302)	(0.238)	(0.249)	(0.241)
log house price			-0.245	-1.771	-1.170	
			(2.426)	(2.258)	(2.403)	
log wage	1.399***	0.573**	-0.397	-1.135	-0.864	-2.904
	(0.018)	(0.201)	(1.181)	(1.092)	(1.145)	(1.599)
Constant	-3.792***	4.778*	11.489	20.478		
	(0.188)	(2.105)	(12.938)	(11.931)		
Adjusted R^2	0.267	0.270	0.300	0.316	0.377	0.502
state FE					O	
county FE						O
# of counties		2854	1196	1196	1196	1158
Observations	21723	21517	15413	15400	15400	15362

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix

A Additional Theoretical Results

A.1 Theoretical Ambiguity the Quantity Elasticity versus Wealth Level

Working with a standard household problem, I show that theoretically rich or poor households could exhibit more procyclical housing ownership. Mainly, I distinguish poor households from rich households by more frequently binding credit constraints, and consider the comparative static in quantities that vary over the cycle.[50]

They are meant to demonstrate the theoretically ambiguous prediction for $\frac{d\log q}{d\log P}$ and are not a comprehensive analysis of the cross section of trading behaviors. In addition to demonstrating the ambiguous theoretical prediction, the model exercise serves two more roles: (1) It explains how the panel data on housing ownership map to the drivers and incidence of business cycles, and (2) the implication of wealth inequality relies on the pattern happening through time - assessing if it is true is partly an empirical exercise, but knowing if forces considered universal can generate the selling pattern is also useful.

I organize the discussion around a simple two-asset consumption-savings model of an individual. Households maximize

$$\sum_{\tau=0}^{\infty} \beta^{\tau} u\left(C_{t+\tau}\right)$$

subject to a standard two-asset budget constraint and a borrowing constraint on the risk-free asset:

$$Y_t + \underbrace{B_t + H_t\left(P_t + D_t\right)}_{W_t} = C_t + H_{t+1}P_t + B_{t+1}Q_t$$

$$B_{t+1} \geq \underline{B}$$

Define wealth W as cum-dividend wealth. Let ξ_t denote the multiplier on the borrowing constraint, scaled by marginal utility and bond price, that is, $\xi_t \equiv \frac{\text{multiplier on borrowing constraint}_t}{Q_t\lambda_t} \geq 0$, because both

[50] Households of different wealth levels have many possible differences that could generate the observed transacting behaviors. For example, (1) except for the top 2%, lower-income households have more cyclical income (Guvenen et al. (2014)), (2) credit-supply fluctuations disproportionately affect lower-income households (Mian et al. (2017a); Mian et al. (2017b)), (3) lower-income households may be more myopic or extrapolative, or (4) lower-income households may lack market knowledge and importantly timing skill, among other possible differences. Here, I show that even with just the difference in the tightness of credit constraint, theoretical predictions are ambiguous.

$Q_t > 0$ and $\lambda_t > 0$ always. ξ_t is the scaled shadow cost of the borrowing constraint and only enters the Euler equation for the risk-free asset. Derivations of key equations closely follow Viceira (2001) and Campbell et al. (2002),[51] and can be found in the Appendix.

First consider the Merton benchmark. Define the housing share as $\theta_t \equiv \frac{H_{t+1}P_t}{H_{t+1}P_t + B_{t+1}Q_t}$. To keep θ_t at a fixed level, those with a higher level of θ_t will buy in boom. Their W_t increases disproportionately more, and thus they have to acquire more H_{t+1}. Because the average risky share is higher for richer households, this simplest benchmark shows clearly that from valuation shocks alone, higher-wealth households should be the ones exhibiting more procyclical net purchases.

Combining the two Euler equations yields

$$\gamma Cov\left(c_{t+1} - c_t, r_{1,t+1}\right) \approx E\left[r_{1,t+1} - \left(r_{f,t} + \xi_t\right)\right] + \frac{1}{2}Var\left(r_{1,t+1}\right)$$

where $r_{1,t+1} \equiv \frac{P_{t+1}+D_{t+1}}{P_t}$ for housing and $r_{f,t} \equiv \frac{1}{Q_t}$ for the risk-free bond.[52]

A reduction in ξ_t from the relaxed borrowing constraint acts like a reduction in the risk-free rate for portfolio choice.[53] This lower effective risk-free rate is a price effect that applies to all risky assets. Note that this lower effective risk-free rate due to reduction in ξ_t is different from a shock to β_t, which would affect all assets equally. An effect on portfolio allocation via the effect on consumption would arise, but the direct price-effect channel is gone. This distinction is important: Papers in macroeconomics often take β_t shocks as a reduced-form way to model incomplete market demand shocks (Werning (2015); Beraja et al. (2016)). With multiple assets, this reduced form would miss the price-effect and portfolio-choice channel. In fact, most macroeconomic models would miss this price-effect, and only capture the forces related to consumption (i.e., the direct effect of β and consumption-income co-movement).

Constraints would also affect the consumption rule. The Euler equation for consumption growth in terms of portfolio return is

$$\left[\log\beta + (1-\theta_t)\,\xi_t\right] - \gamma E_t\left[c_{t+1} - c_t\right] + E_t\left[r_{p,t+1}\right] + \frac{1}{2}Var\left[r_{p,t+1} - \gamma(c_{t+1} - c_t)\right] \approx 0$$

[51] As opposed to the more common numerical solution route, these papers use Taylor approximations to solve for approximate solutions, while taking care to expand properly inside expectations to preserve the precautionary and asset-pricing forces. The positive probability of a zero income and a permanent retirement state allows for strictly positive non-human wealth always. Strictly positive non-human wealth allows for log-linearization. Because I primarily study homeowners, I make a similar assumption of a permanent zero-income possibility and study $W > 0$.

[52] Note that unlike w_t and y_t, and like c_t, ξ_t too is an endogenous policy. These equations are not full solutions.

[53] This effect is similar with a collateral constraint as well, because a borrowing limit scales risky asset holdings by a loan-to-value ratio given by some parameter less than 1, that is, $\phi < 1$.

where $\theta_t \equiv \frac{H_{t+1}P_t}{H_{t+1}P_t+B_{t+1}Q_t}$. With higher ξ_t, the agent simultaneously chooses a lower c_t and lower θ_t, and consequently a lower $E_t r_{p,t+1}$.

For the clearest comparison, I make the following simplification: Assume $\xi_{it} = \xi_i$ for household type i. Denote "poorer households" as those households with higher ξ_i. With this simplification, the solution takes a form similar to that in Viceira (2001). With this simplification, household consumption c_t and portfolio share in risky asset θ_t are given by

$$c_t - y_t = b_0 + b_1 (w_t - y_t)$$
$$\theta_t = \frac{\mu - (r_f + \xi) + \frac{1}{2}\sigma_u^2 - (1 - b_1)\sigma_{yu}}{b_1 \sigma_u^2}$$

where $b_1 = \frac{\rho_w - 1}{\rho_c} > 0$ and $b_0 = -\frac{\text{constant}+(1-\theta)\xi}{b_1 \rho_c}$ with $b_0 < 0$ and $\frac{db_1}{d\xi} > 0$.

Note that b_1 behaves like the marginal propensity to consume (MPC) out of liquidity in the consumption-savings literature. Shocks to labor income y_t are all permanent, so y_t represents the permanent income from which households would borrow. w_t is liquid assets that they can access immediately. Hence, the hedging term implies that if MPC is high, permanent income shocks translate less to consumption growth ($(1 - b_1)$ is lower), and so hedging demand is lower. With higher ξ, b_0 is even more negative and consumption is suppressed. Hence, $\frac{\rho_c}{\rho_w} \downarrow$, and $b_1 \uparrow$. We find that the elasticity of consumption to liquid wealth is higher (again, w_t is like cash on hand in this set-up, because y_t shocks are permanent). Higher b_1 has two opposite effects. On the one hand, human capital is worth less because one cannot borrow against it (i.e., the b_1 in denominator). On the other hand, innovations to permanent income "matter less" in that they translate less to consumption growth, so hedging demand is lower. Given that income shocks are positively correlated with housing returns, this force increases the demand for housing.

Credit supply affects demand for housing via three channels. Taking comparative statics with respect to ξ,

$$\frac{d\theta}{d\xi} = \underbrace{-\frac{1}{b_1\sigma_u^2}}_{<0} + \underbrace{\frac{db_1}{d\xi}}_{>0} \left[\underbrace{-\frac{\theta}{b_1^2\sigma_u^2}}_{<0} + \underbrace{\frac{\sigma_{yu}}{b_1\sigma_u^2}}_{>0} \right]$$

First, note in the combination of the two Euler equations that $\xi_t > 0$ acts as if the risk-free return is higher for those households; a pure price effect would lead to allocation toward more positive risk-free allocation and toward less of all risky assets. Second, with higher ξ and hence higher b_1, future labor

income is less valuable, and thus households behave as if they are more risk averse (i.e., the b_1 in the denominator). Third, because future labor income is less valuable, its covariance with asset prices (i.e., background risk) is also less of a concern. The last force moves in the opposite direction of the other two forces. These three forces apply to all risky assets.[54]

To understand who buys in booms and busts, I consider comparative statics on θ_t, because I have already assumed common price dynamics. Given a one-time permanent shock to a parameter of the model, I consider whose θ_t would increase more.

First, consider a procyclical credit supply. Suppose a looser credit constraint lowers ξ_i more for households with higher levels of ξ_i. With σ_{yu} not too high, lower ξ would translate to higher demand for housing for those households for whom constraints had bound more. Therefore, procyclical credit supply naturally leads to more procyclical net purchase behavior for poorer households who are closer to borrowing and collateral constraints. The relaxation of credit supply in booms can come from market forces or from government policy.

Next, consider a procyclical perceived expected return on risky asset $\tilde{\mu}$. This comparative applies to factually higher return in cycle downturns, as well as to extrapolated higher return in booms in behavioral models. Comparative static yields:

$$\frac{\partial \theta}{\partial \tilde{\mu}} = \frac{1}{b_1 \sigma_u^2}$$

A perceived return would predict that either no transfer of housing would occur along the wealth distribution or a transfer would occur from poor to rich households, who have better capacity to capitalize on the higher return expectation.[55][56]

[54]There are additional forces not in this simple example that would apply specifically to housing. First, see from the portfolio return Euler equation that $\xi_t >$ acts as if β is higher. Household would tilt towards savings from consumption, and consequently owner-occupied housing demand would be lower. This force is specific to durable consumption, of which housing is the primary example. Second, illiquid assets are discounted due to illiquidity, and more difficult self-insurance would increase the cost of illiquidity (Longstaff (2001)). This force is specific to housing as an illiquid asset (i.e., asset with large transaction costs).

[55]Even with an expected housing return shock, an increase in house prices will loosen borrowing/collateral constraints. This effect would be omitted in partial equilibrium. However, even with a relaxation in some credit constraint from a house-price increase, the dominant effect is still the rise in house price (a cost) from the perspective of a buyer. Given a typical loan-to-value constraint $b \geq \phi ph$, for example, an increase in p would increase b only by $\phi < 1$ fraction. This general-equilibrium effect cannot overturn the comparative static. In fact, this discussion highlights a key reason for looking at net housing as opposed to gross borrowing. A causal mechanism going from house price to a looser constraint can increase borrowing by the poor, yet will not get housing to transfer to them on net.

[56]This implication is also seen in Kaplan et al. (2017), in which a price-belief shock by itself reduces homeownership, because owning becomes more expensive than renting. That prices rise so that poor households switch out of homeownership is equivalent to saying poor households experience a lower demand increase than richer households. Endogenous price movement inherits directly the excess demand in a partial-equilibrium set-up.

Therefore, depending on which forces are stronger and what changes are happening at cyclical frequencies, in theory rich or poor households could be holding housing procyclically.

B Derivations

B.1 Wealth Inequality

Suppose

$$\log \mu_t^k = a^k + b^k \log \frac{D_t^k}{P_t^k}$$
$$\log \theta_{it}^k = \log P_t^k + \log Q_{it}^k - \log W_{it}$$

Then

$$\begin{aligned}
cov\left(\theta_{it}^k, \mu_{it}^k\right) =& E\left[\theta_{it}^k \mu_{it}^k\right] - E\left[\theta_{it}^k\right] E\left[\mu_{it}^k\right] \\
=& E\left[\theta_{it}^k\right] E\left[\mu_{it}^k\right] \left[\exp\left(cov\left(\log \theta_{it}^k, \log \mu_t^k\right)\right) - 1\right] \\
\approx& E\left[\theta_{it}^k\right] E\left[\mu_{it}^k\right] cov\left(\log \theta_{it}^k, \log \mu_t^k\right) \\
=& E\left[\theta_{it}^k\right] E\left[\mu_{it}^k\right] \left[cov\left(\log P_t^k, \log \mu_t^k\right) + cov\left(\log Q_{it}^k, \log \mu_t^k\right) - cov\left(\log W_{it}, \log \mu_t^k\right)\right]
\end{aligned}$$

Focus on middle term:

$$\begin{aligned}
acov\left(\theta_{it}^k, \mu_{it}^k\right) \equiv& E\left[\theta_{it}^k\right] E\left[\mu_{it}^k\right] cov\left(\log Q_{it}^k, \log \mu_t^k\right) \\
=& -b^k E\left[\theta_{it}^k\right] E\left[\mu_{it}^k\right] cov\left(\log Q_{it}^k, \log P_t^k\right) \\
=& -b^k E\left[\theta_{it}^k\right] E\left[\mu_{it}^k\right] \frac{cov\left(\log Q_{it}^k, \log P_t^k\right)}{var\left(\log P_t^k\right)} var\left(\log P_t^k\right) \\
\approx& -\left(0.1987424\right)\left(1\right)\left(1.123208\right)\left(.1578269\right)^2 \frac{cov\left(\log Q_{it}^k, \log P_t^k\right)}{var\left(\log P_t^k\right)} \\
\approx& 0.0056 \frac{cov\left(\log Q_{it}^k, \log P_t^k\right)}{var\left(\log P_t^k\right)}
\end{aligned}$$

The assumptions are: (1) process for $x_t^k \equiv \frac{D_t^k}{P_t^k}$, (2) lognormal distributions, (3) approximation around covarinace of 0.

B.1.1 Back-of-envelope

$$\frac{\frac{W_i}{\overline{W}}}{\frac{Y_i}{\overline{Y}}} \approx \frac{1}{1-\left(\mu^i-\mu\right)\frac{1}{\left(1-c_y\right)\frac{Y}{W}}}$$
$$\approx 1+\left(\mu^i-\mu\right)\frac{W}{Y}\frac{1}{1-c_y}$$

where the last approximation is a Taylor expansion around 0.

Plug

$$-b^k E\left[\theta_{it}^k\right] E\left[\mu_{it}^k\right] \frac{cov\left(\log Q_{it}^k, \log P_t^k\right)}{var\left(\log P_t^k\right)} var\left(\log P_t^k\right)$$

into $\left(\mu^i-\mu\right)$.

B.1.2 Estimating behavioral consumption rule

Estimating consumption rule in levels is difficult, so I modify log-linearization method used in macroeconomics. Start from

$$C_i = c_w W_i + c_y Y_i$$

The log deviation of some household's consumption from $\overline{C} = c_w\overline{W} + c_y\overline{Y}$, is $\log C_i - \log\overline{C}$. First-order approximation gives,

$$\log C_i \approx \log\overline{C} + \frac{C_i}{\overline{C}} - 1$$

Plug into the level consumption policy:

$$\log C_i \approx \left(c_w\frac{\overline{W}}{\overline{C}}\right)\log W_i + \left(c_y\frac{\overline{Y}}{\overline{C}}\right)\log Y_i + K$$

where K is a constant given by $K = \log\overline{C} - c_w\frac{\overline{W}}{\overline{C}}\log\overline{W} - c_y\frac{\overline{Y}}{\overline{C}}\log\overline{Y}$.

Based on national account numbers, labor income is roughly 70% of output and consumption is roughly 70% of output, so $c_y\frac{\overline{Y}}{\overline{C}} \approx c_y$.

Estimation of the following regression in the PSID,

$$\log C_{it} = \hat{\gamma}_y \log Y_{it} + \hat{\gamma}_w \log W_{it} + \varepsilon_{it}$$

yields $\hat{\gamma}_y \approx 0.25$ and $\hat{\gamma}_w \approx 0.05$. Note that this estimation omits individuals with non-positive income or wealth. Note also that the estimated coefficients along with the consumption rule are not internally consistent. This part is a rough approximation.

B.2 Micro-foundation

Then, log Euler equations for the two assets,

$$0 = \log\beta - \gamma E_t\left[c_{t+1} - c_t\right] + E_t\left[r_{1,t+1}\right] + \frac{1}{2}Var\left[r_{1,t+1} - \gamma(c_{t+1} - c_t)\right]$$
$$0 = \log\beta - \log\left(1-\xi_t\right) - \gamma E_t\left[c_{t+1} - c_t\right] + r_{f,t} + \frac{1}{2}Var\left[-\gamma(c_{t+1} - c_t)\right]$$

where $r_{1,t+1} \equiv \frac{P_{t+1}+D_{t+1}}{P_t}$ for housing and $r_{f,t} \equiv \frac{1}{Q_t}$ for risk-free. Taking difference:

$$\begin{aligned}\gamma Cov\left(c_{t+1} - c_t, r_{1,t+1}\right) &= E\left[r_{1,t+1} - r_{f,t}\right] + \frac{1}{2}Var\left(r_{1,t+1}\right) + \log\left(1-\xi_t\right)\\ &\approx E\left[r_{1,t+1} - r_{f,t}\right] + \frac{1}{2}Var\left(r_{1,t+1}\right) - \xi_t\\ &= E\left[r_{1,t+1} - \left(r_{f,t} + \xi_t\right)\right] + \frac{1}{2}Var\left(r_{1,t+1}\right)\end{aligned}$$

for $\xi_t \geq 0$.

Log budget constraint

$$w_{t+1} - y_{t+1} \approx k + \rho_w\left(w_t - y_t\right) - \rho_c\left(c_t - y_t\right) - \Delta y_{t+1} + r_{p,t+1}$$
$$r_{p,t+1} = \theta_t\left(r_{1,t+1} - r_{f,t}\right) + r_{f,t} + \frac{1}{2}\theta_t\left(1-\theta_t\right)Var\left(r_{1,t+1}\right)$$

where $\theta_t \equiv \frac{H_{t+1}P_t}{H_{t+1}P_t + B_{t+1}Q_t}$.

Euler with portfolio return:

$$\begin{aligned}\log\beta - \gamma E_t\left[c_{t+1} - c_t\right] + E_t\left[r_{p,t+1}\right] + \frac{1}{2}Var\left[r_{p,t+1} - \gamma(c_{t+1} - c_t)\right] &= \log\left(1 - \xi_t\left(1-\theta_t\right)\right)\\ &\approx -\left(1-\theta_t\right)\xi_t\end{aligned}$$

Simplifications: Assume a fixed $\overline{\xi}$ and stationary environment with fixed expected return as in Viceira (2001). Comparative static on $\overline{\xi}$. (For notational convenience, just stick with ξ with no t subscript.) With actual time-varying expected return, have to keep track of another state variable, as opposed to the one-state set-up (in $w_t - y_t$). Comparative static should be clearest in terms of highlighting the forces. Switch to numerical sooner. For now, just assume shock to μ, i.e., higher expected return in bust as well as higher perceived house price appreciation in early 2000s are all just MIT shocks to $\mu = E\left[r_{p,t+1}\right]$. With fixed ξ, solutions are entirely the same, with a few changes in the coefficients.

Solution still takes the form:

$$c_t - y_t = b_0 + b_1\left(w_t - y_t\right)$$

$$\theta = \frac{\mu - \left(r_f + \xi\right) + \frac{1}{2}\sigma_u^2}{b_1 \sigma_u^2} - \frac{(1-b_1)}{b_1}\frac{\sigma_{yu}}{\sigma_u^2}$$

Note that b_1 is like the MPC out of liquidity we are used to in the consumption-savings literature. Shocks to labor income y_t are all permanent, so y_t represents permanent income that households would borrow from. w_t is liquid assets that they can immediately access. Hence, the hedging term implies: if high MPC, permanent income shocks translate less to consumption growth ($(1-b_1)$ is lower), and so there is lower hedging demand.

Doing the same trick with trivial inequality: $c_{t+1} - c_t = (c_{t+1} - y_{t+1}) + (y_{t+1} - y_t) - (c_t - y_t)$, and for $\gamma = 1$, arrive at

$$\begin{aligned} & b_0 + b_1 k + (1-b_1) g + b_1 E\left[r_{p,t+1}\right] - (b_1\rho_c + 1) b_0 + \left(b_1\rho_w - b_1^2\rho_c - b_1\right)(w_t - y_t) \\ & = \log\beta + (1-\theta)\xi + E\left[r_{p,t+1}\right] + \frac{1}{2}Var\left(r_{p,t+1} - (c_{t+1} - c_t)\right) \end{aligned}$$

where $g = E(y_{t+1} - y_t)$. Setting coefficients to zero,

$$b_1 = \frac{\rho_w - 1}{\rho_c}$$

$$b_0 = -\frac{\text{constant} + (1-\theta)\xi}{b_1\rho_c}$$

with $b_0 < 0$.

With higher ξ, b_0 is even more negative, and consumption is suppressed. Hence $\frac{\rho_c}{\rho_w} \downarrow$, and $b_1 \uparrow$. We get that there is higher elasticity of consumption to liquid wealth (again, w_t is like cash-on-hand in this set-up, since y_t shocks are entirely permanent). Higher b_1 has two opposite effects. On the one hand, human capital is worth less since cannot borrow against it (i.e., the b_1 in denominator). On the other hand, innovations to permanent income "matter less" in that they translate less to consumption growth, so there is lower hedging demand. Given that income shocks are positively correlated with housing returns, this force actually increases the demand for housing.

FIGURE A.1: CoreLogic samples

These figures describe the two samples from CoreLogic used in this paper. In Panel (a), counties colored in blue are included as a balanced panel in the 1998-2013 sample. These areas cover roughly 60% of the US population. In Panel (b), counties colored in blue are included as a balanced panel in the 1988-2013 sample. These areas cover roughly 25% of the US population. In Panel (c), the blue line (with scale on left axis) plots the number of counties that would be included in a balanced, consistent panel of counties in a sample that starts from a given year in the x-axis and ends in 2013. The red line plots the fraction of the US population that would be covered in each sample using concurrent population for each year, while the green line uses the 1990 county population to calculate the population share. The sample that starts in 1988 only includes a few counties but still covers roughly a quarter of the US population. The counties that appear earlier in CoreLogic are not representative also along other dimensions. Panel (d) plots the average house-price index for counties that are in the 1998-2013 CoreLogic sample (red line) and those that are not (blue line). The sample counties had bigger house price boom and bust.

(A) Counties in the 1998-2013 sample

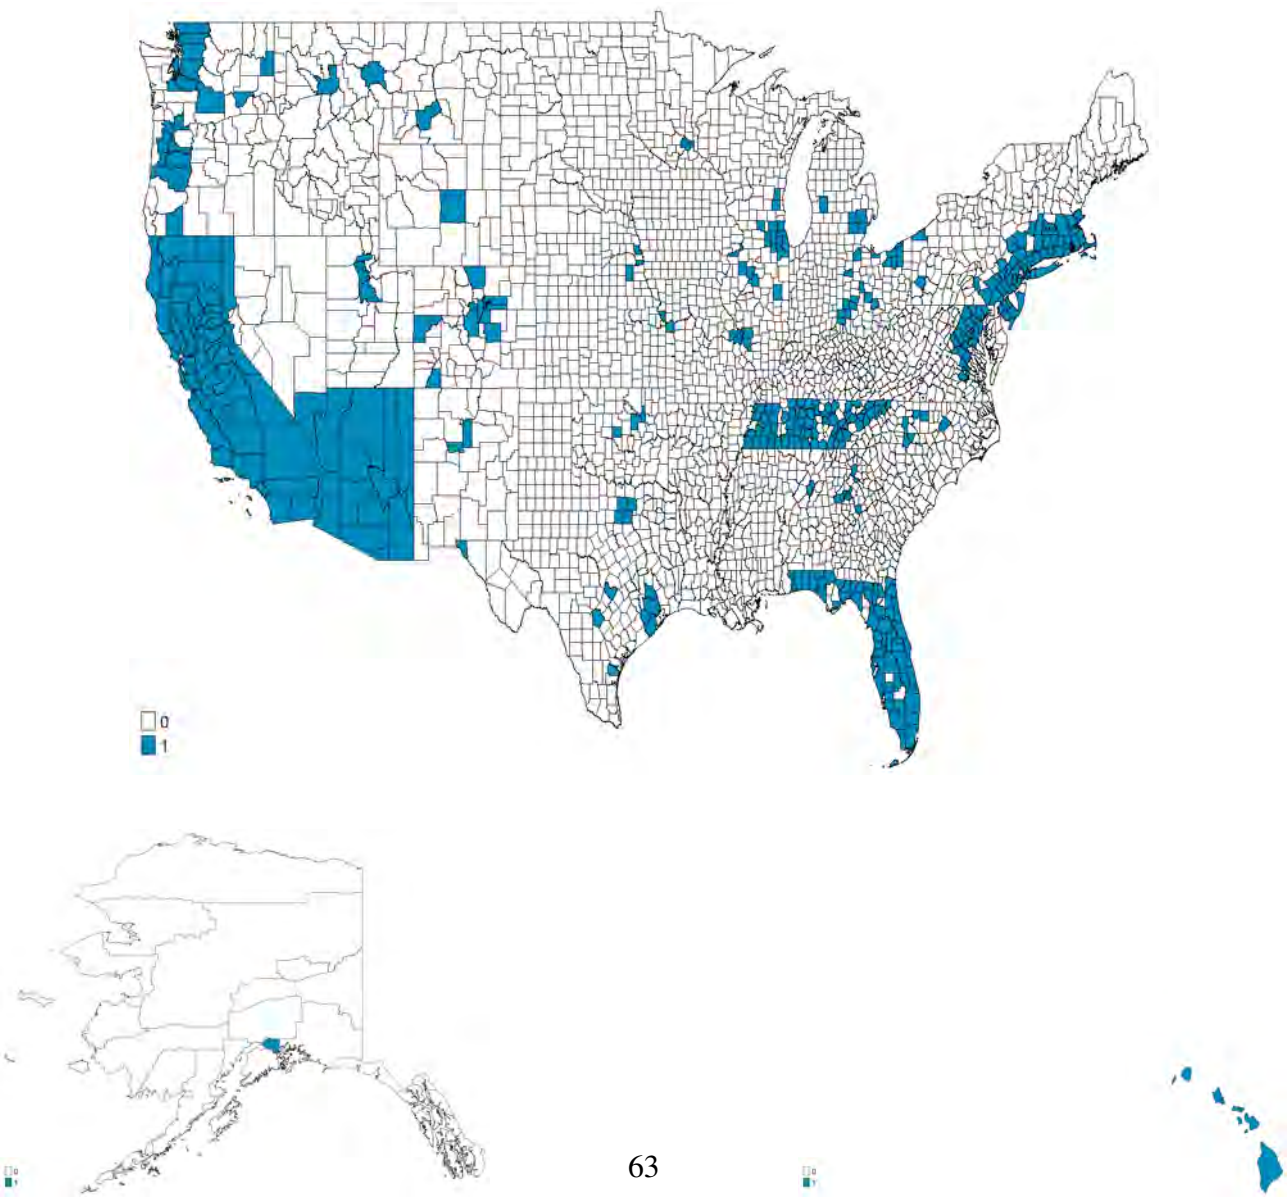

FIGURE A.1: CoreLogic samples (continued)

(B) Counties in the 1988-2013 sample

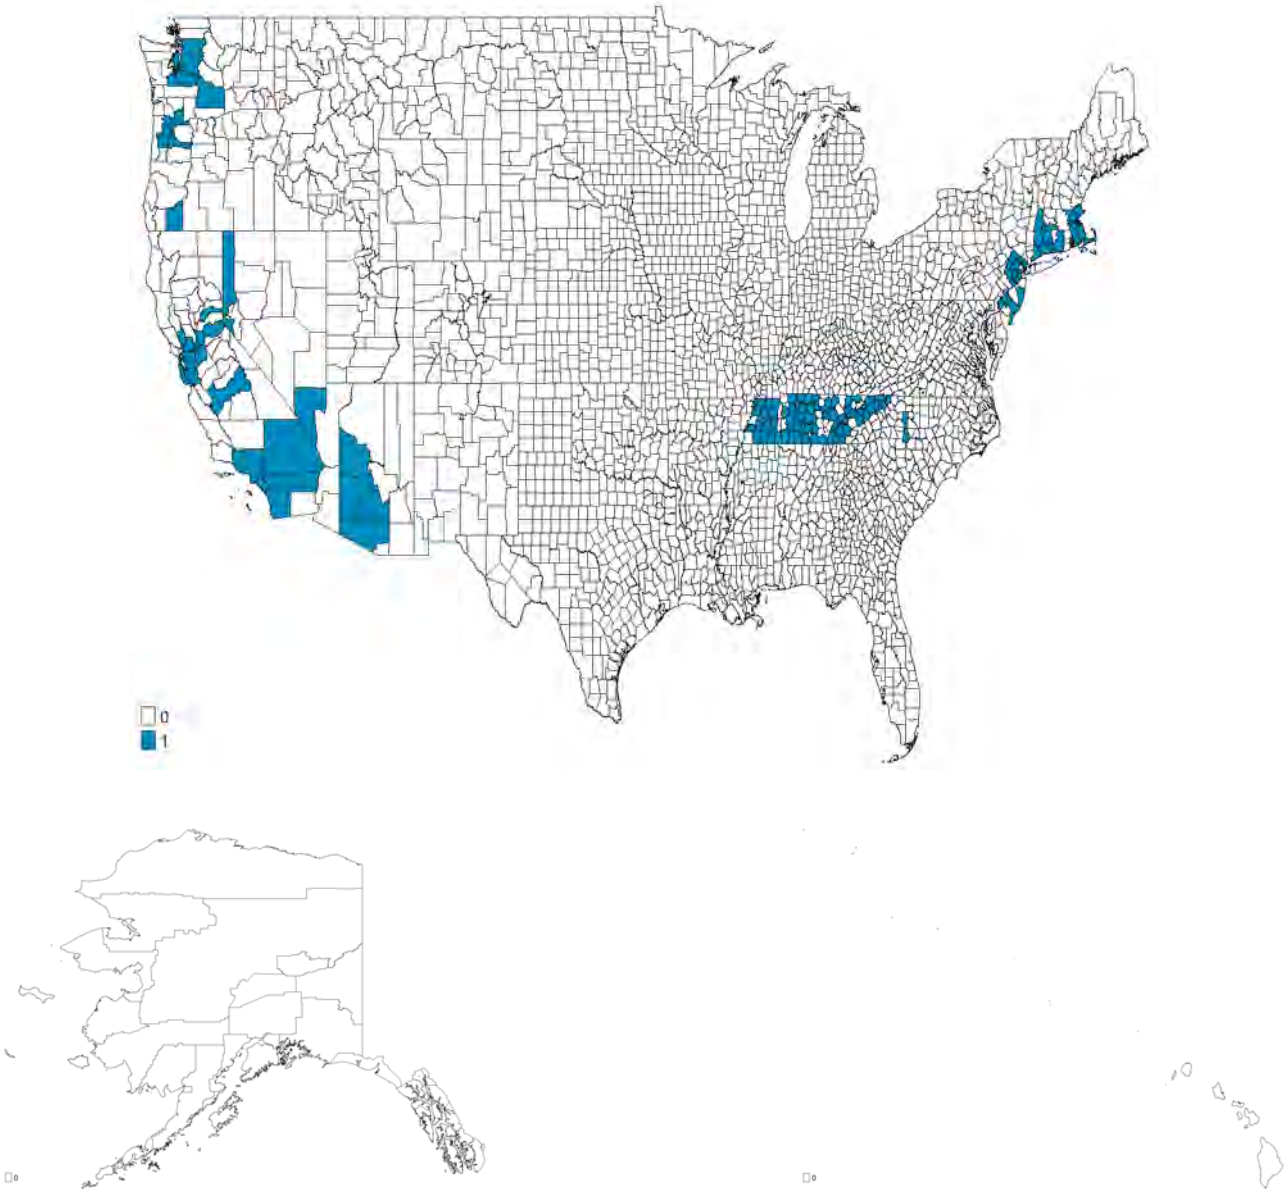

FIGURE A.1: CoreLogic samples (continued)

(C) Fraction of the US population included

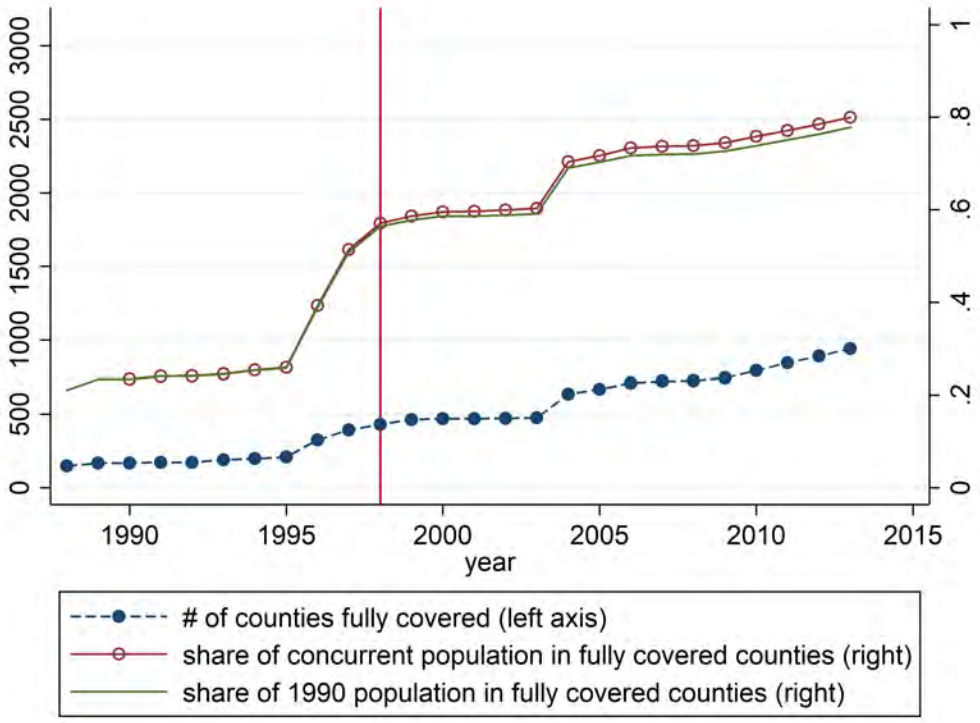

(D) House price of in-sample counties for 1998-2013

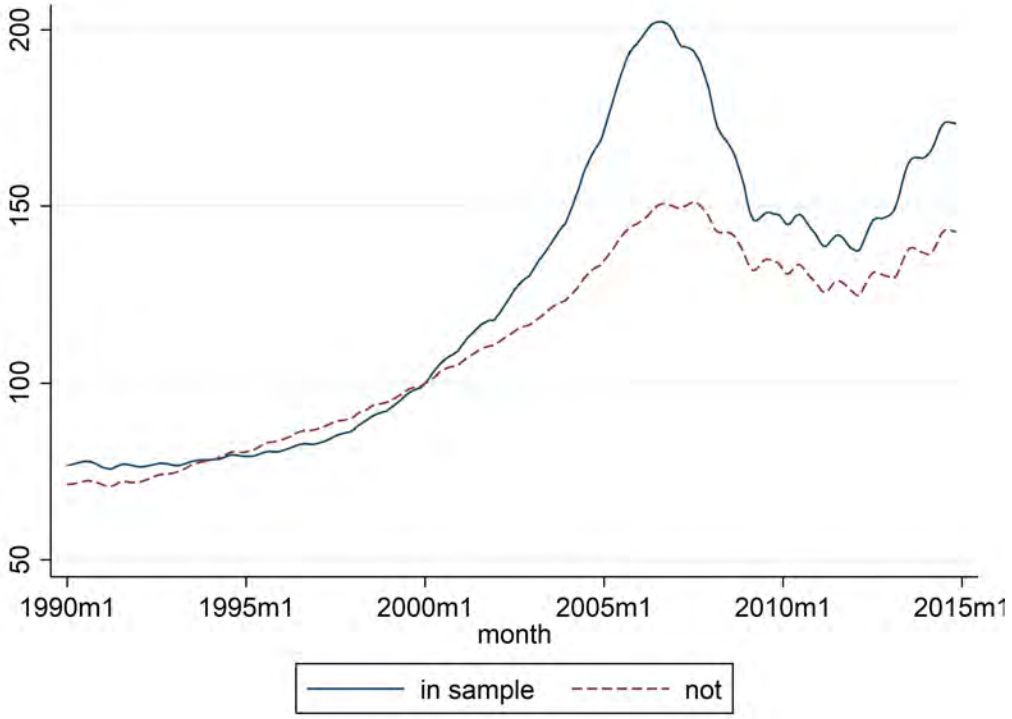

FIGURE A.1: CoreLogic samples (continued)

(E) Aggregate (Fed)

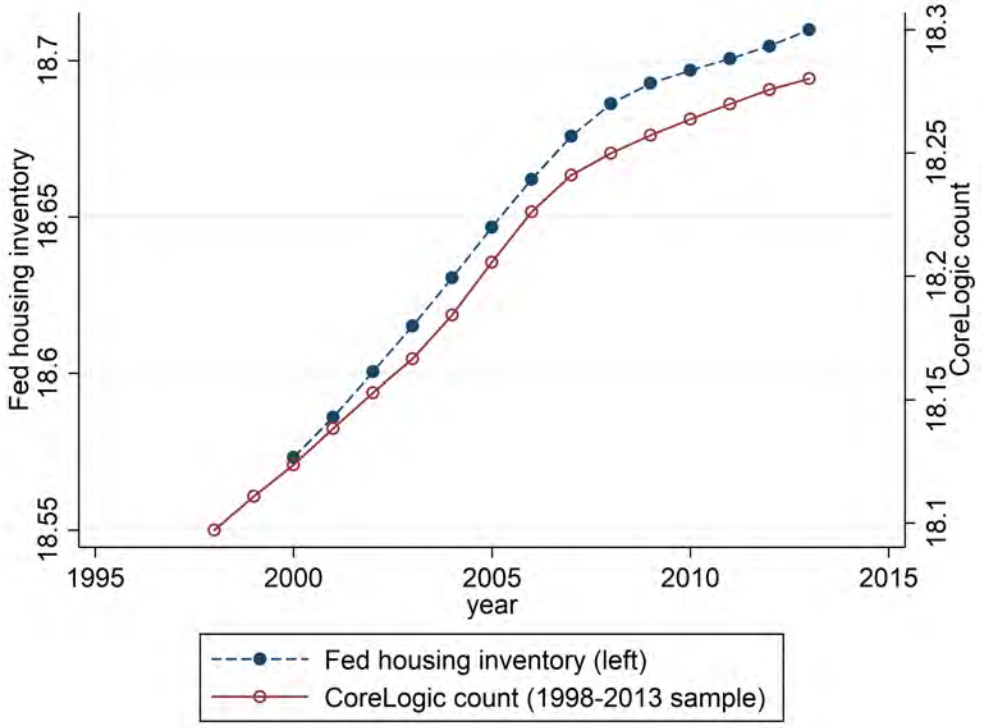

(F) Aggregate (Census)

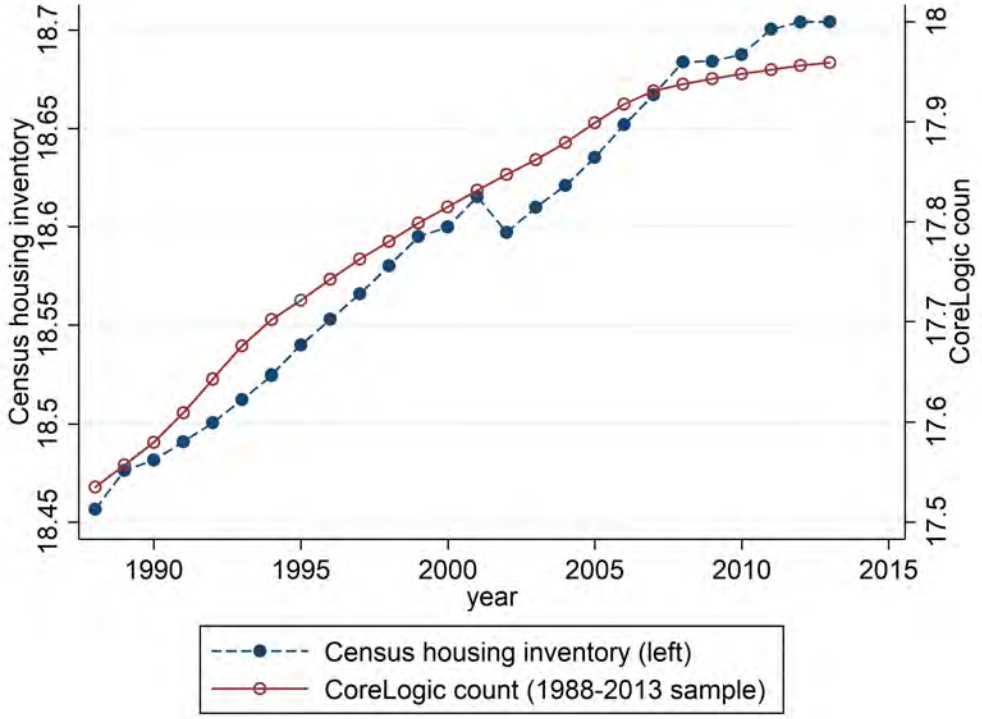

FIGURE A.2: Raw-data pattern, sorting by residence Census income

These figures show raw-data patterns, sorting owners by surnames (and associated 1940 income) on the left column, and sorting owners by the mailing addresses' Census block group (and associated 200 income, within county) on the right column.
Panels (a) and (b) plot the average value of primary residence conditional on owning in 2012-2013, for CoreLogic's assessor record, which covers almost the entire US population in a single year. Panels (c) and (d) plot the per capita holdings of any real estate asset (i.e., count), for selected decile groups, relative to the 1998 levels. Panels (e) and (f) plot the residuals ε_{it} for the same set of selected deciles from the regression:

$$\log(q_{it}) = \alpha_i + \alpha_t + \gamma_i t + \varepsilon_{it}$$

where the regression is weighted by the number of individuals in each decile group, and q_{it} is the holdings of all real estate by number of property by members of the decile group in a given year. For comparison, panel (g) plots the same residuals for CoreLogic national house-price index, i.e., ε_t from

$$\log(P_t) = \gamma_0 t + \varepsilon_t$$

where P_t is the house-price index. The vertical red line indicates 2007.

(A) By block group: primary-residence value

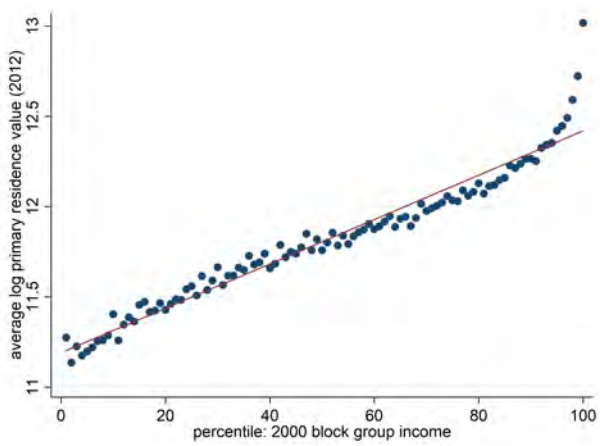

(B) By block group: Per capita holding vs. 1998

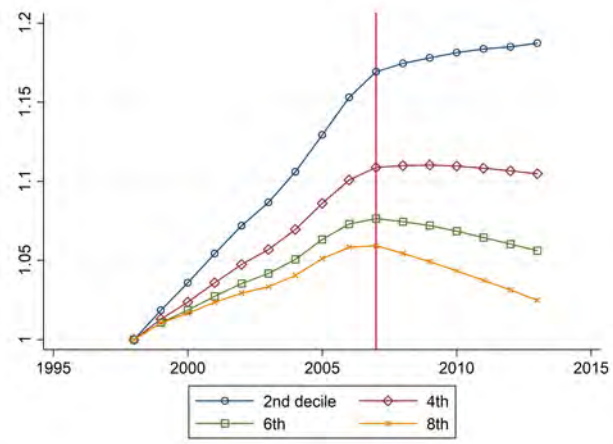

(C) By block group: Detrended log residual

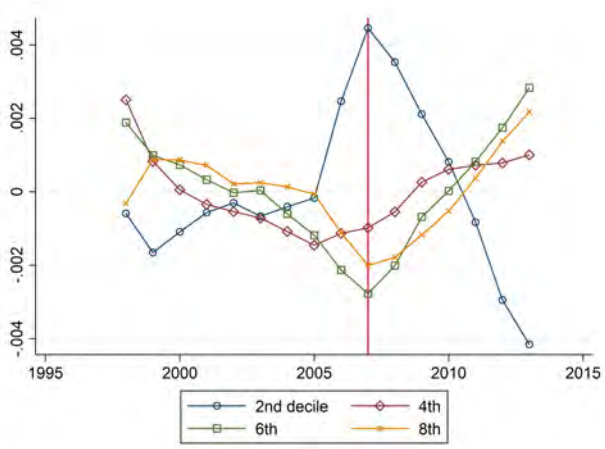

FIGURE A.3: Wealth inequality vs. income inequality

(A) White / Black

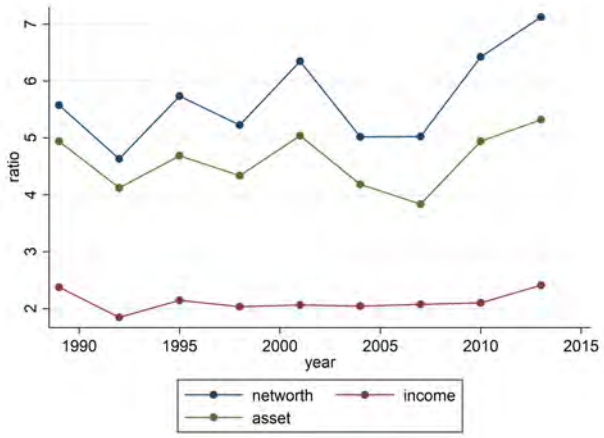

(B) White / Hispanic

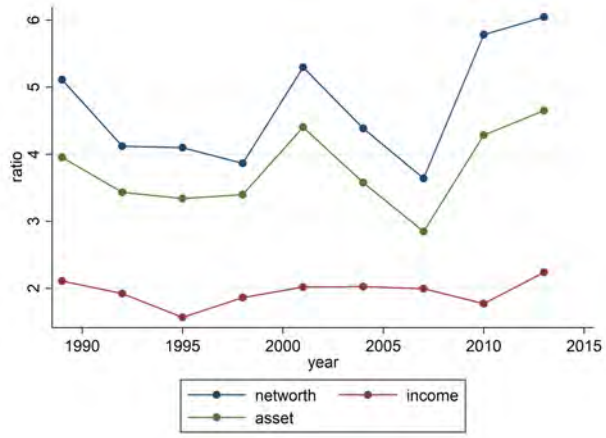

(C) Above median income / below

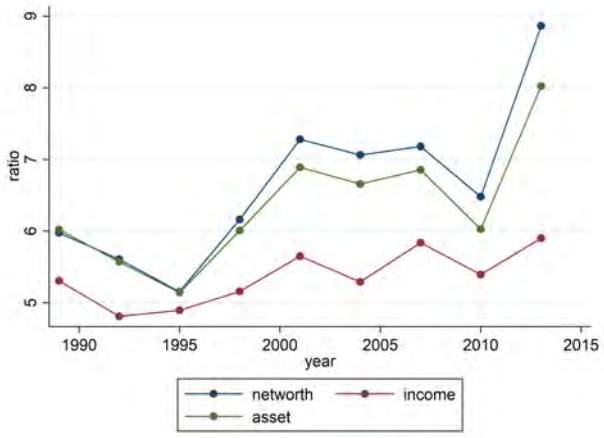

(D) College / no college

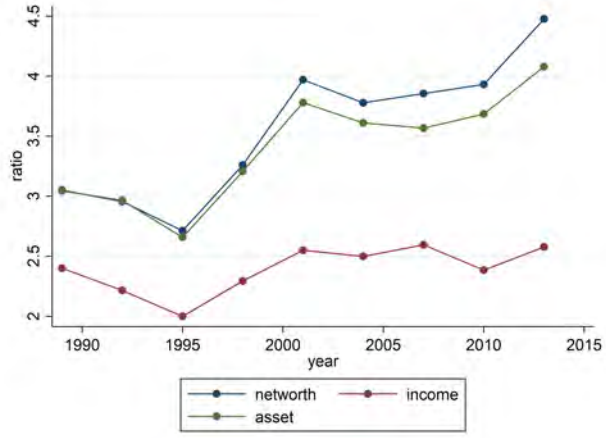

FIGURE A.4: Wealth inequality level: Geographical variation

Panel (a) plots Core-based Statistical Area (CBSA)-level coefficient of variation of net worth in 2012. Net worth coefficient of variation has been calculated for CBSAs using zip code-level variation. Zip code-level net worth has been imputed using capital income from the IRS Statistics of Income, capital income capitalization factors from Saez and Zucman (2016), housing ownership from CoreLogic assessor records, and zip code-level debt stocks from Equifax. Panel (b) plots CBSA-level coefficient of variation of wage income in 2012, calculated using zip code-level wage information. Wage comes directly from "Salaries and wages" in the IRS Statistics of Income. Panel (c) plots business cycle income loadings, calculated at the county-level by regressing changes in county-level log per-capita income on changes in aggregate log per-capita income, using data from the Bureau of Economic Analysis 1969-2015.

(A) Net worth coefficient of variation by CBSA (2012)

FIGURE A.4: Wealth inequality level: Geographical variation (continued)

(B) Wage coefficient of variation by CBSA (2012)

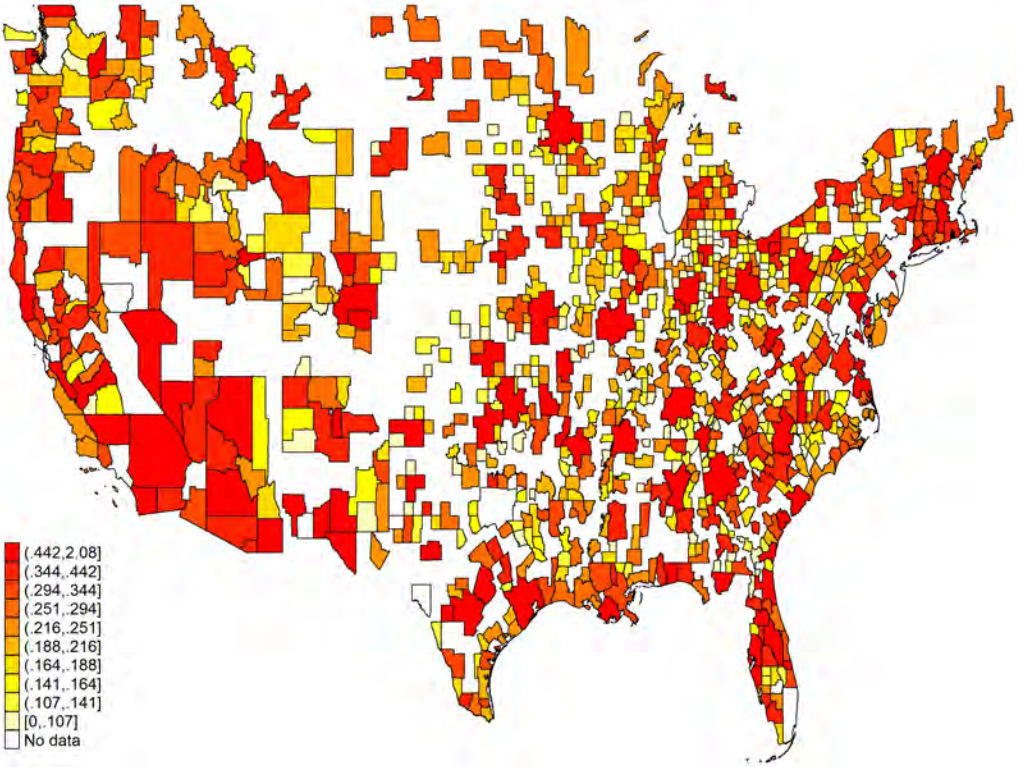

(C) Business cycle loading by county (1969-2015)

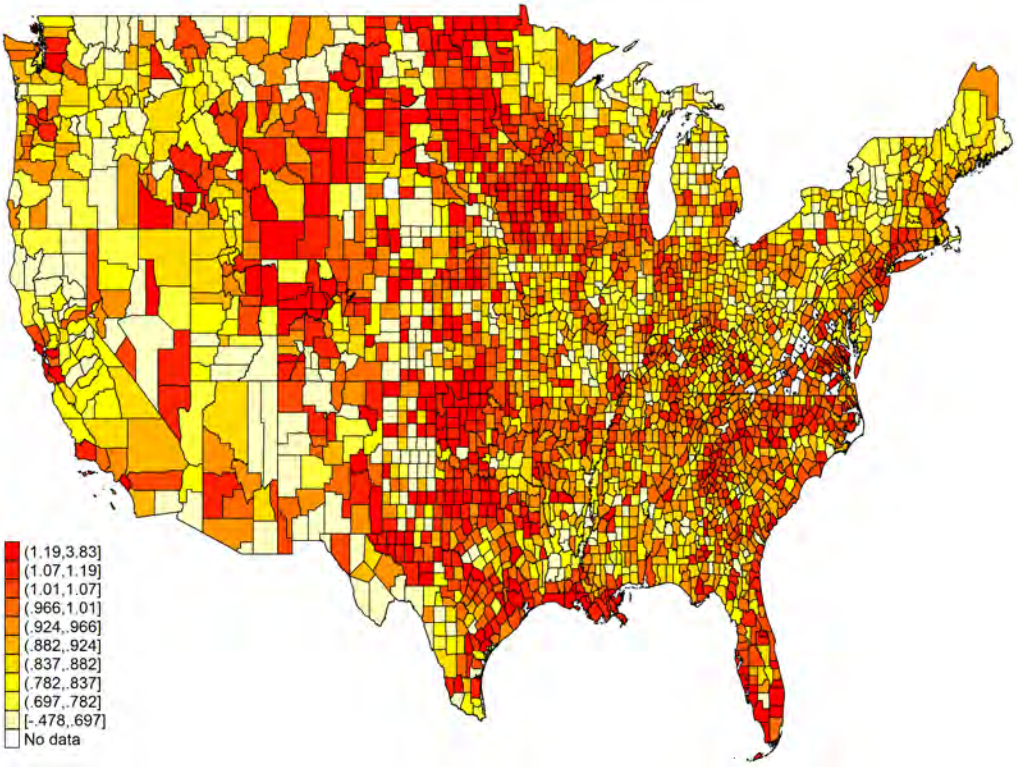

FIGURE A.5: Wealth vs. income (SCF)

From the Survey of Consumer Finances for 2013, panel (a) plots log net worth against log household income, and panel (b) plots log asset against log income. Panel (a) shows a slope of 1.397 of log net worth against log income. Panel (b) shows a slope of 1.530 of log asset against log income.

(A) Log net worth vs. log income (2013)

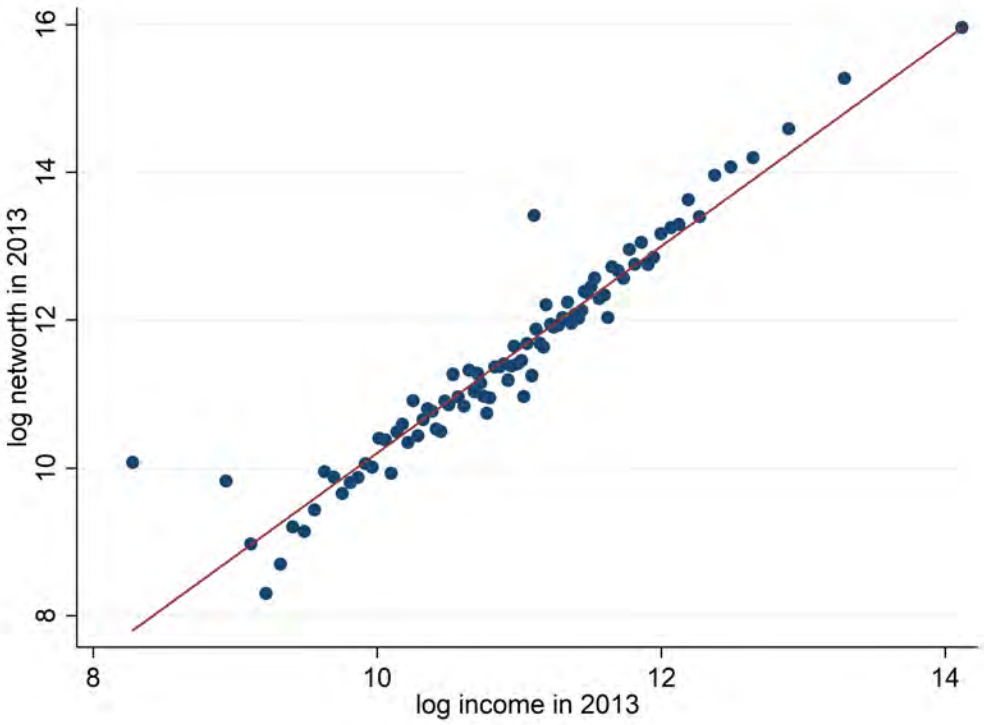

(B) Log asset vs. log income (2013)

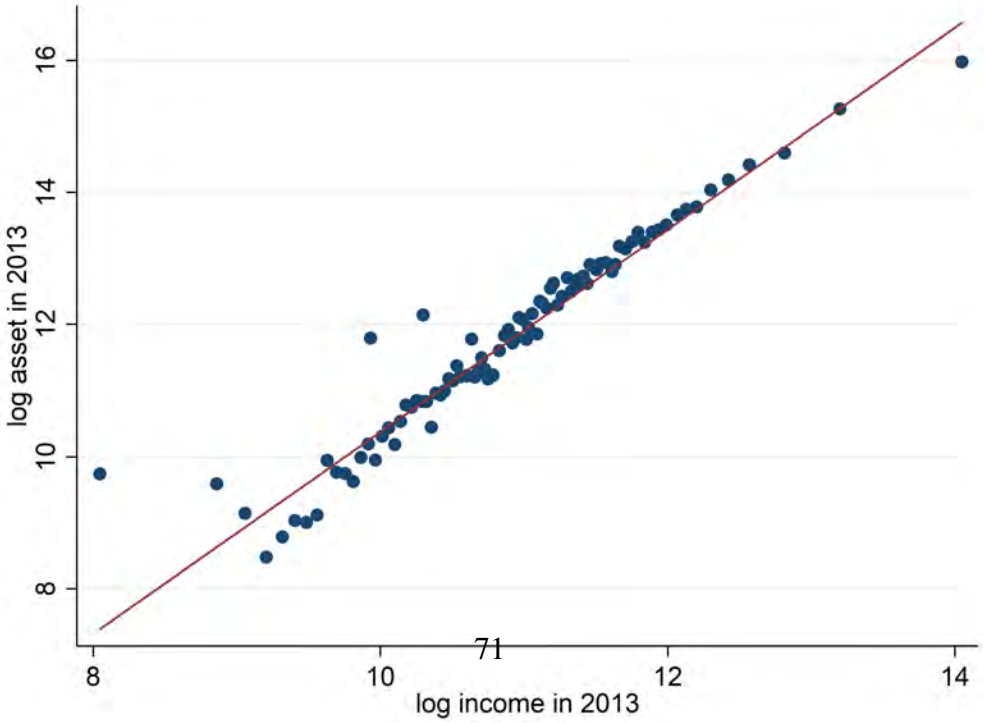

TABLE A.1: Level of wealth inequality: CBSA coefficient of variation (continued)

(A) Asset

	asset c.v. (2012)			with avg cap income 2003-2012	
	(1)	(2)	(3)	(4)	(5)
beta: income per cap	0.557***	0.377***	0.603***	0.347*	0.482**
	(0.108)	(0.077)	(0.122)	(0.165)	(0.166)
wage c.v.		1.202***	1.105***	1.104**	1.106**
		(0.284)	(0.302)	(0.410)	(0.350)
log population			0.043	0.246	0.217*
			(0.030)	(0.129)	(0.095)
log house price			-0.120	0.675	0.651
			(0.103)	(0.524)	(0.515)
Constant	0.256**	0.009	-0.038	-5.624	
	(0.081)	(0.095)	(0.494)	(3.756)	
Adjusted R^2	0.034	0.375	0.387	0.232	0.333
state FE					O
# of CBSA	926	926	670	670	670
Observations	1759	1759	1114	1114	1113

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

TABLE A.2: Level of wealth inequality: Zip code wealth-wage elasticity (continued)

(A) Asset

	log asset per capita ('12)			with avg cap income '03-'12		
	(1)	(2)	(3)	(4)	(5)	(6)
income beta × wage		0.555***	0.609***	0.557***	0.456**	0.600***
		(0.142)	(0.166)	(0.161)	(0.141)	(0.128)
population × wage			0.041**	0.046***	0.045***	0.057***
			(0.013)	(0.013)	(0.013)	(0.013)
house price × wage			-0.083	-0.036	-0.215	0.071
			(0.122)	(0.119)	(0.125)	(0.173)
income beta		-5.688***	-6.368***	-5.816***	-4.569**	
		(1.497)	(1.782)	(1.727)	(1.515)	
log population			-0.372*	-0.458**	-0.468**	-0.581***
			(0.145)	(0.142)	(0.143)	(0.135)
log house price			1.453	0.916	2.533	
			(1.332)	(1.302)	(1.366)	
log wage	1.487***	0.918***	0.720	0.429	1.407*	-0.122
	(0.012)	(0.148)	(0.609)	(0.596)	(0.631)	(0.877)
Constant	-3.980***	1.857	0.936	4.513		
	(0.124)	(1.558)	(6.670)	(6.499)		
Adjusted R^2	0.473	0.474	0.542	0.527	0.614	0.741
state FE					O	
county FE						O
# of counties		3059	1231	1231	1231	1220
Observations	27643	27383	17820	17820	17820	17809

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$